POLITECNICO DI MILANO

Department of Electronics, Information, and Bioengineering

DOCTORAL PROGRAM IN INFORMATION TECHNOLOGY

Agent–based water resources management in complex decision–making contexts

Doctoral Dissertation of:
Matteo Giuliani

Advisor:
Prof. Rodolfo Soncini-Sessa

Co-Advisor:
Prof. Andrea Castelletti

Tutor:
Prof. Carlo Piccardi

Supervisor of the Doctoral Program:
Prof. Carlo Fiorini

2013 - XXVI

Abstract

The presence of multiple, institutionally independent but physically interconnected actors is a distinctive feature of the complexity characterizing most of the decision-making problems in environmental contexts. As dealing with many conflicting stakeholders requires to replace the concept of optimality by that of Pareto efficiency, the presence of many decision makers requires some kind of coordination and a cooperative attitude of the involved parties, who agree on adopting a fully coordinated strategy to maximize the system-level performance. These assumptions are often unpracticable in real world contexts, where the decision makers generally belong to different institutions or countries. In these situations, they independently pursue local interests and produce negative externalities leading to a low system-wide efficiency. Game theory and simulation-based approaches are generally used to analyze these issues from a descriptive standpoint, while their prescriptive use in decision support systems to design coordination mechanisms between the originally self-interested decision-making actors remains a challenge.
This thesis contributes a novel decision analytic framework based on multi-agent system (MAS) to study water resources planning and management problems in complex decision-making contexts. The aim of the proposed decision analytic framework is to combine descriptive and prescriptive methods, which provide informative tools to represent the actual decision-making context as well as decision support procedures to recommend proper coordination mechanisms. The adoption of an agent-based framework naturally allows the representation of a set of decision makers or stakeholders (agents), which act in the same environment and need to coordinate to maximize the system-wide efficiency in the use of the available water. This agent-based perspective aims to move beyond the traditional centralized approach to water resources management and to explore different levels of cooperation, from fully coordinated strategies and full information exchange to completely uncoordinated practices. Moreover, coupling the agent-based modeling scheme with state-of-the-art Control Theory techniques allows a better understanding of the feedbacks between agents objectives, agents decisions, and the environment they share, as well as the description of their co-evolution and co-adaptation under change.

The proposed framework is demonstrated in different problems characterized by distinctive decision-making contexts, which require to adopt different tools and methodologies. The framework is first applied in a hypothetical water allocation planning problem to discuss the issue of balancing system-wide efficiency and solutions practicability. MAS methods based on distributed constraint optimization problems are used to support a watershed authority in evaluating different levels of coordination. Then, the advantages of coordination mechanisms based on the exchange of information are estimated in two real world case studies. The first one assesses the value of cooperation and information exchange in transboundary river basins. The second one shows the potential of co-adapting water demand and supply in agricultural water systems, under current and projected hydroclimatic conditions. Finally, in the last application, the framework combines tools to identify and refine the current operation of the Conowingo reservoir in the Lower Susquehanna system, for balancing evolving demands and system uncertainties. It also introduces a novel method based on input variable selection techniques to support the identification of effective policy mechanisms for environmental protection.

Part of this thesis' contributions has appeared (or has to appear) in the following main publications:

- M. Giuliani and A. Castelletti (2013), Assessing the value of cooperation and information exchange in large water resources systems by agent-based optimization, *Water Resources Research*, 49, 3912-3916.

- M. Giuliani, A. Castelletti, F. Amigoni, and X. Cai (2013), Agent-based distributed optimization as a tool to balance efficiency and practicability in watershed management, *Journal of Water Resources Planning and Management*, (under review).

- M. Giuliani, J.D. Herman, A. Castelletti, and P.M. Reed (2013), Many-Objective Reservoir Policy Identification and Refinement to Reduce Institutional Myopia in Water Management, *Water Resources Research*, (under review).

- M. Giuliani, Y. Li, A. Castelletti, and C. Gandolfi (2013), Co-adapting water demand and supply to changing climate in agricultural water systems, *Global Environmental Change*, (in preparation).

Riassunto

La presenza di diversi attori decisionali, istituzionalmente indipendenti ma fisicamente interconnessi, è uno degli aspetti che rendono i problemi decisionali in campo ambientale particolarmente complessi. La presenza di molti interessi conflittuali richiede infatti di sostituire al concetto di ottimalità quello di efficienza Paretiana. Inoltre, la presenza di molti decisori richiede di assumere un'attitudine cooperativa delle parti coinvolte al fine di massimizzare le performance a livello di sistema attraverso l'adozione di soluzioni completamente coordinate. Tuttavia, queste ipotesi risultano essere inapplicabili in molti contesti reali, dove i decisori spesso appartengono a istituzioni o Paesi diversi. In questi casi, le loro strategie spesso considerano solamente interessi locali, producendo esternalità negative che riducono i benefici a livello di sistema. Questa problematica viene solitamente studiata nel campo della teoria dei giochi con approcci di tipo descrittivo. Al contrario risulta generalmente più difficile adottare approcci prescrittivi di supporto alle decisioni che suggeriscano meccanismi di coordinamento tra gli attori decisionali originalmente indipendenti.

La ricerca presentata in questa tesi introduce un nuovo framework analitico e decisionale basato sulla teoria dei sistemi multi-agente per studiare la pianificazione e gestione delle risorse idriche in contesti decisionali complessi. Tale framework si propone di combinare metodi descrittivi e prescrittivi con lo scopo di fornire strumenti informativi che rappresentino il reale contesto decisionale e, allo stesso tempo, procedure di supporto alle decisioni che suggeriscano efficaci meccanismi di coordinamento. L'utilizzo di un framework di tipo multi-agente permette infatti la rappresentazione di un insieme di agenti (decisori o portatori di interesse) che agiscono all'interno dello stesso ambiente e devono coordinarsi per massimizzare l'efficienza a livello di sistema nell'utilizzo delle risorse idriche disponibili. Questa prospettiva a livello di agente introduce inoltre la possibilità di esplorare diversi livelli di cooperazione tra gli agenti, superando le limitazioni dell'approccio tradizionale di tipo centralizzato. La combinazione di modelli multi-agente con tecniche di ottimizzazione sviluppate all'interno della teoria del controllo permette infine sia di analizzare i feedback esistenti tra obiettivi e decisioni degli agenti e il sistema nel quale agiscono, sia di simularne il co-adattamento e la co-evoluzione

in condizioni socio-economiche e climatiche in rapida evoluzione.
Le potenzialità dell'approccio proposto sono dimostrate applicando il framework a diversi problemi caratterizzati da contesti decisionali diversi, che quindi richiedono strumenti e metodologie differenti. La prima applicazione riguarda un ipotetico problema di allocazione delle risorse idriche in cui è necessario bilanciare l'efficienza delle soluzioni a livello di sistema e la loro applicabilità a livello di agente. L'utilizzo di metodi sviluppati all'interno della teoria dei sistemi multi-agente e basati sulla formalizzazione di problemi di ottimizzazione distribuita permette di supportare l'autorità di bacino nella valutazione di diversi livello di coordinamento. Il secondo e terzo caso di studio si concentrano sull'analisi dei benefici ottenibili attraverso metodi di coordinamento basati sullo scambio di informazione. Nel primo caso viene fornita una stima del valore economico della cooperazione e dello scambio di informazione in bacini internazionali. Nel secondo caso viene dimostrato il potenziale di strategie basate sul co-adattamento dei sistemi di water supply e di water demand nella gestione delle risorse idriche in campo agricolo, sia in condizioni climatiche attuali che future. Nell'ultima applicazione del framework, l'attuale gestione del serbatoio idrico di Conowingo, situato lungo il Susquehanna River, viene prima identificata e successivamente modificata attraverso la proposta di soluzioni in grado di bilanciare tutti gli interessi conflittuali coinvolti e che risultino robuste rispetto alle incertezze legate alle condizioni idro-climatiche. In questo contesto viene anche proposto un nuovo metodo basato su tecniche di input variable selection per l'identificazione di meccanismi di protezione ambientale.

Parte dei contributi di questa ricerca è stata presentata (o è in fase di preparazione) nelle seguenti pubblicazioni:

- M. Giuliani and A. Castelletti (2013), Assessing the value of cooperation and information exchange in large water resources systems by agent-based optimization, *Water Resources Research*, 49, 3912-3916.

- M. Giuliani, A. Castelletti, F. Amigoni, and X. Cai (2013), Agent-based distributed optimization as a tool to balance efficiency and practicability in watershed management, *Journal of Water Resources Planning and Management*, (under review).

- M. Giuliani, J.D. Herman, A. Castelletti, and P.M. Reed (2013), Many-Objective Reservoir Policy Identification and Refinement to Reduce Institutional Myopia in Water Management, *Water Resources Research*, (under review).

- M. Giuliani, Y. Li, A. Castelletti, and C. Gandolfi (2013), Co-adapting water demand and supply to changing climate in agricultural water systems, *Global Environmental Change*, (in preparation).

Contents

List of Figures

List of Tables

1 Introduction

1.1 Complexity in water resources systems

The use of mathematical models to support decision-making in water resources systems is rapidly expanding in the last years due to the recent advances in terms of scientific knowledge of the natural processes, efficiency of the optimization techniques, and availability of computational resources (*Washington et al.*, 2009). However, the problem remains challenging because of three main sources of complexity: *i*) the complexity of most of the natural processes, further complicated by the incoming climate change impacts, which requires using more and more realistic process-based models that are unsuitable to support decision-making processes, as a large numbers of model evaluations is required to run sensitivity analysis, scenario analysis, and optimization algorithms (e.g., *Washington and Parkinson*, 2005; *Carnevale et al.*, 2009); *ii*) the complexity of the socio-economic framework, which generally involves a myriad of conflicting and non-commensurable objectives, representing the interests of multiple stakeholders such as domestic and irrigation supply, flood protection, hydropower production, with additional challenges due to environmental protection, water quality targets, recreational interests, and energy markets (e.g., *Chaves and Kojiri*, 2007; *Soncini-Sessa et al.*, 2007b; *Fernandez et al.*, 2012); *iii*) the complexity of the decision-making institutional context due to the presence of several institutionally independent, but physically interconnected, decision makers both within national jurisdictions and in transboundary contexts (e.g., *Yoffe et al.*, 2003; *Wolf et al.*, 2006; *Zeitoun and Mirumachi*, 2008).
This thesis contributes a novel agent-based decision analytic framework (see Section 1.3.2) focused on the challenges related to the complexity of the decision-making institutional context. The proposed framework aims to support sustainable water resources planning and management in complex environmental systems, characterized by the presence of multiple decision makers and many conflicting stakeholders. Depending on the specific application, the complexity of the natural processes and of the socio-economic context are also addressed. In such multi-objective contexts, the traditional concept of optimality is replaced by that of Pareto efficiency (*Pareto*, 1964): a solution is Pareto optimal (or non-

dominated) if there not exists any other solution which gives a better value for one objective without worsening the performance in at least one other objective. According to classic Operational Research and Control Theory principles, the traditional approach to water resources management adopts a centralized perspective and aims at the maximization of the system-wide efficiency, possibly defined by multiple criteria. To attain such an efficient use of the available resources, the presence of many decision makers requires some kind of coordination and a cooperative attitude of the involved parties, who agree on adopting a fully coordinated strategy on water allocation and distribution in time and space, as well as full knowledge of the current system conditions. These assumptions are often unpracticable in real world contexts. The centralized management, though interesting from a conceptual point of view as it allows to quantify the best achievable performance and to obtain insights on strategies to foster cooperation (e.g., *Anghileri et al.*, 2012), turns out to be of limited practical meaning given the real political and institutional setting (e.g., *Waterbury*, 1987; *Whittington et al.*, 2005; *Wu and Whittington*, 2006). Totally uncoordinated practices, where all the decision makers independently pursue their local objectives, represent a more realistic picture of most real institutional contexts, where these individualistic behaviors generally induce negative externalities leading to a low system-wide efficiency (e.g., *Hardin*, 1968; *Madani*, 2010).
This situation has a Prisoner's Dilemma structure. Table 1.1 shows the Prisoner's Dilemma in a matrix form, with the payoffs representing the utility for each player (i.e., player-1 on the rows and player-2 on the columns) in case of cooperation (C) or non-cooperation (NC). Although the players can gain more from cooperation (i.e., (C, C) represents the best solution at the system-level), non-cooperation is the strictly dominant strategy, meaning that no matter if the other player selects to cooperate or not, it is always better to non-cooperate (i.e., 4>3 and 2>1). The outcome (NC, NC) is a Nash Equilibrium (*Nash*, 1951): given the options of other players, no player can improve his payoff by changing his strategy alone. Nash stability differs from Pareto-optimality. The former is about what is good for an individual without considering what is good for the whole system and the latter is about what is good for the system without considering the interests of the single individuals. The problem structure and outcome might change when other players are involved, such as river basin authorities promoting negotiation processes to incorporate, at different levels, cooperative agreements or coordination mechanisms (e.g., *Madani and Lund*, 2012). These challenges are widely recognized in many research fields, such as economics and game theory, which, however, mainly provide descriptive tools based on what-if or

Table 1.1: Prisoner's Dilemma.

	cooperation (C)	non-cooperation (NC)
cooperation (C)	3,3	1,4
non-cooperation (NC)	4,1	2,2

scenario analyses. An extended analysis of the impacts of self-interested decision makers in the design of decision support systems is still missing. The proposed agent-based decision analytic framework represents a novel approach which combines descriptive and prescriptive methods in order to provide informative tools representing the actual decision-making context (e.g., removing the simplifying assumption of fully cooperative decision makers), as well as decision support procedures, which recommend proper coordination mechanisms between the originally self-interested decision-making actors.

1.2 Planning and management problems

Decision-making processes in water resources systems comprise two different types of problems, namely planning and management, which can be solved using mathematical models and Systems Analysis methods (e.g., *Loucks et al.*, 1980; *Soncini-Sessa et al.*, 2007a, and references therein). In planning problems the decision variables $\mathbf{u}^p$ represents "una tantum" actions, meaning decisions that are taken once, without considering how they might influence analogous decisions in the future. A typical example is the construction of a canal which requires to design its maximum capacity. Planning does not necessarily mean infrastructural interventions. Other common planning problems concern the definition of minimum environmental flow requirements, the allocation of water rights, or the design of reservoir rule-curves. Note that the absence of dynamics in this type of decision does not imply, however, that also the system has to be non-dynamic. In management problems, instead, it is required to take sequential decisions $\mathbf{u}_t$ on the basis of the current system conditions as described by the state vector $\mathbf{x}_t$. The optimal solution of a management problem is a feedback operating policy p defined as a sequence of operating rules providing, step-by-step, the vector of management decisions $\mathbf{u}_t = m_t(\mathbf{x}_t)$. Typical examples of management problems are the design of optimal operating policies for water reservoirs or groundwater pumping systems, where the optimal policy provides at

each time step the optimal release (or pumping) decision as a function of the reservoir (groundwater) level.
The mathematical formulation of a generic planning problem is given in eq. (1.1), while eq. (1.2) defines a generic management problem. Both the problems are formulated with respect to a vector $\mathbf{J}(\cdot)$ of q operating objectives $\left[J^1, \ldots, J^q\right]$ and are subject to appropriate constraints $c_1(\cdot), \ldots, c_r(\cdot) \leq 0$. The problems are also dynamically constrained by the evolution of the system, which can be modeled as $\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}^p, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1})$, where $\boldsymbol{\varepsilon}_{t+1}$ is a vector of stochastic disturbances. Given the multi-objective nature of the problems, their resolution does not yield a single optimal solution, but a set of Pareto-optimal solutions. The image in the objective space of the Pareto-optimal solutions is the Pareto front.

$$\mathbf{u}^{p*} = \arg\min_{\mathbf{u}^p} \mathbf{J}(\cdot) \tag{1.1}$$

$$p^* = \arg\min_{p} \mathbf{J}(\cdot) \tag{1.2}$$

The k-th objective expresses the cost paid by k-th stakeholder over the time horizon $[0, H]$ (see, for more details, *Castelletti et al.* (2008a)):

$$J^k = \lim_{H\to\infty} \Psi_{\boldsymbol{\varepsilon}_1,\ldots,\boldsymbol{\varepsilon}_H} \left[\Phi_{t=0}^{H-1} \left(g_{t+1}^k(\mathbf{x}_t, \mathbf{u}^p, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1})\right)\right] \tag{1.3}$$

where $g_{t+1}^k(\cdot)$ is the k-th immediate cost function associated to each system transition, Φ is an operator for aggregation over time, and Ψ is a filtering criterion to deal with the uncertainties generated by the disturbances $\boldsymbol{\varepsilon}_{t+1}$. Assuming that Φ computes the average value over time and Ψ the expected value over the realizations of the stochastic disturbances, eq. (1.3) represents a measure of vulnerability. Conversely, if Ψ filters the effects of the disturbances by considering the worst-case realization, eq. (1.3) represents a measure of robustness. These formulations of planning and management problems are general and can represent almost all the operational water science problems.

1.3 Agent-based water resources management

1.3.1 Literature review

Distributed Artificial Intelligence

The theory of multi-agent system (MAS) has emerged from a sub-field of researchers in artificial intelligence (AI), called distributed artificial intel-

ligence (DAI). DAI community started forming in the early 1980s when the complexity of AI problems exceeds the capabilities (i.e., knowledge, computing resources, perspective) of a single entity (agent). Particularly complex, large, or unpredictable problems require indeed to develop multiple, specialized, and modular components (agents), which are also able to interoperate via techniques based on negotiation or cooperation (*O'Hare and Jennings*, 1996; *Jennings et al.*, 1998).
The DAI concept of MAS is relatively simple and general: an agent can be defined as a computer system, capable of independent actions to meet its design objectives, and a multi-agent system consists of a number of agents which interact in the same shared environment (*Wooldridge*, 2009). Figure 1.1 provides an abstract view of a MAS, where each agent can be associated to a portion of the system that the agent observes and controls. Given its observation along with the information provided through the interactions with the other agents, each agent executes one of the actions available in its repertoire and modifies the environment. Depending on the characteristics of the environment, the specific structure of the agents as well as the methods defining their behaviors can be different. *Russell and Norvig* (1995) suggest four major distinguishing properties for characterizing MAS environments:

- *deterministic or stochastic*: an environment is deterministic if any action has a guaranteed outcome, and stochastic if there is uncertainty on the effects of the action;

- *static or dynamic*: a static environment can be assumed to remain unchanged (except by the effects of agents actions), while a dynamic environment is characterized by multiple processes evolving in time and beyond agents control;

- *discrete or continuous*: an environment is discrete if it has a finite number of possible states, and continuous if there are infinite states;

- *accessible or inaccessible*: an environment is accessible if each agent can obtain complete, accurate, and up-to-date information about the environment state, and inaccessible otherwise.

The most complex class of environments, also referred to as "open environments" (*Hewitt*, 1986), includes environments that are stochastic, dynamic, continuous, and inaccessible. Interestingly, these properties can be easily mapped into typical features of water systems, which are *i*) *non-deterministic*, due to the presence of stochastic external drivers (e.g., rainfall, temperature, etc.); *ii*) *dynamic*, as they evolve in time; *iii*) *continuous* and spatially distributed; *iv*) *accessible* or *inaccessible*

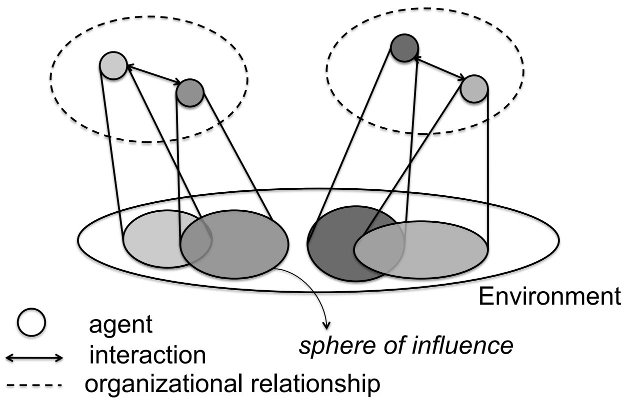

Figure 1.1: Multi-agent system representation.

depending on the degree of information sharing among the agents.
These characteristics then influence the complexity of the agent design process, which tries to find an effective balance between the following features (*Wooldridge and Jennings*, 1995):

- *reactivity*: agents are able to perceive the environment and timely respond to its changes;

- *proactiveness*: agents are able to exhibit goal-directed behaviors;

- *social ability*: agents are capable of interacting with other agents (e.g., from data and information sharing to social activities such as coordination or negotiation).

Purely reactive agents act using a stimulus-response type of behavior and respond to the present state of the environment (*Sycara*, 1998). They do not look at history or plan their strategy over the future. This characteristic allows the design of simple "if-then" agents behaviors. Yet, they make decisions based on local information only and do not predict the effect of their decisions on global behavior. This myopic behaviors can lead to unpredictable and unstable system situations (*Thomas and Sycara*, 1998).
Economics-based mechanisms have been instead utilized to model proactive agents as utility maximizers (*Shoham and Leyton-Brown*, 2009). These approaches, such as market mechanisms, are becoming increasingly attractive as they allow a flexible MAS design according to ready available underlying economics models, with well established roots in game theory and artificial intelligence (*Vidal*, 2009). According to this

approach, it is assumed that the agent's preferences are captured by a utility function, which defines a map from the states of the environment to a real number. Cooperative agents select their actions in order to maximize the total utility at the system-level. A self-interested agent instead chooses a course of action that maximizes its own utility. In a society of self-interested agents, it is desired that if each agent maximizes its local utility, the whole society exhibits good behavior (i.e., good local behavior implies good global behavior). However, in many situations the self-interested agents generally overuse and congest the shared resources leading to the "The Tragedy of the Commons" described by *Hardin* (1968). In these contexts, *Sycara* (1998) defines the goal of MAS research as the design of mechanisms for self-interested agents such that the overall system behavior will be acceptable, which is called mechanism design (*Maskin*, 2008).

Environmental Applications

From the first works developed in DAI, the applications of MAS have rapidly covered a variety of domains, ranging from manufacturing to process control, air-traffic and robot control, or information management. MAS approaches have become a widely used tool in several environmental modeling contexts (e.g., *Athanasiadis*, 2005). The primary goal of most of these studies, also referred to as multi-agent simulations (for a review, see *Bousquet and Le Page*, 2004, and references therein), is to simulate complex systems in order to evaluate macro-level properties emerging from lower-level interactions among the agents.

Agent-based modeling offers several advantages with respect to other approaches (*Bonabeau*, 2002; *Bousquet and Le Page*, 2004): *i*) it provides a more natural description of a system; *ii*) it relaxes the hypothesis of homogeneity in a population of actually heterogeneous individuals; *iii*) it allows an explicit representation of spatial variability; *iv*) it captures emergent global behaviors resulting from local interactions. Purely reactive agents (see the previous section) are largely adopted to define behavioral rules which react to environmental changes (*Le et al.*, 2012). However, the prescriptive use of MAS models in decision support systems remains a challenge due to their mathematical complexity, which requires to shift toward a descriptive standpoint (*Galán et al.*, 2009), developing what-if analyses with respect to a limited number of management alternatives and modeling simple decision mechanisms based on linear programming (*Schreinemachers and Berger*, 2011).

Ecology was the first environmental discipline that adopted MAS approaches due to their similarity with individual-based models (e.g., *Hus-*

ton et al., 1988). The aim of these works is to study the behavior of a population, where the processes involved are modeled at the agent level to take into account the heterogeneity of the individuals (*Bousquet et al.*, 1999). MAS tools are also used to study land use change and agricultural systems (*Berger*, 2001; *Berger and Ringler*, 2002; *Naivinit et al.*, 2010; *Schreinemachers and Berger*, 2011; *Le et al.*, 2012; *Ralha et al.*, 2013) according to the "representative independent farm approach" (*Hanf*, 1989), which consists of a number of independent farm models added up to compose a sector result. Agent-based models are used to represent human-landscape systems where a set of human agents interact with each other and with the environment. The impacts at the system level of changes in agricultural technologies, marked dynamics, and policy intervention are estimated via scenario-based analysis (*Schreinemachers and Berger*, 2011).

In the water resources literature, MAS approaches have been used to model watersheds as spatially distributed, coupled natural-human systems, where the water is shared between the multiple water users (agents). *Berger et al.* (2007) illustrate the potentiality of MAS as a planning tool to evaluate ex-ante different policy scenarios and estimate their effects for different groups of water users. Similarly, scenario-based experiments are used to study domestic water management problems (e.g., *Galán et al.*, 2009) and water distribution for irrigation purposes (e.g., *van Oel et al.*, 2010). MAS tools have been also used to support negotiation (e.g., *Becu et al.*, 2003; *Feuillette et al.*, 2003) in order to facilitate agreements between conflicting stakeholders, focusing more on the representation of the social aspects of stakeholders interactions than on the policy mechanisms. The combination of MAS and role-playing games (e.g., *Barreteau et al.*, 2001; *Guyot and Honiden*, 2006) was demonstrated to be particularly effective to directly involve the stakeholders in participated decision-making processes. In a role-playing game, each player is asked to choose step-by-step the action that the agent has to perform and the results of the model simulation then facilitates discussion and negotiation.

The descriptive standpoint generally adopted in multi-agent simulations is similar to the one used in game theory (e.g., *Madani*, 2010, and references therein), and is more focused on the properties of the outcome situations instead of the mechanisms leading to better solutions, as it would be typical in a prescriptive, optimization-based approach based on proactive MAS. In the water resources field, the first contribution adopting proactive MAS was presented in *Yang et al.* (2009), where the multiple distributed water users are modeled as self-interested agents acting in a distributed decision process to solve a water allocation problem by means of penalty-based decentralized optimization method. A similar approach was then adopted by *Ng et al.* (2011) to optimize pre-season farmers decisions about crop types and best management practices under different scenarios of markets for conventional crops, carbon allowance, and second generation biofuel crops, according to predictions of future prices, costs, yields, and weather. *Huskova and Harou* (2012) propose an agent-based model to estimate the effectiveness of an optimization-driven water market in mitigating increasing water scarcity in England. Agent-based modeling is also adopted by *Nguyen et al.* (2013) to assess the outcomes of different market structures on emission trading. Agents are heterogeneous in their costs and relative size, have imperfect information about their own costs and the costs of others, which can be diminished but not eliminated by investments to improve their information, and are boundedly rational. By solving stochastic optimization problems, the agents decide whether to participate in trading and, if they do, their bidding strategies.

1.3.2 A novel agent-based decision analytic framework

This thesis proposes a novel decision analytic framework based on multi-agent systems to overcome the difficulties related to the complexity of the institutional decision-making contexts in water resources systems. A precise definition of MAS can be difficult due to the many competing, mutually inconsistent answers offered in different disciplines (*Shoham and Leyton-Brown*, 2009). Consequently, in this thesis MAS is considered more as a mindset than a technology (*Bonabeau*, 2002), which relies on the general definitions given in *Wooldridge* (2009) and *Shoham and Leyton-Brown* (2009): *an agent is defined as a computer system situated in some environment and capable of autonomous actions to meet its design objectives; multi-agent systems are those systems that include multiple autonomous entities with either diverging information or diverging interests, or both.*

The selection of a framework based on MAS naturally allows the representation of multiple decision makers or stakeholders (agents), which act in the same environment, thus influencing each other, and need to coordinate to maximize the system-wide efficiency in the use of the available water. The adoption of this agent-based perspective aims to move beyond the traditional centralized approach to water resources management and to analyze different levels of cooperation (e.g., *Watkins*, 2006), from fully coordinated strategies and full information exchange to completely uncoordinated practices where all the decision makers independently pursue their local objectives.

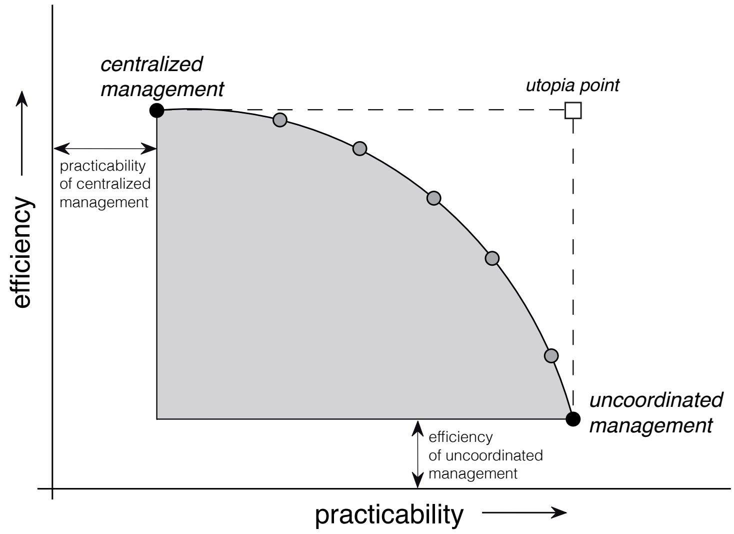

Figure 1.2: Representation of the conflict between the system-wide efficiency and the practicability of watershed management strategies.

As suggested in Figure 1.2, a trade-off exists between system-wide efficiency and the associated practicability in real decision-making processes: the fully cooperative centralized approach aims at the maximization of the former while neglecting the latter, thus leading to solutions which are the most efficient but impracticable, for example the failure of international initiatives such as the ZAMCON Protocol for the management of the Zambezi River (*Tilmant and Kinzelbach*, 2012). At the other extreme, real world management practices are often non-coordinated and, therefore, they usually result in very low efficiency in the utilization of water resources. It is worth noting that the utopia point (i.e., the absolute optima of both efficiency and practicability represented by the white square in Figure 1.2) is usually not feasible, except for rare cases where the two objective functions are actually orthogonal and not conflicting. Yet, there exists room between these two extremes to design intermediate, distributed solutions (i.e., the grey points), more efficient than the uncoordinated ones and more realistic, and thus more practical, than the centralized ones. Different policy mechanisms such as regulatory constraints or economic incentives are usually applied to explore this space (*Pannell*, 2008). However, these mechanisms mainly rely on the empirical experience of the authority in charge to promote

negotiated solutions, while mathematical and technological tools to identify such distributed solutions are nearly undeveloped and, therefore, the applicability of them has been so far limited.
The aim of the proposed agent-based framework is to provide tools to describe the decisions made by each agent as well as insights for the design of proper coordination mechanisms driving the originally uncoordinated decision-making structure toward a system-wide efficient situation. Moreover, coupling the agent-based modeling scheme with state-of-the-art Control Theory techniques allows a better understanding of the feedbacks between agents objectives, agents decisions, and the environment they share, as well as the description of their coevolution and coadaptation under change.
The main components of the proposed framework are illustrated in Figure 1.3. Given a complex water system as the one represented in the figure, which includes multiple reservoirs and four conflicting objectives, the agent-based model of the system is firstly developed. The two decision makers operating the two reservoirs are modeled as active agents, while passive agents represent the stakeholders who do not make decisions but represent other interests affected by the operation of the reservoirs (i.e., flooding issues and domestic water supply). According to the definition of proactive agents, an objective function is associated to both active and passive agents. Three alternative scenarios (represented on the right side of the figure) are then considered: *i*) an upper-bound alternative, where the agents are fully cooperative, exchange complete information, and solve a four-objective optimization problem to attain the maximum of the system-wide efficiency (i.e., the four-objective Pareto front); *ii*) a lower-bound alternative, where the agents are completely self-interested and the active agents act according to local objectives only, thus producing negative externalities on the interests of the passive agents and resulting in uncoordinated solutions at the sytem-level; *iii*) intermediate solutions obtained through mechanism design strategies which preserve the originally uncoordinated decision-making structure of the problem (i.e., the active agents consider their local objectives only) and introduce some coordination mechanisms (e.g., information exchange, normative constraints, or economic incentives) aimed to improve the system-wide efficiency. Finally, the upper and lower bound alternatives provide a reference to evaluate the performance of the mechanism design solutions in terms of system-wide efficiency as well as agent-level practicability.

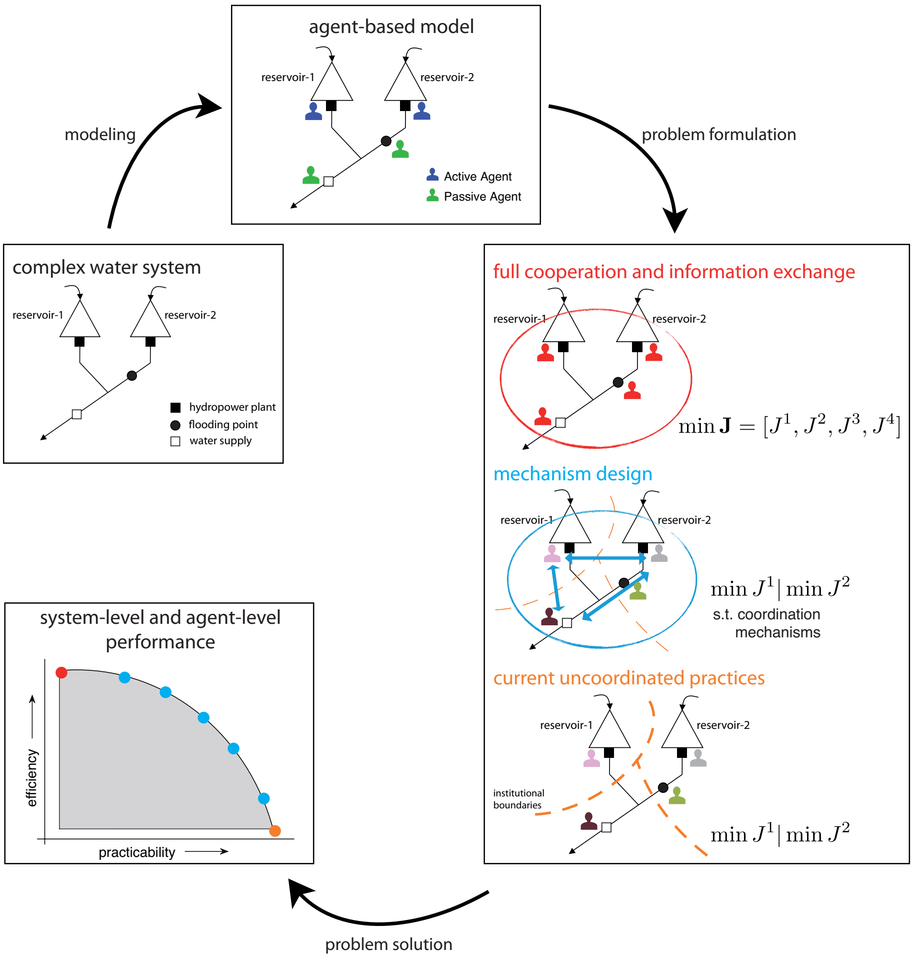

Figure 1.3: Illustration of the agent-based decision analytic framework.

1.4 Overview of the chapters

1.4.1 Chapter 2

Chapter 2 provides the reader with background information and the novel contributions in terms of tools and methodologies used in the applications of the agent-based decision analytic framework presented in this thesis. Section 2.1 frames the proposed framework in the multi-objective decision-making literature, highlighting the advantages of adopting an a posteriori decision-making method. Section 2.2 describes the recent advances in visual analytics which support effective exploration and understanding of complex, high dimensional decision and objective spaces.

Section 2.3 introduces two MAS methods (i.e., distributed constraint satisfaction and distributed constraint optimization) to solve distributed optimal planning problems. They have been developed within the distributed artificial intelligence community and their application in environmental contexts is still missing. Section 2.4 reviews the state-of-the-art methods to solve water management problems, focusing on two approaches (i.e., on-line control and direct policy search) able to scale to large-scale systems and many-objective optimization problems. These features are suitable for the application of the agent-based framework in complex water management problems, where MAS tools are nearly unexplored. Finally, Section 2.5 introduces the direct policy conditioning method to support the identification of effective mechanism design alternatives.

1.4.2 Chapter 3

Chapter 3 introduces the problem of balancing system-wide efficiency and solutions practicability working on a hypothetical water allocation planning problem. Two extreme situations, corresponding to the centralized and uncoordinated strategies in Figure 1.2, are first computed. An extensive analysis of the intrinsic conflict between efficiency and practicability is developed, by highlighting the limitations (i.e., low practicability) of centralized and fully cooperative solutions in systems characterized by the presence of distributed decision making institutions. Then, different levels of coordination are explored by means of MAS methods based on distributed constraint optimization problems. These methods allow the development of constraint-based mechanism design strategies to explore the "gray area" in Figure 1.2. Finally, alternative mechanism design strategies based on economic incentives are also discussed. This chapter is adapted from a paper under review for publication in the Journal of Water Resources Planning and Management (*Giuliani et al.*, 2013a).

1.4.3 Chapter 4

In Chapter 4, the agent-based decision analytic framework is applied in a real world management problem to assess the value of cooperation and information exchange in transboundary systems. The Zambezi River basin is used as a case study to estimate the benefits potentially achievable by the downstream country (i.e., Mozambique), under different levels of cooperation with the two upstream countries (i.e., Zambia and Zimbabwe). The differences in the system-level benefits attained under different scenarios of cooperation provide an estimate of both the economic value of full cooperation, measured as the benefits obtained by full cooperation with respect to the one with coordination only, and the economic value of information exchange, measured as the benefits obtained with coordination with respect to the ones with no cooperation. Chapter 4 represents the first application adopting an agent-based perspective to study real world water management problems in non-cooperative decision-making contexts. Instead of adopting a priori behavioral rules, the agents' behaviors are defined according to an on-line control approach for solving management problems at the agent and at the system-level. Moreover, it contributes a novel procedure to assess the role of information exchange and the benefits that can be achieved with this simple upstream-downstream coordination mechanism. Finally, both the values of full cooperation and information exchange are estimated in a multi-objective context, which requires to work on the Pareto front to a posteriori derive the economic value of the ecological objectives, whose monetization is usually affected by high uncertainty. This chapter is adapted from a journal paper published in Water Resources Research (*Giuliani and Castelletti*, 2013).

1.4.4 Chapter 5

In Chapter 5, the application of the agent-based decision analytic framework shows the potential for a co-adaptation strategy between the agents (i.e., farmers and water managers) decisions to improve the effectiveness of agricultural water management practices. The proposed co-adaptation aims to match the needs of the farmers with the design of water supply management policies, under current and projected climate. The Lake Como serving the Muzza-Bassa Lodigiana irrigation district (Italy) is used to illustrate the methodology on a real world case study. Although many studies have assessed climate change impacts on agricultural practices and water management, most of them assume few scenarios of water demand or water supply separately, while an analysis of their reciprocal feedbacks is still missing. Chapter 5 proposes an integrated procedure to model water supply and demand as coupled human (farmers and water managers) and natural (crops) systems, where people and nature interact reciprocally, form complex feedback loops, and co-evolve under changing conditions. The proposed co-adaptation strategy closes the loop between the two systems (demand and supply) by cross-conditioning farmers and water managers decisions. It allows the farmers to decide the most profitable crop option on the basis of an expected water supply. Knowing the farmers decisions, the water supply strategy (i.e., the Lake Como regulation) is then optimized with respect to the actual irrigation demand of the crops. By iteratively running this procedure, the farmers and the water manager will exchange information until the system converges to an equilibrium. The analysis is performed under current and future climate conditions, in order to assess the potential for the adaptation loop to enhance the efficiency of agricultural water management practices and foster crop production as well as to mitigate climate change adverse impacts. This chapter is adapted from a paper under preparation for Global and Environmental Change (*Giuliani et al.*, 2013d).

1.4.5 Chapter 6

In Chapter 6, the agent-based decision analytic framework combines reservoir policy identification, many-objective optimization under uncertainty, and visual analytics to characterize current water reservoir operations and discover key tradeoffs between alternative policies for balancing evolving demands and system uncertainties. Moreover, this chapter contributes a novel method based on input variable selection techniques to support the identification of effective mechanism design strategies. These tools are demonstrated on the Conowingo Dam, located within the Lower Susquehanna River, USA. The Lower Susquehanna River is an interstate water body that has been subject to intensive water management efforts due to the presence of many stakeholders (agents) affected by the Conowingo Dam operation. The Lower Susquehanna system includes competing demands from urban water supply, atomic power plant cooling, hydropower production, and federally regulated environmental flows. To provide effective support to the Susquehanna River Basin Commission and avoid policy inertia and myopia, the current regulation of Conowingo Dam is identified and refined to balance the conflicting objectives as well as the uncertainties related to the hydroclimatic variability. Then, alternative policy mechanisms are designed by directly constraining the decision space in order to dynamically condition the reservoir operating policy and better balance the primary operating objectives (i.e., guaranteeing the public water supply and maximizing the hydropower revenue) with environmental protection and recreational interests. This chapter is adapted from a paper under review for publication in Water Resources Research (*Giuliani et al.*, 2013c) and a paper in preparation for the 2014 IFAC World Congress (*Giuliani et al.*, 2014).

2 Methods and tools

2.1 Multi-objective decision-making

Contemporary water resources problems often involve multiple, conflicting, and non-commensurable objectives. In such multi-objective contexts, the traditional concept of optimality is replaced by that of Pareto efficiency, which imposes to explicitly consider the preference structure of the parties involved. Yet, the legacy planning strategies that have predominately shaped modern water reservoir operations are termed a priori multi-criteria decision analysis (MCDA) methods (*Cohon and Marks*, 1975). As reviewed by *Chankong and Haimes* (1983), these methods first elicit (or assume) a priority ranking or weighting of objectives, which is then used to reduce the multi-objective operations problem into a single-objective optimization that sought one "compromise" solution. However, these approaches were recognized as making very strong assumptions (linearity, perfect foresight, limited if any uncertainty, convexity, well defined understanding of planning alternatives and preferences, etc.) that could cause severe biases (*Haimes and Hall*, 1977).
Again following the classification of *Cohon and Marks* (1975), an alternative approach to a priori MCDA methods are a posteriori generating techniques, where the full set of Pareto-optimal solutions are generated prior to eliciting the decision makers (DMs) preferences. The underlying benefit of the a posteriori approach is that DMs do not have to state what is preferred in absence of their understanding of what is attainable (assuming a well formulated management problem). The core limitation in this approach is the computational cost of identifying the Pareto front. Classically, these approaches have required similar weighting based methods as used in MCDA a priori methods, with the distinguishing difference that the single-objective optimization is repeated for every Pareto-optimal point generated by adapting the weighting of the objectives (*Chankong and Haimes*, 1983; *Soncini-Sessa et al.*, 2007a; *Coello Coello et al.*, 2007). Interactive approaches (e.g., reference point method (*Wierzbicki*, 1980), Pareto race (*Korhonen and Wallenius*, 1988)) have been developed in order to interactively explore the Pareto space without having to fully compute it in advance, thus mitigating the associate computational burden (e.g., *Deb et al.*, 2006; *Thiele et al.*,

2009). The complexity and high number of questions to be posed to the DM remain an unsolved problem (*Larichev*, 1992).
More recent strategies have transitioned to multi-objective evolutionary algorithms (e.g., *Nicklow et al.*, 2010; *Reed et al.*, 2013) or multi-objective reinforcement learning (e.g., *Castelletti et al.*, 2013b) to avoid the limitation of generating a single Pareto-optimal point per optimization run while also broadening the number and complexity of objectives that can be resolved.
This thesis relies on a posteriori decision-making and the proposed decision analytic framework exploits the recent advances in visual analytics (Section 2.2) to allow the exploration and understanding of complex, high dimensional information through highly interactive visual tools. In what follows, it is always assumed to deal with multi-objective problems, whose solution is a Pareto front.

2.2 Visual analytics

The solution of multi-objective planning and management problems yield a set of Pareto-optimal solutions. Adopting the a posteriori decision-making approach described in the previous section, the DM selects one or more preferred solutions only after the generation of the entire Pareto-optimal set by analyzing the corresponding Pareto front. However, effective tools to explore these large objective spaces are necessary to support more informed decisions as well as effective stakeholders negotiations.
Recent advances in visual analytics (e.g., *Thomas and Cook*, 2005; *Keim et al.*, 2006; *Kollat and Reed*, 2007; *Lotov and Miettinen*, 2008) allow the exploration of large multidimensional spaces, thus favoring the adoption of a posteriori decision-making approaches. The idea of searching and visualizing the Pareto front for two objectives was firstly introduced by *Gass and Saaty* (1955). Then, *Louie et al.* (1984) and *Haimes et al.* (1990) proposed to use decision maps to show three-objectives tradeoffs as collections of two-objective tradeoff curves with different values of the third objective.
The concept of interactive decision maps (*Lotov et al.*, 1998, 2004) was then introduced to inform the DM in problem with more than three objectives. Several slices of the multi-dimensional Pareto front displayed in different colors are superimposed in the same two-objective plane. Animations of the figure for different values of the other objectives allow the analysis of the influence of the remaining objectives. Figure 2.1 shows an example of a five-dimensional interactive decision map.
More recently, the framework proposed by *Kollat and Reed* (2007) al-

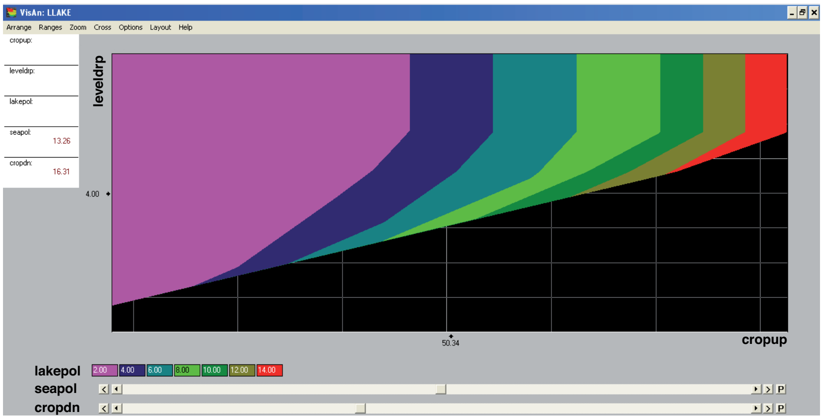

Figure 2.1: Example of interactive decision maps.

lows the user to visualize up to seven design objectives simultaneously using scatter and glyph plots with the AeroVis software. Up to three objectives can be plotted on the spatial coordinate axes. Glyphs color, size, orientation, and transparency allow the visualization of additional objectives. The software allows multiple interactive analyses, such as *i*) quickly change which variables are plotted; *ii*) specify thresholds of interest for each objective and "brush out" solutions that fail to meet DMs requirements; *iii*) perform non-domination sorting with respect to specified objectives; *iv*) mark interesting solutions across different linked views. Figure 2.2 shows an example of a six-dimensional AeroVis plot.
Parallel axes plot (*Inselberg*, 1997) serve as a complementary visual tool for understanding key tradeoffs. Each Pareto-optimal solution is shown as a line crossing the parallel axes, representing the different objectives, at the values of their corresponding performance. The ideal solution would be a horizontal line running along the top/bottom of all of the axes, with the conflicts designated as diagonal lines between two adjacent axes. Figure 2.3 shows the same solutions of Figure 2.2 using a parallel axes plot.
Three-dimensional virtual reality facilities can be used to visualize high-dimensional Pareto front as well as to analyze (e.g., by means of rotation, zooming, and other navigation possibilities) the corresponding Pareto-optimal set aiding the selection of a particular compromise solution (*Madetoja et al.*, 2008). Figure 2.4 shows an example of a 3D virtual reality system used to analyze a three-objectives Pareto front in a virtual reality laboratory.

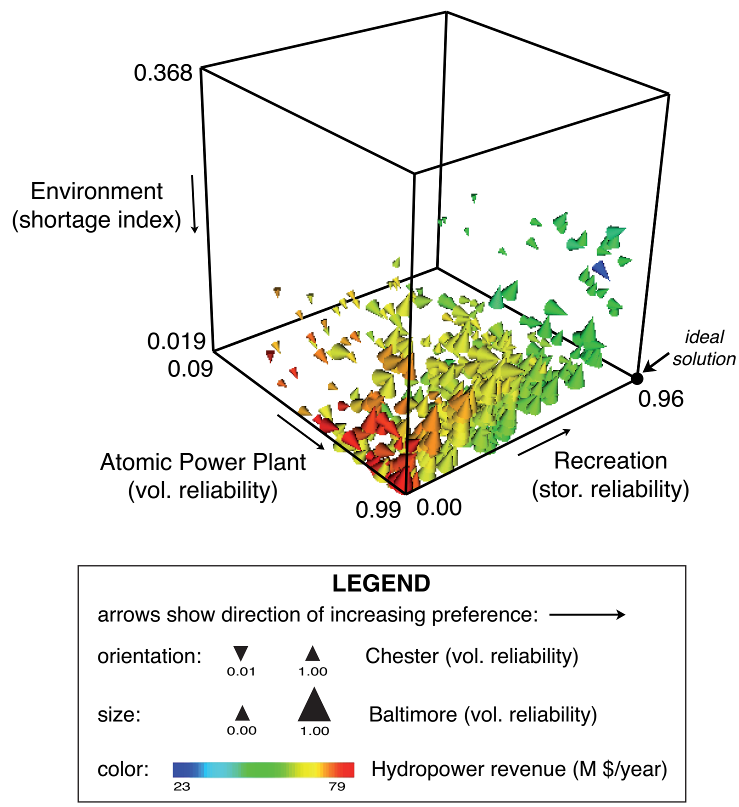

Figure 2.2: Example of AeroVis plot from *Giuliani et al.* (2013c).

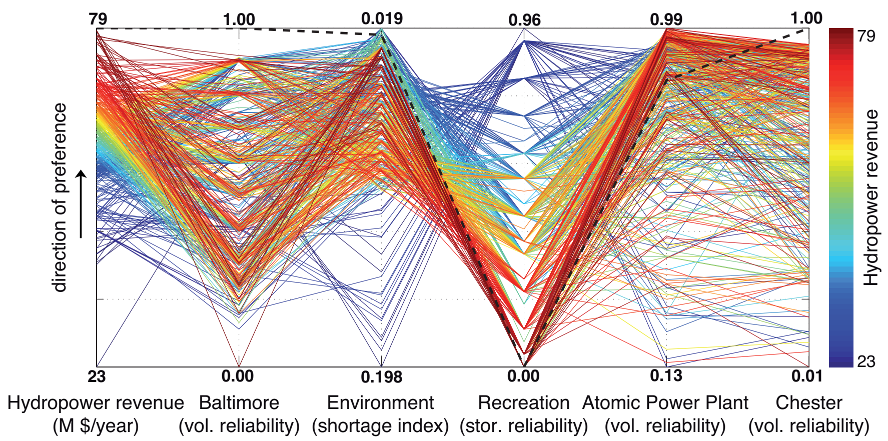

Figure 2.3: Example of parallel axes plot from *Giuliani et al.* (2013c).

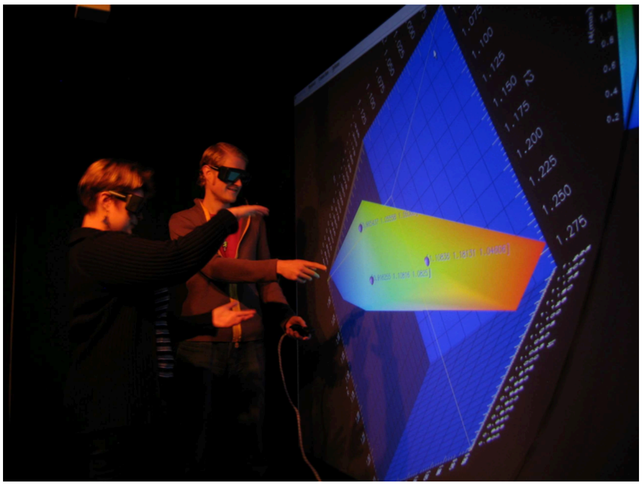

Figure 2.4: Example of 3D virtual reality system (*Madetoja et al.*, 2008).

Transitioning to higher dimensional many-objective formulations may reveal that lower dimensional results represent extreme corners of the objective space that have little interest for DMs (e.g., *Kollat et al.*, 2011; *Kasprzyk et al.*, 2012; *Woodruff et al.*, 2013, and references therein). Moreover, many-objective representations of tradeoffs can help in reducing the negative impacts from two forms of decision bias (*Brill. et al.*, 1990). The first one, *cognitive myopia* (*Hogarth*, 1981), occurs in low dimensional optimization when stakeholders inadvertently ignore alternatives that could strongly influence their decision preferences; the second one, *cognitive hysteresis* (*Gettys and Fisher*, 1979), refers to the challenge that lower dimensional, highly constrained problem representations often reinforce DM's preconceptions.
An example of how adding problem objectives can fundamentally change the interpretation of the results is given in Figure 2.5 (adapted from *Woodruff et al.*, 2013). Figure 2.5a shows the Pareto-optimal solutions of a three-objective problem according to a single-objective perspective (e.g., objective J^1). In this view, a solution minimizing J^1 appears the unique optimum. In Figure 2.5b, a second objective is added and the solution minimizing J^1 corresponds to the minimum also of J^2. If only J^1 and J^2 were considered, a DM would assume that no conflict exists and select the solution that minimizes both objectives (the red point). However, adding J^3 as a third objective disrupts this view of the problem. Figure 2.5c shows the emergence of significant tradeoffs between the three objectives due to the constraint $J^1 + J^2 + J^3 = 1$. For any

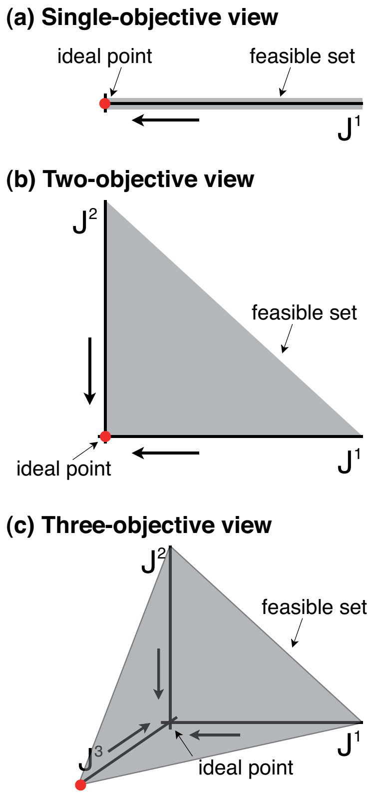

Figure 2.5: Example of how adding problem objectives can fundamentally change the interpretation of the results: in panels (a) and (b) the set of solutions does not show any tradeoff, while in panel (c) a clear tradeoff emerges.

fixed value of J^3, the first two objective are strongly conflicting. The idea that J^1 and J^2 are in agreement represents a misconception and a decision bias which arises from using formulations with fewer objectives. The solution that the DM had thought ideal (the red point corresponding to $J^1 = 0, J^2 = 0$) is actually an extreme point of the Pareto front ($J^1 = 0, J^2 = 0, J^3 = 1$), and a compromise solution that might be potentially interesting for the DM is probably far from there.

2.3 MAS for optimal planning

A wide variety of MAS methods have been developed in the distributed artificial intelligence (DAI) community for planning problems (see Section 1.2), such as distributed scheduling (e.g., *Sycara et al.*, 1991), distributed planning (e.g., *Hirayama and Yokoo*, 1997), distributed sensor and robots networks (e.g., *Modi et al.*, 2001; *Shen et al.*, 2002), satellite constellation (e.g., *Barrett*, 1999), or distributed resources allocation (e.g., *Conry et al.*, 1991). Conversely, as anticipated in Section 1.3.1, most of the MAS applications in the environmental community adopts reactive agents and performs scenario analysis, while proactive agents solving optimization problems are nearly unexplored. In the water resources field, *Yang et al.* (2009) adopt a MAS approach to solve a water allocation problem among multiple, spatially distributed water users. According to an agent-based distributed formulation and assuming one-to-one correspondence between agents and objectives, Problem (1.1) can be reformulated as a sequence of q local problems, each one defined as:

$$\mathbf{u}_i^{p*} = \arg\min_{\mathbf{u}_i^p} J^i(\mathbf{u}_i^p, \mathbf{u}_{-i}^p) \tag{2.1a}$$

subject to

$$c_{i,1}(\mathbf{u}_i^p, \mathbf{u}_{-i}^p), \ldots, c_{i,r}(\mathbf{u}_i^p, \mathbf{u}_{-i}^p) \leq 0 \tag{2.1b}$$

$$\mathbf{u}_i^p \in \mathcal{D}_i \tag{2.1c}$$

where $i = 1, \ldots, q$, $\mathbf{u}_i^p$ are the decision variables of the i–th agent and $\mathbf{u}_{-i}^p$ the ones of all the other agents except i. The i–th agent's decisions are optimized with respect to its local objective function $J^i(\cdot)$, which however depends also on the decisions of the other agents. It is worth noting that also the constraints in eq. (2.1b) depend on both the $\mathbf{u}_i^p$ and $\mathbf{u}_{-i}^p$.

In such distributed problems, the agents look at their local objectives only, without considering the potentially negative externalities that their decisions produce for the others. The two main approaches developed within the DAI field for investigating how agents can instead coordinate their decisions are distributed constraint satisfaction problems (DCSPs, see *Yokoo and Hirayama* (2000)) and distributed constraint optimization problems (DCOPs, see *Modi et al.* (2005)). The DCSP formulation defines a distributed feasibility problem with Boolean constraints, which can be only satisfied or unsatisfied, while the DCOP is a distributed optimization problem dealing with objective functions represented as a weighted sum of costs or valued constraints. According to these formulations, each agent solves a local planning problem with respect to its objective, subject to a set of constraints conditioning its decisions. The constraints represent either physical constraints (e.g., canal capacity) or normative constraints (e.g., minimum environmental flow). The aim of these latter is to condition the individualistic, local decisions of the agents in order to increase the efficiency of the uncoordinated solution. In practice, normative constraints impose some cost on combinations of values selected by agents which produce negative externalities, so that agents are "forced" to select combinations of values that minimize that cost (i.e., that are good from a global standpoint). Moreover, the adoption of DCSP- and DCOP-based methods ensures global solution quality operating efficiently in a distributed search process which attempts to narrow the number of exchanged messages. These characteristics have an important practical consequence, namely the possibility of defining a weak coordination mechanism between the agents with limited transaction costs associated to the agent interactions. This represents a significant advantage in real world application, as ex post assessments of the role of transaction costs showed that they are ubiquitous and nontrivial (*McCann and Easter*, 1999; *McCann et al.*, 2005).

2.3.1 Distributed constraint satisfaction problems

Formally, a DCSP consists of n decision variables, $\mathbf{u}^p = [u_1^p, \dots, u_n^p]$, each assigned to a different agent, where the values of the variables are taken from finite, discrete domains $\mathcal{D}_1, \dots, \mathcal{D}_n$, and of a number of Boolean constraints over the values of these variables $c_1(\mathbf{u}^p), \dots, c_r(\mathbf{u}^p)$, where $c_j(\mathbf{u}^p) \in \{\text{true}, \text{false}\} \quad \forall j$. According to the planning nature of the problem, the decision variables do not change in time (e.g., the assignment of water rights is done once and can be modified only by reformulating a new planning problem).

A solution of a DCSP is an assignment of values to all the variables such that all the constraints are satisfied. DCSPs are solved by employing distributed algorithms, like those surveyed in *Yokoo and Hirayama* (2000), which assume that the constraints are binary (i.e., each constraint involves only two variables) and agents can reliably communicate to exchange the values they select for their variables. Similar to the case of centralized CSPs, these distributed algorithms can be divided in two classes: backtracking algorithms and iterative improvement algorithms. In a backtracking algorithm, such as the asynchronous backtracking algorithm (*Yokoo et al.*, 1992), a value assignment to a subset of decision variables that satisfies all of the constraints within the subset is first constructed. This value assignment is called a partial solution. A partial solution is then expanded by adding new decision variables one by one, until it becomes a complete solution. When no values satisfying all of the constraints with the partial solution are available, the value of the most recently added variable to the partial solution is modified. This operation is called *backtracking*. In iterative improvement algorithms, such as the distributed breakout algorithm (*Hirayama and Yokoo*, 2005), a tentative initial value is assigned to all the decision variables and no partial solution is constructed. Then, a flawed solution is revised according to some heuristic process (e.g., a variable value is changed so that the number of constraint violations is minimized). Iterative improvement algorithms are usually efficient as they do not require an exhaustive search in revising a flawed solution, but they can be incomplete, meaning that they do not guarantee to find a feasible solution; algorithms based on backtracking are instead often complete.

2.3.2 Distributed constraint optimization problems

A DCOP consists of n decision variables, $\mathbf{u}^p = [u_1^p, \ldots, u_n^p]$, each assigned to a different agent, where the values of the variables are taken from finite, discrete domains $\mathcal{D}_1, \ldots, \mathcal{D}_n$, and of a number of valued constraints over the values of these variables $c_1(\mathbf{u}^p), \ldots, c_r(\mathbf{u}^p)$, where $c_j(\mathbf{u}^p) \in \mathbb{R} \quad \forall j$. A solution of a DCOP is an assignment of values to all the variables such that a given objective function g is maximized or minimized. Usually, the objective function is a weighted sum of the functions representing the costs for constraint violations and is minimized, namely $\min_{\mathbf{u}^p} g = \min_{\mathbf{u}^p} \sum_{j=1}^{r} w_j \cdot c_j(\mathbf{u}^p)$, where w_j is the weight of $c_j(\mathbf{u}^p)$.
Probably, the most known (though not the most efficient for all problems) distributed algorithm for solving DCOPs is Adopt (*Modi et al.*, 2005), which performs a global search of the solution space and provides theoretical guarantees on the global solution quality (see Section 2.3.3). Also in the case of Adopt, agents are responsible for choosing the values of their variables. Extensions of the basic formulations of DCSP and DCOP allow each agents to control multiple variables as well as to consider agents solving multi-objective optimal planning problems. Note that a DCSP is a special case of a DCOP, where the constraints are Boolean and, therefore, they can be only satisfied or unsatisfied. A solution of a DCSP can be only feasible or unfeasible. Conversely, a DCOP allows the identification of solutions with a certain degree of quality or cost, depending on the value of the objective function g.
To better clarify the difference between DCSP and DCOP, let consider the simple problem represented in Figure 2.6, where each variable u_i^p $(i = 1, 2, 3)$ is assigned to an agent and has the same domain $\mathcal{D}_i = \{0, 1\}$. The objective function of each agent is $J^i = 10$ if $u_i^p = 1$ and $J^i = 0$ if $u_i^p = 0$. The two links in the figure represent two inequality constraints defined as $u_1^p \neq u_3^p$ and $u_2^p \neq u_3^p$. In the case of the DCSP,

(a)

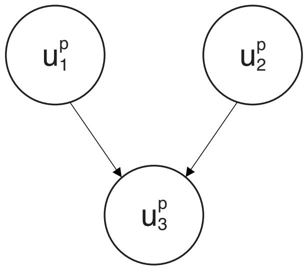

(b)

u_i^p	u_j^p	$c(u_i^p, u_j^p)$
0	0	∞
1	1	∞
0	1	1
1	0	0

Figure 2.6: Example of a DCSP/DCOP problem. Panel (a) shows the constraint graph, where each link represents an inequality constraint. Panel (b) reports the values assigned to the constraints in the DCOP formulation.

the constraints are Boolean. The problem has two feasible solutions $\{u_1^p = 1, u_2^p = 1, u_3^p = 0\}$ and $\{u_1^p = 0, u_2^p = 0, u_3^p = 1\}$, which are considered equally good as they both satisfy all the constraints and no measure of quality is defined. Conversely, the DCOP formulation allows the definition of valued constraints to differentiate solutions with different degrees of quality. Assuming the values of the constraints are defined as in Figure 2.6b, the DCOP has one optimal solutions only, namely $\{u_1^p = 1, u_2^p = 1, u_3^p = 0\}$, which corresponds to a cost for constraint violations $g(u_1^p, u_2^p, u_3^p) = c(u_1^p, u_3^p) + c(u_2^p, u_3^p) = 0$ and to the maximum of system-level benefit (i.e., $J^1 + J^2 + J^3 = 20$). Note that, the DCOP problem has another feasible solution $\{u_1^p = 0, u_2^p = 0, u_3^p = 1\}$, with a higher cost for constraint violations $g(u_1^p, u_2^p, u_3^p) = 2$ and a lower system-level performance (i.e., $J^1 + J^2 + J^3 = 10$). Finally, the DCOP has other two solutions, which are the ones considered unfeasible in the DCSP formulation, which assume an infinite cost in the DCOP formulation and, consequently, will never be considered. The DCOP formulation is therefore more flexible as it allows to discard the same unfeasible solutions as in the DCSP problem and also to distinguish between the feasible solutions depending on their system-level performance.

2.3.3 Asynchronous distributed optimization algorithm

Adopt (Asynchronous Distributed OPTimization) is probably the most known distributed algorithm for DCOP problems (*Modi et al.*, 2005). It can find either an optimal solution or a bounded-error approximate one using only asynchronous and local communication, meaning that agents

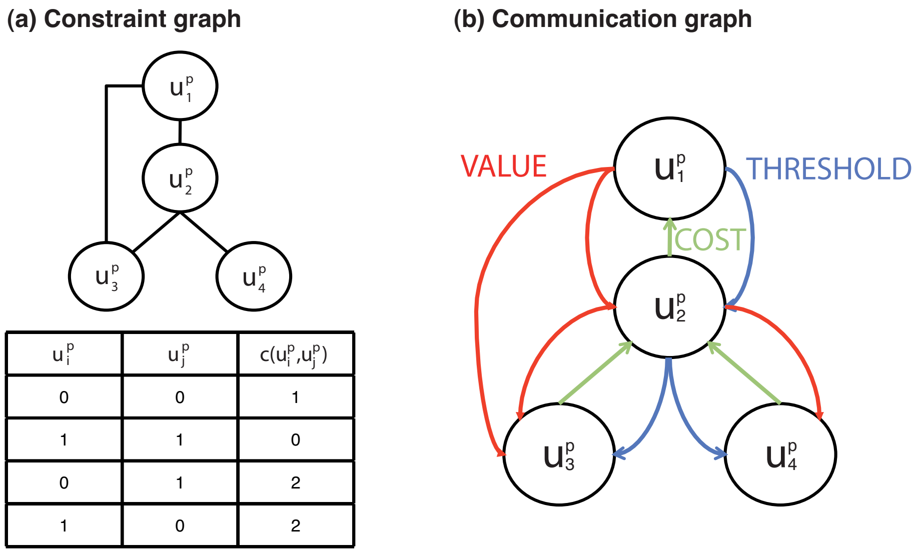

u_i^p	u_j^p	$c(u_i^p, u_j^p)$
0	0	1
1	1	0
0	1	2
1	0	2

Figure 2.7: Example of Adopt distributed search: panel (a) represents the constraint graph as well as the values associated to each constraint; panel (b) shows the communication graph, where the agents are prioritized in a depth-first search tree.

do not broadcast messages to every agent but only to neighboring agents in the attempt to limit the number of exchanged messages.

The main feature of Adopt is to perform a distributed backtrack search (see Section 2.3.1) using an "opportunistic" best-first search strategy, which allows each agent to change variable value whenever it detects there is a possibility that some other solution may be better than the one currently under investigation. This search strategy increases asynchrony because an agent does not need global information to make its local decisions. Moreover, it allows partial solutions to be abandoned before sub-optimality is proved. Consequently, Adopt needs to efficiently reconstruct previously considered partial solutions through the use of backtrack threshold. These two key ideas yield efficient asynchronous search for optimal solutions. Finally, Adopt provides a built-in termination detection mechanism based on a bound interval related to the cost of the optimal solution. When the lower bound equals the upper bound, the cost of the optimal solution has been determined and the agents can terminate the search. Bound intervals are used to search for bounded-error approximation of the optimal solution.

Figure 2.7 reports an example of Adopt distributed search for a problem involving four agents. The agents are organized in a constraint graph, whose nodes are the agents and edges connect agents that share a constraint, which are called neighboring agents (Figure 2.7a). The objective is to find an assignment to all the variables such that the total cost function $g = \sum_{j=1}^{4} c_j$ is minimized. In this example, the optimal solution is $\mathbf{u}^p = [1, 1, 1, 1]$, corresponding to a total cost $g = 0$. Adopt assumes that constraints are at most binary (i.e., they involve one or two variables/agents). No restrictions are placed on the topology of the constraints, so loops are allowed. However, Adopt requires a preprocessing phase to prioritize the agents in a depth-first search tree (e.g. *Korf*, 1985; *Collin and Dolev*, 1994), in which each agent has a single parent and multiple childrens. The result is a communication graph (Figure 2.7b) where the constraints are defined between an agent and any of its ancestors or descendents, but no constraints are allowed between agents in different branches of the DFS tree. When the algorithm starts, all the agents choose their variable values concurrently. Variable values are sent down through constraint edges via VALUE messages. An agent sends VALUE messages only to neighbors lower in the DFS tree and receives VALUE messages only from neighbors higher in the DFS tree. A second type of message, named THRESHOLD message, is sent only from parent to child. A THRESHOLD message contains a backtrack threshold, initially zero. Upon receipt of any type of message, an agent *i*) calculates cost and possibly changes variable value and/or modifies its backtrack threshold; *ii*) sends VALUE messages to its lower neighbors and THRESHOLD messages to its children; *iii*) sends a third type of message, named COST message, to its parent. A COST message contains the cost calculated by the children agent plus any cost reported to this latter by its children. To summarize: variable value assignments (VALUE messages) are sent down the DFS tree, while cost feedback (COST messages) percolate back up the DFS tree. THRESHOLD messages are sent down the tree to reduce redundant search.

The dimensionality of the problem can limit the applicability of Adopt to attain optimal solutions because, although the number of messages exchanged by the agents grows approximately linearly with the number of agents, its worst-case time complexity is exponential in the number of agents. However, there exist approximated algorithms able to find quasi-optimal solutions when the number of agents increases, such as the distributed breakout algorithm (*Hirayama and Yokoo*, 2005).

2.4 MAS for optimal management

So far, optimization-based MAS approaches have been mainly adopted in planning problems (e.g., *Yang et al.*, 2009). However, several water-related applications, such as water reservoir regulation or groundwater pumping systems management, require to take sequential decisions given the current conditions of the system and can be formulated as management problems (see Section 1.2). For example, the optimal reservoir operation can be designed solving Problem (1.2) and deriving an operating policy p, defined as the sequence of operating rules which step-by-step provide the release decisions $\mathbf{u}_t = m_t(\mathbf{x}_t)$ for each dam outlet over the time interval $[t, t+1)$ depending on the state vector $\mathbf{x}_t$ (e.g., reservoir storage and inflow forecast). The use of MAS to design water management policies is nearly unexplored. In this thesis, state-of-the-art Control Theory methods are combined within the proposed agent-based framework to support water resources management.
Stochastic dynamic programming (SDP, see *Bellman* (1957)) is the most flexible and accurate method for solving Problem (1.2). However, SDP and dynamic programming (DP) family methods (e.g., *Powell*, 2007) suffers from a dual curse which prevents its practicable application in large-scale complex systems: *i*) the curse of dimensionality (*Bellman*, 1957), namely the computational cost of dynamic programming grows exponentially with state, decision, and disturbance vectors and would be inapplicable with medium-to-high order dynamical models (e.g., water reservoir networks with more than 2/3 storage units); *ii*) the curse of modeling (*Tsitsiklis and Van Roy*, 1996), meaning the use of in-line model-based computations that make impossible the direct, model-free use of exogenous information into the controller and the use of process-based simulation models (e.g., hydrodynamic and ecological). DP-based approaches have another limitation as they are single-objective methods. They can be thus used to solve multi-objective problems only by reformulating them as a family of parametric single-objective problems and reiteratively running a single-objective optimization for different values of the parameter to explore the Pareto front. This remarkably affects the computational requirements as the number of single-objective problems to solve grows exponentially with the objectives number. The two most common scalarization techniques are the *weighted sum* and *ε-constraint methods* (*Gass and Saaty*, 1955; *Haimes et al.*, 1971). The former is based on a linear combination of the objectives. With the latter, the conversion to a set of single-objective problems is obtained by transforming all the objectives, but one, into constraints. The weighting method is usually preferred as it ensures to find Pareto efficient solutions only,

while the ε-constraint method might provide semi-dominated solutions and, moreover, it does not satisfy the time-separability requirements of SDP. According to the weighting method, Problem (1.2) can be reformulated as:

$$p^* = \arg\min_p \boldsymbol{\lambda} \cdot \mathbf{J}(p) \tag{2.2}$$

where $\boldsymbol{\lambda}$ is the vector of weights, such that they sum to one (i.e., $\sum_{k=1}^{q} \lambda_k = 1$) and are non-negative (i.e., $\lambda_k \geq 0 \quad \forall k$). By solving Problem (2.2) for different values of the weights $\boldsymbol{\lambda}$, a finite subset of the generally infinite Pareto-optimal policy set is obtained. However, the accuracy in the approximation of the Pareto front might be scarce, with a limited solution diversity due to the non-linear relationships between the values of the weights and the corresponding objectives values.
Many approaches have been proposed to overcome the limits of DP family methods depending on the specific features of the problem under study, such as approximate dynamic programming, implicit stochastic optimization, direct policy search, on-line approaches (for a review see *Labadie*, 2004; *Castelletti et al.*, 2008a; *Celeste and Billib*, 2009, and reference therein). The adoption of on-line approaches (see Section 2.4.1) can be particularly suitable in large-scale systems where the curse of dimensionality is the main limiting factor and the optimization problem involves few objectives. Conversely, when the problem takes a many-objective nature, namely it involves more than four or five objectives (*Farina and Amato*, 2002; *Fleming et al.*, 2005), direct policy search approaches coupled with multi-objective evolutionary algorithm (see Section 2.4.2) is preferable. However, these approaches generally do not provide any theoretical guarantee on the optimality of the resulting solutions. Interesting multi-objective approximate dynamic programming methods that ensure some anticipated favorable properties of the policy obtained are recently proposed by *Castelletti et al.* (2013b) and *Giuliani et al.* (2013b).

2.4.1 On-line control

Model predictive control (MPC, see *Bertsekas* (2005) and references therein), also referred to as naive feedback control (NFC, see *Bertsekas* (1976)), is an on-line control approach based on the sequential resolution of multiple, single-objective, open-loop control problems defined over a finite, receding time horizon (*Mayne et al.*, 2000; *Castelletti et al.*, 2008b), which allows to overcome the SDP limits related to the curse

of dimensionality. However, their extension to multi-objective problems (e.g., *Pianosi and Soncini-Sessa*, 2009; *Galelli et al.*, 2012) relies on the same scalarization techniques adopted for DP family methods (see the previous section), which can be prohibitive when the number of objectives increases.
The fundamental idea of on-line control is to exploit, at each time t, a forecast over the finite horizon $[t, t+h]$ of the external drivers (e.g., the inflow), called nominal value. The forecast is obtained by a model predictor that uses all the information available at time t (e.g., precipitation, inflow at previous time). The corresponding sequence of optimal decisions $\mathbf{u}_t^*, \dots, \mathbf{u}_{t+h-1}^*$ is then obtained by solving a mathematical programming problem assuming that the realization of the disturbances will be equal to the predicted nominal value. However, only the control $\mathbf{u}_t^*$ is actually applied. At time $t+1$, a new problem is formulated over the horizon $[t+1, t+1+h]$ on the basis of the updated information available, namely the state of the system at time $t+1$ as well as updated forecasts of the external drivers (*Castelletti et al.*, 2008b). This feedback can partially compensate the effects of the disturbances, as it is very unlikely that the actual realization of the disturbances is equal to the predicted nominal values, with the system actually not evolving as expected. The availability of good forecast reduces the distance between expected and actual conditions, thus allowing to obtain decisions close to optimality.
However, if the quality of the predictions is low, also the obtained decisions will have poor performance as the realizations of the disturbances are very different from the predicted nominal values and the system is actually not evolving as expected. In such a case, there is the possibility of explicitly taking into account the stochasticity of the disturbances by replacing the NFC approach with an open-loop feedback control (OLFC). The OLFC formulation of the control problem does not rely on predicted nominal values but describes the disturbances according to their probability distribution and computes the objectives through some functions to filter the disturbances (e.g., the expected values). As with the NFC approach, the problem is solved over a finite horizon $[t, t+h]$ and, at the next time instant, the problem is reformulated over the horizon $[t+1, t+h+1]$ according to the current condition of the system and, possibly, with the probability distributions of the disturbances conditioned on exogenous information available at that time. Finally, it is possible to further improve the performance of the OLFC policy by adopting a partial open-loop feedback control (POLFC) formulation (e.g., *Pianosi and Soncini-Sessa*, 2009), which explicitly assumes that in the future the state of the system will be measured and a new problem will be reformulated. The POLFC problem, therefore, computes at each time step the

optimal control for the first step and the optimal policy for the following steps. POLFC problems are functional problems and so they have a greater computational complexity than OLFC and NFC problems. Yet, this increased complexity is usually reflected in a better performance. Indeed, usually the performance of these three approaches are ordered as CL $\leq$ POLFC $\leq$ OLFC $\leq$ NFC (*Bertsekas*, 1976), where CL is the Closed-Loop solution of Problem (2.2).
Adopting an MPC approach, Problem (2.2) can be reformulated as:

$$\mathbf{u}^* = \arg \min_{\mathbf{u}_t,\dots,\mathbf{u}_{t+h-1}} \boldsymbol{\lambda} \cdot \mathbf{J}(\cdot) \tag{2.3a}$$

subject to

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1}) \tag{2.3b}$$

$$\mathbf{u}_t \in \mathcal{U}_t \tag{2.3c}$$

Though being largely adopted in process-engineering problems (see *Scattolini*, 2009, and references therein), MPC is less applied in the water resources management (e.g., *Niewiadomska-Szynkiewicz et al.*, 1996; *Pianosi and Soncini-Sessa*, 2009; *Anand et al.*, 2011; *Galelli et al.*, 2012) due to the difficulties of obtaining accurate forecasts of the several disturbances affecting the natural system. Moreover, on-line approaches are mainly used to overcome the limitation of DP family methods and maintain the traditional centralized perspective adopted in the water management literature, while their application in non-cooperative settings such as the ones considered in this thesis is nearly unexplored.

2.4.2 Direct policy search

Direct policy search (DPS, see *Schmidhuber* (2001); *Rosenstein and Barto* (2001)), also known as parameterization-simulation-optimization in the water resources literature (*Koutsoyiannis and Economou*, 2003), is a simulation-based approach where the operating policy is first parameterized within a given family of functions (e.g., linear or piece-wise linear) and then the parameters optimized with respect to the operating objectives (see also *Oliveira and Loucks*, 1997; *Momtahen and Dariane*, 2007; *Celeste and Billib*, 2009; *Pianosi et al.*, 2011; *Ostadrahimi et al.*, 2012; *Guo et al.*, 2013). Following *Nalbantis and Koutsoyiannis* (1997), DPS can be seen as an optimization-based generalization for multi-objective problems of well known simulation-based, single-purpose heuristic operating rules (for a review, see *Lund and Guzman*, 1999). Prior studies (e.g., *Baglietto et al.*, 2006; *Momtahen and Dariane*, 2009; *Castelletti et al.*, 2012a) have adopted the DPS approach mainly to overcome the computational and dimensional limitations of DP family of methods, without considering the realities of the current systems' operations and have thus rarely been adopted. However, beside the computational advantages, DPS results particularly useful in studying existing system by describing the current system operation in terms of a parametrized operating policy. This latter can be then refined according to an enlarged many-objective perspective. This policy identification and refinement procedure aims to design improved solutions, which better address the tradeoffs between original and potentially new objectives given significant hydroclimatic uncertainties and maintain some features of the current operation (e.g., satisfaction of regulatory constraints or the operators' current preference structure) to increase their practicability in current water system operations.

Policy identification

Although real world decision makers (DMs) generally do not accept the validity of sophisticated decision support tools (e.g., *Yeh*, 1985; *Labadie*, 2004; *Castelletti et al.*, 2008a, and references therein), in making their decisions they necessarily look at the current or expected systems conditions (e.g., current water level, forecasted inflow) when they close the loop between their operating decisions and the system's conditions (*Soncini-Sessa et al.*, 2007a). For example, in the case of water reservoir operators, this is done either implicitly, while they are tracking the reservoir rule curve, or explicitly, when their operation relies on empirical operating rules (e.g., those reviewed by *Lund and Guzman* (1999)). In both cases, their decision mechanism can be formalized as an operating policy p. Modeling the DM behavior means identifying this policy p by assuming that the operating rules $\mathbf{u}_t = m_t(\mathbf{x}_t)$ belong to a given class of functions, namely $\mathbf{u}_t = m(\mathbf{x}_t, \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ is a vector of unknown time-varying parameters. The values of $\boldsymbol{\theta}$ can be determined by looking, when available, at the historical system operation, which in the case of reservoir operation is given by the time series of levels and associated releases. Hence, the historical policy can be identified (parametrized) via regression by estimating the parameters $\boldsymbol{\theta}$ that minimize some distance metrics between historical releases and modeled ones (*Guariso et al.*, 1986; *Corani et al.*, 2009). This explicit policy identification approach can be adopted only when the historical time series are available. A more general procedure, which does not require historical data, is based on the assumption that the DMs are rational proactive agents acting to maximize their benefit, which can be expressed by a specific objective function. Optimizing the rule parameters with respect to this objective function yields a policy that implicitly captures the actual DM behavior. In the water reservoir literature, a number of parameterizations of operating rules have been proposed, such as the New York City rule (*Clark*, 1950), the well known spill-minimizing "space rule" (*Clark*, 1956; *Johnson et al.*, 1991), or the Standard Operating Policy (*Draper and Lund*, 2004). However, many rules in practice are based largely on empirical or experimental successes and they were designed, mostly via simulation, for single-purpose reservoirs (*Lund and Guzman*, 1999). In complex many-objective problems, a priori knowledge can be counterproductive, since it might restrict the search for the optimal policy to a subspace of the decision space that might not include the optimal solution. The adoption of universal approximators such as artificial neural networks or basis functions (e.g. *Barron*, 1993; *Kurková and Sanguineti*, 2001; *Zoppoli et al.*, 2002) partially overcomes this limitation by providing flexibility to the shape of the operating rule. In this work, gaussian radial basis functions (RBFs) are selected to model the operating rule as they are capable of representing functions for a large class of problems (*Tsitsiklis and Van Roy*, 1996; *Menache et al.*, 2005; *Busoniu et al.*, 2011). With RBFs, the k-th release decision in the vector $\mathbf{u}_t$ (with $k = 1, \ldots, n$) is defined as:

$$u^k = \sum_{i=1}^{N} w_i^k \varphi_i \tag{2.4}$$

where N is the number of RBFs and w_i is the weight of the i-th RBF (φ_i). The weights are formulated such that they sum to one (i.e., $\sum_{i=1}^{n} w_i = 1$) and are non-negative (i.e., $w_i \geq 0 \quad \forall i$). The single RBF is defined as follows:

$$\varphi_i(\mathbf{x}) = \exp\left[-\sum_{j=1}^{M} \frac{(x_j - c_{j,i})^2}{b_{j,i}^2} \right] \tag{2.5}$$

where M is the number of input variables ($\mathbf{x}$) and $\mathbf{c}_i, \mathbf{b}_i$ are the M-dimensional center and radius vectors of the i-th RBF, respectively. The centers of the RBF must lie within the bounded input space and the radii must strictly be positive (i.e., using standardize variables, $\mathbf{c}_i \in [-1, 1]$ and $\mathbf{b}_i \in (0, 1]$). The parameter vector $\boldsymbol{\theta}$ is therefore defined as $\boldsymbol{\theta} = [(c_1, \ldots, c_M), (b_1, \ldots, b_M), (w_1, \ldots, w_n)]_1^N$.

Policy refinement

Once the historical operating policy has been identified, the same family of functions can be used to refine (optimize) it in a multi-objective perspective, thus exploring the original operation for different tradeoffs. Technically, the policy parameters ($\boldsymbol{\theta}$) are determined by solving the following multi-objective problem:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathbf{J}(\boldsymbol{\theta}) \tag{2.6a}$$

subject to

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1}) \tag{2.6b}$$

$$\mathbf{u}_t = m(\mathbf{x}_t, \boldsymbol{\theta}) \tag{2.6c}$$

$$\boldsymbol{\theta} \in \Theta \tag{2.6d}$$

DPS methods search for the optimal policy directly in the policy space, with the operating objectives that are optimized by moving the values of the policy parameters (e.g., *Rückstiess et al.*, 2010; *Kormushev and Caldwell*, 2012), as opposed to dynamic programming family methods that evolve in the objective space. In addition to overcome the curse of dimensionality and the curse of modeling limiting DP family methods, DPS offers other advantages as it does not require the system to be a discrete automaton, the objective function to be separable in time and the disturbances uncorrelated in time discretization. However, it does not provide any theoretical guarantee on the optimality of the resulting operating policies, which are strongly dependent on the choice of the class of functions to which they belong and on the ability of the optimization algorithm to deal with non-linear models and objectives functions, complex and highly constrained decision spaces, and many conflicting objectives.

Stochastic hydrology generation

The DPS approach designs Pareto-optimal policies by solving Problem (2.6) where the values of the objectives are obtained via simulation of the system under the corresponding parametrized policy. This simulation-based optimization approach requires in principle to simulate the system over a wide range of system conditions. Consequently, the use of observed records to perform simulation over historical conditions tends to bias the resulting policies. Conversely, the optimization performed over a broad

ensemble of synthetic hydroclimatic variables derives more robust solutions.
A large number of methods for synthetic hydroclimatic data generation has been proposed in literature (e.g., *Box and Jenkins*, 1970; *Lall and Sharma*, 1996; *Yates et al.*, 2003). According to *Rajagopalan et al.* (2010), these methods can be classified as parametric approaches, which assume a standard functional form for the observed data, and nonparametric approaches, which instead define empirical distributions. In this thesis, the nonparametric K-Nearest Neighbor resampling method proposed by *Nowak et al.* (2010) is used. This data-driven method captures the observed statistics, is consistent with the lag correlation structures in the observed data, and ensures summabililty and continuity across the daily time scale.
The procedure is based on the generation of annual data and their disaggregation to daily values. The synthetic annual data Z are generated through an autoregressive model of the first order calibrated over the historical time series. The historical data are also used to compute the proportion matrix P, which contains the proportion of the annual data occurring in each day of the year. Then, K nearest neighbors of one generated annual value Z are identified, with $K = \sqrt{N_y}$ being N_y the number of years in the historical time series. A weight is assigned to each of the K-nearest neighbors as follows:

$$W(i) = \frac{\left(\frac{1}{i}\right)}{\left(\sum_{i=1}^{K} \frac{1}{i}\right)} \tag{2.7}$$

where i is the neighbor index, with $i = 1$ identifying the nearest neighbor. According to the probabilities defined by their weights, one of the K-nearest neighbors is randomly selected. Finally, the proportion vector corresponding to the selected year y is used to disaggregate the generated annual flow Z to obtain daily data $d = Z \cdot p_y$. The procedure is iterated to generate an ensemble of daily streamflows for each Z, and then repeated for multiple synthetic annual data.

Multi-objective evolutionary algorithms

Multi-objective evolutionary algorithms (MOEAs) are iterative search algorithms that evolve a Pareto-approximate set of solutions by mimicking the randomized mating, selection, and mutation operations that occur in nature (*Goldberg*, 1989; *Back et al.*, 2000; *Coello Coello et al.*, 2007). These mechanisms allow MOEAs to deal with challenging multi-objective problems characterized by multi-modality, nonlinearity, and

discreteness (see *Nicklow et al.* (2010) for an extensive review of MOEA applications in water resources).
The self-adaptive Borg MOEA (*Hadka and Reed*, 2013) is used in this thesis. It employs multiple search operators that are adaptively selected during the optimization based on their demonstrated probability of generating quality solutions. The Borg MOEA has been shown to be highly robust across a diverse suite of challenging multi-objective problems, where it met or exceeded the performance of other state-of-the-art MOEAs (*Hadka and Reed*, 2012; *Reed et al.*, 2013). In addition to adaptive operator selection, the Borg MOEA assimilates several other recent advances in the field of MOEAs, including an ε-dominance archiving with internal algorithmic operators to detect search stagnation, and randomized restarts to escape local optima. The flexibility of the Borg MOEA to adapt to challenging, diverse problems makes it particularly useful for addressing DPS problems, where the shape of the operating rules and the parameters values are problem-specific and completely unknown a priori.
According to the DPS approach, the Borg MOEA starts with a population of N individuals, representing N randomly generated parameter vectors $\boldsymbol{\theta}$. The algorithm evaluates the fitness of each individual by simulating the system according to the operating policy defined by the corresponding value of $\boldsymbol{\theta}$ and evaluating the objective vector $\mathbf{J}(\boldsymbol{\theta})$. Then, a new population is generated by selection, crossover and mutation with respect to the best individuals (i.e., the ones obtaining the highest values of fitness) according to the Pareto dominance criterion. This process is then repeated for a given number of iterations until a good approximation of the Pareto front is obtained.

2.5 Direct policy conditioning to support mechanism design

The aim of mechanism design is the identification of mechanisms to condition the decisions of self-interested agents such that the overall system behavior will be acceptable (*Sycara*, 1998; *Maskin*, 2008). Most of the mechanism design strategies is generally based on regulatory constraints or economic incentives (*Pannell*, 2008). Yet, these strategies mainly rely on the empirical experiences of the institutions in charge to promote negotiated solutions, while mathematical and technological tools to support the definition of these mechanisms are nearly undeveloped. A typical example related to water reservoir operation is the definition of minimum environmental flow (MEF) constraints to guarantee adequate

conditions downstream of artificial reservoirs, possibly with seasonal or monthly varying values. Most major reservoirs have had their rule curves defined in prior decades (*U.S. Army Corps of Engineers*, 1977), considering one of four primary uses: generation of power, flood risk reduction, irrigation or drinking water supply. In more recent years, bilateral negotiation processes have been adopted to include new operational targets (e.g., environmental protection, recreation and transportation). Traditionally, the four primary uses take precedent over all other concerns and these secondary objectives are optimized assuming the four primary uses form constraints. However, the resulting reservoir operations are often inefficient as they fail to explore the full set of tradeoffs between evolving multi-sector objectives and preferences in the water system.
This thesis contributes a novel method, called direct policy conditioning (DPC), to define effective mechanism design strategies for complex water management problems by directly constraining the operating policy space. This mechanism allows the design of dynamic constraints that, like the operating policy, are able to exploit the feedback provided by the system conditions. Conversely, the MEF constraint does not consider the system conditions (e.g., the current reservoir storage) and provide a static (i.e., fixed) constraint on the reservoir release. The DPC method relies on the direct policy search (DPS) approach described in Section 2.4.2. In multi-objective problems, the DPS approach searches the set of Pareto-optimal solutions of the problem, designed explicitly considering a suite of conflicting operating objectives. This set, therefore, represents the best achievable performance at the system-level. The DPC method, instead, maintains the original preference structure and designs the control policy considering only the primary objectives. The Pareto-optimal set is used to formulate a set of constraints for the DPC policy design problem in order to account for the secondary objectives. These constraints are defined in the decision space of the policy parameters and, therefore, directly influence the operating policy space (Figure 2.8). Conversely, traditional MEF constraints are defined with respect to the release decisions and does not directly act on the real decision space.
The DPC method is based on a three-step procedure: *i*) multi-objective problem formulation (including all the operating objectives, both primary and secondary) and solution via multi-objective DPS to design a set of Pareto-optimal policies representing the best ideal achievable performance and the associated policy parameterizations (see Section 2.4.2); *ii*) identification via input variable selection techniques of the subset of optimal control policy parameters that are more related to the secondary objectives, actually excluded in the design of the reservoir regulation; *iii*) design of a control policy considering only the primary objectives, subject

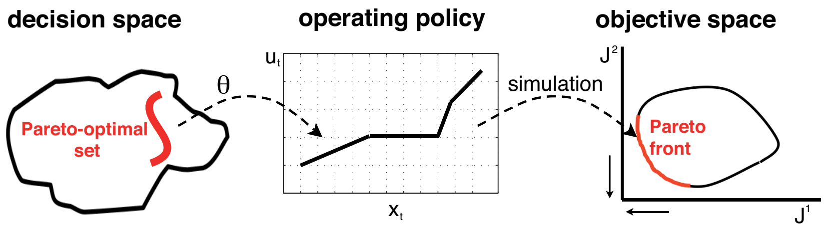

Figure 2.8: Illustration of the difference between decision space, operating policy, release decision, and objective space. The proposed direct policy conditioning acts in the decision space, while traditional MEF constraints are defined with respect to the release decision u_t.

to a set of dynamic constraints accounting for the secondary ones.

2.5.1 Iterative Input Selection algorithm

Generally, IVS problems arise every time a variable of interest has to be modeled as a function of a subset of potential explanatory variables, or predictors, but there is uncertainty about which subset to use among a number, usually large, of candidate sets available (*George*, 2000). The goal of IVS is threefold (*Guyon and Elisseeff*, 2003): *i*) improving model performance by avoiding the interference of redundant information and more effectively exploiting the data available for model calibration; *ii*) providing faster and more cost effective models; *iii*) assisting in the interpretation of the underlying process by enabling a more parsimonious and compact representation of the observational data. In this thesis, IVS techniques are used to characterize the complex relationships and feedbacks between operating objectives, management decisions, and the system dynamic evolution.

Many approaches have been proposed in literature, especially to support emulation modeling (for a review, see *Castelletti et al.*, 2012b). According to *Maier et al.* (2010), two families of methods have been developed depending on whether they implicitly or explicitly assume an underlying model in the selection process. Model-free methods (e.g., *Peng et al.*, 2005; *Bowden et al.*, 2005; *Hejazi and Cai*, 2009) do not explicitly use any model and rank the candidate input variables with respect to a statistical measure representing how strong the relationship is between each input and the output under investigation. The most significant variables can then be selected according to some pre-defined criterion. Model-

based methods (e.g., *Das*, 2001; *Guyon and Elisseeff*, 2003) are instead explicitly based on models, whose performance is evaluated for each variable that is added or removed from the selected set.
The model-free, forward-selection, iterative input selection (IIS) algorithm (*Galelli and Castelletti*, 2013) is used in this thesis. Given the output variables v^o and the set of candidate inputs $\mathbf{v}^i$, the IIS algorithm first ranks these latter with respect to a statistical measure of significance and adds only the best performing input v^* to the set of selected variables $\mathcal{V}$. This operation aims to avoid the inclusion of redundant variables, as one an input is selected, all the inputs highly correlated with it may become useless in the next iterations. Then, the algorithm estimates a model of v^o with input $\mathcal{V}$, such as $v^0 = \hat{m}(\mathcal{V})$, and estimates the model performance with a suitable distance metric D (e.g., the coefficient of determination) as well as the model residuals, which become the new output at the next iteration. The algorithm stops when the best variable returned by the rank is already in the set $\mathcal{V}$, or when over-fitting conditions are reached. Among the many alternative model classes, IIS relies on extremely randomize trees (Extra-Trees), a tree-based method proposed by *Geurts et al.* (2006) that was empirically demonstrated to outperform other models in terms of modeling flexibility, computational efficiency, and scalability with respect to the input dimensionality. Moreover, the Extra-Trees structures can be exploited to infer the relative importance of the input variables and, therefore, they can be used as an input ranking procedure (*Wehenkel*, 1998; *Fonteneau et al.*, 2008; *Castelletti et al.*, 2011b). A tabular version of the IIS algorithm is given in Algorithm 1.

2.5.2 Direct policy conditioning formulation

The aim of DPC is to effectively condition the operating policies designed only considering the primary objectives (e.g., hydropower production, flood prevention, water supply) in order to take into account the secondary objectives (e.g., environment, recreation). This conditioning is defined by exploiting the reference provided by the Pareto-optimal solutions of Problem, namely the set of optimal operating policy parameters $\boldsymbol{\theta}^*$. To guarantee the flexibility of the resulting policy, the conditioning is applied to a subvector of parameters $\tilde{\boldsymbol{\theta}} \subseteq \boldsymbol{\theta}$, which is identified through the iterative input selection (IIS) algorithm described in the previous section. Given a dataset of candidate inputs (i.e., policy parameters) and the corresponding output variables (i.e., the secondary objectives), the IIS algorithm identifies the set of input variables that are more related to the selected outputs. The dataset required for IIS experiments

Algorithm 1 Iterative Input Selection

Inputs: a dataset $\mathcal{F}$ of candidate inputs $\mathbf{v}^i$ and the output variable to explain v^o.
Initialization:
Set $\mathcal{V} \leftarrow 0, \hat{v}^o \leftarrow v^o, D_{old} \leftarrow 0$
Iterations: repeat until stopping conditions are met
- select the most relevant input $v^* \in \mathbf{v}^i$ to explain $\hat{v}^o$
- **if** $v^* \in \mathcal{V}$, **return** $\mathcal{V}$ **end if**
- $\mathcal{V} \leftarrow \mathcal{V} \cup v^*$
- $\hat{m}(\cdot) \leftarrow$ Extra-Trees$(\mathcal{F}, v^o, \mathcal{V})$
- $\hat{v}^o \leftarrow v^o - \hat{m}(\cdot)$
- $\Delta D \leftarrow D(v^o, \hat{m}(\cdot)) - D_{old}$
- $D_{old} \leftarrow D(v^o, \hat{m}(\cdot))$
- **until** $\Delta D < \varepsilon_D$

return $\mathcal{V}$

can be generated via random sampling of the policy parameters' space or, alternatively, the Pareto-optimal set can be directly used. In the first case (i.e., random sampling), a set of policy parameters vectors is randomly sampled in Θ. Each vector defines a randomly generated operating policy. The system is hence simulated following the randomly generated policy to compute the operating objectives. The combination of the randomly sampled policy parameters and their performance in terms of the secondary objectives forms the IIS dataset. In the second case (i.e., Pareto-optimal set), the dataset used is composed by the set of Pareto-optimal policy parameterizations, obtained solving the multi-objective problem with respect to all the operating objectives, and their corresponding performance in the secondary objectives.
Given the Pareto-optimal set $\boldsymbol{\theta}^*$ and the extracted subvector $\tilde{\boldsymbol{\theta}}$, the DPC problem can be formulated as:

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathbf{J}^p(\boldsymbol{\theta}) \tag{2.8a}$$

s.t.

$$\boldsymbol{\theta} \in \Theta' \subseteq \Theta \tag{2.8b}$$

$$\Theta' = \{\boldsymbol{\theta} \in \Theta : \theta_i = \theta_i^* \pm \gamma, \forall \theta_i \in \tilde{\boldsymbol{\theta}}\} \tag{2.8c}$$

where the objective function vector $\mathbf{J}^p$ includes only the primary objectives and the subvector $\tilde{\boldsymbol{\theta}}$ is constrained to take values in a neighborhood of size γ of the optimal values $\boldsymbol{\theta}^*$. These constraints impose the optimal values of the policy parameters according to the signature identified by the IIS algorithm. The aim of this conditioning is to partially consider the secondary objectives in the design of the control policy targeted only to the primary objectives.

3 Agent-based mechanism design test problem

The purpose of this chapter is to apply the agent-based decision analytic framework (Section 1.3.2) to evaluate different levels of coordination in a simple, non-dynamic water allocation problem. The Y-shaped hypothetical water system described in *Yang et al.* (2009) is used to illustrate the methodology. In the considered case study, six agents represent six conflicting water users sharing the same river. Although the illusory simplicity of this planning problem, the considered case study actually includes multiple sources of complexity characterizing many real world problems, such as the upstream-downstream asymmetry, the presence of agents deciding in parallel and in series, the difference between primary objectives associated to real decisions (e.g., water supply demands driving the amount of water to divert from the river or hydropower production defining the releases from the dam) and secondary environmental concerns. In this chapter, multi-agent system (MAS) are not adopted only as a modeling tool, but MAS distributed optimization methods are explicitly employed. The two situations corresponding to a centralized solution where all the agents cooperatively maximize the total benefit of the system (i.e., the sum of the benefit of each agent) and, on the other extreme, an uncoordinated solution where the agents consider their local objective functions only are first considered. Then, a constraint-based mechanism design strategy based is proposed, where a watershed authority is in charge of imposing soft (normative) constraints on the originally self-interested agents' decisions to drive the solutions in the space in-between these two extremes. The resulting problem is a distributed, constrained planning problem which can be effectively managed through MAS methods (Section 2.3) based on distributed constraint satisfaction problems (DCSP, see *Yokoo and Hirayama* (2000)) and distributed constraint optimization problems (DCOP, see *Modi et al.* (2005)). Adopt algorithm is used to solve DCSP and DCOP problems. Moreover, alternative mechanism design strategies based on economic incentives are discussed. Finally, the agent-based framework allows the analysis of the consequences of individualistic behaviors by the agents on the system benefit, in order to identify potentially critical situations in case of wa-

ter scarcity.
This chapter is adapted from a journal paper under review for publication in the Journal of Water Resources Planning and Management (*Giuliani et al.*, 2013a).

3.1 Problem formulation

The proposed mechanism design strategy is tested on a numerical case study firstly introduced by *Yang et al.* (2009) and then further analyzed by *Giuliani et al.* (2012a). The system is composed by a Y-shaped river with one mainstream and one main tributary, see Figure 3.1. The mainstream provides water to a city for municipal and industrial uses and to an irrigation district downstream. Moreover, the river is dammed for creating an artificial reservoir (just downstream with respect to the city water supply diversion) in order to produce hydropower energy. A second agricultural district diverts water for irrigation purposes from the tributary. Finally, two river stretches just downstream with respect to the irrigation diversions are particularly interesting from an ecological point of view as they are identified as primary fish habitats. The agent-based model of the system therefore comprises six agents representing the different water-related interests:

A_1: municipal water supply to the city;

A_2: hydropower production;

A_3: irrigation water supply to the agricultural district on the tributary;

A_4: irrigation water supply to the agricultural district on the lower mainstream;

A_5: ecological preservation in the tributary;

A_6: ecological preservation in the mainstream.

A quadratic concave objective function is associated to each agent, which preserves the nonlinear characteristics of a real objectives: $J^i = a_i u_i^2 + b_i u_i + c_i$ (the values of the parameters are reported in Table 3.1). For the sake of readability, in this chapter the decision variables of the planning problem are denoted as $\mathbf{u}$ instead of $\mathbf{u}^p$ as in Problem (1.1). The quadratic formulation of the objective functions allows to obtain negative benefits for values of the decision variables very different from the optimal ones. These negative benefits represent potential costs the agents may have to pay in extreme situations. The six agents have different

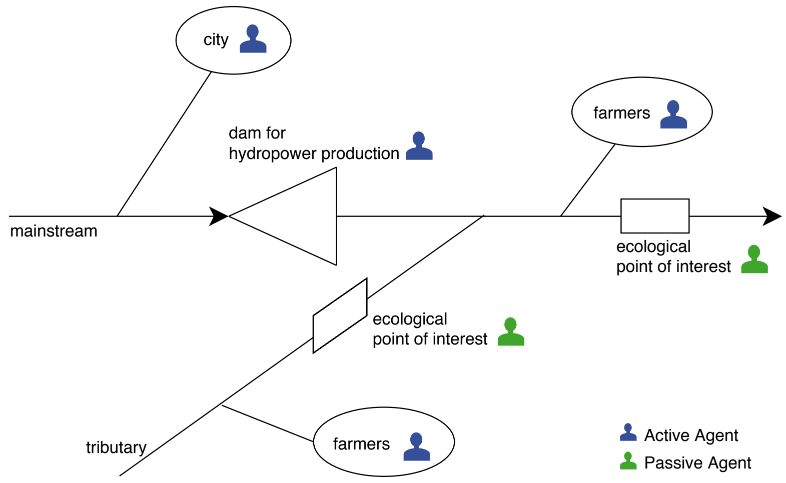

Figure 3.1: Schematic map of the system.

nature: a first group (i.e., A_1-A_2-A_3-A_4) includes *active agents* (shown in blue in Figure 3.1), who really make decisions about the amount of water to divert from the river or to be released from the dam in order to explicitly maximize the corresponding objective function J^i (with $i = 1, \dots, 4$); agents A_5-A_6 are instead defined as *passive agents* (shown in green in Figure 3.1), who do not make decisions but represent the ecological interests through the functions J^5 and J^6, which are explicitly optimized only in the centralized case.

Table 3.1: Values of parameters defining agents objective functions (*Yang et al.*, 2009). Coefficients a_i, b_i and c_i are dimensionless.

Parameter	Value	Parameter	Value	Parameter	Value
a_1	-0.20	b_1	6	c_1	-5
a_2	-0.06	b_2	2.5	c_2	0
a_3	-0.13	b_3	6	c_3	-6
a_4	-0.15	b_4	7.6	c_4	-15
a_5	-0.29	b_5	6.28	c_5	-3
a_6	-0.056	b_6	3.74	c_6	-23

Assuming for simplicity a non-dynamic situation (all the variables, both flows and reservoir storage, are expressed as volumes [L^3]), the watershed optimization problem, subject to hard (physical) constraints, can be formulated as:

$$\max_{u_1} J^1(u_1) \quad s.t. \quad u_1 \leq Q_1 \tag{3.1a}$$

$$\max_{u_2} J^2(u_2) \quad s.t. \quad u_2 \leq S + Q_1 - u_1 \tag{3.1b}$$

$$\max_{u_3} J^3(u_3) \quad s.t. \quad u_3 \leq Q_2 \tag{3.1c}$$

$$\max_{u_4} J^4(u_4) \quad s.t. \quad u_4 \leq u_2 + Q_2 - u_3 \tag{3.1d}$$

$$J^5(u_5) \quad s.t. \quad u_5 = Q_2 - u_3 \tag{3.1e}$$

$$J^6(u_6) \quad s.t. \quad u_6 = u_2 + u_5 \tag{3.1f}$$

where Q_1 is the mainstream inflow, Q_2 the tributary inflow, and S the reservoir storage. The constraints expressed above are all physical constraints. Three hydrological scenarios are defined representing different water availability situations, namely high, medium, and low flow conditions, see Table 3.2. In the first case (i.e., high flow scenario), the water available allows each active agent to achieve its optimal solution; in the medium flow scenario, instead, the water available in the system is insufficient to satisfy all the agents demands, thus producing upstream-downstream water sharing interactions, which are further tightened up in the low flow scenario.

Table 3.2: High, medium, and low flow scenarios.

Hydrological variable	High flow	Medium flow	Low flow
Q_1 [L^3]	80	40	15
Q_2 [L^3]	40	20	8
S [L^3]	10	8	3

The fully cooperative centralized solution of Problem (3.1) is compared to three different distributed alternatives:

- an uncoordinated solution where each active agent acts independently considering its objective only. The upstream agents are in a favorable condition as they can decide what is the best for themselves while the downstream agents can use only the water remaining. The objective of the passive agents are not considered.

- A DCSP solution where the active agents try to maximize their objective functions, but the assignment of values to the decision variables has to be feasible (i.e., all the constraints imposed on the problem have to be satisfied). The interests of the passive agents are partially considered by means of the soft (normative) constraints that a watershed authority might impose on the agents decisions.

- A DCOP solution where the active agents decisions, which aim to optimize their objective functions, have to satisfy the hard (physical) constraints and are influenced by soft (normative) constraints aiming to protect the interests of the passive agents. The soft constraints may be violated. However, it is guaranteed that the solution found minimizes the sum of soft constraints cost violations.

3.2 Results

The proposed agent-based decision analytic framework requires first to compute the two extreme solutions of Problem (3.1), namely the centralized and the uncoordinated one, for each hydrological scenario. It is assumed that the goal of a centralized management strategy is to maximize the total system benefit, meaning the sum of the benefit of the six agents explicitly considering also the ecological objectives and assuming the same importance for each agent. This ideal solution, which represents the most efficient management strategy, is then compared with the other extreme situation where the agents act independently by optimizing their local objective functions only, without considering the potentially negative externalities that their decisions produce on the other agents' objectives, in particular with respect to the ecological passive agents. In this scenario, Problem (3.1) is solved as a sequence of optimization problems from the upstream agents to the downstream ones. The comparison between the system benefit for these two extreme solutions is represented in Figure 3.2. Not surprisingly, the centralized solution (red bars) produces higher benefits than the uncoordinated one (blue bars) in all the considered scenarios, and the gap between the two solutions increases when water availability decreases and the conflicts among different water users become stronger and stronger.
Given these two extreme reference solutions and according to the proposed coordination mechanism, the watershed authority imposes the following set of normative constraints in order to protect the environment, especially in the medium and low flow scenarios in which upstream agents

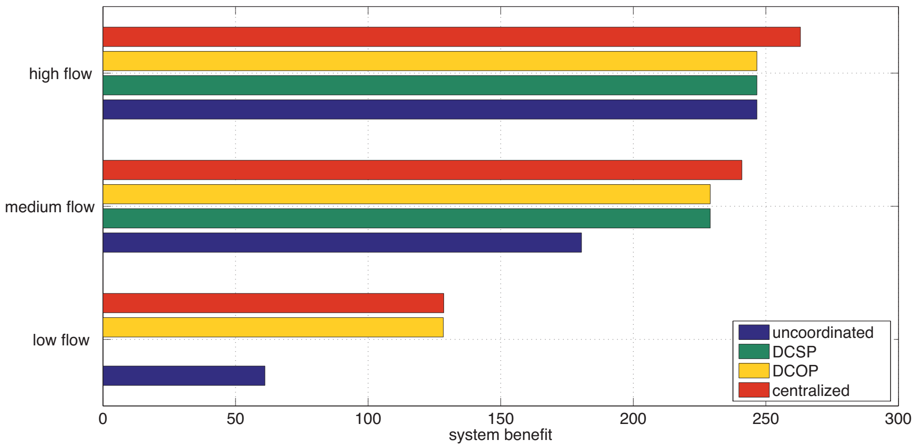

Figure 3.2: Comparison of system benefits for centralized, uncoordinated, and DCSP/DCOP solutions.

overuse the available water producing externalities over the downstream agents suffering water shortage:

$$\begin{aligned}
&\alpha_1 - u_1 \leq 0\\
&\alpha_2 - Q_1 + u_1 \leq 0\\
&\alpha_3 - u_3 \leq 0\\
&\alpha_4 - u_4 \leq 0\\
&\alpha_5 - Q_2 + u_3 \leq 0\\
&\alpha_6 - u_2 - (Q_2 - u_3) + u_4 \leq 0
\end{aligned} \tag{3.2}$$

where $\alpha_1 = 12$ is the minimum water demand of the city, $\alpha_2 = 10$ is the minimum flow requirement for hydropower production, $\alpha_3 = 8$ and $\alpha_4 = 15$ are the minimum water demands of the farmers on the tributary and on the lower mainstream, respectively, $\alpha_5 = 6$ and $\alpha_6 = 10$ are the flow requirements for the protection of the fish habitats on the tributary and on the lower mainstream, respectively. The hard (physical) constraints in Problem (3.1) are obviously non-violable, while the normative constraints defined in eqs. (3.2) may be violated by either some self-interested agents or the nature in case of very low flow conditions (e.g., in the low flow scenario, the tributary flow is equal to 8, the minimum demand of the farmers is 8 and the environmental flow requirement is 6: in such a case, if agent A_3 diverts only 2 in order not to violate the environmental constraint, there is still a violation of the

farmers minimum demand as $u_3 \geq 8$, being $u_3 = 2$).
It is important to point out that there is a significant difference between the agent-based solutions and the centralized one: this latter, indeed, assumes full cooperation and coordination between the agents, who are no longer decision entities but actually become actuators of the decisions made by a centralized authority. By imposing normative constraints, the distributed nature of the decision process is instead preserved in the agent-based solutions. The development of mechanism design strategies aims to obtain an approximation of the ideal centralized solution, which remains different from the institutional (agents vs. actuators) as well as from the mathematical formulation (four single-objective problems vs. a unique six-objective optimization problem) points of view.
The new distributed and constrained optimization problem formulation can be effectively managed through DCSP and DCOP algorithms described in Section 2.3 in order to obtain distributed solutions in-between the two extreme situations so far described. In both DCSP and DCOP formulations, the distributed optimization problem is subject to a new set of constraints, comprising the hard (physical) ones already defined in Problem (3.1) along with the soft (normative) constraints introduced in eqs. (3.2). Both DCSP and DCOP problems are solved using the implementation of Adopt (see Section 2.3.3) provided by FRODO, an open-source framework for DCOP (*Léauté et al.*, 2009). The problem is solved as a sequence of optimization problems from upstream to downstream, where each active agent is considering its local objective function only or, possibly, according to a depth-first search tree, in which each agent has a single parent and multiple children. In the DCSP formulation, where the constraints are Boolean, all the constraints have to be satisfied and, therefore, there is no difference between physical and normative constraints. On the other hand, the DCOP formulation deals with physical and normative constraints in different ways, and violations of these latter are allowed. In particular, for DCOP, the cost of violation of the soft constraint $u_1 - \alpha_1 \leq 0$ is 0, when actually $u_1 - \alpha_1 \leq 0$, and $u_1 - \alpha_1$, otherwise. In a similar way, the costs for the other soft constraints can be calculated. In the computation of the weighted sum of the constraints violation costs, which is minimized in the DCOP solutions, equal weights are assumed for each soft constraint (namely, $w_j = 1 \quad \forall j$).
A graphical comparison of the system benefit for the centralized and the uncoordinated solutions with respect to the results obtained solving the new distributed constrained problem adopting the proposed DCSP- and DCOP-based approaches is represented in Figure 3.2 by green and yellow bars, respectively. In the high flow scenario, the uncoordinated, DCSP, and DCOP solutions are all equivalent because in this scenario, charac-

terized by high water availability, each active agent is actually able to maximize its objective function; yet the centralized solution outperforms all the other solutions because it explicitly optimizes the ecological objective functions (J^5 and J^6) which are included in the computation of the system benefit but not optimized in the other cases. In the medium flow scenario, the DCSP and DCOP solutions have a system benefit that is higher than the uncoordinated solution and they are equivalent because it is possible to find out a solution that does not violate any normative constraint. Finally, in the low flow scenario, the DCSP solution does not exist because it is impossible to satisfy all the normative constraints due to very low water availability; the DCOP solution, instead, largely outperforms the uncoordinated one and it is almost equivalent to the centralized one.
More details are provided in Figure 3.3, where the benefit of each agent is represented separately. The centralized solution, which is globally better than all the others in the high flow scenario, is able to effectively deal with the upstream-downstream relationships between the different agents: the benefits of A_2, A_3 and A_4 are slightly lower than in other solutions in order to generate more significant benefits for the ecological agents A_5 and A_6. The medium flow scenario provides the most interesting results, as it emphasizes the improvement in the ecological agents' benefits in the DCSP and DCOP solutions (which are equivalent in this scenario) with respect to the uncoordinated solution. The imposition of the normative constraints defining minimum water requirements for each agent, indeed, allows to partially take into account also the interests of the passive agents, thus leading to a higher system benefit. Finally, in the low flow scenario, where even the centralized solution is not able to guarantee a positive benefit to the ecological agents, the DCOP solution (the DCSP one does not exist) is able to balance at least the benefits of the active agents avoiding the upstream overuse of water (e.g., see the differences between the couples of upstream-downstream agents (A_1, A_2) and $(A_3, A_4\ )$ in the uncoordinated solution and in the DCOP one).
Given the results in Figure 3.3 and assuming that the benefit is measured in monetary terms, the possibility of developing mechanism design strategies based on economic incentives can be also assessed. Let consider again the medium flow scenario: by comparing the agents decisions in the uncoordinated ($\mathbf{u} = [15, 21, 20, 21]$) and in the centralized ($\mathbf{u} = [15, 33, 12, 21]$) scenarios, it is evident that the different benefits are due to the decisions of A_2 and A_3, which produce negative externalities on the benefits of A_5 and A_6. Hence, there exists the chance to push agents A_2 and A_3 to change their decisions in order to favor A_5 and A_6 by compensating their losses. Let first focus on the tributary sub-

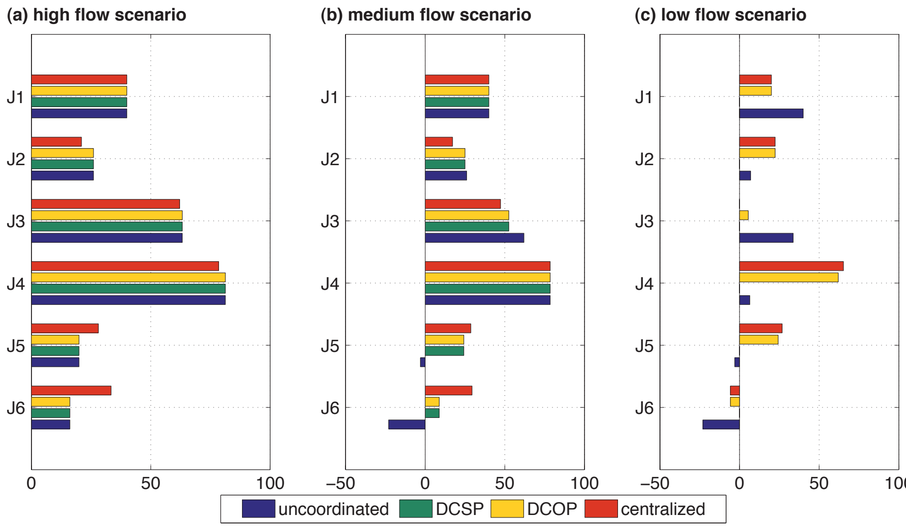

Figure 3.3: Agents benefits in the three considered scenarios.

problem, where A_3 is diverting $u_3 = 20$ in the uncoordinated scenario, corresponding to a local benefit $J^3 = 62.00$ and an environmental benefit $J^5 = -3$, while in the centralized scenario $u_3 = 12$, corresponding to $J^3 = 47.28$ and $J^5 = 28.68$. In this case, an incentives-based solution might require A_3 to decrease its decision in order to increase the benefit for A_5 and, then, compensate this decision, and the corresponding loss, by mean of a monetary payment. A potential solution might be the following: as in the DCSP and DCOP solutions A_3 chooses $u_3 = 14$, with a local benefit $J^3 = 52.52$, and is compensated by an economic incentive equal to its loss (i.e., 62.00-52.52=9.48) in order not to decrease its final benefit. Then, if the incentive was paid by the watershed authority, the benefit for A_5 would be equal to 24.24. Another possibility is that agent A_5 has to directly compensate A_3 without any external action. In such a case, the final benefit for A_5 would be 24.24-9.48=14.76, which is still higher than the benefit obtained in the uncoordinated scenario. In the same way, agent A_2 may be pushed to increase its release decision to $u_2 = 25$ by compensating its losses with respect to benefit achievable in the uncoordinated scenario (i.e., 26.04-25=1.04). Again, if the compensation was paid by the watershed authority, the benefit for agent A_6 would be 8.80, while if A_6 had to compensate A_2, its final benefit would be 8.80-1.04=7.76. Yet, even if this mechanism may appear easy to be implemented, two major concerns have to be pointed out: first, these compensation measures require to clearly define who has to compensate the upstream agents that change their decisions. This problem is very complex and goes beyond the scope of this application. Second, the quantification of a "sufficient compensation" is another complex issue and, sometimes, the compensation of the losses might be impracticable (see the conflict between hydropower revenue and irrigation supply described in *Anghileri et al.* (2012)).

The results presented in Figure 3.2 consider the imposition of a single set of normative constraints by the watershed authority. The solutions obtained with the DCOP approach improve the uncoordinated solutions increasing the system benefit. However, some agents might consider the imposed constraints as a too restrictive decision by the watershed authority and, consequently, they might decide not to consider these constraints in order to increase their local benefits. The effects of these individualistic behaviors then involve the entire system and, usually, tend to decrease the benefits of the other agents. In particular, the consequences of the individualistic behaviors in the low flow scenario is analyzed because it is assumed that this kind of strategies is more likely to be adopted in water shortage conditions.

According to several environmental MAS applications (e.g., *Charness and Rabin*, 2002; *Jager and Janssen*, 2003; *Yang et al.*, 2009; *Poteete*, 2010), the consequences at a system level produced by individualistic behaviors of the agents can be assessed by modifying Problem (3.1) introducing a parameter β_i which multiplies the objective function of the i-th agent ($\beta_i = [0, 10]$). These coefficients represent the "selfishness" of the i-th agent, meaning his preference for his local benefit against the total benefit of the system. According to this formulation, the original problem is defined by setting $\beta_i = 1 \quad \forall i$. On the other hand, it is possible to represent an individualistic behavior of the i-th agent by increasing the value of β_i. Results are represented in Figure 3.4 where the benefit of the individualistic agent is reported on the x axis, while the benefit of all the remaining agents is on the y axis. The obvious optimum would be a solution that maximizes both. However, it is evident looking at Figure 3.4 that there exists a trade-off between the maximization of the local objective functions through individualistic behaviors and the maximization of the system benefit: the solutions obtained by moving β_i represent a set of Pareto efficient alternatives. The knowledge of this set is particularly relevant as it overcomes the difficulties limiting the a priori calibration of β_i. This process is indeed not straightforward and requires additional, subjective preference information, further biased by the existence of non-unique preference representation. Moreover, looking at the trade-offs existing between the Pareto efficient alternatives, the watershed authority might estimate the marginal costs of individ-

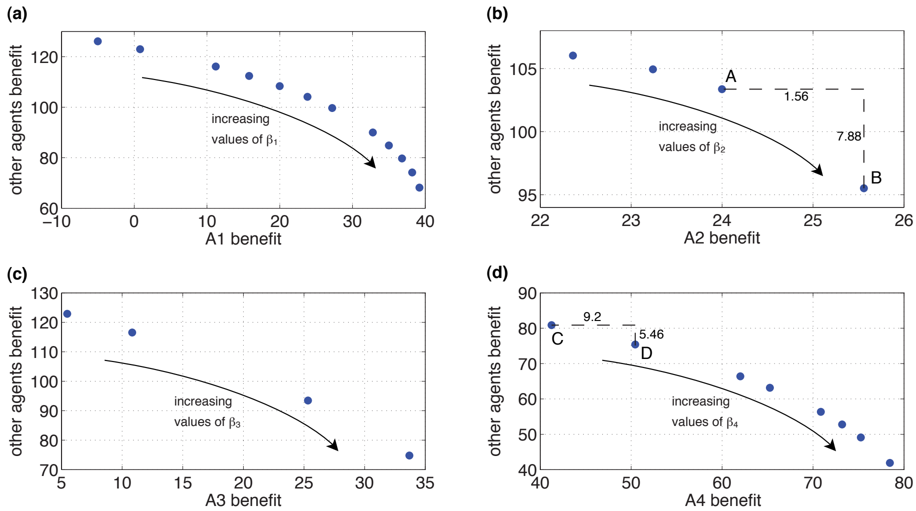

Figure 3.4: Individualistic behaviors of the active agents for different values of β_i.

ualistic behaviors for the system benefit in order to identify critical as well as tolerable cases. An example of critical behavior is represented by solutions A and B in Figure 3.4b, where the individualistic behavior of agent A_2 produces a limited increase of the local benefit (around 1.5) corresponding to a higher decrease of the other agents benefits (almost 8). On the other side, an individualistic behavior that might be tolerated is represented in Figure 3.4d by solutions C and D, where the increase in the local benefit (almost 10) is higher than the negative effect on the other agents (around 5).

In order to prevent the individualistic behaviors just analyzed, the watershed authority might try to modify the set of normative constraints imposed to the agents in order to better explore the space between the two extreme centralized and uncoordinated solutions, looking for another compromise between system efficiency and practicability. The system benefits obtained solving Problem (3.1) in the medium-flow scenario with the different sets of normative constraints reported in Table 3.3 are represented in Figure 3.5: it can be observed that more strict/weak constraints yield to solutions that are closer to the centralized/uncoordinated one. In particular, the original set of constraints α^1 allows to obtain a solution almost equivalent to the centralized one. Yet, the rigidity of these constraints might induce some individualistic behaviors. More flexible constraints slightly decrease the system benefit and the solutions move

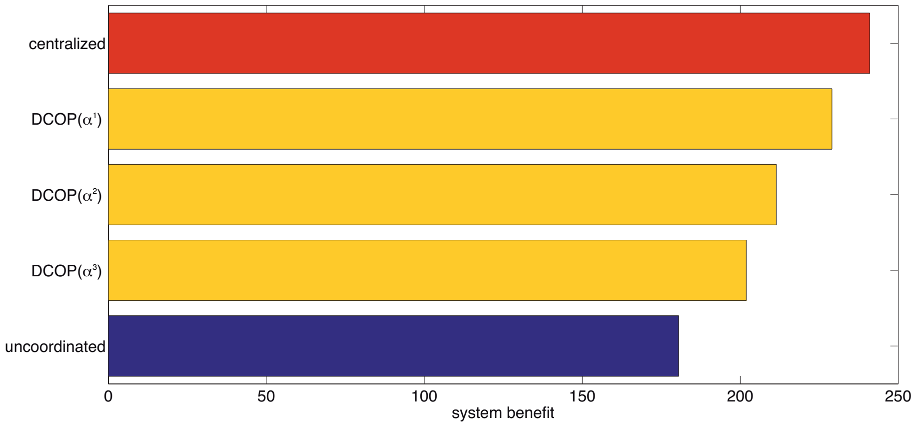

Figure 3.5: Effects of different set of normative constraints on system benefit.

towards the uncoordinated one, which is the most preferable for the active agents.

Table 3.3: Different sets of normative constraints α_i (expressed in [L^3]).

Parameter	Set 1 (original)	Set 2	Set 3
α_1	12	6	3
α_2	10	5	3
α_3	8	4	2
α_4	15	7	4
α_5	6	3	2
α_6	10	5	3

A conflict therefore arises in the mechanism design process because the definition of strict constraints, allowing to approach the centralized solution with good performance for the passive ecological agents and with respect to the system-wide efficiency, produces a decrease in the benefits of the active agents who can decide to adopt individualistic behaviors. Conversely, weak constraints which should be accepted by the active agents yield to practicable solutions with lower benefits for the passive ecological agents and, consequently, lower system-wide efficiency. To analyze such a conflict, two indices measuring system-wide efficiency and practicability are necessary. Efficiency is measured in terms of the resulting total system benefit (i.e., the sum of the benefits of the six agents), while practicability is defined as the percentage of available water which is not constrained by the normative constraints (e.g., in the medium flow scenario and soft constraints α^1, the total available water is $Q_1 + Q_2 + S$ = 68, the constrained water is equal to $\alpha_1 + \alpha_2 + \alpha_3 + \alpha_4 + \alpha_5 + \alpha_6$ = 61, hence practicability is (68-61)/68=7/68=0.1029). In the centralized solution, all the decisions are made by the centralized authority and imposed to the agents, meaning that all the water is constrained and the practicability is equal to 0. In the uncoordinated management, there are no soft constraints and, consequently, practicability is 100%. Looking at Figure 3.6, representing the solutions obtained by varying the set of normative constraints in the space of efficiency and practicability (measured by the system benefit and the percentage of not-constrained water, respectively), it is evident that it is not possible to simultaneously maximize both efficiency and practicability. However, the analysis of this trade-off curve coupled with the assessment of the individualistic behaviors effects can support the watershed authority in designing an effective coordination mechanism. Finally, note that the two extremes of the trade-off curve in Figure 3.6 correspond to the centralized and uncoordinated solutions (as in Figure 1.2) which are probably the less interesting strategies because, usually, the best compromise solution is set in the middle of the Pareto front.

3.3 Discussion and final remarks

In this chapter, a simple application of the agent-based decision analytic framework is used to show the key features characterizing decision-making problems that involve multiple and self-interested decision makers. In particular, the framework is used to develop constraint-based mechanism design strategies to drive the inefficient uncoordinated practices toward solutions that are balanced with respect to efficiency and practicability. The approach is tested on a hypothetical non-dynamic problem, characterized by the presence of several human and ecological agents.

Results show that it is possible to identify distributed solutions by applying constraint-based mechanism design strategies adopting the DCSP- and DCOP-based approaches. These coordinated solutions are more efficient than the uncoordinated ones and more realistic and politically practicable in real decision-making processes than the centralized management. Yet, the ideal, fully cooperative centralized solution remains the most efficient alternative (i.e., it produces the highest system benefit),

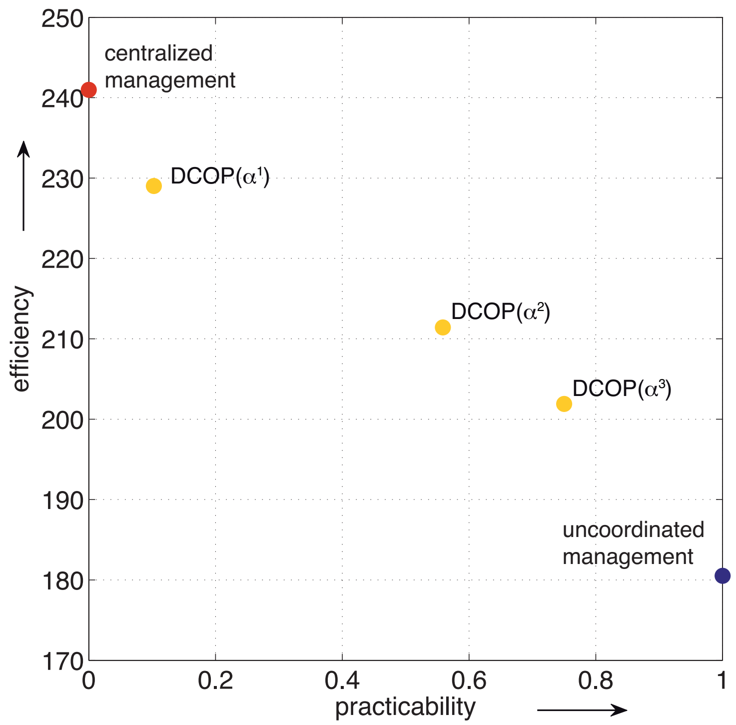

Figure 3.6: Representation of the conflict between system-wide efficiency and practicability, measured by the system benefit and the percentage of available water which is not constrained by the normative constraints, respectively.

but the improvement of the DCSP/DCOP solutions with respect to the uncoordinated management is substantial. The proposed framework allows also the identification of potential incentives-based solutions, which have to be further analyzed in order to establish if an economic compensation is feasible and who has to compensate the upstream agents.
Moreover, the considered case study allows the analysis of individualistic behaviors by the agents. Especially in situations of water scarcity, it is possible that some agents might consider not to comply with the normative constraints imposed by the watershed authority, thus producing negative externalities on the benefits of the other agents. The analysis of the trade-offs between increasing the local benefit and the negative effects for the remaining agents allows to identify critical situations which have to be carefully considered in designing the normative constraints.
Finally, different sets of normative constraints are considered to estimate

the trade-off curve between efficiency (i.e., the total system benefit) and practicability (i.e., the percentage of available water which is not constrained by the normative constraints) and the corresponding solutions confirm that the proposed DCSP- and DCOP-based approaches successfully identify coordination mechanisms which produce solutions in the space in-between the two extreme situations of centralized and uncoordinated management.
In conclusion, it is worth analyzing the scalability of the proposed methods with respect to the dimensionality of the problem as well as of the number of the agents. In general, the dimensionality of the problem is not a limiting factor in the uncoordinated scenario, which requires to solve a sequence of optimization problems, one for each active agent. Some limitations arise in the centralized, fully cooperative solution. However, the most advanced optimization techniques, such as Borg MOEA (*Hadka and Reed*, 2013; *Reed et al.*, 2013) described in Section 2.4.2, are able to find optimal solutions for challenging problems characterized by many-objective formulations, multi-modality, nonlinearity, discreteness, severe constraints, stochastic objectives. For DCSP/DCOP the dimensionality of the problem can limit the applicability of Adopt algorithm because, although the number of messages exchanged by the agents grows approximately linearly with the number of agents, its worst-case time complexity is exponential in the number of agents. However, there exist approximated algorithms able to find quasi-optimal solutions when the number of agents increases, as distributed breakout algorithm (*Hirayama and Yokoo*, 2005).

4 Upstream-downstream coordination in transboundary systems

The purpose of this chapter is to apply the agent-based decision analytic framework (Section 1.3.2) to estimate the value of cooperation and information exchange in large-scale, transboundary river basins characterized by multiple and originally non-cooperative decision makers (agents). The Zambezi River basin is used to illustrate the methodology. It represents one of the largest river basins in Africa and is shared by eight countries: Angola, Botswana, Malawi, Mozambique, Namibia, Tanzania, Zambia and Zimbabwe. The four largest reservoirs in the basin (Ithezhithezhi, Kafue-Gorge, Kariba, and Cahora Bassa) are mainly operated for maximizing the economic revenue from hydropower energy production, with considerable negative effects on the aquatic ecosystem in the Zambezi delta due to the alteration of the natural flow regime (e.g., *Beilfuss and dos Santos*, 2001; *Tilmant et al.*, 2010). Currently, the four reservoirs are managed by independent and non-cooperative decision makers. This uncoordinated situation is comparatively analyzed with respect to a first level of cooperation (i.e., coordination) by introducing information sharing among the agents. Moreover, the ideal, fully cooperative centralized solution with complete information exchange among the agents is also analyzed. The optimization of the agents' decisions is done according to a model predictive control scheme (MPC, see Section 2.4.1), which is particularly suitable for large scale systems. In particular, given the increasing ability in forecasting techniques producing accurate forecast of the inflows processes in the Zambezi River basin, MPC seems particularly promising for this specific problem.

This chapter is adapted from a journal paper published in Water Resources Research (*Giuliani and Castelletti*, 2013).

4.1 Water management in transboundary systems

The presence of multiple, institutionally independent but physically interconnected decision makers is a distinctive feature of many water resources systems, especially in transboundary river basins: water flows and, by moving, creates a hydrological interdependency between basin users (*Alam et al.*, 2009). More than 250 rivers in the world belong to trans-national basins, accounting for about 60% of freshwater worldwide (*UNEP*, 2002). Yet, disputes arise over contested use of water at different spatial and institutional scales. While control over water resources might not be achieved through water wars (*Zeitoun and Warner*, 2006), subnational disputes often occur (*Wolf*, 1998; *Julien*, 2012) and, whether violent or not, they increase the pressure on water resources along with climate-change induced water scarcity, increasing water demand, and decreasing quality.
Most of the world's largest and disputed trans-national river basins has been studied assuming a centralized decision-making framework, exploring the potential for a more efficient water management at the system-wide scale: the Nile (e.g., *Guariso and Whittington*, 1987; *Wu and Whittington*, 2006; *Block and Strzepek*, 2010)), the Zambezi (e.g., *Gandolfi et al.*, 1997; *Tilmant et al.*, 2010, 2012)), the Euphrates-Tigris (e.g., *Kibaroglu and Unver*, 2000; *Altinbilek*, 2004)). Yet, centralized management assumes a cooperative attitude and full information exchange by the involved parties, and these rarely correspond to the actual socio-political setting in a river basin. When decision makers belong to different countries or institutions, or even sectors, it is very likely that they act considering only their local objectives, producing global externalities that negatively impact on other objectives (*Bernauer*, 2002).
Similarly, although data and information sharing at the basin level is considered a precondition to achieve cooperation (*Gerlak et al.*, 2011), a full information exchange practice is far from being applied in most of transboundary basins. Especially upstream countries have a tendency to restrict information exchange, as it is not in their local interest to give full access to the available information, thus ensuring a favorable position with respect to the countries with no access to information (*Timmerman and Langaas*, 2005).
According to *Watkins* (2006), three levels of cooperation can be identified in a water resources system: *coordination* by information sharing, *collaboration* by developing adaptable national plans, and *full cooperation* by developing joint ownership of infrastructure assets. Given the

lower bound in terms of system-level efficiency represented by the case of completely independent and non-cooperative agents, the role of information exchange is analyzed by introducing coordination as first level of cooperation. The advantage of this basic level of cooperation is its intrinsic feasibility, as the independent, local optimal strategy of each agent is guaranteed. Indeed the practicability of cooperative solutions requires that the benefits to any participant are at least equal to what that participant would obtain by acting unilaterally (*Wu and Whittington*, 2006), and the glue for any stable cooperation among multiple decision makers has to be the self-interest of each participating actor (*Waterbury*, 1997). Conversely, more sophisticated cooperation strategies require to design compensation measures (e.g., international incentives), which modify the benefits of the agents to satisfy this condition of individual rationality (*Sheikhmohammady et al.*, 2011; *Madani and Lund*, 2012). Non-cooperation and coordination are then comparatively analyzed with respect to a reference scenario of full cooperation and complete information exchange among the agents, which is equivalent to the ideal centralized solution of the problem. The differences in the system-level benefits achievable under the different scenarios allow to estimate both the economic value of full cooperation (VFC), measured as the benefits obtained by full cooperation with respect to the ones with coordination only, and the economic value of information exchange (VIE), measured as the benefits obtained with coordination with respect to the ones with no cooperation.

4.2 Problem formulation

4.2.1 The Zambezi River basin case study

The Zambezi River basin (Figure 4.1) is located in south-east Africa. It is shared by eight countries and drains a catchment area of 1.39 million km^2. The river flows eastwards for 2,750 km, from the headwaters, in the Kalene hills in north-west Zambia, to the Indian Ocean in Mozambique. The river consists of three sections: *upper* from the sources to Victoria Falls; *middle*, from the Falls to Cahora Bassa, including the junction with Kafue River tributary; *lower*, from Cahora Bassa to the Indian Ocean, including the Shire River tributary flowing from Lake Malawi. The northern part of the basin belongs to the tropical summer rainfall zone, while moving south the climate becomes more arid.

Because of the high runoff generated in the upper parts of the basin, combined with a fall of more than 1,000 meters during its course to the ocean, the river provides a good opportunity for hydropower energy

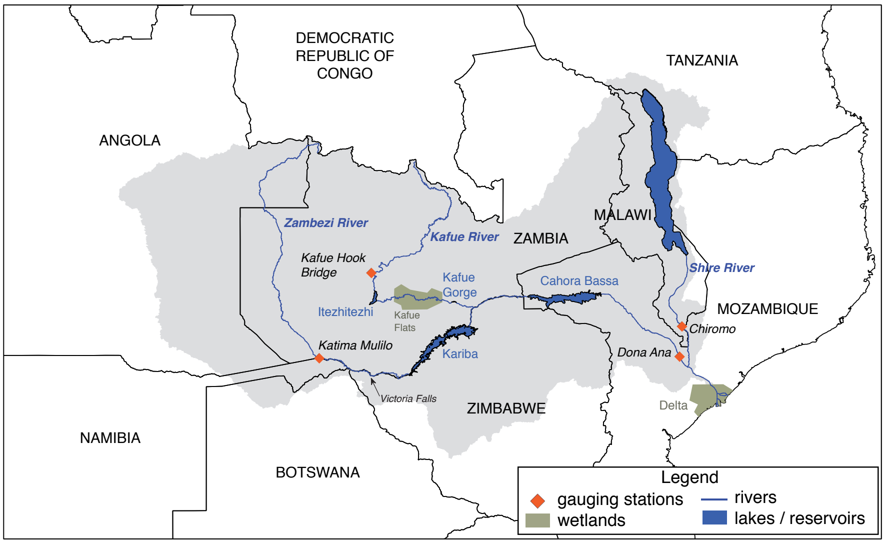

Figure 4.1: The Zambezi River basin.

production. Two dams are located on the tributary Kafue River, which belongs entirely to Zambia and has a catchment area of 155,000 km^2. The Government of Zambia built the Kafue Gorge dam in 1972 and, then, coupled it with the Itezhitezhi dam in 1977, which integrates the insufficient storage at Kafue Gorge due to the high losses for evaporation caused by the presence of the Kafue Flats, an extensive floodplain area where the river flows slowly for 250 km (it takes about two months from Itezhitezhi dam to reach the Kafue Gorge one) with an average gradient of 2.7 cm/km. The operation of Itezhitezhi dam is governed by hydropower generation needs at Kafue Gorge Dam, except for an ecological constraint in March imposing a minimum flow of 300 m^3/s to preserve the Kafue Flats ecosystem (*Beilfuss and dos Santos*, 2001). In the lower Kafue catchment there are several agricultural users with an associated water demand of nearly 15 m^3/s; however, more than half of this flow returns to the system (*Beilfuss and dos Santos*, 2001) and, therefore, it is assumed as negligible in this work. Kariba and Cahora Bassa are the two largest reservoirs on the mainstream. Kariba dam, divided in North and South Banks, belonging respectively to Zambia and Zimbabwe, was completed in 1959; the two countries jointly manage the reservoir through the Zambezi River Authority comprising ministries from both Zambia and Zimbabwe. Cahora Bassa dam was filled in 1974 in Mozambique and controls a large portion of the flow in the lower sec-

tion.
Kariba and Itezhitezhi-Kafue Gorge reservoirs regulate almost 90% of the flow in the middle Zambezi, while Cahora Bassa controls a large portion of the flow in the Zambezi delta. According to *Beilfuss and dos Santos* (2001), the flows to the delta observed after the completion of Cahora Bassa (i.e., 1974-1999) have been reduced during the entire flooding season with respect to the condition before the construction of the dam (i.e., 1930-1974), including a 64% reduction in the mean monthly flow during February-April. The Zambezi runoff measured at Muturara (Dona Ana gauging station) decreased from 3,200 m^3/s between 1930 and 1958 to 2,200 m^3/s over the past 25 years. Many works have recently studied the ecology of the Zambezi River basin and the effect on these large storage operation on the ecosystem (e.g., *Timberlake*, 2000; *Beilfuss*, 2001; *Beilfuss and Brown*, 2010; *Tilmant et al.*, 2010)). As in *Tilmant et al.* (2010), this work focuses on the delta region where a target pulse of 7,000 m^3/s during the peak flow season in February and March was established to restore part of the natural seasonal flow regime. The river discharge in the delta can be actually disaggregated into the releases from Cahora Bassa reservoir and the contribution from the lower Zambezi catchment tributaries, among which the Shire River is the largest one. The Shire River is the outflow of Lake Malawi, the only natural lake in the Zambezi basin, and its average annual runoff at Chiromo station is 483 m^3/s. Since 1960, Lake Malawi outflow has been partially regulated by the Liwonde Barrage aiming to maintain dry season flows in the Shire River for run-of-river hydropower generation. The regulation of Liwonde Barrage in not considered in this application and the Shire River is considered as an inflow.

4.2.2 The agent-based model

The agent-based model of the Zambezi River system (Figure 4.2) comprises six agents representing the six water-related interests in the system, namely the five hydropower plants and the environment in the delta. The five agents associated to the power plants (Itezhitezhi (A_{ITT}), Kafue Gorge (A_{KG}), Kariba North Bank (A_{KAn}), Kariba South Bank (A_{KAs}), and Cahora Bassa (A_{CB})) are *active-controller agents* (shown in blue in Figure 4.2), who operate on a dynamic portion of the system having an internal state (the reservoir storage) and, therefore, decide according to a closed-loop control scheme. The agent in the delta (A_E) is, instead, a *passive agent* (shown in green in Figure 4.2), who does not make any decision but represents the ecological interest in the delta region. According to the real political and institutional setting in the basin, the

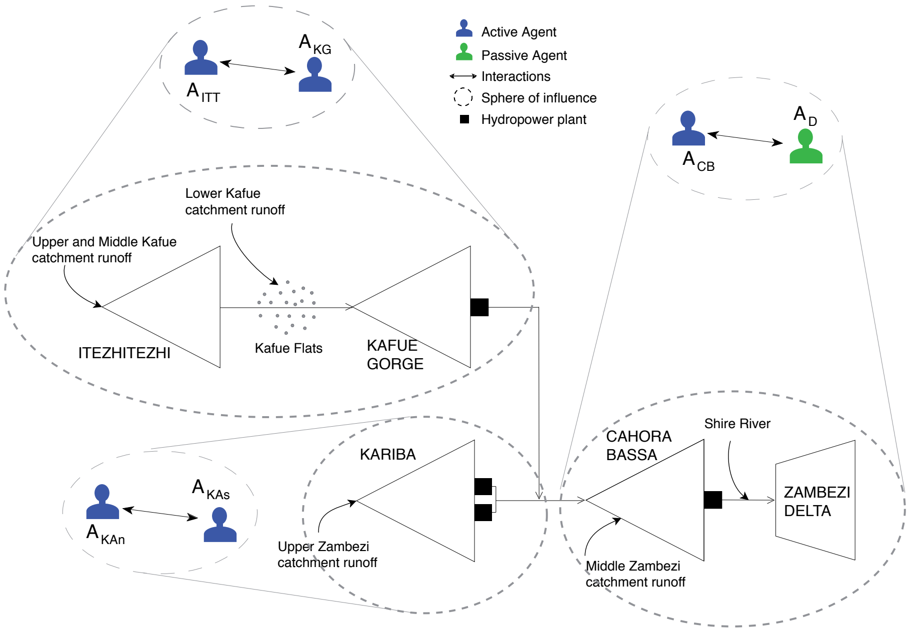

Figure 4.2: Agent-based model of the Zambezi system.

agents are grouped in three coalitions of two agents sharing a common strategy: Itezhitezhi and Kafue Gorge dams require to be jointly operated by agents A_{ITT} and A_{KG} for hydropower production at Kafue Gorge; the regulation of Kariba North and South Banks for hydropower production is established by the Zambezi River Authority, represented by agents A_{KAn} and A_{KAs}; the regulation of Cahora Bassa is designed by the last coalition including agents A_{CB} and A_E and has to consider both hydropower production and ecological preservation of the delta. Only the last coalition is interested in the protection of the environment in the delta as, according to the agent-based model of the system, this coalition represents the interests of Mozambique country.
The water system dynamics is described by the mass balance equations of the reservoirs storages with a monthly time step, as follows:

$$
\begin{aligned}
s_{t+1}^{ITT} &= s_t^{ITT} + (\varepsilon_{t+1}^{ITT} - r_{t+1}^{ITT}) \cdot \Delta - e_t^{ITT} \\
s_{t+1}^{KG} &= s_t^{KG} + (\varepsilon_{t+1}^{KG} + r_{t+3}^{ITT} - r_{t+1}^{KG}) \cdot \Delta - e_t^{KG} \\
s_{t+1}^{KA} &= s_t^{KA} + (\varepsilon_{t+1}^{KA} - r_{t+1}^{KA}) \cdot \Delta - e_t^{KA} \\
s_{t+1}^{CB} &= s_t^{CB} + (\varepsilon_{t+1}^{CB} + r_{t+1}^{KA} + r_{t+1}^{KG} - r_{t+1}^{CB}) \cdot \Delta - e_t^{CB}
\end{aligned} \tag{4.1}
$$

where ε^i_{t+1} (m^3/month) (i=ITT;KG;KA;CB) is the inflow to the i-th reservoir in the interval $[t, t+1)$; Δ is the integration time-step; the release r^i_{t+1} (m^3/month) is given by the release function $r^i_{t+1} = f(s^i_t, u^i_t, \varepsilon^i_{t+1})$, where u^i_t (m^3/month) is the decision (control) and $r^i_{t+1}(\cdot)$ is a non-linear function describing the stochastic relation between the decision u_t and the actual release r_{t+1} (*Piccardi and Soncini-Sessa*, 1991); finally, e^i_t (m^3/month) is the mean monthly losses for evaporation. In particular, the evaporation at Kafue Gorge is calibrated (*Gandolfi et al.*, 1997) to take care also of the significant evaporation losses in the Kafue Flats. According to the monthly time step adopted, river branches are modeled as plug-flow canals with negligible travel time, except for the Itezhitezhi-Kafue Gorge connection which requires two months.

At the basin-wide level, reservoir operation aims at satisfying four different objectives: to maximize the hydropower production (TWh/year) at Kafue Gorge ($J^{H,KG}$, associated to the agents A_{ITT}-A_{KG}), at Kariba North and South Banks ($J^{H,KA}$, related to agents A_{KAn}-A_{KAs}), at Cahora Bassa ($J^{H,CB}$) and to protect the ecosystem in the Zambezi delta (J^E), both associated to agents A_{CB}-A_E. Previous studies considered the environmental requirements in the delta region as an additional constraint (*Gandolfi et al.*, 1997), or as an economic objective by monetizing river flows through a marginal benefit function (*Tilmant et al.*, 2010). The environmental objective is instead represented by a specific objective function J^E (associated to agent A_E) defined as the average water deficit in the delta with respect to the target peak flow of 7,000 m^3/s in February and March. The cooperative solution of the problem at the basin-wide level yields a set of Pareto-optimal or trade-off solutions between the two main objectives, namely total hydropower energy production $J^{H,tot}$ and ecological preservation J^E, and enables a trade-off analysis. On the other hand, in both the coordinated and non-cooperative scenarios the multi-objective problem is defined at the agent (actually coalition of agents) level. The first coalition (agents A_{ITT} and A_{KG}) operates Itezhitezhi and Kafue Gorge and solves a single-objective problem with respect to $J^{H,KG}$ (i.e., maximization of hydropower production at Kafue Gorge). The second one (agents A_{KAn} and A_{KAs}) operates the two power plants at Kariba solving again a single-objective problem with respect to $J^{H,KA}$ (i.e., maximization of hydropower production at Kariba). The third coalition, representing the Mozambique interests (agents A_{CB} and A_E), operates Cahora Bassa dam and solves a two-objective problem related to the maximization of hydropower production at Cahora Bassa $J^{H,CB}$ as well as the protection of the Zambezi delta J^E. In sum, two single-objective optimization problems are formulated for the upstream agents and a two-objective optimization problem for the downstream agents, yielding to a Pareto front in the downstream objectives space $J^{H,CB}$ and J^{E}.

4.2.3 Cooperation and information exchange

Three different scenarios of cooperation/information-exchange are comparatively analyzed: *i*) *fully cooperative and informative* agents, who agree to establish a joint action strategy and exchange all the information; *ii*) *coordinated* agents, where the agents exchange full information (i.e., hydrological data and management objectives) on their sub-systems but do not cooperate to find a globally optimal solution; *iii*) *non-cooperative and non-informative* agents, where agents are completely individualistic and do not share any information. Realistically, there exists another situation where the agents partially exchange information: they may share hydrological data but do not reveal their management objectives (*Giuliani et al.*, 2012b). Depending on the scenario adopted, the formulations of the optimization problems solved by the agents to design their optimal decisions are different as described next.

Fully cooperative and informative scenario

Under this scenario, the agents, beyond sharing information, agree on sharing also the operating objectives in order to establish a joint management strategy for the entire system. This is equivalent to the centralized approach commonly adopted in most of the water resources literature, in which the management problem is formulated as a q-objective, stochastic, periodic, non-linear, closed-loop optimal control problem as described in Section 2.4.
To overcome stochastic dynamic programming (SDP) curse of dimensionality (*Bellman*, 1957), which prevents the use of SDP in the large-scale system of the Zambezi River, a model predictive control (MPC, see Section 2.4.1) approach is adopted. The fully cooperative problem can be therefore formalized as an open-loop control problem defined over a receding horizon h equal to 3 months as follows:

$$\mathbf{u}^* = \arg \min_{\mathbf{u}_t,\ldots,\mathbf{u}_{t+h-1}} \boldsymbol{\lambda} \cdot \mathbf{J}(\cdot) \tag{4.2a}$$

subject to

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \boldsymbol{\varepsilon}_{t+1}) \tag{4.2b}$$

$$\mathbf{u}_t = [u_t^{ITT}, u_t^{KG}, u_t^{KAn}, u_t^{KAs}, u_t^{CB}] \in \mathcal{U}_t \tag{4.2c}$$

The inflows prediction needed to solve Problem (4.2) can be obtained in different ways. Among the set of alternative modeling approaches (autoregressive models, neural networks, extremely randomize trees, etc.), linear periodic PAR(1) models (e.g., *Salas et al.*, 1980) were identified because they are able to provide good forecasts of the inflow processes according to the adopted monthly time step, see the explained variance in Table 4.1. In particular, the interannual variability characterizing the inflows in the Zambezi system (e.g., *Rocha and Simmonds*, 1997; *Mazvimavi and Wolski*, 2006) has limited impacts on the considered objectives and can be captured by autoregressive models.

Table 4.1: Model prediction performance measured with the Nash-Suttcliffe efficiency index over the validation period 1974-1980 (Cahora Bassa inflows are synthetically generated).

Inflow	R^2
Itezhitezhi	0.772
Kafue Gorge	0.772
Kariba	0.770
Cahora Bassa	0.992
Shire River	0.768

Non-cooperative scenario

Under this scenario, the agents look at their local objectives and do not share information on the respective sub-systems. Problem (4.2) is therefore solved as a sequence of local problems. The i-th agent's problem (or the problem of the i-th coalition of agents) becomes:

$$\mathbf{u}^{*i} = \arg \min_{\mathbf{u}_t^i, \dots, \mathbf{u}_{t+h-1}^i} \boldsymbol{\lambda}^i \cdot \mathbf{J}^i(\cdot) \tag{4.3a}$$

subject to

$$\mathbf{x}_{t+1}^i = f_t(\mathbf{x}_t^i, \mathbf{u}_t^i, \boldsymbol{\varepsilon}_{t+1}^i) \tag{4.3b}$$

$$\mathbf{u}_t^i \in \mathcal{U}_t^i \subset \mathcal{U}_t \tag{4.3c}$$

where the i-th agent considers local objective functions $\mathbf{J}^i$ only and the decisions are limited to a subset $\mathcal{U}_t^i$ of the entire feasible decision set $\mathcal{U}_t$.

Moreover, the i-th agent is able to observe only the state variables $\mathbf{x}_t^i$ belonging to its portion of the system. In this scenario the upstream agents are in a favorable condition as they can independently decide what is the best for themselves, while the downstream agents will be affected by and have to adapt to these decisions. This two-step approach is the most common in real-world problems (e.g., *Goor et al.*, 2007; *Anghileri et al.*, 2012). Without information exchange, the information available to the i-th agent comprises only the variables directly observed and controlled by the agent himself. Everything else is a stochastic driver for its sub-system. Conversely, the introduction of information exchange enlarges the information system on which the decisions of the downstream agents are made as described in the next section.

Coordinated scenario

Under this scenario, the agents look at their local objectives but now also share information on the respective sub-systems. On the basis of this information, which comprises hydrological data and the other agents' operating objectives, the downstream agents can develop a model of the upstream sub-system $\psi(\cdot)$ equivalent to the one used by the upstream agents to optimize their decisions. According to this model, the downstream agents can simulate off-line the upstream optimal behavior and exploit an enlarged information system to make more informed decisions. According to this scenario, the i-th agent's problem (or the problem of the i-th coalition of agents) becomes:

$$\mathbf{u}^{*i} = \arg \min_{\mathbf{u}_t^i, \ldots, \mathbf{u}_{t+h-1}^i} \boldsymbol{\lambda}^i \cdot \mathbf{J}^i(\cdot) \tag{4.4a}$$

subject to

$$\mathbf{x}_{t+1}^i = f_t(\mathbf{x}_t^i, \mathbf{u}_t^i, \mathbf{w}_t^i, \boldsymbol{\varepsilon}_{t+1}^i) \tag{4.4b}$$

$$\mathbf{u}_t^i \in \mathcal{U}^{\mathbf{i}}{}_t \subset \mathcal{U}_t \tag{4.4c}$$

$$\mathbf{w}_t^i = \psi(\mathbf{I}_t^{-i}) \tag{4.4d}$$

where $\mathbf{w}_t^i$ are the variables affecting the i-th agent sub-system dependent on the decisions of the other agents, which can be described by the model ψ based on the exchanged information $\mathbf{I}_t^{-i}$ related to the sub-system observed by the other agents.

4.3 Results

The fully cooperative, coordinated, and non-cooperative management of the Zambezi River systems are evaluated by simulating the agent-based model (Section 4.2.2) over the historical time series of inflows on the period 1974-1980. Apart from the Middle Zambezi Catchment that required a process of synthetic time series generation, the data used in this work, represented in Figure 4.3, were provided by the Global Runoff Data Centre (http://www.bafg.de/).

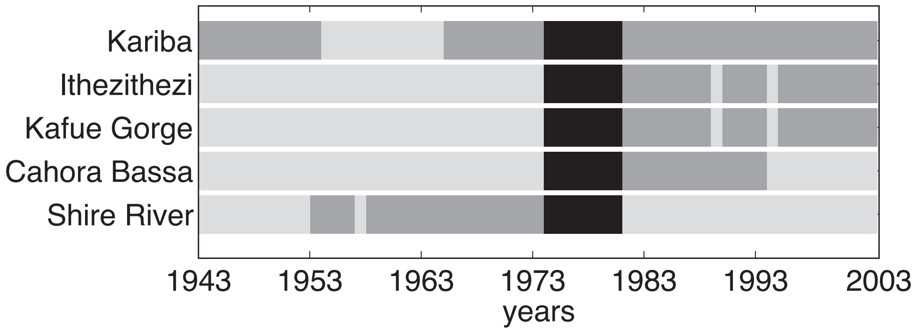

Figure 4.3: Available inflows data (source: Global Runoff Data Centre): the validation-simulation period is in black, the data used in the calibration of the AR models are in dark grey, missing data are in light grey

4.3.1 Fully cooperative solutions

In order to evaluate the tradeoff relationships between hydropower production and ecological preservation of the Zambezi delta, the Pareto front in the objectives space $J^{H,tot}$ and J^E was approximated by solving Problem (4.2) for different combinations of weights $\boldsymbol{\lambda}$. Figure 4.4 shows the obtained Pareto efficient solutions: with solution S^H, which considers only the hydropower objective, the energy production is equal to 34.03 TWh/year and the corresponding average flow deficit in February and March is equal to 3,143 m^3/s. This result is comparable with the value obtained by *Gandolfi et al.* (1997) on the basis of the models were calibrated. The resulting energy production seems instead overestimated, in particular at Kafue Gorge and Kariba, with respect to the one obtained by *Tilmant et al.* (2010) using models calibrated on historical productions. Adopting the same correction, an energy production equal to 30.3 TWh/year is estimated, which is consistent with the hi-

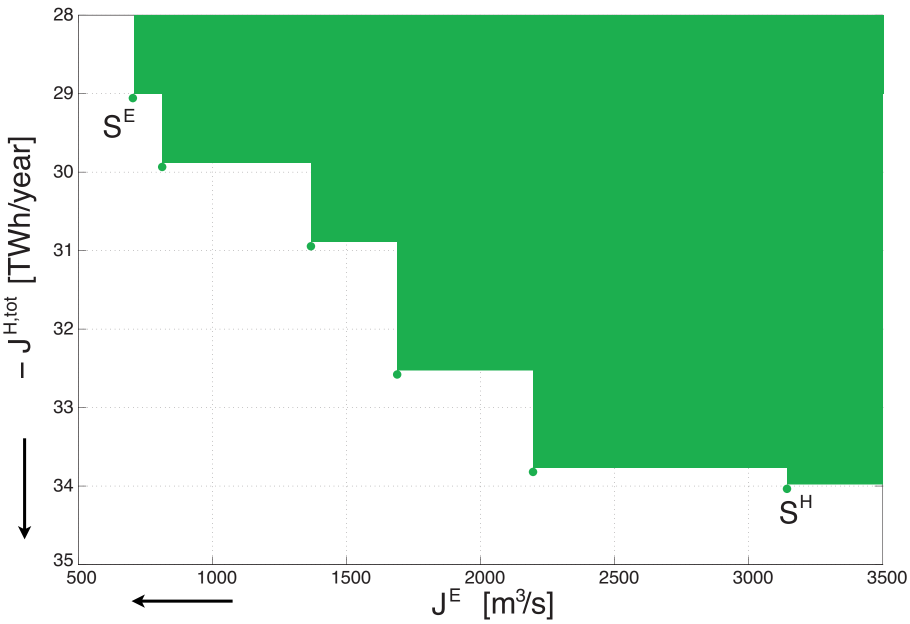

Figure 4.4: Fully cooperative (centralized) Pareto front in the objectives space $J^{H,tot}$ and J^E.

storical productions. However, the real decision makers acting in the system do not look at hydropower production only, but implicitly consider and weight also other minor objectives (e.g., irrigation supply in the Kafue catchment). Moreover, the decision makers can also store water when the power plants are unavailable deciding to turbine later with a higher head. Hence, it seems reasonable that the historical production is lower than the obtained values. The introduction of the environmental objective alters the decision makers objectives balance. For this reason, although the values used by *Gandolfi et al.* (1997) probably overestimate the hydropower production, they provide a solid ground to project the decision makers behavior in a multi-objective context. Moreover, the primary aim of this work is not to improve the efficiency of hydropower energy production system in the Zambezi River basin, but rather to assess the differences in the objective values due to different scenarios of cooperation and information exchange as explained next.

The hydropower production for solution S^E, which considers only the ecological objective, is lower than that for S^H and equals 29.05 TWh/year, with the flow deficit decreased to 703 m^3/s. The comparison of the average monthly flows in the Zambezi delta for these two solutions (Figure

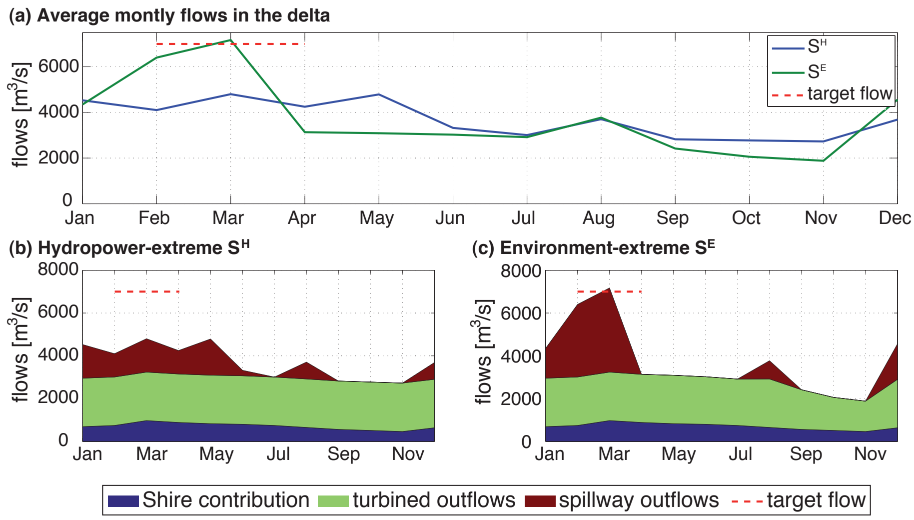

Figure 4.5: Comparison of average monthly flows in the delta (panel (a)) along with their different contributions for the two extremes solutions of the fully cooperative Pareto front S^H in panel (b) and S^E in panel (c).

4.5a) shows a distinctive difference in the flows in February and March due to the ecological objective. Although solution S^E, aiming at minimizing the difference with respect to the target peak flow of 7,000 m^3/s in February and March, does not obtain a null deficit, it is able to restore a peak flow close to the natural condition. However, this release strategy reduces water availability in the rest of the year, while the hydropower-based regulation of Cahora Bassa dam produces almost constant flows in the delta. It is also worth noting that the water released for environmental purposes is actually wasted with respect to hydropower production, as shown in Figure 4.5b-c, where the different contributions to the flows in the delta are separated: S^E spills a significant volume of water in February and March trying to guarantee the target flow, while spillway outflows are obviously more limited for S^H.

The tradeoff between the two objectives can be estimated by pair comparison of the solutions in the objectives space: the improvement in one objective is compensated by the worsening of the other. As an example, the difference in hydropower production between S^H and S^E (34.03 - 29.05 = 4.98 TWh/year) is balanced by the decrease of average water deficit in the delta in February and March (from 3,143 m^3/s to 703 m^3/s). Note that the curvature of the tradeoff curve is almost constant and does not significantly change near the extremes. This unexpected result can be explained by the limited length of the peak flow season when the environmental flow target has to be satisfied: the aggregation of the hydropower production along the entire year mitigates the negative impacts of the ‘unproductive water’ released in the peak flow season. The curvature of the Pareto front obtained by computing the two objectives only in February and March changes moving along the front.

The economic cost of the water deficit in the delta can be estimated by considering the hydropower revenue instead of the energy production. According to *Tilmant et al.* (2010) and *Whittington et al.* (2005), a constant price of electricity equal to 80 US$/MWh is assumed, which is independent from the actual power production as a detailed model of the entire South African Power Pool (SAPP) would be too complex and, moreover, the energy produced in the Zambezi system is low compared to that of South-Africa. By the analysis of the last prices available on the SAPP website, it seems that the proposed value is slightly overestimated. However, the reference price found in the literature is used to enable comparison with previous works. The Pareto-optimality concept guarantees that the decrease in the hydropower revenue associated to a given solution with respect to solution S^H is compensated by the corresponding reduction of the environmental costs in the delta. Therefore, the cost of the water deficit $\nu^{\boldsymbol{\lambda}}$ of the solution $S^{\boldsymbol{\lambda}}$ associated to the vector of weights $\boldsymbol{\lambda}$ in the objectives space can be estimated as the slope of the line connecting that solution to solution S^H (*Wu et al.*, 2010), as follows

$$\nu^{\boldsymbol{\lambda}} = \frac{|J^{H,tot}(S^H) - J^{H,tot}(S^{\boldsymbol{\lambda}})| \times 10^6 \times 80}{|J^E(S^H) - J^E(S^{\boldsymbol{\lambda}})|} \quad \frac{US\$}{m^3/s} \tag{4.5}$$

Comparing solution S^H with the other extreme of the Pareto front S^E (obtained with $\boldsymbol{\lambda} = [.0001, .9999]$ in order not to consider multiple optima), the yearly economic cost of the water deficit in the delta is $\nu^E = 4.98 \times 10^6 \times 80 \ / \ 2{,}440 = 1.6319 \times 10^5$ US$/(m^3/s). The value of $\nu^{\boldsymbol{\lambda}}$ is actually different for each point of the Pareto front as each solution corresponds to a different balance of the two objectives and its range of variability is between 0.1805×10^5 and 1.6319×10^5 US$/(m^3/s), while the mean value is $\bar{\nu} = 1.0825 \times 10^5$ US$/(m^3/s).

It is interesting to observe that both the average value $\bar{\nu}$ and the whole range of variability of $\nu^{\boldsymbol{\lambda}}$ are very close to the upper bound of the marginal-benefit functions proposed in literature to monetize the benefits associated with the flows in the delta, namely 1.7143×10^5 US$/(m^3/s) corresponding to 1000 US$/ha (*Woodward and Wui*, 2001; *Brander et al.*, 2006; *Tilmant et al.*, 2010). This result suggests that the explicit consideration of an environmental objective is equivalent to assigning a high economic value to the ecological preservation of the aquatic ecosystems. From a practical point of view, this result also confirms the findings of *Whittington et al.* (2005): water resources management does not sufficiently consider the environmental requirements which assume secondary importance. Indeed, the range of wetland economic valuation spans from this upper bound to a lower bound of only 100 US$/ha, corresponding to 1.7143×10^4 US$/(m^3/s) that is much lower than the obtained $\nu^{\boldsymbol{\lambda}}$.

4.3.2 Coordinated and non-cooperative solutions

In order to assess the role of information exchange, the system was then simulated under different levels of cooperation, in particular assuming the non-cooperative scenario and the coordinated one, which basically differ in the degree of information sharing (see Problems (4.3)-(4.4)). The agents' objectives are then evaluated one by one. Figure 4.6 shows a graphical representation of the objectives, comparing the agent-based solutions in these two scenarios with a fully cooperative compromise solution. This latter is shown for demonstration purposes and was selected among the set of Pareto optimal solutions obtained in the cooperative scenario by adopting the criterion of the minimum distance with respect to the Utopia point (*Eschenauer et al.*, 1990), which identifies the absolute optima of all the objectives and is usually an unfeasible solution.

Not surprisingly, all the agent-based solutions produce the same performance in the upstream objectives $J^{H,KG}$ and $J^{H,KA}$ in every scenario (the blue and red circles in Figure 4.6 have the same size), because the upstream agents act independently. Moreover, these solutions outperform the centralized one for the upstream objectives (blue and red circles are bigger than the green one). On the other side, looking at the downstream objectives $J^{H,CB}$ and J^E, the fully cooperative solution outperforms the others (the green point is placed in the bottom-left part of the figure): the joint management is clearly able to better exploit the upstream-downstream relationships in order to guarantee a good compromise among all the agents' objectives.

For both the coordinated and non-cooperative scenarios, a set of Pareto optimal solutions (Figure 4.7) between the downstream objectives $J^{H,CB}$ and J^E are obtained by changing the vector of weights $\boldsymbol{\lambda}^i$ in the optimization problem of agents A_{CB} and A_E. The interesting evidence is that increasing the degree of information exchange improves the performance of the downstream agents in both the objectives: the non-cooperative Pareto front obtained with non-information sharing (blue

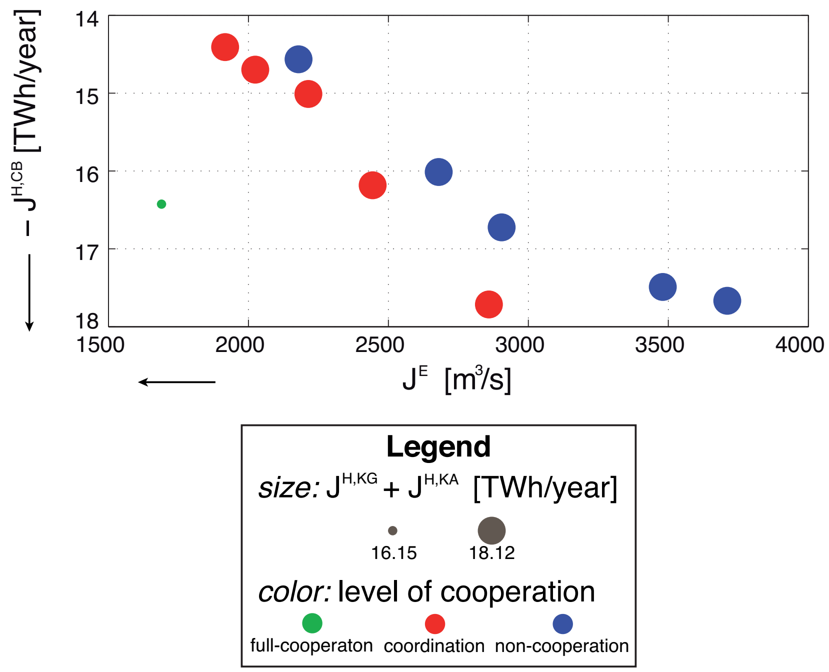

Figure 4.6: Projection of the 4D Pareto front in the objectives plan $J^{H,CB}$ and J^{E}. Circle size represents the upstream agents' objectives $J^{H,KG} + J^{H,KA}$ (the bigger the circle, the better is the solution). The agent-based solutions in the non-cooperative (blue circles) and in the coordinated (red circles) scenarios are compared with a fully cooperative compromise solution (the green circle).

points) is dominated by the coordinated one, which assumes complete information sharing (red points). To quantify the role of information on the whole Pareto front, it is possible to look at the values of the hypervolume indicator (*Zitzler et al.*, 2003) in the two scenarios. The hypervolume indicator measures the volume of objective space dominated by a given Pareto front, thus allowing the comparison of Pareto fronts obtained with different methods respecting the dominance relationships (*Knowles and Corne*, 2002). High values of this indicator are obtained for Pareto fronts that are both converged and diverse. The hypervolume indicator is equal to 0.426 in the case of non-cooperation, while it is equal to 0.495 for coordinated agents. These results confirm, from a quantitative point of view, the general superiority of the solutions obtained by simple cooperation through full information exchange: the

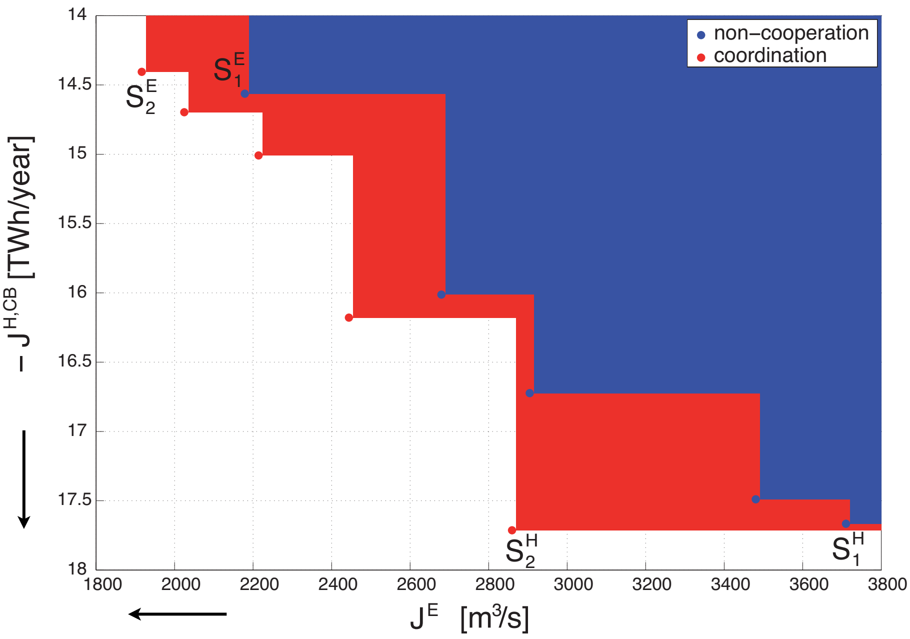

Figure 4.7: Pareto front in the downstream agents' objectives space $J^{H,CB}$ and J^E for different scenario of information sharing (i.e., non-cooperation and coordination).

non-information sharing solutions are indeed Pareto-dominated by the coordinated ones and information exchange allows to improve both the considered objectives. Moreover, these advantages for the downstream agents are obtained without affecting the upstream decisions and thus the corresponding benefit. Finally, observe that the improvement by information exchange varies along the Pareto front, meaning that the role of information sharing might depend upon the objective considered. Specifically, the improvements in the two extreme solutions are equal to 0.27% (hydropower only) and 12.04% (environment only), thus suggesting that the information exchange, and consequently the cooperation, has a higher marginal value for the environmental objective.

4.3.3 Value of full cooperation and information exchange

The above analysis demonstrates that even a basic level of cooperation, only based on the exchange of information among the agents, might increase both the benefits for the downstream agents and the overall utility at the basin-wide level with respect to a non-cooperative setting.

An economic evaluation of the improvements potentially achievable by cooperation and information exchange might represent an effective basis to identify the most suitable policy mechanisms (e.g., economic incentives) to be implemented in order to favor a more cooperative attitude (*Whittington et al.*, 2005). From the comparative analysis of the solutions obtained respectively with the fully-cooperative and the coordination scenarios against the non-cooperative one, it is possible to infer the economic value of full cooperation (VFC) and of information exchange (VIE) associated to the differences in the performance at the basin-wide level.

To this end, the physical objectives (energy production and flow rate deficit) must be converted into monetary values. In the case of the hydropower energy production, the conversion is made by assuming again the African energy price equal to 80 US$/MWh (*Tilmant et al.*, 2010; *Whittington et al.*, 2005). For the water deficit in the delta, the cost can be computed using the values $\nu^{\boldsymbol{\lambda}}$ previously estimated on the fully cooperative Pareto front. Since the marginal value of the information varies along the Pareto front and depends upon the objective considered, the two extreme solutions of the Pareto front are first considered and, subsequently, a compromise solution is analyzed.

For the hydropower extreme of the Pareto front (points S_1^H and S_2^H in Figure 4.7), the annual energy production of the coordinated strategy is 0.0471 TWh/year more than the one in the non-cooperative case. This difference corresponds to 3.76×10^6 US$/year, which is the economic VIE for hydropower energy production. Dually, considering the other extreme which only minimizes objective J^E (points S_1^E and S_2^E in Figure 4.7), the reduction of the flow deficit is equal to 262.428 m^3/s. Hence, the VIE for an ecological management of Cahora Bassa is estimated adopting the value of ν^E previously identified for solution S^E (1.6319×10^5 US$/(m^3/s)) and equals 42.82×10^6 US$/year.

Finally, it is interesting to estimate the role and the economic value of cooperation by comparing one solution for each of the three considered scenarios. Since the fully cooperative solutions were obtained by solving Problem (4.2) in a centralized way, while the agent-based solutions are derived assuming different levels of cooperation in the resolution of Problems (4.3)-(4.4), the selection of which solutions to compare is not straightforward. Therefore, for each scenario, the "most interesting" solution is again identified according to the criterion of the minimum distance from the Utopia point (*Eschenauer et al.*, 1990). The three selected solutions are reported in Table 4.2. It is evident that the fully cooperative solution is worse for the upstream agents, especially for the Kariba hydropower production, but enables improved performance for the downstream agents' objectives that are more damaged in the other scenarios by the individualistic strategies adopted upstream, as already shown in Figure 4.6. The economic values of full cooperation and information exchange were then estimated by looking at the monetary gain corresponding to the four objectives in the three considered solutions (assuming the average cost of the water deficit $\bar{\nu}$): the fully cooperative solution guarantees a monetary gain equal to 2.42×10^9 US\$/year, which is significantly higher than 2.39×10^9 US\$/year and 2.37×10^9 US\$/year given by the coordinated and non-cooperative solutions, respectively. These results allow to estimate a VIE and a VFC equal to 15.7×10^6 US\$/year and 28.2×10^6 US\$/year, respectively. On the basis of this information, the agents (and hence the real decision makers) might reconsider their behaviors and introduce a complete information exchange which produces advantages to the downstream agents without affecting the benefits for the upstream ones. Moreover, given the reference of the fully cooperative solutions along with the VFC, the knowledge of the best performance ideally achievable can be used to get insights on strategies to foster more sophisticated cooperation and negotiation strategies.

Table 4.2: Solutions selected in the Pareto space according to the minimum distance from Utopia point criterion.

Scenario	$J^{H,KG}$ [TWh/year]	$J^{H,KA}$ [TWh/year]	$J^{H,CB}$ [TWh/year]	J^E [m^3/s]
Non-cooperation	7.67	10.45	14.56	2,178
Coordination	7.67	10.45	14.40	1,916
Full cooperation	7.63	8.51	16.42	1,689

4.4 Discussion and final remarks

In this chapter, the agent-based decision analytic framework is used to study different degrees of cooperation and information exchange among multiple decision makers and/or stakeholders in the large-scale, transboundary Zambezi River basin. The framework also allows to economically quantify the value of full cooperation and the of information exchange, which might be a fundamental information to set up a negotiation process.

The agent-based model of the Zambezi River comprises five active agents, representing the managers of the five main hydropower plants in the sys-

tem, and a passive agent modeling the ecosystem of the river delta. Three different scenarios of cooperation and information exchange among the agents have been evaluated: *i*) fully cooperative and informative agents, who agree to act coordinately and exchange all the information, which is equivalent to assume a centralized decision maker who manages simultaneously the entire system; *ii*) coordinated agents who exchange information, without, however, actively cooperate to find a globally optimal solution; *iii*) non-cooperative and non-informative agents, who are completely individualistic and do not share any information.
Results show that it is possible to improve the conditions of the downstream agents representing the Mozambique interests (i.e., hydropower production at Cahora Bassa and protection of the ecosystem in the Zambezi delta) with respect to the actual non-cooperative setting by introducing coordination among the agents. This simple mechanism allows the downstream agents to more effectively adapt to the upstream strategies only because they know these strategies. Pareto front solutions obtained in the coordinated scenario outperform the corresponding solutions with no-cooperation at all. According to the concept of Pareto dominance, coordination is therefore worthwhile in large water resources systems as it allows to obtain solutions that are better (at least for one objective) with respect to the non-cooperative one.
The multi-objective nature of the analysis allowed to estimate the tradeoffs between hydropower energy production and the protection of the Zambezi delta. The explicit consideration of an environmental objective corresponds to assign a high economic value to the ecological preservation of the aquatic ecosystems in the delta, close to the upper limit of the range of wetland valuation proposed in literature. Moreover, it is interesting to observe that the role as well as the economic value of information exchange changes according to the considered Pareto-efficient solution and, in particular, the marginal value is higher for the environmental objective.
More in general, information exchange might have a primary role in rebalancing the upstream/downstream asymmetry in the Zambezi River basin, as it allows the downstream agents to better adapt to the upstream management strategies, with no consequence for these latter. Compared to ideal centralized solution (full cooperation and information exchange), the stability of this coordinated solution is guaranteed as each agent cannot improve its benefit acting unilaterally.
Finally, the economic Value of Full Cooperation and Value of Information Exchange are estimated by comparing one solution for each scenario: the economic gain achievable at the system-level by moving from coordination to the ideal full cooperation is 28.2×10^6 US\$/year, while the

introduction of the information exchange as a first level of cooperation, i.e. coordination, with respect to the non-cooperative scenario produces an economic gain equal to 15.7×10^6 US$/year. Adopting the historical production based calibration of the models as in *Tilmant et al.* (2010), the energy production is lower, resulting in different VIE and VFC. In particular, the fully cooperative, coordinated and non-cooperative solutions guarantee monetary gains equal to 2.14×10^9 US$/year, 2.087×10^9 US$/year and 2.073×10^9 US$/year, respectively. The corresponding VIE is equal to 14.7×10^6 US$/year, while the VFC is 58.4×10^6 US$/year, as the fully cooperative solution is better than the coordinated one in the downstream objectives, which do not vary, and is worse in the upstream objectives, which are reduced. Consequently, the gap between the coordinated and the fully cooperative solutions increases. Note that, in reality, VIE depends on two factors: first, the agents (institutions) have a constrained capacity, which limits their ability in the exploitation of the shared information, depending on the accuracy of the models they use. The costs to develop these models is assumed as negligible with respect to the benefits produced by the use of the models to optimize the operation in the system. Furthermore, VIE tends to decrease in time, because by acquiring experience over time, the agents have more data available for the identification of more and more effective forecast models. The availability of better forecasts allows the agents in the non-cooperative scenario to obtain a performance closer to the ones in the coordinated scenario, with the gap decreasing in time, meaning that the value of information exchange decreases as well. The accuracy of the forecast models is however limited and VIE does not tend to zero.
In conclusion, it is worth summarizing the limitations of the proposed approach, in particular with respect to the computational costs and the scalability for larger systems . The centralized solution is by far the more expensive, while the non-cooperative one is the more flexible. As the problem scales up, a number of algorithms can be adopted to mitigate the computational burden, ranging from approximate dynamic programming (e.g., *Powell*, 2007) for relatively small systems to simulation-based optimization methods (e.g., *Koutsoyiannis and Economou*, 2003). For each algorithm using the weighting method, the associated computational costs grow exponentially with the number of objectives considered. The number of weights combinations required to accurately approximate the real continuos Pareto front might change from problem to problem. Clearly, the higher the accuracy the higher the computational costs. The interaction with the decision maker usually helps in refining the approximation (adopting a more dense sampling of the weights space) in the region of interest for the decision maker. Without such a direct interac-

tion, the weights can be selected to ensure that the shape of the front was reasonably represented.

5 Co-adaptation in agricultural water systems

The purpose of this chapter is to apply the agent-based decision analytic framework (Section 1.3.2) to demonstrate the potential of co-adapting water demand and supply in agricultural water systems. According to *IPCC* (2001), *adaptation to climate change can be defined as an adjustment in ecological or socio-economic system in response to observed or expected changes in climatic stimuli and their effects, in order to mitigate adverse impacts of change or take advantage of new opportunities.* Effective adaptation strategies hence involves decisions made across scales, from individuals to governments at local, regional, and national level, and international agencies (*Adger et al.*, 2005). In this chapter, the Lake Como serving the Muzza-Bassa Lodigiana irrigation district (Italy) is used to illustrate the methodology. A distributed-parameter, dynamic model of the irrigation district allows the simulation of crop growth and yield over a range of hydroclimatic conditions, irrigation strategies, and water-related stresses. The proposed co-adaptation strategy aims to cross-condition the decisions of farmers and water managers: the farmers decide the most profitable crop option on the basis of an expected water supply; knowing the farmers decisions, the water supply strategy (i.e., the regulation of Lake Como) is optimized with respect to the actual irrigation demand of the crops. By iteratively running this procedure, the farmers and the water manager will exchange information until the system converges to an equilibrium, where water supply and demand are considered as coupled human (farmers and water managers) and natural (crops) systems (*Liu et al.*, 2007). The proposed approach is tested in two different climatic scenarios, namely current and projected conditions, to assess the potential for the co-adaptation to enhance the efficiency of agricultural water management practices and foster crop production, as well as to mitigate climate change adverse impacts.
This chapter is adapted from a paper under preparation for Global and Environmental Change (*Giuliani et al.*, 2013d).

5.1 Agricultural water management adaptation to climate change

Agriculture is the main land use in the world and also the sector characterized by the highest water demand (*FAO*, 2003; *Jury and Vaux*, 2007; *Phalan et al.*, 2013), with 24% of the total harvested cropland that is under irrigation (*Portmann et al.*, 2010). To meet projected growth in human population and per-capita food demand, agricultural production will have to significantly increase in the next decades along with the corresponding irrigation consumptions (e.g., *Tubiello et al.*, 2007; *de Fraiture and Wichelns*, 2010). Yet, water availability, which is often a key factor in determining crop productivity, is expected to decrease over the next century due to climate change impacts (*IPCC*, 2007), with 3.1 billion urban dwellers expected to experience water shortages by 2050 (*McDonald et al.*, 2011). Squeezing more crop out of the same drop will be one of the biggest challenges of this century (*Marris*, 2008) in order to guarantee adequate irrigation supply in cropping regions where precipitations are expected to decline. If current crops, such as rice, cotton, or sugar cane, were not irrigated, their production would indeed decrease by up to 60% (*Siebert and Döll*, 2010). Consequently, irrigated areas are expected to expand over the next years (*Neumann et al.*, 2011), producing a further increase of the related water consumptions. This means that competition for water between agriculture and the other sectors is destined to increase (*Falkenmark*, 2013).
The impacts of climate change on water resource systems has been intensively studied in the literature (see the review by *Leavesley*, 1994, and references therein), mostly focusing on the hydrological cycle and the underlying natural processes to analyze the changes in the distribution of river flows and groundwater recharge over space and time, according to projected changes in temperature and precipitation (e.g., *Di Baldassarre et al.*, 2011; *Escaler et al.*, 2012). The effects on the natural, economical, and social sphere are also studied (e.g., *Schaefli et al.*, 2007; *Ajami et al.*, 2008; *Hingray et al.*, 2007). Climate variability has been demonstrated the most significant factors influencing year to year crop production (e.g., *Mall et al.*, 2006; *Kang et al.*, 2009), because crops' growth is strongly dependent on water availability and temperature. However, crops' growth processes are also strongly affected by farmers practices as well as water supply management strategies. Agricultural water systems, therefore, comprise a natural (crops) component coupled with a human (farmers and water managers) component, and both are subjected to changing climatic conditions. As a consequence, these coupled systems

are co-evolving under change. However, although several studies have assessed climate change impacts on agricultural water systems, most of them considers the two problems separately. Moreover, current practices are generally established according to historical agreements and normative constraints and, in the absence of dramatic failures, the shift toward a more efficient integrated water management is not easily achievable.

5.1.1 Water demand adaptation

Farmers practices are composed by many different activities which are planned and scheduled periodically, most of them at the beginning of the agricultural season. Since water is becoming more and more scarce, farmers are called to modify these practices to maximize production and farm revenue (e.g., *Marques et al.*, 2005; *John et al.*, 2005; *Ng et al.*, 2011). Many options are potentially available for marginal modifications of existing agricultural systems (e.g., *Seo et al.*, 2005; *Mall et al.*, 2006; *Deressa et al.*, 2009): *i*) the change of cultivar types and rotations by replacing crops sensitive to climate and water stresses with more resistant ones (*Howden et al.*, 2007); *ii*) the shift of sowing and harvesting dates to match crops' temperature requirements, such as early planting to avoid heat and drought stresses in the late summer (*Rosenzweig and Tubiello*, 2007); *iii*) the adoption of high efficiency irrigation techniques (*Cai and Rosegrant*, 2004) along with the implementation of other practices (e.g., conservation tillage) to conserve soil moisture and face drought conditions (*Bindi and Howden*, 2004). According to *McCarl* (2006), the implementation of adaptation strategies would allow a substantial increase in the crop yields with respect to a no-adaptation baseline.

5.1.2 Water supply adaptation

Farmers and water demand adaptation is only one part of the equation, because changes in the water supply management strategies significantly affect crop productivity, particularly in irrigated agricultural systems. Different operations of the water supply system, especially in the case of water reservoirs, can significantly alter the amount of water available for the farmers both in time and space. On-demand regulation of water reservoirs and irrigation canals is a promising option to improve the efficiency in the utilization of the available water (*Mareels et al.*, 2005; *Galelli and Soncini-Sessa*, 2010). The adaptation of the supply system to the undergoing change is also a key factor in mitigating projected climate impacts (*Anghileri et al.*, 2011). However, prior studies in this area (e.g., *van Oel et al.*, 2010; *Ng et al.*, 2011; *Garcia-Vila and Fereres*,

2012) assume irrigation as an external input and evaluate few scenarios of water availability, while an analysis of the feedbacks between water supply and demand is still missing.

5.2 Problem formulation

5.2.1 The Lake Como system case study

Lake Como is a regulated lake in Northern Italy (Figure 5.1). Its storage is about 254 Mm^3 and it is fed by a 3,500 km^2 catchment, characterized by the typical Alpine hydrological regime with scarce discharge in winter and summer, and peaks in late spring and autumn due to snowmelt and rainfall. The lake inflow and effluent is the Adda River, serving eight run-of-river hydroelectric power plants and feeding a dense network of irrigation canals, which supports five irrigation districts with a total surface of 1,400 km^2. The regulation of the lake aims also to prevent flooding along the lake shores, especially in Como city.
The Muzza-Bassa Lodigiana district is one of the irrigation districts served by the Adda River, located south-east of the city of Milan in the Pianura Padana region (Figure 5.1). It has an area of about 700 km^2 and irrigation is practiced with the border method (or free-surface flooding). Major crops are cereals (especially corn) and permanent grass. The district is divided in 66 irrigation units, which represent the decision-making authority selecting the crop to grow. The irrigation supply is provided by the Muzza main canal, which originates from the River Adda, and is hence controlled by the regulation of the Lake Como.

The model of the Lake Como

The Lake Como is modeled focusing on its storage capacity and the downstream system, with an approximated representation of the upstream catchment and the small artificial reservoirs operated for hydropower production. The lake dynamics is defined by a mass balance equation as follows:

$$s_{t+1} = s_t + n_{t+1} - r_{t+1} \tag{5.1}$$

where s_t is the lake storage, n_{t+1} and r_{t+1} are the inflows and the outflows in the time interval $[t, t+1)$, respectively. In particular, the release is given by the release function $r_{t+1} = R_t(s_t, u_t, n_{t+1})$ which accounts for any possible deviation of the actual release from the decision u_t due to unintentional spills or any other physical or legal constraints (*Piccardi*

Figure 5.1: Schematic map of the system: Lake Como, Adda River, and the Muzza-Bassa Lodigiana irrigation district.

and Soncini-Sessa, 1991). Lake Como management is driven by multiple objectives, such as flood prevention, hydropower production, irrigation supply. This work focuses only on irrigation supply. The water supply objective J^S is defined as the quadratic daily average water deficit over the simulation horizon H with respect to the irrigation demand w, as follows:

$$J_S = \frac{1}{H} \sum_{t=1}^{H} \beta_t \left(\max(w_t - q_{t+1}^M, 0)^2 \right) \tag{5.2}$$

where β_t is a time-varying coefficient taking into consideration the different relevance of the water deficit in different periods of the years and q_{t+1}^M is the flow diverted from the Adda River in the Muzza main canal, namely $q_{t+1}^M = \min(\alpha^M \cdot r_{t+1}, q^{max})$, with α^M representing the water allowance of the Muzza district and $q^{max} = 110$ m^3/s the capacity of the canal. This quadratic formulation aims to penalize severe deficits in a single time step, while allowing for more frequent, small shortages (*Hashimoto et al.*, 1982).

The model of the Muzza district

The dynamic model of the Muzza district includes three main modules, devoted to specific tasks: *i*) a distributed-parameter water balance module which simulates water sources, conveyance, distribution, and soil-crop water balance (*Facchi et al.*, 2004) ; *ii*) a heat units module which simulates the sequence of growth stages as a function of the accumulation of heat units, according to the growing degrees principle (*Neitsch et al.*, 2011); *iii*) a crop growth module which estimates the optimal and actual yields depending on possible stresses experienced during the agricultural season due to insufficient water supply from rainfall and irrigation (*Steduto et al.*, 2009; *Raes et al.*, 2009).
The water balance module partitions the irrigation district with a regular mesh of cells with a side length of 250 m, which allows the representation of the space variability of crops, soil types, meteorological inputs, and irrigation distribution. Each individual cell identifies a soil volume which extends from the soil surface to the lower limit of the root zone. This soil volume is subdivided into two layers, modeled as two non-linear reservoirs in cascade (Figure 5.2): the upper one (evaporative layer) represents the upper 15 cm of the soil; the bottom one (transpirative layer) represents the root zone and has a time-varying depth. The water percolating out of the bottom layer constitutes the recharge to the groundwater system. The dynamics of the water contents $\theta_{1,t}^{(i)}$ and $\theta_{2,t}^{(i)}$ in the evaporative and transpirative layers for each cell i are described by the following mass balance equations:

$$\theta_{1,t+1}^{(i)} = \theta_{1,t}^{(i)} + (R_{t+1}^{(i)} - C_{t+1}^{(i)}) - Q_{r,t+1}^{(i)} - E_{t+1}^{(i)} - Q_{u,t+1}^{(i)} + I_{t+1}^{(i)} \quad (5.3a)$$

$$\theta_{2,t+1}^{(i)} = \theta_{2,t}^{(i)} + Q_{u,t+1}^{(i)} - Tr_{t+1}^{(i)} - Q_{g,t+1}^{(i)} \quad (5.3b)$$

where R_{t+1} is the rainfall, C_{t+1} the canopy interception, $Q_{r,t+1}$ the surface runoff, E_{t+1} the evaporation, $Q_{p,t+1}$ the percolation to the transpirative layer, I_{t+1} the irrigation supply, Tr_{t+1} the transpiration, $Q_{g,t+1}$ the outflow to the groundwater system, all in the time interval $[t, t+1)$. In particular, canopy interception and surface runoff are computed through the *Braden* (1985) formula and the Curve Number method (*USDA-SCS*, 1972), respectively. Evaporation and Transpiration are computed using

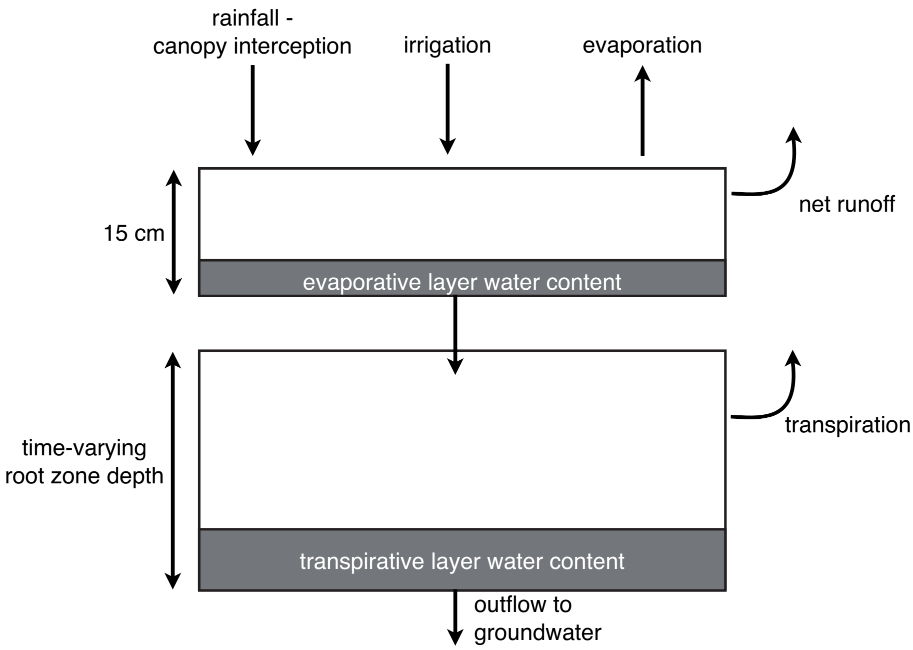

Figure 5.2: Representation of the soil volume along with the processes modeled for each cell of the distributed-parameter water balance module.

the FAO-56 dual crop coefficient method (*Allen et al.*, 1998). The former is defined as a function of the reference crop evapotranspiration $ET0_{t+1}$ and $\theta_{1,t}$, while the latter depends on $ET0_{t+1}$, $\theta_{2,t}$ and the basal coefficient K_{cb}, which is strongly related to the crop growth stage. Drainage discharges $Q_{p,t+1}$ and $Q_{g,t+1}$ are determined considering a simplified model based on a Darcian-type gravity flow in the unsaturated soil (*Gandolfi et al.*, 2006).

The heat units module defines the relationships between the temperature and some variables and parameters related to the crop growth stage (e.g., root length, basal coefficient, leaf area index), which also influence the water balance module. According to the heat units theory, crops growth stage at time t in the i-cell is defined as a function of the cumulated heat units ($HU_t^{(i)}$). A range is defined for each crop: the minimum is the base temperature T_b, which determines the day of sowing (i.e., when $HU_t^{(i)} > T_b$), and the maximum is the cutoff temperature, over which the heat units are no longer cumulated.

Finally, the crop growth module first estimates the maximum yield achievable in optimal conditions and, then, reduces it to take into account the stresses due to insufficient water supply from rainfall and irrigation experienced during the agricultural season. The yield response to water stresses is estimated in each cell i according to the empirical function proposed in the AquaCrop model (*Steduto et al.*, 2009; *Raes et al.*, 2009) and based on the approach proposed in the FAO Irrigation & Drainage Paper n. 33 (*Doorenbos et al.*, 1979):

$$1 - \frac{Y^{(i)}_{real}}{Y^{(i)}_{opt}} = k_y \left(1 - \frac{Tr^{(i)}_{real,tot}}{Tr0^{(i)}_{tot}}\right) \tag{5.4}$$

where Y_{real} and Y_{opt} are the actual and optimal yield, $Tr_{real,tot}$ and $Tr0_{tot}$ the actual and optimal transpiration during the whole growth period, and k_y is a crop-specific coefficient relating yield decline and water stress. In particular, the optimal yield is computed as the product of the crop optimal Harvest Index (HI_{opt}) and the optimal biomass produced at the end of the agricultural season ($B^{(i)}_{opt}$), assuming that no nutrient stresses happened:

$$Y^{(i)}_{opt} = HI_{opt} \cdot B^{(i)}_{opt} \tag{5.5}$$

where the biomass B_{opt} is a function of the crop Water Productivity, the transpiration Tr_{t+1}, the reference evapotranspiration $ET0_{t+1}$, and an adimensional stress coefficient taking into account the negative effects of cold conditions.

The dynamic model just introduced is used by the 66 irrigation units in the Muzza district to optimize the choice of the crop to grow by maximizing the expected economic revenue at the end of the year, as follows:

$$J^k_D = \pi \cdot \sum_{i=1}^{N_k} Y^{(i)}_{real} \qquad k = 1, \ldots, 66 \tag{5.6}$$

where π is the price of the cultivated crop, N_k is the number of cells belonging to the k-th irrigation unit, and $Y^{(i)}_{real}$ the actual yield in the i-th cell. The profitability of farmers crop choices is strongly influenced by crop prices (e.g., *Marques et al.*, 2005). However, a detailed description of market prices' dynamics goes beyond the scope of this work, as it would require to model local as well as global factors (e.g., *Kantanantha et al.*, 2010). Fixed crops prices are assumed as in *Paudel and Hatch* (2012), using the values published online by EUROSTAT.

Hydroclimatic Scenarios

In this work, the effectiveness of the co-adaptation strategy is evaluated in two hydroclimatic scenarios, namely current and projected conditions. The generation of the projected time series of hydroclimatic variables in climate change conditions is based on the application of a cascade of models: the worst emission scenario provided by the Intergovernmental Panel on Climate Change (*IPCC*, 2000) is used as input for a general circulation model (GCM), which provides the boundary conditions for a regional circulation model (RCM). In particular, the HadAM3H model (*Pope et al.*, 2000) is used as GCM and the RACMo model (*Lenderink et al.*, 2003) as RCM. Since the spatial resolution of RACMo is too rough to provide representative climatic scenarios at the basin scale, a statistical downscaling method based on quantile mapping (*Déqué*, 2007; *Boé et al.*, 2007) was applied to correct RCM outputs as in *Anghileri et al.* (2011).
This statistical downscaling method compares historical measured data with the outputs obtained from RACMo simulations over the historical (backcast) period (i.e., 1961-1990), in order to estimate a site-specific quantile-quantile correction function. By assuming that this relationship will not change in the future, it can be applied to the RACMo output over the projected (forecast) period (i.e., 2071-2100). The downscaled variables are then used as inputs for a standard rainfall-runoff model to obtain the projected time series of inflows. The data used for the downscaling were obtained from the PRUDENCE project (*Christensen and Christensen*, 2007). Figure 5.3 shows the difference of the downscaled hydroclimatic variables under projected climate change conditions with respect to their current values. Although the annual volume of water available (Figure 5.3a) does not significantly change, its different distribution during the year is expected to negatively impact on the agricultural water supply. The irrigation demand is indeed high in the late-spring and summer period, when projected water availability is low. Moreover, the projected climate change impacts show an increase of the temperature in the Muzza irrigation district (Figure 5.3b), producing higher evapotranspiration rates and, consequently, an increase of the irrigation water demand.

5.2.2 The co-adaptation problem

The proposed co-adaptation strategy aims to coordinate farmers and water managers adaptation options, see Sections 5.1.1-5.1.2. In this section, the two optimization problems associated to the coupled water supply

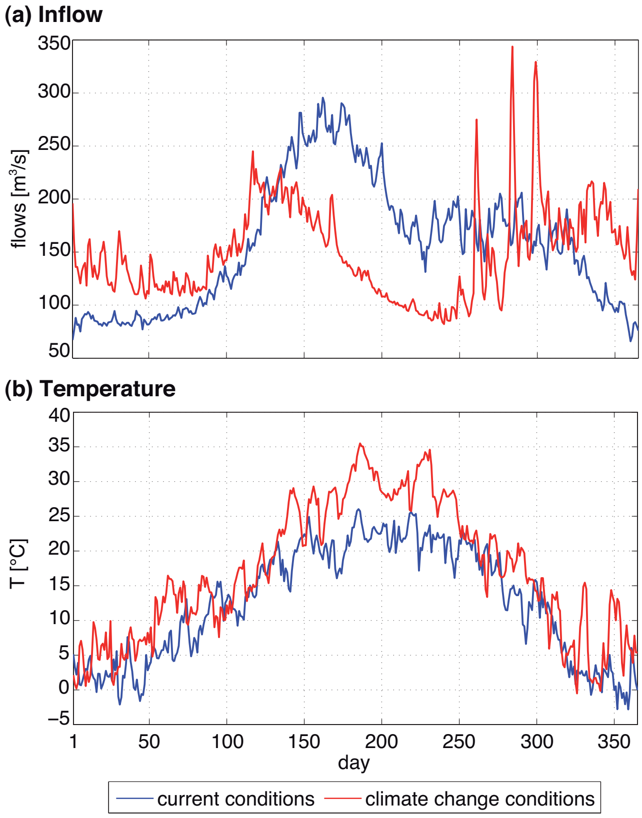

Figure 5.3: Effects of climate change on the mean inflows to the Lake Como (panel (a)) and mean temperature in the Muzza irrigation district (panel (b)).

and demand systems are formulated. The water supply manager is modeled as an active agent acting according to a daily operating policy p which, given the current storage of the lake s_t, provides the volume $u_t = m_t(s_t)$ to be released over the time interval $[t, t+1)$ (i.e., the next 24 hours). The optimal operating policy p^* is designed by formulating and solving a stochastic, periodic, non-linear, closed-loop optimal control problem (see *Castelletti et al.*, 2008a, and references therein) of a dynamic system which evolves according to the model defined in eq. (5.1):

$$p^* = \arg\min_{p} J_S(p, w) \tag{5.7}$$

in which $J_S(p, w)$ is the water supply objective as defined in eq. (5.2), which depends on the operating policy p and the irrigation demand w. Since the focus of this work is the irrigation supply, the problem associated to the water supply agent is a single-objective problem. The extension to multipurpose water reservoirs considering more objectives (e.g., flood protection, hydropower production or environmental preservation) is however straightforward. Among the set of optimization methods available to solve the management problem formulated in eq. (5.7), see Section 2.4, the optimal operating policy of the Lake Como is designed using stochastic dynamic programming (SDP, see *Bellman* (1957)), as it is the most adopted and accurate method for solving optimal control problems, offering performance guarantee and proof of convergence.
The water demand problem regards the pre-season decisions of the farmers on the crops pattern $\mathbf{u}_{crop}$. Each irrigation unit, which represents the decision-making authority in charge of selecting the crop to grow, is modeled as an active agent. The optimal crop pattern can be obtained by solving 66 non-linear optimization problems (one for each agent) based on the dynamic model of the Muzza district described in Section 5.2.1:

$$u_{crop}^{k*} = \arg\max_{u_{crop}^k} J_D^k(u_{crop}^k, p) \qquad k = 1, \dots, 66 \tag{5.8}$$

where $J_D^k(u_{crop}^k)$ is the objective function of the k-th agent as defined in eq. (5.6), which depends on the agent crop choice $u_{crop}^k \in \mathcal{U}_{crop}$ and the operating policy p adopted by the water supply agent. The set $\mathcal{U}_{crop}$ comprises five different crops representing the most commonly grown in the Pianura Padana agricultural system, namely tomato, grass, corn, soybean, and rice. On the basis of the solution of Problem (5.8) for each agent, representing the optimal crop pattern $\mathbf{u}_{crop}^*$ for the entire district, it is possible to derive the actual irrigation demand $w = w(\mathbf{u}_{crop}^*)$ that the water supply system has to satisfy. Due to the complexity of producing accurate long-term hydroclimatic forecast required by the model to solve Problem (5.8) at the beginning of the agricultural season, it is assumed that the farmers have a perfect forecast of the future hydroclimatic conditions. The resulting performance will therefore represent an upper-bound solution. The introduction of forecast errors may result in suboptimal farmers' decisions and performance degradation.
The aim of the proposed co-adaptation strategy is to cross-condition the decision-making problems of the agents (i.e., water manager and farmers), as shown in Figure 5.4. The procedure starts by deriving the optimal operating policy of the Lake Como p^0 on the basis of an a priori demand w^0. Then, the water supply system is simulated over a horizon

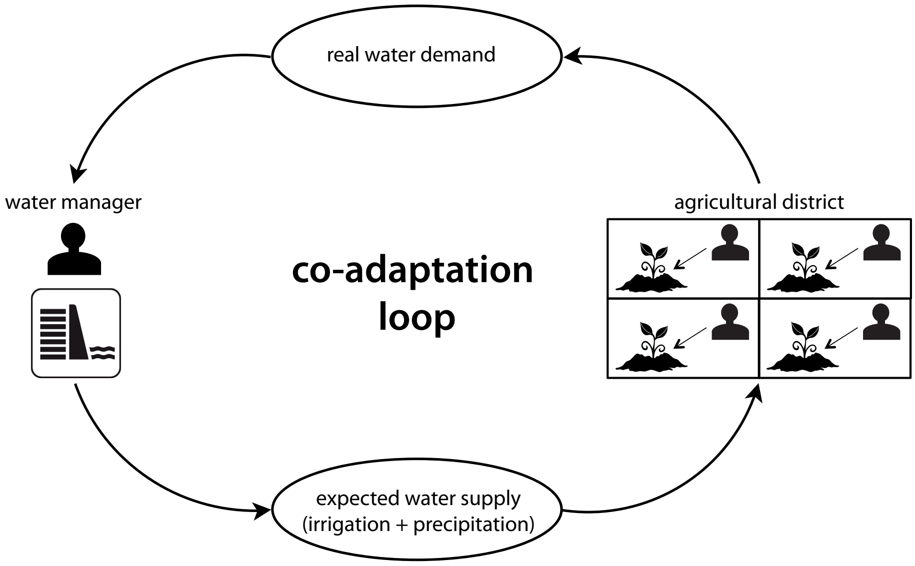

Figure 5.4: Illustration of the co-adaptation loop between water supply and demand.

of one year according to p^0 to obtain the trajectories of the expected releases and the water available for irrigation q^0_{irr}. This latter is used as input for the model of the water demand subsystem (i.e., the dynamic model of the Muzza district). This model allows the simulation of the irrigation water distribution, the computation of the hydrologic balance in the root zone, and the estimation of optimal and actual crop growth on a daily basis. For a given q^0_{irr}, the farmers optimize the choice of the best crop pattern $\mathbf{u}^*_{crop0}$ for the coming year (i.e., the one producing the highest revenue). Moreover, the model estimates the water requirements of the crops from which the real water demand of the whole district can be derived. On the basis of this new water demand (i.e., $w^1 = w^1(\mathbf{u}^*_{crop0})$), the release policy of the lake can be re-designed. The procedure is then iterated, with the optimization and simulation of the new policy p^1, the farmers' optimization $\mathbf{u}^*_{crop1}$ and the estimation of a new demand w^2. The iterations are stopped when the system converges to an equilibrium. It is assumed that convergence is obtained when the number of farmers' decisions changing between two consecutive iterations is lower than a desired threshold (i.e., 13 agents, corresponding to 20% of the Muzza irrigation units).

5.3 Results

The effectiveness of the proposed co-adaptation loop is tested in current and projected hydroclimatic scenarios, namely the years 1999 and 2099, which are assumed as representative of average conditions. Beside the hydroclimatic conditions, the co-adaptation and co-evolution of agricultural water management is influenced by a number of exogenous variables dependent on the socio-economic context. Since most of the agricultural districts in the Pianura Padana region has currently to grow mainly corn and grass for sustaining livestocks, a constraint is defined according to historical land use data to limit the agents' decisions and guarantee to grow corn and grass in a fixed number of cells of the Muzza district. More detailed analysis on the role of the socio-economic context are beyond the scope of this work and will be the subject of future research.

5.3.1 Co-adaptation in current conditions

In order to assess the effectiveness of the proposed co-adaptation loop, the performance of the current practices has to be estimated. On the supply side, without any interaction between the farmers and the water manager, the operation of Lake Como is designed by the water supply agent trying to minimize the water supply objective J_S, see eq. (5.2), according to the a priori trajectory of irrigation demand illustrated in Figure 5.5.

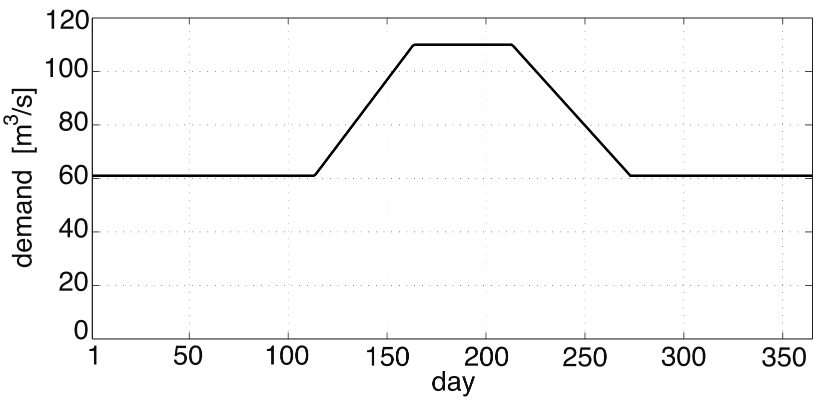

Figure 5.5: A priori irrigation demand of the Muzza district.

On the demand side, historical land use data show that the farmers in the Muzza district mainly grow corn and grass. A baseline alternative

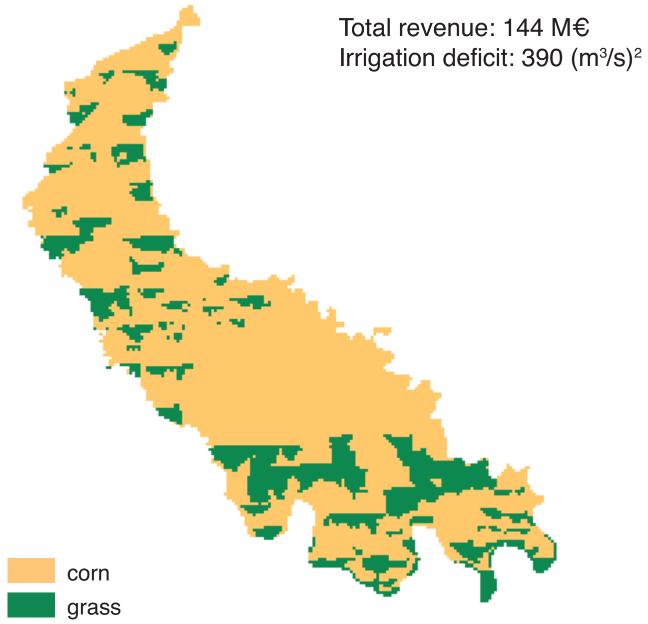

Figure 5.6: Baseline alternative: current crop pattern is defined by a mix of corn and grass.

(Figure 5.6) is therefore designed by assigning corn and grass to around 78% and 22% of the cultivated area, respectively. Under this alternative, the water supply objective (i.e., the daily average squared water deficit) is equal to 390 $(m^3/s)^2$ and the total economic revenue at the end of the agricultural season is equal to 144 million of €.
In the baseline alternative, both the water supply and demand practices represent an approximation of the current condition with no guarantee of being the optimal solution for the integrated agricultural water system as there is no interaction between agents' decisions. The adoption of the co-adaptation loop directly connects the decision-making problems of the agents to promote coordinated practices aimed to improve the system-level efficiency. The proposed co-adaptation strategy allows the selection by the farmers of the crop with the highest expected economic revenue on the basis of an expected water availability, which is dependent on the water supply management. Given the crop choices, the water supply agent can tune the management policy with respect to the actual water demand of the planted crops. Results show that, in current hydroclimatic conditions, tomato is selected by 59 agents, with only 7 agents deciding to grow rice (Figure 5.7). The estimated total economic revenue is equal

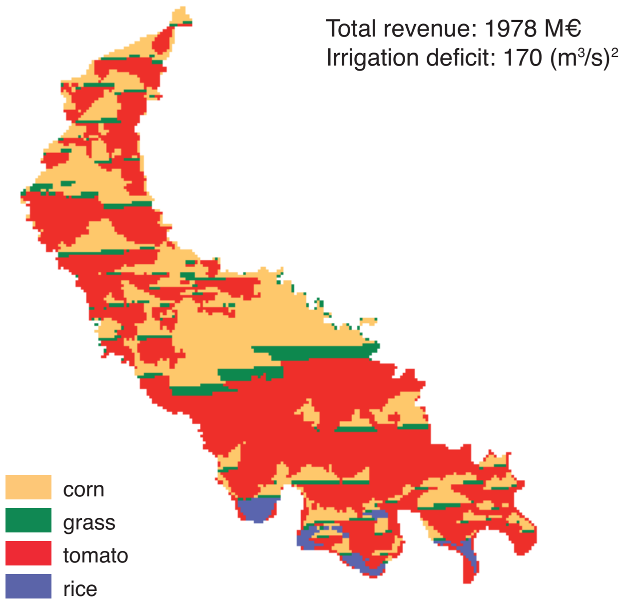

Figure 5.7: Co-adaptation alternative in current hydroclimatic conditions.

to 1,978 million of €, almost ten times higher than the one attained in the baseline alternative. This significant improvement is produced by two factors: *i*) most of the agents selects to grow tomato, which is the best crop in terms of productivity (i.e., 78,400 kg/ha) as well as the one characterized by the highest price (i.e., 0.6 €/kg); *ii*) the accurate tuning of the irrigation management by the water supply agent on the basis of the actual water demand of the selected crop, which produces a significant decrease of the irrigation deficit (i.e., from 390 $(m^3/s)^2$ to 170 $(m^3/s)^2$) and, consequently, of the water-related stress. It is worth noting that the low value of water deficit along with the agents' decisions to grow tomato, which is a crop characterized by a high water demand and particularly sensitive to water stresses (i.e., its parameter k_y describing the yield response to water stress in eq. (5.4) is equal to 1.35), suggests that the amount of water available in the current hydroclimatic conditions is not a limiting factor, especially assuming to regulate the lake for irrigation purposes only. However, this shift from corn and grass toward tomato will require significant costs for converting the farmers' equipments, with a largely reduced final profit for the farmers. The medium-long term profitability of this solution therefore requires further

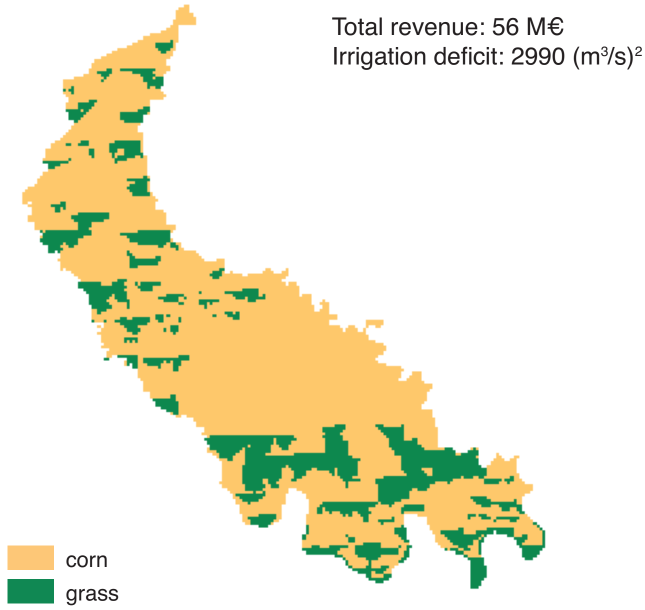

Figure 5.8: Baseline alternative in climate change conditions.

analysis, including the projected profitability of tomato over the next years affected by climate change conditions.

5.3.2 Co-adaptation in climate change conditions

The previous section shows the effectiveness of the proposed co-adaptation loop to enhance the efficiency of agricultural water management practices and foster crop production. The benefits achievable with the proposed co-adaptation strategy are expected to be more relevant in climate change conditions, characterized by less water available and increased temperature, which might negatively impact on a priori historical practices.

To estimate the potential worsening of the current situation, the baseline alternative (i.e., a mix of corn and grass, with a water supply policy defined with respect to the a priori irrigation demand) is re-evaluated via simulation over the projected hydroclimatic conditions in 2099. Results are reported in Figure 5.8. With no-adaptation of either the crop pattern or the water supply, the performance of the baseline significantly degrades. The total economic revenue reduces to 56 million of €, with an irrigation deficit equal to 2990 $(m^3/s)^2$. These results demonstrate

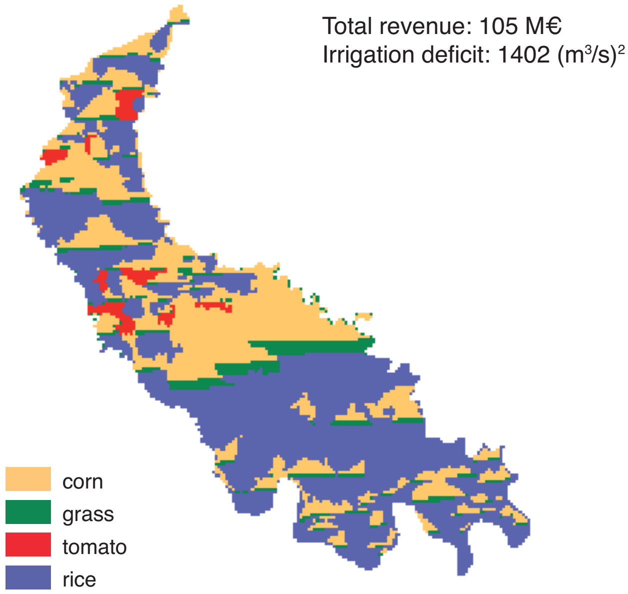

Figure 5.9: Co-adaptation alternative in climate change conditions.

that climate change impacts represent significant concerns for the Lake Como system and adaptation strategies are essential to sustain adequate agricultural productions.

The activation of the co-adaptation loop can contribute to mitigate the negative impacts of climate change on both the water supply and demand sides. The results of the co-adaptation alternative in projected conditions are represented in Figure 5.9. A clear trend is evident with respect to the co-adaptation alternative in current conditions, with most of the farmers selecting to grow rice (56 agents) instead of tomato (10 agents). This behavior can be explained by the lower sensitivity of rice to water stresses (i.e., its k_y parameter is equal to 0.6) with respect to that of tomato. This property allows high levels of production, yielding to a total revenue equal to 105 million of €, although a high irrigation deficit (i.e., 1402 $(\text{m}^3/\text{s})^2$). The results support the possibility to effectively adopt deficit irrigation strategies in future water scarcity conditions as suggested by *Zhang and Oweis* (1999), meaning that the supply of less water than the full irrigation requirements actually produces limited yield reductions. Moreover, deficit irrigation might be also unavoidable due to the expected decrease of water availability in the summer months combined with the projected increase of the tempera-

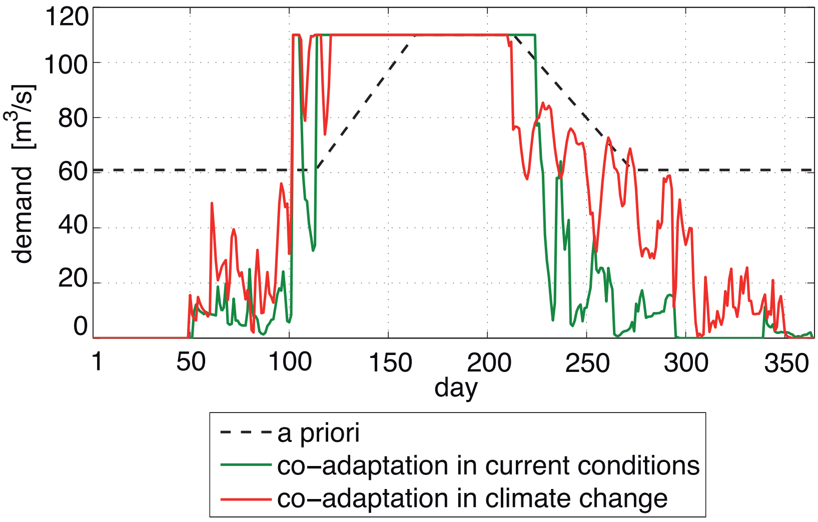

Figure 5.10: Comparison of the a priori historical demand of the Muzza irrigation district with the actual irrigation demands estimated with the co-adaptation alternatives in current and climate change conditions.

ture (Figure 5.3). Finally, warmer temperatures are expected to induce an extension of the agricultural season as illustrated in Figure 5.10, which shows a trend of sowing dates moved up in March and harvesting delayed in October and November.

5.4 Discussion and final remarks

In this chapter, the agent-based decision analytic framework is used to estimate the potential for a co-adaptation strategy of water supply and demand in the Lake Como system. The analysis aims to assess the benefit achievable in current as well as projected hydroclimatic conditions, affected by climate change impacts. Both the water manager regulating the Lake Como and the farmers selecting the crop pattern to grow are modeled as agents, sharing the same environment subject to the same changing conditions. The proposed co-adaptation strategy explicitly connects the decision-making problems of the agents by cross-conditioning their decisions: the farmers select the crop to grow as the one with the highest revenue on the basis of an expectation of irrigation water supply; knowing the selected crop pattern, the water manager optimizes the ma-

nagement of the irrigation supply system to match the actual demand of the crops.
Results show that it is possible to significantly foster the crop production by co-adapting the agents decisions, with an increase of the total revenue from 144 to 1,977 million of €. This gain is obtained by growing tomato (i.e., the crop with the highest productivity and price) and accurately supply water to limit water-related stresses. The irrigation deficit is equal to 170 $(m^3/s)^2$ with co-adaptation, significantly lower than the one in the baseline alternative (i.e., 390 $(m^3/s)^2$). However, the costs required for converting the farmers' equipments might significantly reduce the final profit in the short-term. Under climate change conditions, the performance of the baseline and of the co-adaptation alternative degrade. Without any form of adaptation, the current practices represented by the baseline alternative attain a total revenue equal to 56 million of €, representing a decrease of 60% with respect to the revenue in current conditions. By co-adapting the agents' decisions, the negative impacts of climate change are mitigated. Most of the farmers grows rice, obtaining a total revenue equal to 105 million of €, which corresponds to an increase of 87% with respect to the no-adaptation baseline. In this case, the conversion of farmers' equipment might have a lower impacts due to the long-term horizon of the investment.
The proposed co-adaptation strategy represents a promising step toward the enhancement of water resources exploitation in agricultural systems. However, many aspects of the proposed approach require further investigations. The results illustrated in this chapter are obtained by simulating the system over a single year for each hydroclimatic scenario, considered as representative of normal conditions, assuming a perfect-forecast of the hydroclimatic conditions in the coming year for the selection of the best crop to grow. The introduction of non-perfect forecasts and the repetition of the experiments over multiple years will allow the validation of the results over a broader set of uncertain conditions, depending on the natural hydroclimatic variability. Moreover, the study is focused on the management of the Lake Como for irrigation supply only. Potentially negative effects, for example in terms of flooding in the Como city, will have to be considered. The adoption of a multi-objective water supply management is expected to further degrade the performance in terms of agricultural production. The flexibility of the co-adaptation loop to conflicting water supply interests will be analyzed. Finally, the representation of the socio-economic system is very simple. Future efforts will concentrate on introducing dynamic crop prices and different scenarios of flexibility of the socio-economic framework.

6 Mechanisms design in complex river basin management problems

The purpose of this chapter is to apply the agent-based decision analytic framework (Section 1.3.2) to support water reservoir operations in managing growing water demands as well as hydroclimatic uncertainties. The procedure is demonstrated on the Conowingo reservoir, an interstate water body shared by Pennsylvania (PA) and Maryland (MD) in the Lower Susquehanna River, characterized by the presence of many conflicting stakeholders (agents). Currently that dam provides water supply to Chester (PA) and Baltimore (MD), cooling water for the Peach Bottom atomic power plant, and minimum regulated flows as defined by the Federal Energy Regulatory Commission (FERC) to protect fishery resources. In low flow conditions, FERC requirements tend to drawdown storage levels, increasing the conflict between the other stakeholders' objectives and reducing the recreational value (e.g., boating and fishing activities) of the system. This system represents a complex river basin management problem, where the decision-making process is not spatially distributed as in the other applications (i.e., there is one water reservoir). However, the adoption of the informative and decision-making supportive tools of the agent-based decision analytic framework allows the characterization of the complex interactions between the stakeholders affected by the Conowingo dam operation. The combination of reservoir policy identification and many-objective optimization under uncertainty (Sections 2.4.2) supported by visual analytics techniques (Section 2.2) successfully captures current reservoir operations and discovers key tradeoffs between alternative policies. Moreover, this chapter contributes a novel method, called direct policy conditioning, to design policy mechanisms for environmental protection through input variable selection techniques (Section 2.5).

Part of this chapter is adapted from a journal paper under review for publication in Water Resources Research (*Giuliani et al.*, 2013c) and a paper for the 2014 IFAC World Congress describing the direct policy conditioning method is under preparation (*Giuliani et al.*, 2014).

6.1 Policy inertia and myopia in water reservoir regulation

River basin management has traditionally been challenged by multiple competing water demands, including domestic and irrigation supply, flood protection, and hydropower production. Additional challenges arise with environmental regulations for flows, water quality targets, recreational interests, and energy markets (e.g., *Brown and Carriquiry*, 2007; *Fernandez et al.*, 2012), emphasizing the need to rethink the way freshwater resources are distributed, managed, and used (*Gleick*, 2002). Concerning water storage systems, such a paradigm shift is not easily achievable: the possibility of re-designing water reservoir regulation is strongly limited by historical agreements and regulatory constraints (*Fernandez et al.*, 2013). The limited flexibility of water laws, for example in the United States, creates policy inertia, where water institutions are highly unlikely to change their current practices in absence of a dramatic failure or water conflict (*Sheer*, 2010). Yet, no guarantee exists that historical management policies will not fail in coming years, especially as water managers face growing water demands and increasingly uncertain hydrologic regimes (*Milly et al.*, 2008). There is a significant need to better understand the consequences of current reservoir operations while discovering alternative policies that better balance competing objectives and performance uncertainties.
Prior studies in this area have often neglected the challenging realities of reservoir operations, assuming complete flexibility when designing optimal operation via optimization models (e.g., *Yeh*, 1985; *Labadie*, 2004; *Castelletti et al.*, 2008a, and references therein), and mostly focusing either on improving system-wide performance, by including hydroclimatic information to better condition the decisions (e.g., *Klemeš*, 1977; *Tejada-Guibert et al.*, 1995; *Hejazi et al.*, 2008), or on solving increasingly larger systems, by addressing the associated curse of dimensionality (e.g., *Cervellera et al.*, 2006; *Castelletti et al.*, 2010a, and references therein). Reservoir operators generally reject the validity of using optimization models to directly inform actual real-time operations, in particular when they include uncertainty explicitly (*Celeste and Billib*, 2009). Consequently, these tools are rarely employed in real operational contexts (*Teegavarapu and Simonovic*, 2001). Instead, operators prefer simpler tools, such as rule curves (*Loucks and Sigvaldason*, 1982; *Loucks et al.*, 2005), even though these tools are not able to adapt release decisions when the system deviates from the "normal" hydroclimatic conditions assumed in the design of the rule (*Maass et al.*, 1962; *Howard*, 1999).

The more uncertain the hydrologic system, the more frequent the deviations from the assumed baseline flow conditions and, accordingly, the lower the effectiveness of rule-curve-based operations. This is particularly critical given that rule curves are rarely redesigned to account for changing hydroclimatic conditions. Closing the loop between real operational decisions and evolving river conditions (e.g., *Soncini-Sessa et al.*, 2007a) will be key to effectively adapt to increasingly variable and extreme hydrologic conditions.
In addition to this inflexibility, traditional rule curves also suffer from myopia: they fail to explore the full set of tradeoffs between evolving multi-sector objectives and preferences in river basins. Most major reservoirs have had their rule curves defined in prior decades, where planning methods required strong a priori assumptions on the preferences (or priorities) of a representative, idealized decision maker across a limited number of operating objectives (*U.S. Army Corps of Engineers*, 1977). Just as a changing hydrological context poses a challenge, evolving objectives and preferences for reservoir operations can be another mode of failure for fixed rule curves. Although these issues have long been recognized, only recently have a posteriori methods coupled with visual analytics emerged to address them for complex engineered river basin systems, see Sections 2.1-2.2.

6.2 Problem formulation

6.2.1 The Lower Susquehanna case study

The Susquehanna River (Figure 6.1) is the longest river on the eastern United States, draining a catchment area of about 71,000 km^2 through New York, Pennsylvania, and Maryland, ultimately contributing 50% of the freshwater flowing into the Chesapeake Bay. The Conowingo reservoir is an interstate water body shared by Pennsylvania and Maryland in the Lower Susquehanna, about 16 km from the Susquehanna River mouth. The dam, which was completed in 1928 for hydropower generation purposes, is the largest non-federal dam in the U.S. regulating a large share of the flow in the Lower Susquehanna with substantial impacts on multiple stakeholders. The Conowingo reservoir contributes to the water supply of Chester (PA) and Baltimore (MD). Conowingo releases are also critical for cooling the Peach Bottom atomic power plant and downstream releases are subject to minimum flow requirements defined by the Federal Energy Regulatory Commission (FERC) to protect fishery resources. Moreover, in 1968 the reservoir was connected to the Muddy Run Pumped Storage Hydroelectric Facility, which cycles water

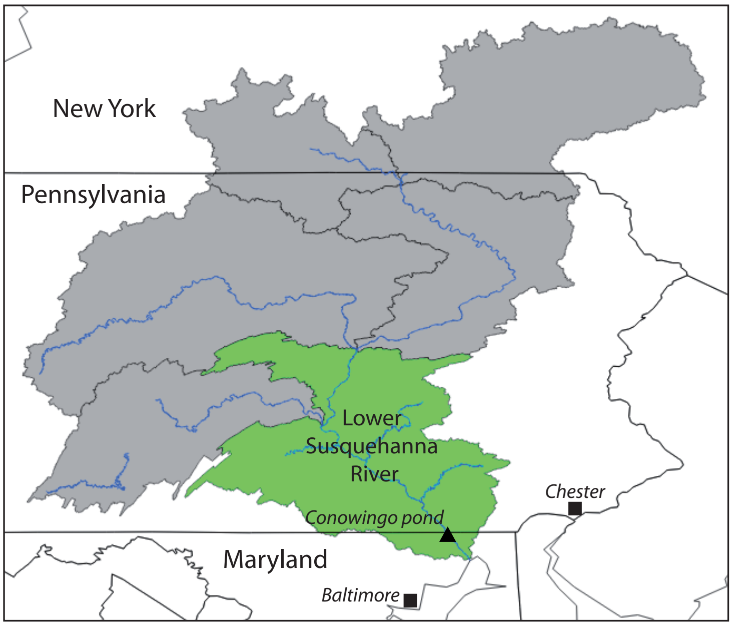

Figure 6.1: Map of the Susquehanna River basin.

back and forth from Conowingo for additional power generation. Finally, the Conowingo reservoir provides valuable recreational and ecosystem services.

The FERC minimum flow requirements introduced in 1988 protect fishery resources threatened by the hydropower management of the dam. The Conowingo reservoir is unique as a high valued river basin system that is being adaptively managed by the SRBC in collaboration with its core service constituencies (*Federal Energy Regulatory Commission*, 1989). In average flow conditions, water availability is generally sufficient to maintain hydroelectric operations, water supply, meet environmental flow requirements, and sustain recreational activities. Yet, in low flow conditions challenging tradeoffs emerge for Conowingo operations to supply water to Baltimore, Chester, and the Peach Bottom atomic power plant, while seeking to minimize negative impacts on the recreational and touristic interests. The normal level of the Conowingo reservoir along with the critical levels for water supply and the target level for recreation are reported in Table 6.1.

The interstate Susquehanna River Basin Commission (SRBC) actively coordinates conflicting water demands and water related interests between the basin's stakeholders. However, growing regional water de-

Table 6.1: Reference levels for the Conowingo reservoir (ft).

Normal level	108.5
Touristic weekend recreational level	106.5
Critical level for Peach Bottom atomic power plant	103.5
Critical level for Chester water supply	100.5
Critical level for Baltimore water supply	91.5

mands and climate change are significant concerns for the SRBC. As a recent example, the SRBC coordinated a regional planning effort assessing a set of alternative modifications to the FERC requirements to mitigate the negative impacts of the low reservoir levels (*Swartz*, 2006). The effort represents a substantial participatory negotiation process, with the SRBC promoting a direct involvement of the stakeholders to evaluate the different alternatives with the support of OASIS model simulations (*Randall et al.*, 1997; *Sheer and Dehoff*, 2009), a general purpose water resources model that uses a linear program solver to allocate water to meet multi-sector demands. Given the results of the modeled alternatives, the possibility of including the 800 cfs leakages from the closed dam gates toward meeting the downstream minimum flow requirements has been selected as the most critical action in managing the Conowingo dam during drought periods. The result of this intensive planning effort is the identification of alternative management strategies for implementing the credit for leakages and specifying the hydrologic conditions under which this credit is warranted. According to *Swartz* (2006), the most promising alternatives are summarized below:

- *Baseline*, representing the current situation where the Conowingo reservoir provides public water supply and the downstream releases are regulated by Exelon for hydropower production, subject to the downstream FERC minimum flow requirements (i.e., to release at least the maximum between the minimum environmental flow and the inflow registered at Marietta) without including the credit for the leakage;

- *Automatic Credit*, which proposes to automatically include the credit allowance in meeting the FERC minimum flow requirements. Yet, the credit is never available in April, May, or June, in order to protect fish migration.

- *Critical Level*, which proposes to allow the 800 cfs credit only when the elevation of the Conowingo reservoir drops below a critical level equal to 104.5 ft. That stage was selected because it guarantees operations at Peach Bottom and Muddy Run facility. The credit is never available in April, May, or June.

- *Minimum Flow*, which proposes to consider the FERC minimum flow requirements as absolute constraints independently from the flows at Marietta. The credit for the leakage is always counted in this scenario, but in April, May or June.

All of these alternatives have been designed to manage credit for leakages and the minimum flow requirements only, according to the historical agreements and the regulatory constraints. However, there is currently a limited understanding of the potential benefits achievable by significant modifications of the current Conowingo reservoir operating policy as well as the impacts of uncertainties in the basin's hydrologic regimes. This chapter builds on the adaptive management efforts of the SRBC and contributes a set of candidate policies that could aid the dam operator in balancing its multi-sector demands and hydroclimatic uncertainty. The study is focused on the identification and refinement of the operating policy for the Conowingo dam, while a fixed weakly rule is assumed for the operation of Muddy Run according to the hydropeaking strategy reported in *Swartz* (2006) and illustrated in Figure 6.2.
The model of the system (Figure 6.3) is mainly based on the representation of the dynamics of the two water reservoirs defined by the mass balance equations of the water volume s_t^i stored in each reservoir (i=Conowingo, Muddy Run):

$$\begin{aligned} s_{t+1}^{CO} &= s_t^{CO} + q_{t+1}^{CO} - r_{t+1}^{CO} - E_{t+1}^{CO} - q_{t+1}^{p} + r_{t+1}^{MR} \\ s_{t+1}^{MR} &= s_t^{MR} + q_{t+1}^{MR} - r_{t+1}^{MR} - E_{t+1}^{MR} + q_{t+1}^{p} \end{aligned} \tag{6.1}$$

where q_{t+1}^{CO} is the inflow to the Conowingo reservoir in the interval $[t, t+1)$, which is composed by the flow measured at Marietta gauging station and the lateral contribution between Marietta and the reservoir and q_{t+1}^{MR} is the inflow to Muddy Run. The release r_{t+1}^i is given by the release function $r_{t+1}^i = f(s_t^i, u_t^i, q_{t+1}^i)$, where u_t^i is the release decision and $r_{t+1}^i(\cdot)$ is a non-linear function describing the stochastic relation between the decision u_t^i and the actual release r_{t+1}^i (*Piccardi and Soncini-Sessa*, 1991). The Conowingo release r_{t+1}^{CO} is composed by four different releases, one

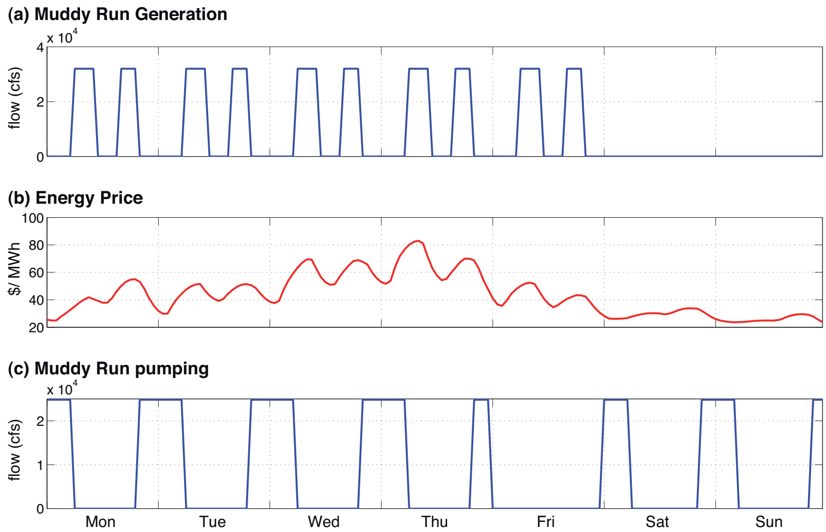

Figure 6.2: Weakly rule defining Muddy Run facility hydropeaking operation.

for each reservoir outlet (i.e., the atomic power plant, Baltimore and Chester water supply, and the downstream releases connected to the Conowingo hydropower plant). The water pumped from Conowingo to Muddy Run is represented by q^p_{t+1}. Finally, E^i_{t+1} is the loss for evaporation in the two reservoirs. The decision time step is equal to 4 hours to balance the need of following the hourly dynamics of the energy prices and the specification of a time step sufficiently long to not be impacted by the mechanics of turbines operation (e.g., cavitation, ramp up).
The multiple stakeholders (agents) interests affected by the Conowingo dam operation are modeled using the following six objectives, computed over a simulation horizon H:

- *Hydropower revenue*: the annual economic revenue from energy production at the Conowingo hydropower plant (to be maximized) defined in eq. (6.2) as the product of the hourly energy production HP_t (MWh) and the hourly energy price ρ_t (US$/MWh). The hydropower production is defined as $HP_t = \left(\eta g \gamma_w \bar{h}_t q^{Turb}_t\right) \cdot 10^{-6}$, where η is the turbine efficiency, $g = 9.81$ (m/s^2) the gravitational acceleration, $\gamma_w = 1000$ (kg/m^3) the water density, $\bar{h}_t$ (m) the net hydraulic head (i.e., reservoir level minus tailwater level), q^{Turb}_t (m^3/s) the turbined flow. According to *Exelon* (2010), the energy prices are defined by the 7-hour moving average of the historical

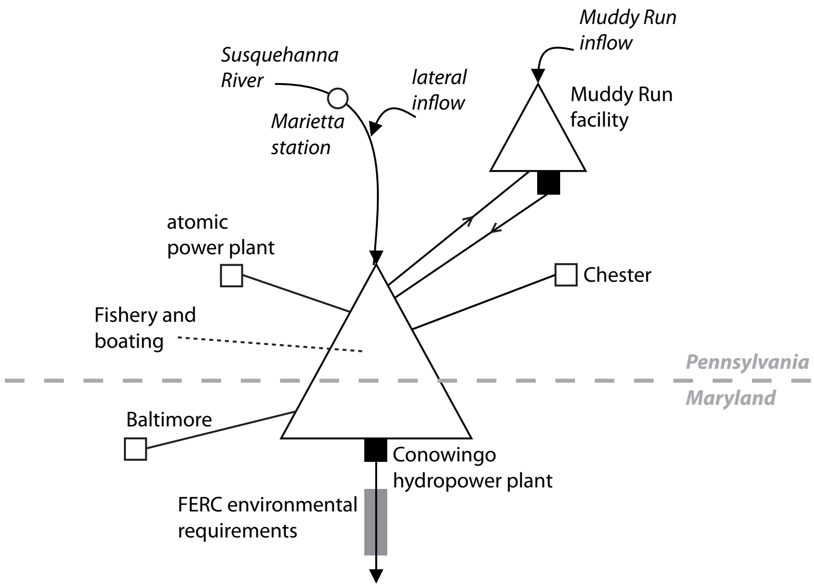

Figure 6.3: Schematic representation of the main components of the Lower Susquehanna model.

energy price trajectory in PJM energy market;

$$J^{hyd} = \sum_{t=1}^{H} (HP_t \cdot \rho_t) \tag{6.2}$$

- *Water supply* to Baltimore, Chester, and the Atomic Power Plant: the daily average volumetric reliability (to be maximized) defined as

$$J^{VR,i} = \frac{1}{H} \sum_{t=1}^{H} \left(Y_t^i / D_t^i \right) \tag{6.3}$$

where Y_t^i (m^3) is the daily delivery, D_t^i (m^3) is the corresponding demand (Figure 6.4a-c), and i = (Baltimore, Chester, Atomic Power Plant);

- *Recreation*: the storage reliability (to be maximized) in the weekends of the touristic season (i.e., from Memorial Day to Labor Day), defined as

$$J^{SR} = 1 - \frac{n_F}{2N_{we}} \tag{6.4}$$

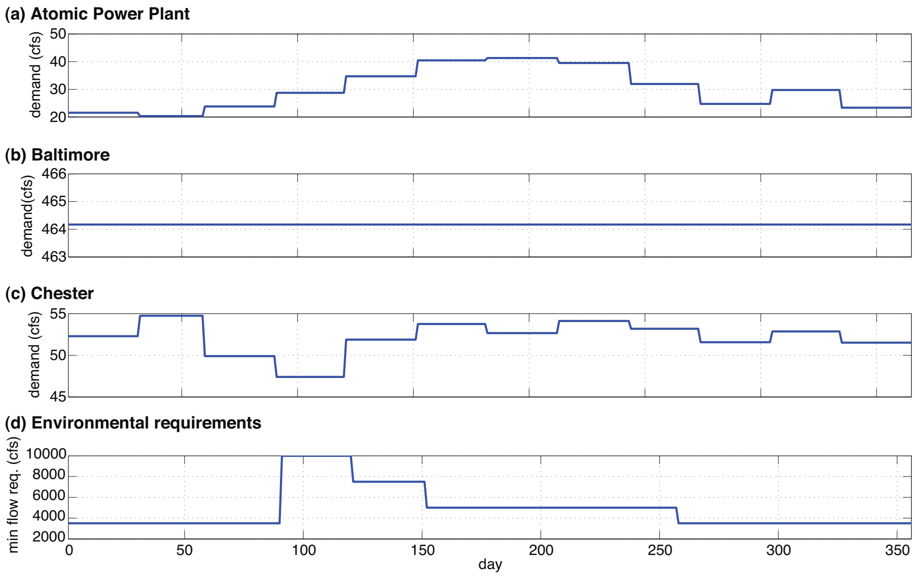

Figure 6.4: Public water supply demands for the Peach Bottom atomic power plant (a), Baltimore (b), and Chester (c), and the FERC minimum flow requirements (d).

where n_F is the number of days during which the reservoir level is below the target level of 106.5 ft (which guarantees boating and recreational activities) and N_{we} is the number of weekends in the touristic season;

- *Environment*: the daily average shortage index with respect to the FERC minimum flow requirements (to be minimized), defined as

$$J^{SI} = \frac{1}{H} \sum_{t=1}^{H} \left(\frac{\max(Z_t - Y_t, 0)}{Z_t} \right)^2 \tag{6.5}$$

where Y_t (m^3) is the daily release and Z_t (m^3) is the corresponding FERC flow requirement (Figure 6.4d). The quadratic formulation aims to penalyze severe deficits in a single time step, while allowing for more frequent, small shortages (*Hashimoto et al.*, 1982).

6.2.2 Many-objective policy identification and refinement under uncertainty

In this chapter, the first goal of the agent-based decision analytic framework is to provide support to the SRBC in re-designing the operation of

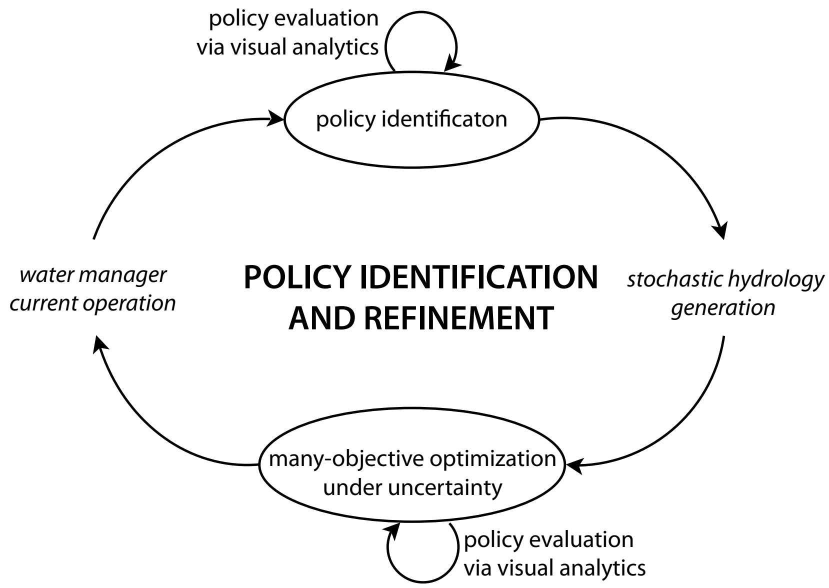

Figure 6.5: Illustration of the proposed procedure combining reservoir policy identification, policy refinement with many-objective optimization under uncertainty, and visual analytics.

the Conowingo reservoir using a combination of reservoir policy identification, policy refinement with many-objective optimization under uncertainty, and visual analytics according to the procedure shown in Figure 6.5.

First, the current baseline operating policy is identified in the form of a radial basis functions (RBFs) policy, mapping relevant information into release decisions. According to the implicit policy identification approach (Section 2.4.2), the dam operator is modeled as a rational agent seeking to maximize primary operating objectives (i.e., guaranteeing the public water supply and maximizing the hydropower revenue), subject to the FERC minimum environmental flow requirements. The baseline policy is hence defined as a multi-input (i.e., time and reservoir level) single-output function (i.e., downstream release decision) with four RBFs, accounting for 20 parameters. The three water supply withdrawals are set equal to the corresponding demands. This hypothesis means that they are always satisfied if the level in the reservoir is sufficiently high to activate the corresponding outlets. The quality of the identified baseline policy is validated by its ability to replicate historical release dynamics (i.e., the flows measured at the USGS gauging station downstream of the Conowingo dam). Subsequently, this baseline policy is refined via many-objective optimization, and the associated tradeoffs visually explored. The Pareto-optimal policies are defined as multi-input multi-output functions with four RBFs (accounting for 32 parameters), which provide the four release decisions, corresponding to the downstream release as well as the ones for the public water supply, as a function of time and reservoir level. Visual analytics plays a key role in the proposed framework by allowing to comparatively analyze the current policy in the context of the full tradeoffs surface as well as in the corresponding operating policy decision space. Two different formulations of the Lower Susquehanna River management problem are considered. The first formulation, which will be termed the historical formulation, is defined in eq. (6.6), where the operating objectives (see eqs. (6.2)-(6.5)) are evaluated over the historical realization of the hydroclimatic variables, namely inflows and evaporation rates.

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathbf{J}(\boldsymbol{\theta}) \tag{6.6}$$

To assess the vulnerability of the solutions to hydroclimatic uncertainties, in the second formulation, which will be termed stochastic formulation, the same objectives are instead evaluated over an ensemble Ξ of stochastic inflows and evaporation rates realizations generated with the K-nearest neighbor resampling methods described in Section 2.4.2. Figure 6.6 illustrates the annual flow duration curves for both the historical flows (1930–2001) as well as the stochastic ensemble at the Marietta gauging station. Although the stochastic ensemble directly models the autocorrelation and variability within historical record, the generated equally plausible water years clearly cover a far broader range of hydroclimatic conditions. This is especially true for the low flow conditions that have the most critical impact on the Conowingo dam's operations. The uncertainty introduced by the stochastic ensemble is then filtered adopting a minimax approach (eq. (6.7)) which minimizes the objectives in the worst-case realization. This approach identifies robust operating policies able to guarantee certain performance. Good solutions must indeed robustly perform for rapidly increasing numbers of Monte Carlo samples during the search process because new, independent samples are used to evaluate the objectives in successive iterations of the algorithm search. If a solution survives to the final generation, it has already been evaluated for a rapidly increasing number of realizations based on its ability to survive and propagate in the search population (*Miller and Goldberg*, 1996).

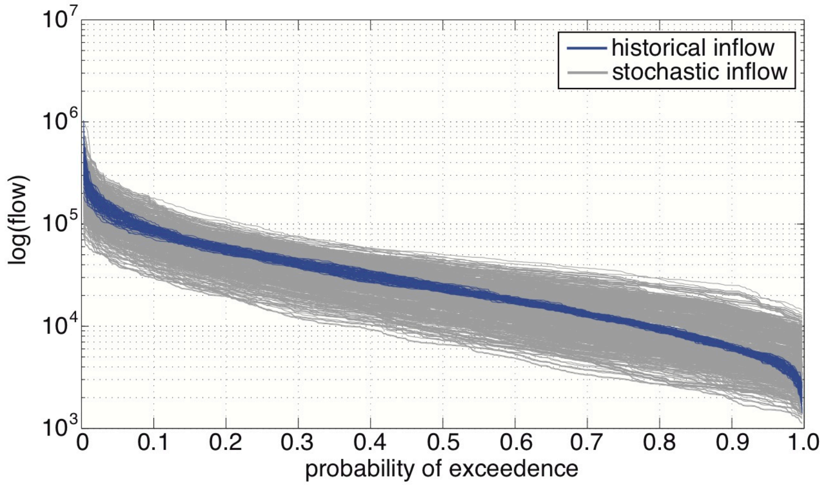

Figure 6.6: Annual Flow Duration Curves of the flows at Marietta gauging station. The historical records (1930–2001) are in blue, the generated stochastic ensemble in gray.

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \max_{\Xi} \mathbf{J}(\boldsymbol{\theta}) \tag{6.7}$$

The proposed policy identification and refinement procedure employs the Borg MOEA to optimize the operating policies via direct policy search (DPS, see Section 2.4.2). The default algorithm parametrization suggested by *Hadka and Reed* (2013) is used as the Borg MOEA has been demonstrated to be relatively insensitive to the choice of parameters, showing a high probability of attaining successful search if the algorithm is run for a sufficient number of iterations (*Hadka and Reed*, 2012; *Reed et al.*, 2013). Epsilon dominance is used to set the resolution of the operating objectives. In this work, epsilon values are equal to 0.5 for hydropower revenue, 0.05 for Baltimore, Chester, and Atomic Power Plant volumetric reliability, 0.05 for recreational storage reliability, 0.001 for the environmental shortage index. The computational requirements for this study were dominated by the optimization under stochastic hydroclimatic conditions. In the stochastic optimization, each function evaluation performed by the Borg MOEA comprises 50 Monte Carlo simulations over a one year horizon. The stochastic optimization was run for 1 million function evaluations. To improve solution diversity and avoid dependence on randomness, the solution set from each formulation is the result of 30 random optimization trials (i.e., 30 seeds with

50 million simulations each yields 1.5 billion simulations in total). The final Pareto-optimal policies are obtained as the set of nondominated solutions identified from the results of all the optimization trials.
The stochastic optimization was performed on the Texas Advanced Computing Center (TACC) Stampede Cluster. The 6,400 nodes of the TACC Stampede system each contain two Intel Xeon E5 processors and one Intel Xeon Phi Coprocessor, for a total of 102,400 processing cores. Each optimization run was parallelized to be run on 4096 processing cores simultaneously. In total, approximately 200,000 computing hours were required to complete the study, ensuring the best possible approximation to the Pareto-optimal solution set within the limits of computational tractability. It should be noted that the computational experiment is more rigorous than would be necessary in practice. Parallel search is used to maximize the exploration of the problem's decision space while minimizing the time required to attain the search results. The Borg MOEA was robust in solving the problem with far fewer function evaluations (100,000-200,000) and showed limited variability across multiple random search trials.

6.2.3 The agent-based model

The policy identification and refinement procedure described in the previous section aims to support the design of improved operating policies able to address many competing demands and system uncertainties. Problem (6.6), and its stochastic extension (eq. 6.7), assumes the point of view of the SRBC and considers all the stakeholders (agents) affected by the Conowingo dam operation as equally important. An objective function is indeed defined for each agent (Section 6.2.1), yielding to the formulation of six-objective management problems. Yet, it is very unlikely that primary objectives such as the public water supply have the same strategic role of the recreational interests. The adoption of an agent-based perspective allows a more realistic representation of the actual decision-making structure of the Lower Susquehanna problem. In the agent-based model illustrated in Figure 6.7, each stakeholder is modeled by one agent, but they are differentiated in active agents (shown in blue) and passive agents (shown in green). Four active agents represent the primary objectives, namely public water supply to Chester and Baltimore, atomic power plant cooling, and hydropower production. These active agents can be also associated to the four controlled releases of the Conowingo reservoir. Two passive agents represent instead the Environment and Recreation objectives. Although the environmental protection is a significant concern for the SRBC, it is modeled by a passive agent

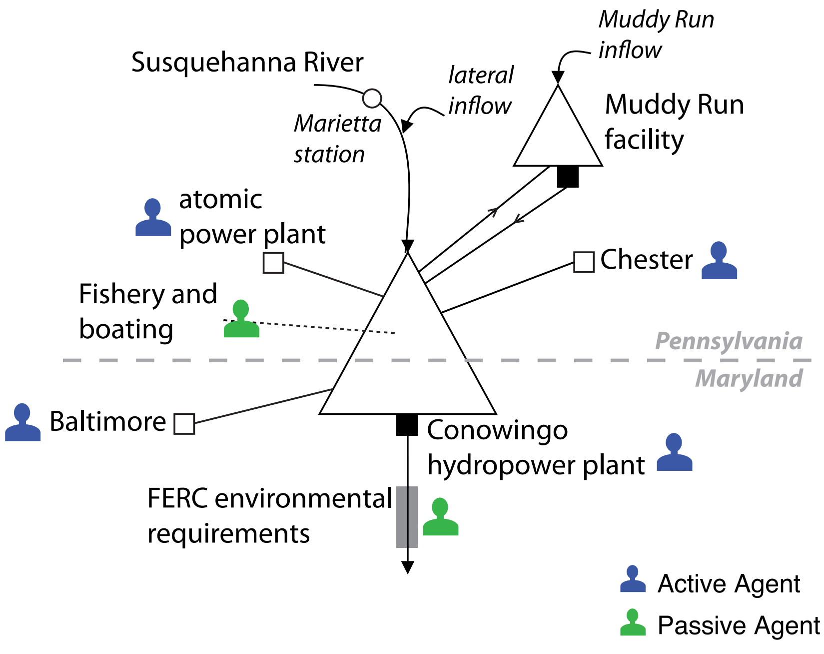

Figure 6.7: Agent-based model of the Lower Susquehanna system.

as the FERC requirements introduced in 1988 represents an important example of mechanism design where a watershed authority, namely the SRBC, tries to protect the environment threatened by the operation of the system, originally driven only by the primary objectives. Also the alternatives recently negotiated by SRBC and described in Section 6.2.1 are alternative mechanism design options to attain a better system-level performance by constraining the Conowingo reservoir operation.

In this chapter, the second goal of the agent-based decision analytic framework is to identify effective mechanism design strategies through the application of the direct policy conditioning method (Section 2.5). The stochastic Pareto-optimal policies designed with the 6-objective problem formulation provide a reference for the best achievable performance. At the other extreme, a lower bound solution is also considered as representative of the conditions before 1988, with no flow requirements. This solution is designed by solving a four-objective problem which considers only the interests of the four active agents. Finally, the current policy and the SRBC negotiated alternatives represent empirical mechanism design solutions to be contrasted with the solution obtained with DPC.

6.3 Results

6.3.1 Identification of the baseline alternative

In order to discover operating policies that could improve the management of the Lower Susquehanna, it is pivotal to accurately model the dynamics and preferences currently guiding the operation of the Conowingo Dam. The current Conowingo operations are identified using the implicit policy identification approach (Section 2.4.2) and validated by running a simulation over historical hydroclimatic conditions (i.e., inflows and evaporation rates) and comparing the resulting releases with respect to the flows measured at the USGS gauging station downstream of the Conowingo dam (USGS gauge 01578310).
The results of the policy identification are reported in Figure 6.8. Figure 6.8a shows the trajectory of the daily Conowingo reservoir level in 1999, with the attained values varying between the minimum and maximum elevations of 101.2 ft and 110.2 ft. This range of storage levels fall within the feasible limits imposed on Conowingo Dam. Year 1999 was selected as it represents a highly challenging dry period, where operations in the system are actively managing the tradeoffs for Conowingo in low flow conditions. The simulated and observed trajectories of releases along with the cumulative releases are shown in Figures 6.8b and 6.8c, respectively. Both figures show that the implemented implicit modeling approach effectively captures the historical operation of the system. Note that the over-estimation of the peaks (spills) in Figure 6.8b can be explained by the 4-hour time step of the model which keeps the spillways open longer than in reality. Figure 6.8d shows the release decisions for different reservoir levels according to the estimated policy: the concave shape serves to maximize hydropower production, which depends on the turbined flow and the net hydraulic head (i.e., reservoir level minus tailwater level), with higher releases in the hours with higher energy prices.

The baseline policy identified and illustrated in Figure 6.8 provides a highly flexible tool for contrasting how the Conowingo Dam's current operations perform relative to alternative operating policies designed via many-objective policy refinement. Figure 6.9 compares the baseline performance with the 6-objective historical Pareto-optimal policies, where Recreation, Atomic Power Plant, and Environment are plotted on the primary axes, with the black arrows identifying the directions of increasing preference. The orientation of the cones represents the reliability of meeting Chester's water supply demands, with the best solutions represented by upward cones. The size of the cones is proportional to reliability of meeting Baltimore's water supply demands, with the best solutions

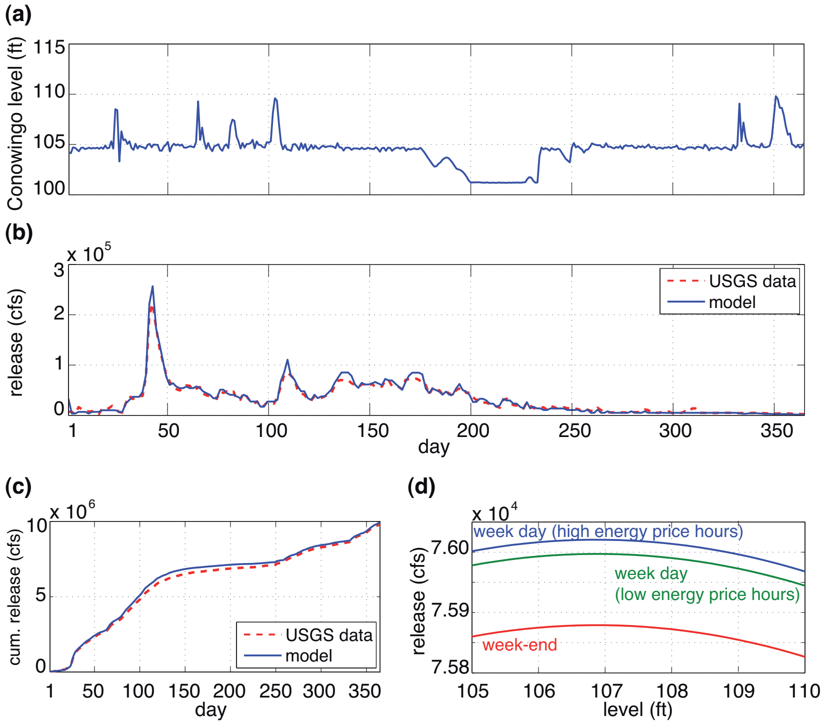

Figure 6.8: Trajectories of Conowingo reservoir level in 1999 under the estimated baseline policy (panel (a)). Comparison of releases and cumulated releases in 1999 (panels (b)-(c), respectively) obtained via simulation of the estimated baseline policy with the ones measured downstream of Conowingo dam. Representation of the estimated baseline policy (panel (d)).

represented by the largest cones. Finally, the hydropower revenue is represented by the color of the cones where maximum revenues are red. So in the figure, the ideal solution is a large red cone oriented upward near the ideal point designated as the black circle in the bottom-right corner of the figure. The baseline policy is identified by the boxed cone. This policy is very good in terms of hydropower production, with a revenue of 79 million US$, and water supply to Baltimore and Chester, with volumetric reliability equal to 1.0 in 1999. It also demonstrates good performance in terms of reliably meeting the Environment objective (i.e., the FERC minimum environmental flow constraint), while it struggles to reliably provide water for cooling the atomic power plant attaining a volumetric reliability of 0.85. The Peach Bottom atomic power plant has the high-

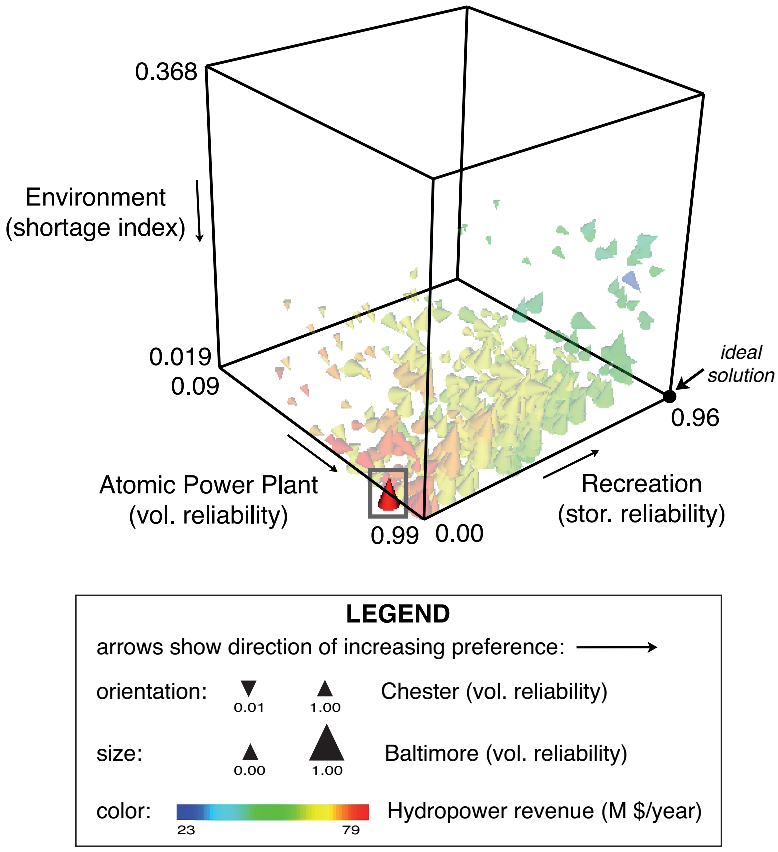

Figure 6.9: Comparison of the performance over historical hydroclimatic conditions of the baseline alternative and the historical Pareto-optimal policies.

est intake (Table 6.1) and therefore suffers water shortages when the reservoir level decreases (see the trajectory in Figure 6.8a). Finally, the baseline policy has a very poor performance in terms of Recreation, with a storage reliability equal to 0.0. The results shown in Figure 6.9 demonstrate that is possible to exploit publicly available historical streamflow observations to discover the implicit structure of preferences driving the historical reservoir operation: hydropower revenue, water supply, and low flow environmental concerns are most strongly emphasized in the baseline operating policy for Conowingo. This can occur because either these concerns are easily satisfied or they are strongly shaping management preferences (or both).

Figure 6.9 also shows the optimized policies that compose the deter-

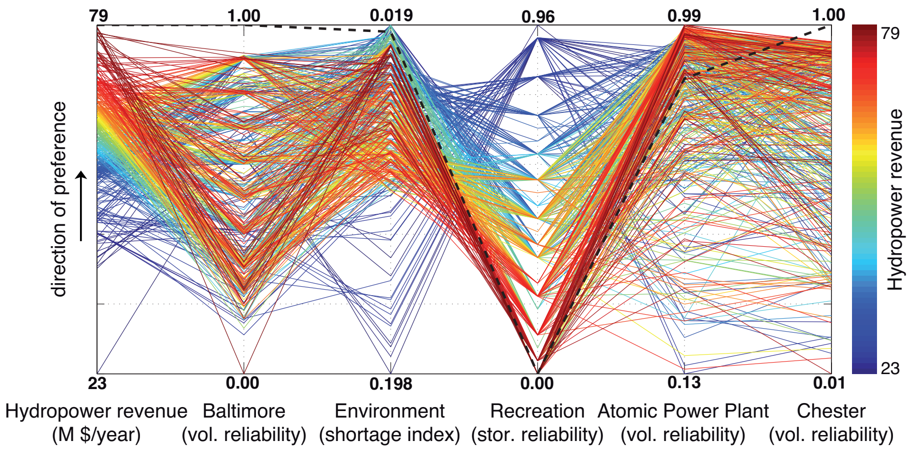

Figure 6.10: Parallel axes plot comparing the performance over historical hydroclimatic conditions of the baseline alternative and the historical Pareto-optimal policies.

ministic historical formulation's Pareto-optimal set. The objective calculations in these results are based on the historical realization of the hydroclimatic variables. The current operating policy performs very well in most objectives relative to the Pareto-optimal alternatives except for the Recreation and Atomic Power Plant objectives. Although Recreation may be viewed as a less critical concern, the reduced reliability of providing cooling water to the Peach Bottom atomic power plant is likely a far more critical concern. However, an increase in this meeting cooling water needs will negatively impact on the reliability of the water supply to Baltimore and Chester (i.e., a strong tradeoff between these objectives). Figure 6.10 presents a parallel axes plot (*Inselberg*, 1997) to serve as another visual tool for understanding key interacting tradeoffs for the Lower Susquehanna. This parallel axes plot representation shows each solution as a line crossing the six axes, representing the six objectives, at the values of their corresponding performance. In the plot, the objective values are normalized between their minimum and maximum values and the axes are oriented so that the direction of preference is always upward. Consequently, the ideal solution would be a horizontal line running along the tops of all of the axes. The conflicts are designated as diagonal lines between two adjacent axes. Figure 6.10 shows clear tradeoffs, especially when seeking to maximize the hydropower revenue represented by red solutions. Attaining high reliability for the atomic power plant cooling water supply strongly conflicts with contributing to Baltimore's water

supply and maintaining sufficiently high reservoir levels for recreation. Baltimore's water supply contributions from Conowingo Dam also face a strong conflict with meeting the Environment objective (or FERC regulations), probably because Baltimore has the highest public water demand (Figure 6.4).
Overall, when evaluated using solely the observed historical record for the Lower Susquehanna system, Conowingo Dam's current baseline policy effectively addresses several of the system's primary operating objectives and raises some concerns about reliably providing cooling to the Peach Bottom atomic power plant. However, these results are evaluated over a single realization of historical hydroclimatic conditions and are reflective of the current approach used to model and manage the basin. A key question still remains unanswered by these results. Are the history-based reliabilities for the multi-sector services provided by Conowingo overconfident and biased by neglecting hydroclimatic uncertainties?

6.3.2 Policy performance over stochastic hydroclimatic conditions

In order to reply to the above question, all of the alternatives illustrated in Figures 6.9-6.10 are re-evaluated via simulation over an ensemble of 50 stochastic hydroclimatic realizations, with the objective values computed according to the minimax approach defined in eq. (6.7). The comparison between the performance differences between the historical and stochastic conditions is presented in Figure 6.11. Note that this performance evaluation is actually biased toward allowing the history-based to maintain high levels of performance given that the re-evaluations use only 50 hydroclimatic scenarios versus the full 10,000 realizations illustrated in Figure 6.6. Consequently, degradations in performance are of significant concern. Figure 6.11 illustrates how the stochastic re-evaluation degrades the prior results. The performance of the baseline solution is significantly worse under stochastic conditions, with substantial degradation in hydropower revenue (from 79 to 39 million US$), environmental shortage index (from 0.023 to 0.106), and reliability of the atomic power plant supply (from 0.85 to 0.63). On the other hand, it maintains high reliability for both Baltimore and Chester, while the Recreation reliability remains equal to zero. The performance of the historical Pareto-optimal policies also strongly degrade when moving to the stochastic simulation. The parallel axes plot compares the historical evaluation (green) with the stochastic re-evaluation (red). Significant reductions in performance occur for hydropower revenue (with the best solutions degrading from 79 to 53 million US$), environment (with the shortage index of the worst

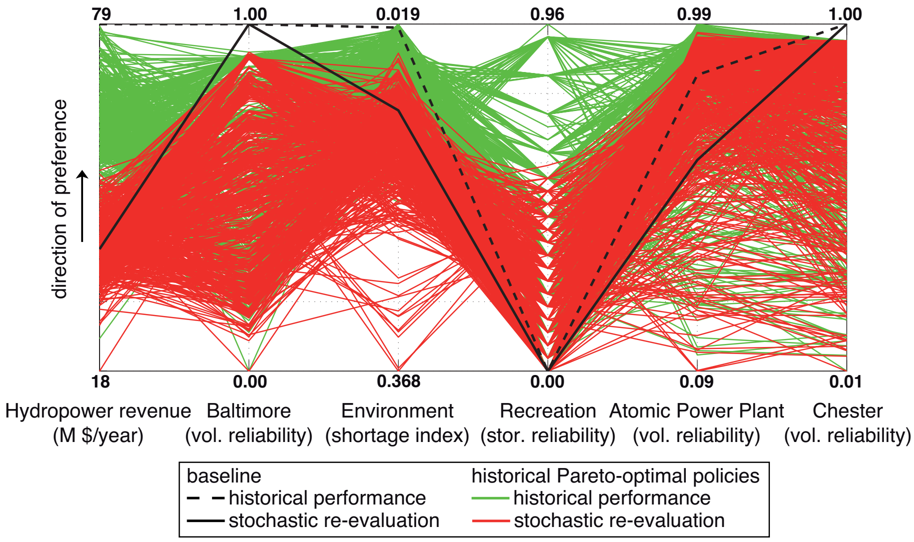

Figure 6.11: Comparison of the historical performance of the baseline policy and the historical Pareto-optimal solutions and their re-evaluation over an ensemble of 50 stochastic realizations.

solutions increasing from 0.198 to 0.368), and recreation (with the highest reliability decreasing from 0.96 to 0.57). These results clearly show that the intrinsic uncertainties in the natural processes strongly impact the policies' performance, with some objectives more sensitive than others to hydroclimatic variability. The analysis over a single realization of historical hydroclimatic conditions is therefore weak, indicating that stochasticity must be explicitly considered in the design of effective water management strategies for the Lower Susquehanna.

6.3.3 Stochastic Pareto-optimal policies

Growing water demands and low flow conditions are significant concerns for the SRBC that shaped recent adaptive management efforts to identify a set of potential modifications to the current baseline operation seeking to better balance the multi-stakeholders demands within the Lower Susquehanna (as summarized in Section 6.2.1). Figure 6.12 compares the relative performance of the baseline policy and these modified alternatives via simulation over the same 50 stochastic hydroclimatic scenarios discussed above. The proposed alternatives do offer significantly different performance across the six objectives illustrated in the parallel axes plot. The critical level alternative (red) has the same performance in

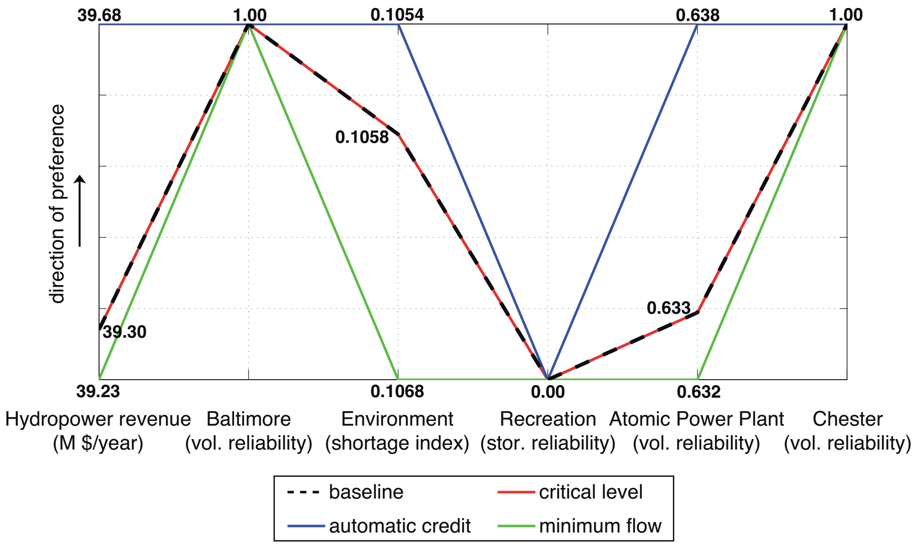

Figure 6.12: Performance of the alternatives negotiated by the SRBC over an ensemble of 50 stochastic hydroclimatic realizations.

all objectives as the baseline solution (dashed black line). Interestingly, the minimum flow alternative (green) produces results that are counter to the goal of the negotiated agreements, obtaining lower performance than under the baseline operation in many objectives and demonstrating that the Conowingo reservoir is unable to meet the environmental requirements in drought conditions. Finally, the automatic credit alternative (blue) Pareto-dominates the baseline solution, making it the most promising alternative. Overall, the results of Figure 6.12 are consistent with the those attained in the SRBC facilitated negotiation (*Swartz*, 2006).

However, this set of solutions has been identified by focusing solely on how to manage the credit for conveyance leakages. Figure 6.13 illustrates the effects of broadening the scope of the analysis to compare the performance of this set of alternatives with those achievable via policy refinement by using many-objective optimization under uncertainty. It shows the 6-dimensional objective space, with the baseline and the three alternatives proposed by the SRBC represented by the opaque boxed cones, while the stochastic Pareto-optimal policies are shown with transparency. Figure 6.13 suggests that the alternatives proposed by the SRBC are essentially equivalent to the baseline policy in this broader scoped problem formulation (i.e., the cones are almost overlapped). These results illus-

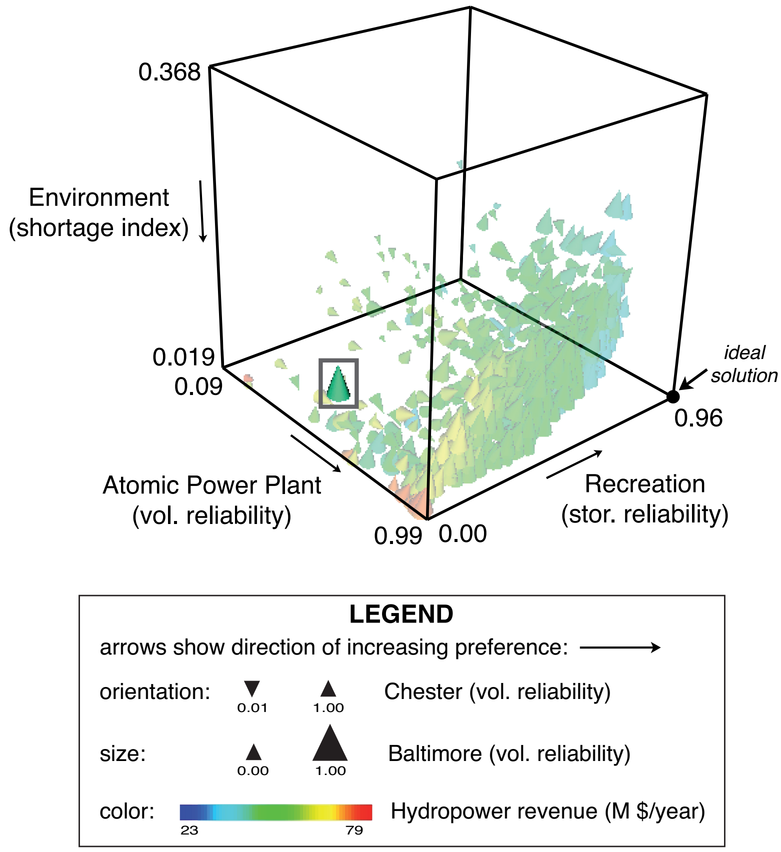

Figure 6.13: Comparison of the performance over stochastic hydroclimatic conditions of the baseline policy, the alternatives negotiated by the SRBC, and the stochastic Pareto-optimal policies.

trate how the policy inertia (i.e., the resistance to changing operating policies) can induce policy myopia. Moreover, Figure 6.13 shows the potentiality of an a posteriori decision-making approach for providing a broader contextual understanding of objective tradeoffs and alternative reservoir policies. The analysis over stochastic hydroclimatic conditions clarifies the potential refinement of the baseline policy, especially in those objectives which are more sensitive to system uncertainties. The performance of the baseline solution can be significantly improved in terms of Hydropower, Atomic Power Plant, Environment, and Recreation at the cost of a small reduction in the reliability for Baltimore and Chester, which have the option to obtain water from other sources. Depending

on the SRBC structure of preference, the hydropower revenue can be increased from 39 to 74 million US$, the reliability of the atomic power plant supply from 0.63 to 0.97, the storage reliability from 0 to 0.85, and the environmental shortage index can be reduced from 0.106 to 0.023.
To better understand the stochastic Pareto-optimal policies and provide useful information for the SRBC and the stakeholders (agents) involved in the problem, the dynamic behavior of the Lower Susquehanna system under these alternative regulations is analyzed. Figure 6.14 shows the trajectories of downstream release, reservoir level, and atomic power plant supply in 1999. The historical trajectory is shown as a thick black line and the stochastic Pareto-optimal solutions are colored for the respective Hydropower revenue, Environment, Recreation, and Atomic Power Plant objectives in panels (a), (b), (c), and (d) respectively. Figure 6.14a shows the downstream release trajectories (in logarithmic scale). The main difference between the current baseline operations for the Conowingo Dam and the high-revenue trajectories (red) is the latter's high releases during the summer, when the reservoir level is generally low and higher releases are needed to maximize energy production. The ability to sustain summer releases is also critical for maintaining high levels of performance for the Environment objective as shown by the red trajectories in Figure 6.14b. Here, the red policies are able to provide higher releases in the summer by allowing small and short duration deficits with respect to the minimum flow requirements (dotted black line) as induced by the quadratic formulation of the shortage index objective (*Hashimoto et al.*, 1982). These results highlight that FERC regulations may strongly reduce the sustainability of the basin's multi-sector services. The adoption of RBFs policies purposefully avoided overly constraining the problem formulations to attain a broad scope of operating policies and their consequent tradeoffs for the Conowingo Dam. Figure 6.14c shows the Conowingo reservoir level trajectories. A clear pattern is evident in summer (the recreation objective is formulated with respect to the touristic season only) with the red policies generating periodic peaks during the weekends. The baseline policy produces consistently draws down and consequently performs very poorly in terms of Recreation. Finally, Figure 6.14d represents the atomic power plant supply trajectories. The red trajectories, especially in summer, are slightly less than the water demand and this conservative strategy avoids the reservoir level drawdown obtained with the baseline regulation (see panel (c)), thus ensuring the possibility of using the outlet located at 103.5 ft. for a longer period.
To illustrate how the agent-based proposed decision analytic framework could be exploited to provide direct recommendations to the SRBC, the

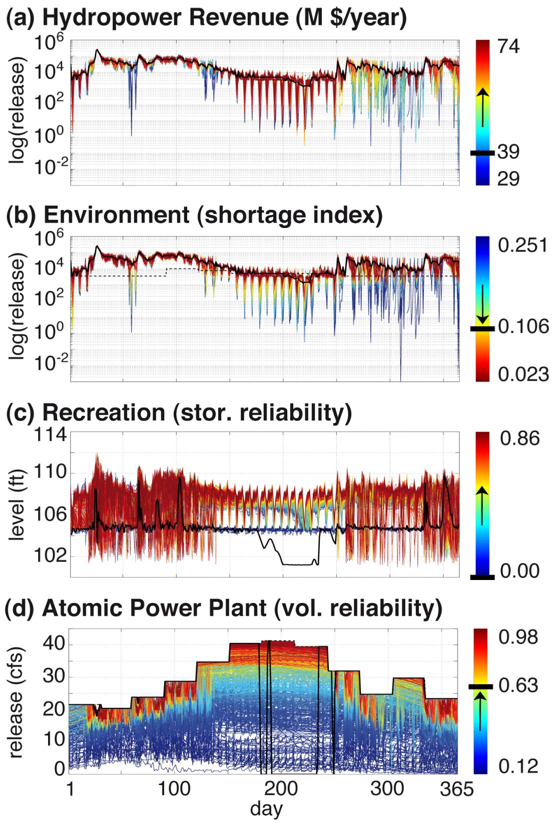

Figure 6.14: Comparison of the trajectories of the downstream release, level, and atomic power plant supply under the baseline alternative (thick black line) and the stochastic Pareto-optimal policies over historical hydrology.

analysis of the trajectories shown in Figure 6.14 is coupled with the investigation of their corresponding reservoir policies as illustrated in Figure 6.15. In particular, since the policies are time-varying, the analysis is focused on the shape of the summer operating rules, in a week day for hydropower and in a weekend for the other objectives. As shown in Figure 6.14, the summer is the most critical period of the year, when most of the challenging tradeoffs emerges. Each line therefore defines the release decisions as a function of the reservoir water levels (for a fixed time instant). Most of the stochastic Pareto-optimal policies representing the downstream releases (Figure 6.15a-c) is more conservative and releases less water than the baseline alternative except for high water

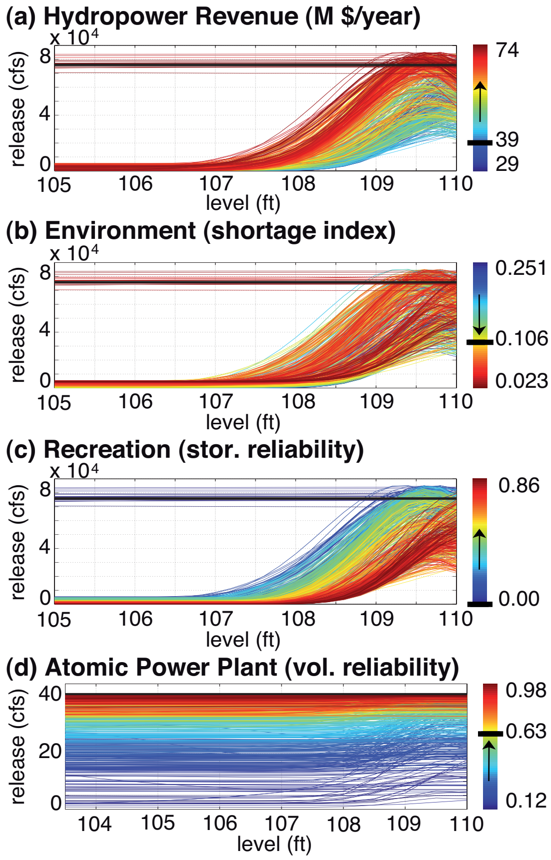

Figure 6.15: Comparison of the baseline alternative (thick black line) with the stochastic Pareto-optimal policies colored with respect to their respective performance.

levels conditions, thus saving water to face droughts. Note the inversion of the colors between panel (a) and panel (c), confirming the strong conflict existing between Hydropower and Recreation. The best policies for this latter (red lines in panel (c)) do not release when the level is below 108 ft to maximize the storage reliability, but they strongly reduce the corresponding hydropower revenue (blue lines in panel (a)). Conversely the best policies in panel (a) are poorly performing in panel (c), with the baseline policies, which is similar to blue stochastic Pareto-optimal solutions, that attains a storage reliability equal to zero. The policies for the atomic power plant release (Figure 6.15d) are more flat. Again, the best policies (dark red) are more conservative that the baseline and releases slightly less than the water demand. This strategy results to be particularly effective in uncertain conditions, and allows significantly improvements in the performance of the baseline policy marked on the colorbars. Moreover, Figure 6.15 shows the full range of options available to the SRBC to achieve a desired level of performance on the multi-objective tradeoff surface.

6.3.4 SRBC recommendations

To further illustrate how the SRBC could refine the current operation of Conowingo Dam, the management preferences identified on the baseline policy are used to eliminate (or "brush out") Pareto-optimal policies that fail to meet the SRBC requirements. The full stochastic Pareto-optimal set contains 1490 solutions. This set of solutions provides a rich context for understanding complex management tradeoffs and dynamics. Figure 6.16 illustrates how to obtain a smaller subset of interesting candidates policies to improve upon the Conowingo Dam's current baseline operations. The following criteria are applied to select the policies shown in Figure 6.16: Hydropower revenue $\geq$ 60 million US\$/year, Atomic Power Plant reliability $\geq$ 0.90, Baltimore reliability $\geq$ 0.85, Chester reliability $\geq$ 0.90, and environmental shortage index $\geq$ 0.10. The underlying idea is to select solutions that outperform the baseline policy, as well as the other alternatives negotiated by the SRBC, at the cost of a small reduction in the reliability for Baltimore and Chester. All the selected alternatives represent potentially interesting compromise solutions which effectively balance the competing multi-sector demands in the Lower Susquehanna system according to the preference structure associated to the baseline operation, which considers Recreation as a secondary objective. Among the set of policies selected in Figure 6.16, one policy is selected to be further analyzed.
Results show that adopting the recommended solution instead of maintaining the current reservoir regulation, the SRBC would potentially attain an increase in the hydropower revenue equal to 22 million US\$/year, 0.07 in recreational storage reliability, 0.29 in the Atomic Power Plant reliability, and a reduction of 0.05 in terms of environmental shortage index. These values correspond to a relative improvement of 56% in Hydropower, 47% in Environment, and 46% in Atomic Power Plant. The increase in Recreation is limited (i.e., from 0 to 0.07) as the current preference structure does not prioritize this objective. However, as shown in Figure 6.15c, there exists a large opportunity for obtaining higher storage reliability by adopting other policies with different balances of the objectives. In general, the recommended solution exhibits the potential

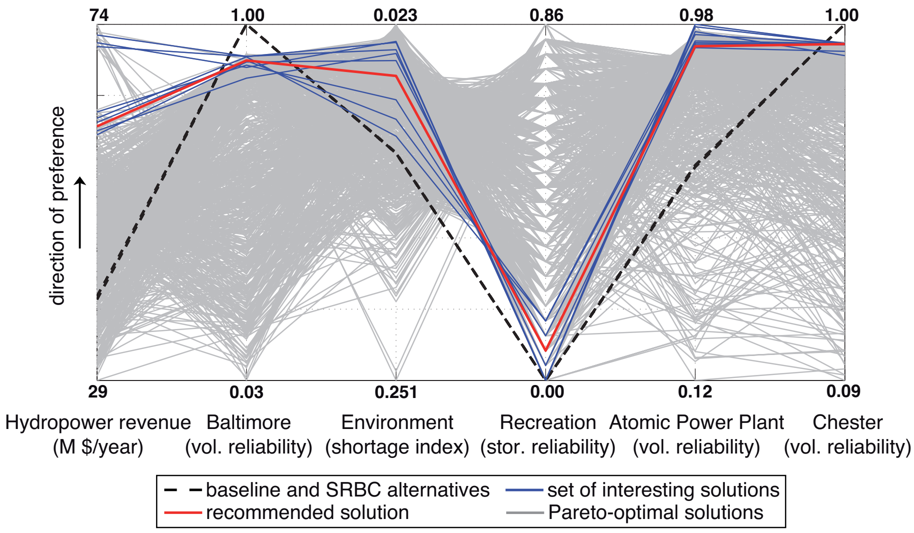

Figure 6.16: Identification of a set of potential candidate policies that might replace the baseline. The criteria adopted are the following: Hydropower revenue ≥ 60 million US\$/year, Atomic Power Plant reliability ≥ 0.90, Baltimore reliability ≥ 0.85, Chester reliability ≥ 0.90, environmental shortage index ≥ 0.10

to significantly outperform the baseline regulation with careful modification of summer releases, producing a policy that is more robust to hydroclimatic uncertainties and also better addresses the tradeoffs across the Conowingo Dam's multiple stakeholders objectives.

6.3.5 Direct policy conditioning

Although the recommendations proposed in the previous section significantly improve the performance of the baseline policy in addressing system uncertainties and objective tradeoffs, the recommended policies are designed by assuming the SRBC effectively controls the Lower Susquehanna system and represents all the stakeholders (agents) involved. The many-objective optimization problem indeed considers six objectives, one for each agent. The stochastic Pareto-optimal policies therefore represents the best achievable performance at the system-level, possibly filtered as in Section 6.3.4, but they do not consider the actual decision-making structure of the Lower Susquehanna problem, which can be better represented by the agent-based model described in Section 6.2.3,

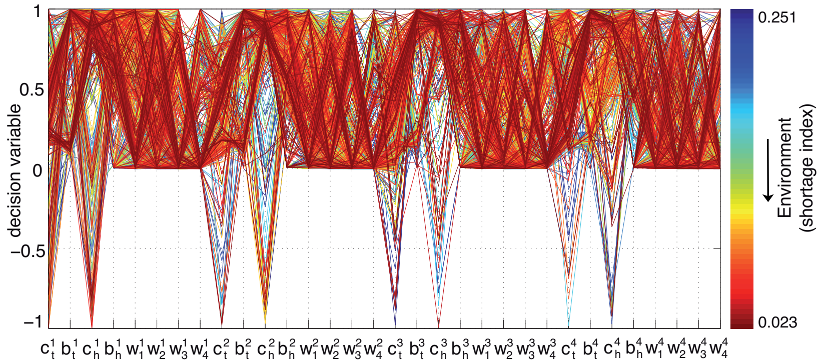

Figure 6.17: Representation of the 32 parameters values defining the stochastic Pareto-optimal policies.

comprising four active agents (i.e., hydropower and public water supplies) and two passive agents (i.e., environment and recreation). The adoption of direct policy conditioning (DPC, see Section 2.5) provides potential mechanism design strategies to better constrain the individualistic policy which reflect only the interests of the active agents in order to protect the environment. The performance in terms of environmental shortage index under stochastic hydroclimatic conditions (Figure 6.13) is indeed a significant concern under the baseline policy as well as the alternatives negotiated by the SRBC, which are all based on the FERC minimum flow requirements.
The decision space of the Lower Susquehanna management problem is shown in Figure 6.17, where each axis represents one parameter of the stochastic Pareto-optimal policies. Each line represents one RBFs operating policy, colored for the respective performance in the Environment objective. Given the complexity of this 32-dimensional decision space, it is difficult to recognize an environmental signature in the values of the parameters, meaning a clear range for each parameter associated to a good performance in terms of environmental shortage index. The adoption of input variable selection techniques aims to support the identification of the subset of parameters that are more related to the Environment objective as described next. The iterative input selection (IIS) algorithm (Section 2.5) is used in this work.

Iterative input selection

In order to identify the most significant decision variables with respect to the Environment objective, the IIS algorithm is first run working

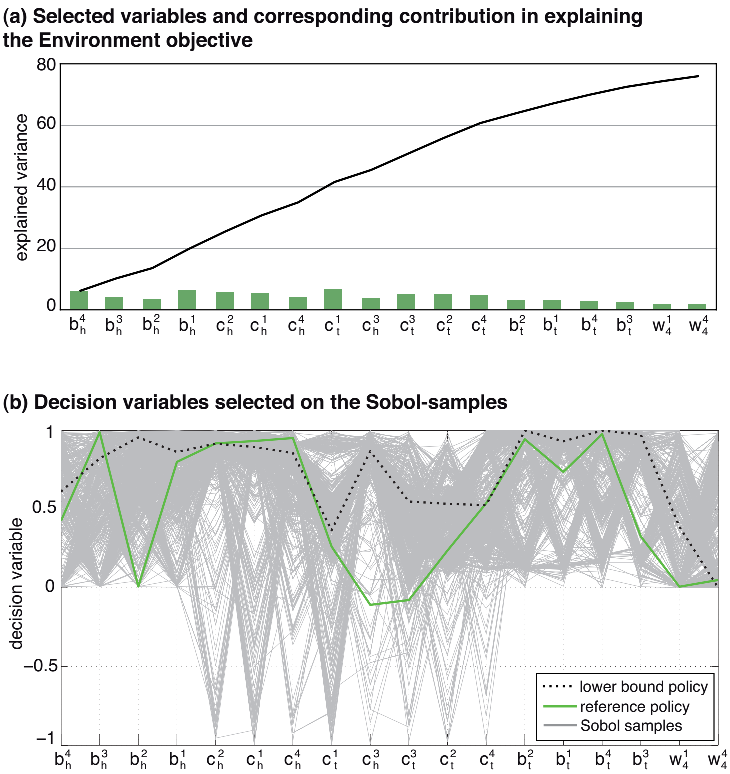

Figure 6.18: Representation of the IIS results on the Sobol samples. Panel (a) shows the set of selected variables and their contributions in explaining the Environment objective. Panel (b) represents the values of the selected decision variables for different policies

with a dataset generated with a design of experiments that samples the RBFs parameter space using the quasi-random Sobol sequence (*Sobol*, 2001) coupled with the cross-sampling method proposed by *Saltelli et al.* (2008). In this work, 10,000 parameter sets are generated from the Sobol sequence, with the cross-sampling method creating a total of 660,000 parametrizations of the RBFs policies. Each policy is then simulated to obtain the corresponding performance. The Environment objective is selected as output to explain, while the candidate input for IIS are the parameters of the RBFs policies. Figure 6.18a shows the 18 parameters selected by the IIS algorithm and their contribution in explaining the output (i.e., Environment objective) variance. Not surprisingly, IIS selects all centers and radii of the RBFs, which are strongly affecting the

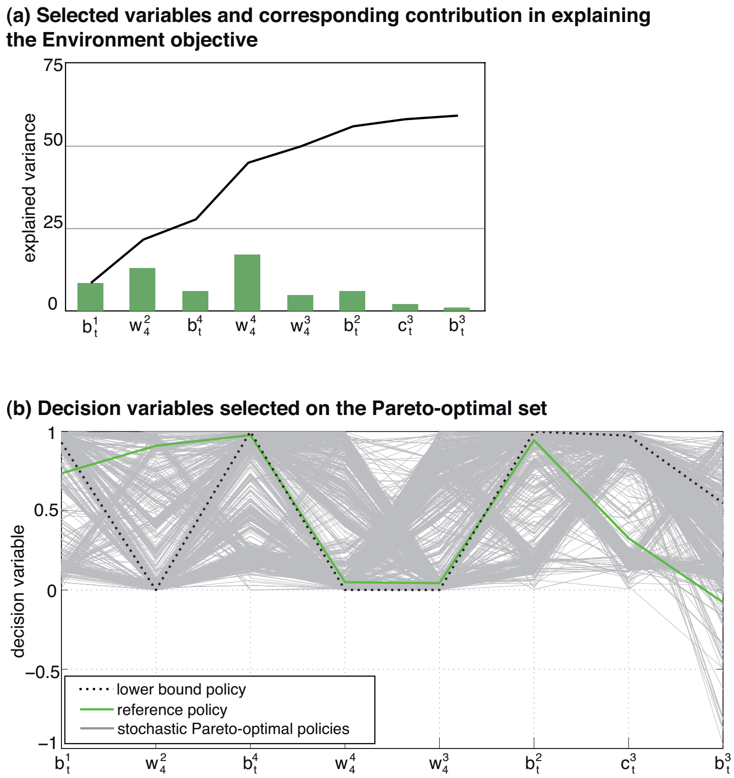

Figure 6.19: Representation of the IIS results on the Pareto-optimal set. Panel (a) shows the selected variables and their contributions in explaining the Environment objective. Panel (b) represents the selected decision variables for different policies.

policy shape and, consequently, the release decisions impacting on the Environment objective. Figure 6.18b shows the values of the selected parameters in the reference policy (i.e., the one with the best performance in terms of environmental shortage index) with respect to the sampled parameter sets. Moreover, the RBFs parameters in the lower bound policy (i.e., the one considering only the active agents objectives) are shown to highlight the main differences with respect to the reference policy.

The IIS algorithm is also run working directly on the parametrizations of the Pareto-optimal policies. The underlying idea is that the information content of the Pareto-optimal set might be more informative than a randomly generated dataset in conditioning the operating policies. Figure 6.19a shows the parameters selected and their contribution in explaining the Environment objective. In this case, less parameters are selected (i.e., only 8), with a lower explained variance (i.e., 60% with respect to 75%). Figure 6.19b compares the values of the selected parameters in the reference and lower bound policies with respect to the entire stochastic Pareto-optimal set. Note that most of the selected parameters are RBFs radii related to the time variable (i.e., b_t^i), while no parameters related to the Conowingo level are selected. These results suggest that the time of the releases is the most significant factor affecting the Environment objective, while the level importance decreases when moving from random parametrizations to the stochastic Pareto-optimal set.

Direct policy conditioning policies

In this section, the effectiveness of DPC in developing mechanism design strategies which preserve the actual decision-making structure and better protect the passive agents objectives (i.e., the Environment) is demonstrated. The historical formulation of the Lower Susquehanna DPS problem given in eq. (6.6) is conditioned by reducing the active agents decision space. The decision variables (i.e., the RBFs parameters) selected by the IIS algorithm are constrained to be close to the values assumed in the reference policy with the best performance in terms of environmental shortage index under stochastic hydroclimatic conditions. Figure 6.20 reports the comparison of DPC policies (shown opaque) with respect to the baseline and the SRBC negotiated alternatives (represented by the almost overlapping boxed cones). The performance of the stochastic Pareto-optimal policies are show with transparency, while the lower bound policy is represented by the dashed circled cone. All the DPC policies overcome the performance of the lower bound policy with respect to the Environment objective. Moreover, most of these policies meet or improve the performance of the baseline and the SRBC negotiated alternatives in terms of environmental shortage index, demonstrating that the constraints imposed in the RBFs parameters decision space represents an effective alternative to the classical minimum environmental flow constraints. Finally, the flexibility of this dynamic conditioning mechanism, which exploits the feedback represented by the system condition by means of the corresponding operating policies, allows the identification of solutions attaining higher performance than the baseline in terms of Hydropower revenue and Atomic Power Plant.
Figure 6.21 illustrates two candidate solutions designed with the proposed DPC approach working on the Sobol samples (cyan line) and on the stochastic Pareto-optimal set (blue line). These policies are compared with the baseline and the SRBC alternatives. The lower bound

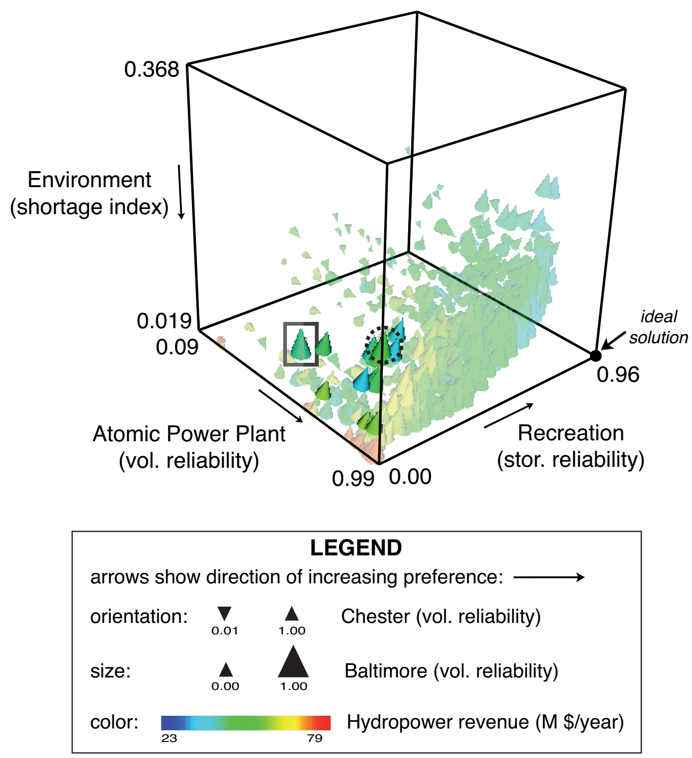

Figure 6.20: Performance of the DPC policies (opaque cones) over stochastic hydroclimatic conditions. The baseline policy and the alternatives negotiated by the SRBC are identified by the boxed cones, the lower bound policy by the dashed circled cone, while the stochastic Pareto-optimal policies with transparency.

policy and the stochastic Pareto-optimal policies are also shown as they represent the two extreme solutions. Results show that adopting the DPC solution derived from the Pareto-optimal set (Sobol samples) instead of maintaining the current reservoir regulation, the SRBC would potentially attain an increase in the hydropower revenue equal to 18.6 (17.6) million US$/year, 0.29 (0.33) in the Atomic Power Plant reliability, and a reduction of 0.038 (0.032) in terms of environmental shortage index. These solutions therefore exhibits the potential to outperform the baseline regulation, producing a robust policy that reflects the actual decision-making structure of the Lower Susquehanna system and successfully protects the environment by means of effective conditioning of the Conowingo reservoir operating policy.

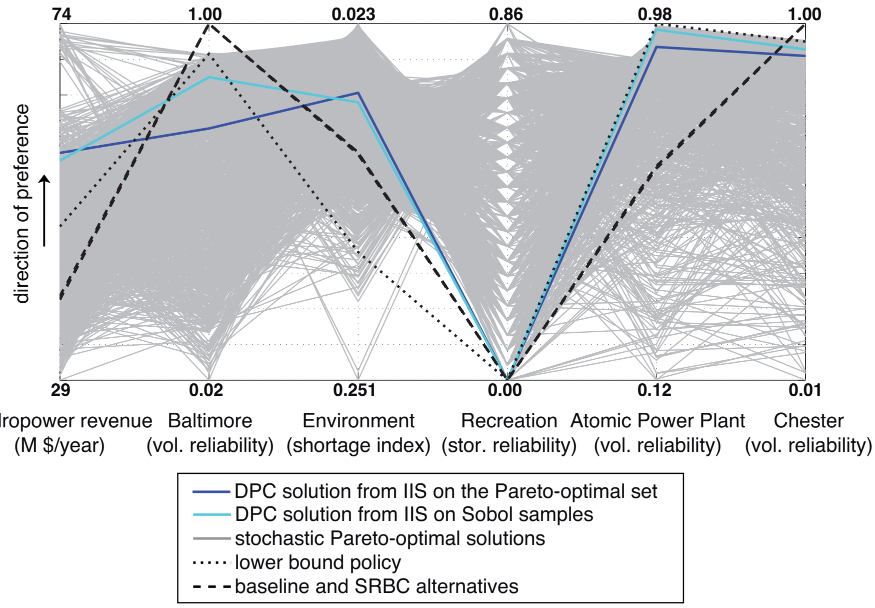

Figure 6.21: Comparison of the performance over stochastic hydroclimatic conditions of the DPC policies, the baseline policy, and the alternatives negotiated by the SRBC.

Figures 6.20-6.21 represent the performance of the operating policies with respect to the six operating objectives, with the different mechanism design strategies aiming to guarantee effective environmental protection. However, these figures are focused on the system-level performance (i.e., the Environment objective) and do not provide any information on the practicability of the illustrated solutions. Let assume that the number of decision variables subject to constraints is a proxy to estimate the practicability of the operating policies, with less constraints corresponding to more practicable (acceptable) solutions for the active agents. The lower bound policy, representing the situation before the introduction of the FERC flow requirements, is therefore the most practicable solution (i.e., zero constraints). At the other extreme, the reference solution (i.e., the stochastic Pareto-optimal policy with the lowest environmental shortage index) is a fully constrained solution with fixed values for all the 32 decision variables. The baseline policy and the other alternatives proposed by the SRBC impose a constraint on the downstream release decision, thus involving 20 RBFs parameters. The DPC solutions are instead defined with more soft conditioning mechanisms (i.e., 18 or 8 constraints in the case of IIS run on Sobol samples or Pareto-optimal set,

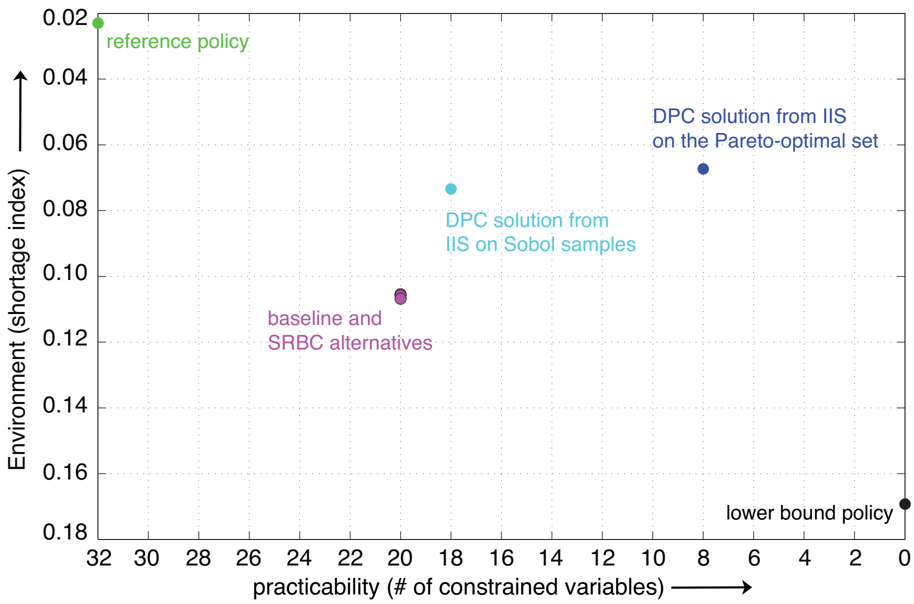

Figure 6.22: Representation of the conflict between system-wide efficiency and practicability, measured by the Environment objective and the number of constrained decision variables, respectively.

respectively). Figure 6.22 represents these policies in the practicability-Environment (i.e., system-wide performance) space and demonstrates a clear tradeoff between these two metrics. These results show that DPC solution designed according to the environmental signature identified on the Pareto-optimal set is able to attain a good environmental performance with limited conditioning, thus representing a potentially highly practicable solution.

6.4 Discussion and final remarks

In this chapter, the agent-based decision analytic framework is applied to support water reservoir operation dealing with many competing demands and hydroclimatic uncertainty in the Lower Susquehanna system. The approach demonstrates the potential to overcome policy inertia and myopia by providing the Susquehanna River Basin Commission (SRBC) with the full range of opportunities to improve the current reservoir operation according to a multi–objective and stochastic perspective.

The adopted implicit policy identification method captures the current operation of the dam and defines the historical policy by fitting radial basis functions to existing system dynamics. After identifying the historical baseline policy, the combination of evolutionary many-objective optimization with visual analytics allows the discovery of improved operating policies. The optimization was performed over stochastic samples of inflows and evaporation rates to ensure robustness of the solutions.
Results show that the system's current history-based operations are negatively biased to overestimate the reliability of the reservoir's multi-sector services. The a posteriori analysis of the stochastic Pareto-optimal solutions and the set of alternatives negotiated by the SRBC shows that these latter are essentially equivalent to the historical policy, with negligible differences in performance compared to the stochastic Pareto-optimal policies. Moreover, the proposed framework has successfully identified a subset of alternative reservoir policies that are robust to hydroclimatic uncertainties while being capable of better addressing the tradeoffs across the Lower Susquehanna competing demands. The comparison of the historical performance with the one that would be attained with the recommended policy provides an estimate of the regret that the SRBC would experience by maintaining the current policy. By adopting the recommended policies, the expected hydropower revenues increase by 22 million US$, with significant advantages also in Environment and Atomic Power Plant objectives. The improvement in Recreation is instead limited as the current preference structure does not prioritize this objective. These advantages are obtained at the cost of a small reliability reduction for Baltimore and Chester reliability, which have the option to obtain water from other sources.
Finally, the adoption of the direct policy conditioning (DPC) method exploits the reference provided by the stochastic Pareto-optimal policies to identify alternative operating policies for the Conowingo reservoir which better represent the actual decision-making context, modeled by four active agents and two passive agents. The federal minimum flow requirements and the alternative negotiated by the SRBC can be considered as examples of mechanism designs, aimed at the protection of the environment, which is a critical concern under stochastic hydroclimatic conditions. Results show that the DPC policies exhibits the potential to outperform the baseline alternative in terms of Environment and also Hydropower revenue, Atomic Power Plant reliability, and Recreation. Moreover, these policies are designed with more soft conditioning than the FERC flow requirements (i.e., they define less constraints) and, therefore, might represent highly practicable solutions.
The results presented in this chapter are obtained assuming stationary

hydroclimatic conditions and only considering the uncertainty in the hydroclimatic variables. Broadly, there are many uncertain factors that can influence the system including shifting objectives, evolving demands, and climate change. Moreover, although the minimax approach used to filter the system uncertainties guarantee certain performance over different hydroclimatic conditions, other filtering criteria might be used depending on the risk aversion of the SRBC, such as the Laplace criterion (*Laplace*, 1951), which looks at the expected performance, the Hurwicz criterion (*Hurwicz*, 1951), which considers a weighted combination of the worst and best case, or the Savage criterion (*Savage*, 1951), which minimizes the regret of adopting a wrong decision. Depending on the adopted filtering criterion, the set of optimal solutions varies.

Future efforts will concentrate on estimating the robustness of the policies under conditions of deep uncertainty (*Kasprzyk et al.*, 2012), such as enlarging the variability of the hydroclimatic variables, introducing non-stationarity and climate change effects, and considering uncertain water demands and energy prices. Furthermore, the potential of DPC in designing policies for different objective tradeoffs can be assessed. Finally, the application of DPC in the definition of coordination mechanisms in multireservoir systems, such as the Zambezi River basin, seems a promising method to support mechanism design strategies in large-scale, transboundary water management problems.

7 Conclusions and future research

The aim of this thesis was to introduce a novel agent-based decision analytic framework to study water resources planning and management problems in complex decision-making contexts, involving multiple decision makers and many conflicting stakeholders. The proposed framework combines descriptive and prescriptive methods in order to provide informative tools, which represent the actual decision-making context, as well as decision support procedures, which recommend proper coordination mechanisms.

The first application of the framework to a hypothetical water allocation planning problem presents an extensive analysis of the intrinsic conflict between system-wide efficiency and solutions practicability, when multiple decision-making actors are involved. Results show that approaches relying on distributed constraint optimization problems are able to support a watershed authority in the identification of mechanism design strategies based on normative constraints or economic incentives. The imposition of different sets of constraints allows the exploration of the trade-off curve between system-wide efficiency and agent-level practicability. The adopted agent-based methods shows the potential to design coordination mechanisms that produce solutions in the space in-between the two extreme situations of centralized and uncoordinated management.

The second application of the framework to the Zambezi River basin management problem demonstrates the effectiveness of a simple coordination mechanism based on the exchange of information. The introduction of full information sharing between Zambia, Zimbabwe, and Mozambique allows the downstream country (i.e., Mozambique) to better adapt to the upstream strategies, with no consequence for these latter. The stability of this coordinated solution is guaranteed, as each country cannot improve its benefit acting unilaterally. The comparison of the system-level performance in three scenarios of cooperation, namely full cooperation, coordination, and non-cooperation, provides an estimate of the value of cooperation and information exchange, which might represent a fundamental information to set up a negotiation process.

The exchange of information was successfully applied also in the Como system to promote the activation of co-adaptation strategies of the water supply management and farmers' practices. The aim of this approach is to cross-condition the decisions of farmers and water managers: the farmers select the crop to grow as the one with the highest revenue on the basis of an expectation of irrigation water supply; knowing the selected crop pattern, the water manager optimizes the management of the irrigation supply system to match the actual demand of the crops. Results show that the proposed approach successfully enhances the efficiency of agricultural water management practices and fosters crop production. Moreover, co-adaptation is demonstrated to be effective in mitigating climate change adverse impacts.
In the last application, the proposed framework effectively combines reservoir policy identification, many-objective optimization under uncertainty, and visual analytics to characterize the current Conowingo reservoir operations and discover key tradeoffs between alternative policies in the Lower Susquehanna system. Results show that the current history-based performance are negatively biased and overestimated. The a posteriori analysis of the stochastic Pareto-optimal solutions allows the identification of a sub-set of alternative reservoir policies that are robust to hydroclimatic uncertainties, while being capable of better addressing the tradeoffs across the Lower Susquehanna competing demands. Moreover, the adoption of the direct policy conditioning method successfully exploits the reference provided by the stochastic Pareto-optimal policies to identify alternative mechanism design options for environmental protection. Results show that the direct policy conditioning solutions are able to meet or exceed the performance of the alternatives based on the classical minimum environmental flow constraint in terms of environmental protection, hydropower revenue, and reliability of the atomic power plant cooling, thus representing a potential compromise solutions with respect to system-level efficiency and practicability.

In conclusion, the proposed agent-based decision analytic framework represents a promising contribution to address the complexity of the decision-making institutional contexts. It allows the exploration of different levels of cooperation between the decision makers through an agent-based modeling approach, where agents' decisions are designed by state-of-the-art Control Theory techniques. Moreover, the proposed coordination mechanisms, based on information exchange and direct policy conditioning, successfully influence the originally uncoordinated agents and increase the system-level performance.
Finally, this thesis suggests several directions for future investigations,

which can be summarized as follows:

- development of new methods and algorithms based on distributed constraint optimization problems to deal with multi-objective and dynamic (management) problems and their application to real world case studies;
- elaboration of more realistic models in terms of description of the processes and agents' decisions, in order to better represent the current baseline conditions as well as to provide more credible support to the decision makers;
- design of effective mechanisms to promote negotiation processes for the identification of improved solutions at the system-level, possibly supported by game-theoretic analysis, which might have practical impact from a policy-making point of view;
- improvement of the interaction with the real decision makers and stakeholders to better capture their interests and more effectively communicate the results of the research works.

Acknowledgements

This PhD thesis was carried out under the supervision of Rodolfo Soncini-Sessa and the co-supervision of Andrea Castelletti. I wish to thank them for the opportunity they gave me and for their teaching and consulting during the fulfillment of this thesis. I'm extremely grateful to Andrea Castelletti who has been always available for discussion and support and has shown me how to develop my skills and perform cutting-edge research.

I sincerely thank all the members and ex-members of the Integrated Water Resources Management group as well as the other members of the Environmental Systems group for their suggestions and inspiring conversations. Thanks to the students I co-supervised during these three years who collaborated with me for the development of some aspects of this research. I would also like to acknowledge all the co-authors of the research presented in this thesis and, in particular, Francesco Amigoni, Ximing Cai, Yu Li, Claudio Gandolfi, Jon Herman, and Patrick Reed. Special thanks go to Patrick Reed, who gave me the possibility of visiting Penn State University (Pennsylvania) and provided me a strong support in developing the research on the Susquehanna River problem. I wish to thank all the PSU group and, in particular, Jon Herman, Alisha Fernandez, Joe Kasprzyk, Matthew Woodruff, and David Hadka.

I am grateful to Prof. Amaury Tilmant (Université Laval) who carefully reviewed this thesis and to my tutor Carlo Piccardi.

This research was partially supported by *Fondazione Fratelli Confalonieri*.

Bibliography

Adger, W., N. Arnell, and E. Tompkins (2005), Successful adaptation to climate change across scales, *Global Environmental Change*, *15*(2), 77–86, doi:http://dx.doi.org/10.1016/j.gloenvcha.2004.12.005.

Ajami, N. K., G. M. Hornberger, and D. L. Sunding (2008), Sustainable water resource management under hydrological uncertainty, *Water Resources Research*, *44*(11), doi:10.1029/2007WR006736.

Alam, U., O. Dione, and P. Jeffrey (2009), The benefit-sharing principle: Implementing sovereignty bargains on water, *Political Geography*, *28*(2), 90–100.

Allen, R., L. Pereira, D. Raes, and M. Smith (1998), Fao irrigation and drainage paper 56, *Tech. rep.*, FAO, Rome, Italy.

Altinbilek, D. (2004), Development and management of the Euphrates-Tigris basin, *International Journal of Water Resources Development*, *20*(1), 15–33.

Anand, A., S. Galelli, M. Giuliani, L. Samavedham, and D. Schwanenberg (2011), Coordinating Multiple Model Predictive Controllers for Water Reservoir Networks Operation, in *Proceedings of the MODSIM 2011 International Congress on Modelling and Simulation*, Perth, Australia.

Anghileri, D., F. Pianosi, and R. Soncini-Sessa (2011), A framework for the quantitative assessment of climate change impacts on water-related activities at the basin scale, *Hydrology and Earth System Sciences*, *15*(6), 2025–2038, doi:10.5194/hess-15-2025-2011.

Anghileri, D., A. Castelletti, F. Pianosi, R. Soncini-Sessa, and E. Weber (2012), Optimizing watershed management by coordinated operation of storing facilities, *Journal of Water Resources Planning and Management*, doi:doi:10.1061/(ASCE)WR.1943-5452.0000313.

Athanasiadis, I. (2005), A review of agent-based systems applied in environmental informatics, in *Proceedings of the MODSIM 2005 International Congress on Modelling and Simulation*, Melbourne, Australia.

Back, T., D. Fogel, and Z. Michalewicz (2000), *Handbook of evolutionary computation*, IOP Publishing Ltd. and Oxford University Press, Bristol, U.K.

Baglietto, M., C. Cervellera, M. Sanguineti, and R. Zoppoli (2006), Water reservoirs management under uncertainty by approximating networks and learning from data, in *Topics on System Analysis and Integrated Water Resource Management*, edited by Elsevier, pp. 117–139.

Barreteau, O., F. Bousquet, and J. Attonaty (2001), Role–playing games for opening the black box of multi-agent systems: method and lessons of its application to Senegal River Valley irrigated systems, *Journal of artificial societies and social simulation*, *4*(2), 5.

Barrett, A. (1999), Autonomy architectures for a constellation of spacecraft, in *Artificial Intelligence, Robotics and Automation in Space*, vol. 440, p. 291.

Barron, A. (1993), Universal approximation bounds for superpositions of a sigmoidal function, *IEEE Transactions on Information Theory*, *39*(3), 930–945.

Becu, N., P. Perez, A. Walker, O. Barreteau, and C. Le Page (2003), Agent based simulation of a small catchment water management in northern Thailand: description of the CATCHSCAPE model, *Ecological Modelling*, *170*(2), 319–331.

Beilfuss, R. (2001), Prescribed flooding and restoration potential in the Zambezi delta, Mozambique, in *Working Paper #3, Program for the sustainable management of Cahora Bassa dam and the lower Zambezi Valley*.

Beilfuss, R., and C. Brown (2010), Assessing environmental flow requirements and trade-offs for the Lower Zambezi River and Delta, Mozambique, *International Journal of River Basin Management*, *8*(2), 127–138.

Beilfuss, R., and D. dos Santos (2001), Patterns of hydrological change in the Zambezi Delta, Mozambique, in *Working Paper #2, Program for the sustainable management of Cahora Bassa dam and the lower Zambezi Valley*.

Bellman, R. (1957), *Dynamic programming*, Princeton University Press, Princeton.

Berger, T. (2001), Agent-based spatial models applied to agriculture: a simulation tool for technology diffusion, resource use changes and policy analysis, *Agricultural Economics*, *25*(2–3), 245–260, doi: 10.1111/j.1574-0862.2001.tb00205.x.

Berger, T., and C. Ringler (2002), Tradeoffs, efficiency gains and technical change-Modeling water management and land use within a multiple-agent framework, *Quarterly Journal of International Agriculture*, *41*(1-2), 119–144.

Berger, T., R. Birner, N. Mccarthy, J. DíAz, and H. Wittmer (2007), Capturing the complexity of water uses and water users within a multi-agent framework, *Water Resources Management*, *21*(1), 129–148.

Bernauer, T. (2002), Explaining success and failure in international river management, *Aquatic Sciences*, *64*(1), 1–19.

Bertsekas, D. (1976), *Dynamic programming and stochastic control*, Academic Press, New York.

Bertsekas, D. (2005), Dynamic programming and suboptimal control: a survey from ADP to MPC, *European Journal of Control*, *11*(4-5).

Bindi, M., and M. Howden (2004), Challenges and opportunities for cropping systems in a changing climate, in *New directions for a diverse planet*, Brisbane, Australia, proceedings of the 4th International Crop Science Congress.

Block, P., and K. Strzepek (2010), Economic Analysis of Large-Scale Upstream River Basin Development on the Blue Nile in Ethiopia Considering Transient Conditions, Climate Variability, and Climate Change, *Journal of Water Resources Planning and Management*, *136*(2), 156–166.

Boé, J., L. Terray, F. Habets, and E. Martin (2007), Statistical and dynamical downscaling of the seine basin climate for hydro-meteorological studies, *International Journal of Climatology*, *27*(12), 1643–1655.

Bonabeau, E. (2002), Agent–based modeling: Methods and techniques for simulating human systems, *Proceedings of the National Academy of Sciences of the United States of America*, *99*(Suppl 3), 7280–7287.

Bousquet, F., and C. Le Page (2004), Multi-agent simulations and ecosystem management: a review, *Ecological Modelling*, *176*(3-4), 313–332.

Bousquet, F., O. Barreteau, C. Le Page, C. Mullon, and J. Weber (1999), An environmental modelling approach: the use of multi-agent simulations, *Advances in environmental and ecological modelling*, *113*, 122.

Bowden, G., G. Dandy, and H. Maier (2005), Input determination for neural network models in water resources applications. Part 1 - background and methodology, *Journal of Hydrology*, *301*(1-4), 75–92.

Box, G., and G. Jenkins (1970), *Time series analysis: Forecasting and control*, Holden Day, San Francisco.

Braden, H. (1985), Ein energiehaushalts und verdunstungsmodell fur wasser und stoffhaushaltsuntersuchungen landwirtschaftlich genutzer einzugsgebiete, *Mitteilungen deutsche bodenkundliche gesselschaft*, *422*, 294–299.

Brander, L., R. Florax, and J. Vermaat (2006), The empirics and wetlands valuation: A comprehensive summary and a meta-analysis of the literature, *Environmental & Resource Economics*, *33*, 223–250.

Brill., E., J. Flach, L. Hopkins, and S. Ranjithan (1990), MGA: A Decision Support System for Complex, Incompletely Defined Problems, *IEEE Transactions on Systems, Man, and Cybernetics*, *20*(4), 745–757.

Brown, C., and M. Carriquiry (2007), Managing hydroclimatological risk to water supply with option contracts and reservoir index insurance, *Water Resources Research*, *43*(11), doi:10.1029/2007WR006093.

Busoniu, L., D. Ernst, B. De Schutter, and R. Babuska (2011), Cross–Entropy Optimization of Control Policies With Adaptive Basis Functions, *IEEE Transactions on systems, man and cybernetics–Part B: cybernetics*, *41*(1), 196–209, doi:10.1109/TSMCB.2010.2050586.

Cai, X., and M. Rosegrant (2004), Irrigation technology choices under hydrologic uncertainty: A case study from Maipo River Basin, Chile, *Water Resources Research*, *40*(4), doi:10.1029/2003WR002810.

Carnevale, C., G. Finzi, E. Pisoni, and M. Volta (2009), Neuro-fuzzy and neural network systems for air quality control, *Atmospheric Environment*, *43*(31), 4811–4821.

Castelletti, A., F. Pianosi, and R. Soncini-Sessa (2008a), Water reservoir control under economic, social and environmental constraints, *Automatica*, *44*(6), 1595–1607, doi:10.1016/j.automatica.2008.03.003.

Castelletti, A., F. Pianosi, and R. Soncini-Sessa (2008b), Receding horizon control for water resources management, *Applied Mathematics and Computation*, *204*(2), 621–631.

Castelletti, A., S. Galelli, M. Restelli, and R. Soncini-Sessa (2010a), Tree-based reinforcement learning for optimal water reservoir operation, *Water Resources Research*, *46*(W09507).

Castelletti, A., S. Galelli, M. Restelli, and R. Soncini-Sessa (2011b), Tree-based feature selection for dimensionality reduction of large-scale control systems, in *Proceedings of the IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning*, Paris, France.

Castelletti, A., F. Pianosi, X. Quach, and R. Soncini-Sessa (2012a), Assessing water reservoirs management and development in Northern Vietnam, *Hydrology and Earth System Sciences*, *16*(1), 189–199, doi:10.5194/hess-16-189-2012.

Castelletti, A., S. Galelli, M. Ratto, R. Soncini-Sessa, and P. Young (2012b), A general framework for dynamic emulation modelling in environmental problems, *Environmental Modelling & Software*, *34*, 5–18.

Castelletti, A., F. Pianosi, and M. Restelli (2013b), A multiobjective reinforcement learning approach to water resources systems operation: Pareto frontier approximation in a single run, *Water Resources Research*, *49*, doi:10.1002/wrcr.20295.

Celeste, A., and M. Billib (2009), Evaluation of stochastic reservoir operation optimization models, *Advances in Water Resources*, *32*(9), 1429–1443, doi:10.1016/j.advwatres.2009.06.008.

Cervellera, C., V. Chen, and A. Wen (2006), Optimization of a large-scale water reservoir network by stochastic dynamic programming with efficient state space discretization, *European Journal of Operational Research*, *171*(3), 1139–1151.

Chankong, V., and Y. Haimes (1983), *Multiobjective decision making: theory and methodology*, North-Holland, New York, NY.

Charness, G., and M. Rabin (2002), Understanding social preferences with simple tests, *The Quaterly Journal of Economics*, *117*(3), 817–869.

Chaves, P., and T. Kojiri (2007), Deriving reservoir operational strategies considering water quantity and quality objectives by stochastic fuzzy neural networks, *Advances in Water Resources*, *30*(5), 1329–1341, doi: 10.1016/j.advwatres.2006.11.011.

Christensen, J. H., and O. B. Christensen (2007), A summary of the PRUDENCE model projections of changes in European climate by the end of this century, *Climatic Change*, *81*, 7–30.

Clark, E. (1950), New York control curves, *Journal of the American Water Works Association*, *42*(9), 823–827.

Clark, E. (1956), Impounding reservoirs, *Journal of the American Water Works Association*, *48*(4), 349–354.

Coello Coello, C., G. Lamont, and D. V. Veldhuizen (2007), *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic Algorithms and Evolutionary Computation)*, 2 ed., Springer, New York.

Cohon, J. L., and D. Marks (1975), A review and evaluation of multiobjective programing techniques, *Water Resources Research*, *11*(2), 208–220.

Collin, Z., and S. Dolev (1994), Self–stabilizing depth first search, *Inform. Process. Lett.*, *49*, 297–301.

Conry, S., K. Kuwabara, V. Lesser, and R. Meyer (1991), Multistage negotiation for distributed constraint satisfaction, *IEEE Transactions on Systems, Man and Cybernetics*, *21*(6), 1462–1477.

Corani, G., A. Rizzoli, A. Salvetti, and M. Zaffalon (2009), Reproducing human decisions in reservoir management: the case of lake Lugano, in *Proceedings of the 4th International ICSC Symposium on Information Technologies in Environmental Engineering*, Thessaloniki, Greece.

Das, S. (2001), Filters, wrappers and a boosting-based hybrid for feature selection, in *Proceedings of the 18th International Conference on Machine Learning*, Williamstown, MA.

de Fraiture, C., and D. Wichelns (2010), Satisfying future water demands for agriculture, *Agricultural Water Management*, *97*(4), 502–511, doi: 10.1016/j.agwat.2009.08.008.

Deb, K., J. Sundar, N. Udaya Bhaskara Rao, and S. Chaudhuri (2006), Reference point based multi–objective optimization using evolutionary algorithms, *International Journal of Computational Intelligence Research*, *2*(3), 273–286.

Déqué, M. (2007), Frequency of precipitation and temperature extremes over France in an anthropogenic scenario: Model results and statistical correction according to observed values, *Global and Planetary Change*, *57*(1), 16–26.

Deressa, T., R. Hassan, C. Ringler, T. Alemu, and M. Yesuf (2009), Determinants of farmers' choice of adaptation methods to climate change in the Nile Basin of Ethiopia, *Global Environmental Change*, *19*(2), 248–255, doi:10.1016/j.gloenvcha.2009.01.002.

Di Baldassarre, G., et al. (2011), Future hydrology and climate in the River Nile basin: a review, *Hydrological Sciences Journal*, *56*(2), 199–211, doi:10.1080/02626667.2011.557378.

Doorenbos, J., A. Kassam, and C. Bentvelsen (1979), Yield response to water, irrigation and drainage. Paper no. 33, *Tech. rep.*, Food and Agriculture Organization, Rome, Italy.

Draper, A., and J. Lund (2004), Optimal Hedging and Carryover Storage Value, *Journal of Water Resources Planning and Management*, *130*(1), 83–87, doi:10.1061/(ASCE)0733-9496(2004)130:1(83).

Escaler, I., A. Cabello, C. Campos, B. Lesjean, C. Herve-Bazin, and M. Velasco (2012), Climate change and the water sector in Europe: a review of research and technology development needs, *Journal of Water and Climate Change*, *3*(2), 79–94, doi:10.2166/wcc.2012.022.

Eschenauer, H., J. Koski, and A. Osyczka (1990), *Multicriteria design optimization: procedures and applications*, Springer, Berlin, Heidelberg, NewYork.

Exelon (2010), Operations modeling calibration report, *Tech. rep.*, Gomez and Sullivan Engineers.

Facchi, A., B. Ortuani, D. Maggi, and C. Gandolfi (2004), Coupled SVAT–groundwater model forwater resources simulation in irrigated alluvial plains, *Environmental Modeling & Software*, *19*(11), 1053–1063.

Falkenmark, M. (2013), Growing water scarcity in agriculture: future challenge to global water security, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *371*(2002), 20120,410.

FAO (2003), Unlocking the Water Potential of Agriculture, *Tech. rep.*, Food and Agriculture Organization, Rome, Italy.

Farina, M., and P. Amato (2002), On the optimal solution definition for many-criteria optimization problems, in *Proceedings of the NAFIPS-FLINT International Conference '2002*, pp. 233–238, IEEE Service Center, Piscataway, Ney Jersey.

Federal Energy Regulatory Commission (1989), Order Approving Settlement Agreement for the Conowingo Hydroelectric Project, *Tech. Rep. Project No. 405-009, Docket No. EL80-38-000.*

Fernandez, A., S. Blumsack, and P. Reed (2012), Evaluating wind–following and ecosystem services for hydroelectric dams in PJM, *Journal of regulatory economics*, *41*(1), 139–154, doi:10.1007/s11149-011-9177-9.

Fernandez, A., S. Blumsack, and P.M. (2013), Operational constraints and hydrologic variability limit hydropower in supporting wind integration, *Environmental Research Letters*, *8*(2), doi:10.1088/1748-9326/8/2/024037.

Feuillette, S., F. Bousquet, and P. Le Goulven (2003), SINUSE: a multi-agent model to negotiate water demand management on a free access water table, *Environmental Modelling & Software*, *18*(5), 413–427.

Fleming, P., R. Purshouse, and R. Lygoe (2005), Many-Objective optimization: an engineering design perspective, in *Proceedings of the Third international conference on Evolutionary Multi-Criterion Optimization*, pp. 14–32, Guanajuato, Mexico.

Fonteneau, R., L. Wehenkel, and D. Ernst (2008), Variable selection for dynamic treatment regimes: a reinforcement learning approach, in *Proceedings of the European Workshop on Reinforcement Learning*, Villeneuve d'Ascq, F.

Galán, J., A. López-Paredes, and R. Del Olmo (2009), An agent-based model for domestic water management in valladolid metropolitan area, *Water Resources Research*, *45*(5).

Galelli, S., and A. Castelletti (2013), Tree-based iterative input variable selection for hydrological modeling, *Water Resources Research*, *49*, 4295–4310, doi:10.1002/wrcr.20339.

Galelli, S., and R. Soncini-Sessa (2010), Combining metamodelling and stochastic dynamic programming for the design of reservoir release policies, *Environmental Modeling & Software*, *25*(2), 209–222, doi:10.1016/j.envsoft.2009.08.001.

Galelli, S., A. Goedbloed, D. Schwanenberg, and P. van Overloop (2012), Optimal Real-Time Operation of Multi-Purpose Urban Reservoirs: A Case Study in Singapore, *Journal of Water Resources Planning and Management.*

Gandolfi, C., G. Guariso, and D. Togni (1997), Optimal Flow Allocation in the Zambezi River System, *Water Resources Management*, *11*, 377–393.

Gandolfi, C., A. Facchi, and D. Maggi (2006), Comparison of 1d models of water floe in unsaturated soils, *Environmental Modelling & Software*, *21*, 1759,Äì–1764.

Garcia-Vila, M., and E. Fereres (2012), Combining the simulation crop model AquaCrop with an economic model for the optimization of irrigation management at farm level, *European Journal of Agronomy*, *36*(1), 21–31, doi:10.1016/j.eja.2011.08.003.

Gass, S., and T. Saaty (1955), Parametric objective function - Part II, *Operations Research*, *3*, 316–319.

George, E. (2000), The variable selection problem, *Journal of the American Statistical Association*, *95*(452), 1304–1308.

Gerlak, A., J. Lautze, and M. Giordano (2011), Water resources data and information exchange in transboundary water treaties, *International Environmental Agreements-Politics Law and Economics*, *11*(2), 179–199.

Gettys, C., and S. Fisher (1979), Hypothesis plausibility and and hypothesis generation, *Organizational Behavior and Human Performance*, *24*(1), 93–110.

Geurts, P., D. Ernst, and L. Wehenkel (2006), Extremely randomized trees, *Machine Learning*, *63*(1), 3–42.

Giuliani, M., and A. Castelletti (2013), Assessing the value of cooperation and information exchange in large water resources systems by agent–based optimization, *Water Resources Research*, *49*, 3912–3916, doi:10.1002/wrcr.20287.

Giuliani, M., A. Castelletti, F. Amigoni, and X. Cai (2012a), Multi-Agent Systems optimization for distributed watershed management, in *Proceedings of the 6th International Congress on Environmental Modeling and Software (iEMSs)*, Leipzig, Germany.

Giuliani, M., A. Castelletti, and R. Soncini-Sessa (2012b), Multi-Agent Water Resources Management: Centralized vs Decentralized Approach, in *Proceedings of the 10th International Conference on Hydroinformatics (HIC)*, Hamburg, Germany.

Giuliani, M., A. Castelletti, F. Amigoni, and X. Cai (2013a), Distributed optimization as a tool to balance efficiency and practicability in watershed management, *Journal of Water Resources Planning and Management*, (under review).

Giuliani, M., S. Galelli, and R. Soncini-Sessa (2013b), A dimensionality reduction approach for Many-Objective Markov Decision Processes: application to a water reservoir operation problem, *Environmental Modeling & Software*, (under review).

Giuliani, M., J. Herman, A. Castelletti, and P. Reed (2013c), Many-Objective Reservoir Policy Identification and Refinement to Reduce Institutional Myopia in Water Management, *Water Resources Research*, (under review).

Giuliani, M., Y. Li, A. Castelletti, and C. Gandolfi (2013d), Co-adapting water demand and supply to changing climate in agricultural water systems, *Global Environmental Change*, (in preparation).

Giuliani, M., A. Castelletti, and P. Reed (2014), Direct policy conditioning for environmental protection, in *Proceedings of the 19th IFAC World Congress*, Cape Town, SA, (submitted).

Gleick, P. (2002), Water management – Soft water paths, *Nature*, *418*(6896), 373, doi:10.1038/418373a.

Goldberg, D. (1989), *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison Wesley, Boston, MA.

Goor, Q., A. Alia, P. Van Der Zaag, and A. Tilmant (2007), Impacts of the southeastern Anatolia Project in Turkey on the performance of the Tabqa dam and hydropower plant in Syria, in *Changes in Water Resources Systems: Methodologies to Maintain Water Security and Ensure Integrated Management – Proceedings of Symposium HS3006 at IUGG2007. IAHS Publ. 315*, pp. 92–99, Perugia, Italy.

Guariso, G., and D. Whittington (1987), Implications of ethiopian water development for Egypt and Sudan, *International Journal of Water Resources Development*, *3*(2), 105–114.

Guariso, G., S. Rinaldi, and R. Soncini-Sessa (1986), The Management of Lake Como: A Multiobjective Analysis, *Water Resources Research*, *22*(2), 109–120, doi:10.1029/WR022i002p00109.

Guo, X., T. Hu, X. Zeng, and X. Li (2013), Extension of parametric rule with the hedging rule for managing multireservoir system during droughts, *Journal of Water Resources Planning and Management*, *139*(2), 139–148, doi:10.1061/(ASCE)WR.1943-5452.0000241.

Guyon, I., and A. Elisseeff (2003), An introduction to variable and feature selection, *Journal of Machine Learning Research*, *3*, 1157–1182.

Guyot, P., and S. Honiden (2006), Agent–based participatory simulations: Merging multi–agent systems and role-playing games, *Journal of Artificial Societies and Social Simulation*, *9*(4).

Hadka, D., and P. Reed (2012), Diagnostic assessment of search controls and failure modes in many–objective evolutionary optimization, *Evolutionary Computation*, *20*(3), 423–452.

Hadka, D., and P. Reed (2013), Borg: An Auto–Adaptive Many–Objective Evolutionary Computing Framework, *Evolutionary Computation*, *21*(2), 231–259.

Haimes, Y., and W. Hall (1977), Sensitivity, responsivity, stability and irreversibility as multiple objectives in civil systems, *Advances in water resources*, *1*(2), 71–81.

Haimes, Y., L. Lasdon, and D. Wismer (1971), On a bicriterion formulation of the problems of integrated system identification and system optimization, *IEEE Transactions on Systems, Man and Cybernetics*, *1*, 296–297.

Haimes, Y., K. Tarvainen, T. Shima, and J. Thadathil (1990), *Hierarchical multiobjective analysis of large–scale systems*, Hemisphere Publishing Corporation New York.

Hanf, C. (1989), *Agricultural sector analysis by linear programming models: approaches, problems and experiences*, 20, Wissenschaftsverlag Vauk.

Hardin, G. (1968), The tragedy of the commons, *Science*, *162*(3859), 1243–1248.

Hashimoto, T., J. Stedinger, and D. Loucks (1982), Reliability, resilience, and vulnerability criteria for water resource system performance evaluation, *Water Resources Research*, *18*(1), 14–20, doi: 10.1029/WR018i001p00014.

Hejazi, M., and X. Cai (2009), Input variable selection for water resources systems using a modified minimum redundancy maximum relevance (mMRMR) algorithm, *Advances in Water Resources*, *32*(4), 582–593.

Hejazi, M., X. Cai, and B. Ruddell (2008), The role of hydrologic information in reservoir operation – learning from historical releases, *Advances in Water Resources*, *31*(12), 1636–1650.

Hewitt, C. (1986), Offices are open systems, *ACM Transactions on Information Systems (TOIS)*, *4*(3), 271–287.

Hingray, B., N. Mouhous, A. Mezghani, K. Bogner, B. Schaefli, and A. Musy (2007), Accounting for global-mean warming and scaling uncertainties in climate change impact studies: application to a regulated lake system, *Hydrology and Earth System Sciences*, *11*(3), 1207–1226.

Hirayama, K., and M. Yokoo (1997), Distributed partial constraint satisfaction problem, in *Proceedings of the Third International Conference on Principles and Practice of Constraint Programming (CP97)*, Linz, Austria.

Hirayama, K., and M. Yokoo (2005), The distributed breakout algorithms, *Artificial Intelligence*, *161*(1–2), 89–115.

Hogarth, R. (1981), Beyond discrete biases: Functional and dysfunctional aspects of judgemental heuristics, *Psychological Bulletin*, *90*(1), 197–217.

Howard, C. (1999), *Death to Rule Curves*, chap. 231, pp. 1–5, doi: 10.1061/40430(1999)232.

Howden, S. M., J. Soussana, F. Tubiello, N. Chhetri, M. Dunlop, and H. Meinke (2007), Adapting agriculture to climate change, *Proceedings of the National Academy of Sciences of the United States of America*, *104*(50), 19,691–19,696, doi:10.1073/pnas.0701890104.

Hurwicz, L. (1951), Optimality criteria for decision making under ignorance, *Cowles commission papers*, *370*.

Huskova, I., and J. Harou (2012), An agent model to simulate water markets, in *Proceedings of the 6^{th} International Congress on Environmental Modeling and Software (iEMSs)*, Leipzig, Germany.

Huston, M., D. DeAngelis, and W. Post (1988), New computer models unify ecological theory, *BioScience*, *38*(10), 682–691.

Inselberg, A. (1997), Multidimensional detective, in *Proceedings of the IEEE Symposium on Information Visualization*, pp. 100–107, IEEE.

IPCC (2000), IPCC: Special Report on Emission Scenarios, *Tech. rep.*, Intergovernmental Panel on Climate Change (IPCC).

IPCC (2001), Climate Change 2001: Impacts, Adaptation and Vulnerability. Summary for Policy Makers, *Tech. rep.*, Intergovernmental Panel on Climate Change (IPCC).

IPCC (2007), IPCC: Fourth Assessment Report, *Tech. rep.*, Intergovernmental Panel on Climate Change (IPCC).

Jager, W., and M. Janssen (2003), The need for and development of behaviourally realistic agents, in *Multi-Agent-Based Simulation II*, *Lecture Notes in Computer Science*, vol. 2581, edited by J. Sichman, F. Bousquet, and P. Davidsson, pp. 36–49, Springer Berlin.

Jennings, N., K. Sycara, and M. Wooldridge (1998), A roadmap of agent research and development, *Autonomous agents and multi–agent systems*, *1*(1), 7–38.

John, M., D. Pannell, and R. Kingwell (2005), Climate Change and the Economics of Farm Management in the Face of Land Degradation: Dryland Salinity in Western Australia, *Canadian Journal of Agricultural Economics*, *53*(4), 443–459, doi:10.1111/j.1744-7976.2005.00029.x.

Johnson, S., J. Stedinger, and K. Staschus (1991), Heuristic operating policies for reservoir system simulation, *Water Resources Research*, *27*(5), 673–685.

Julien, F. (2012), Hydropolitics is what societies make of it (or why we need a constructivist approach to the geopolitics of water), *International Journal of Sustainable Society*, *4*(1), 45–71.

Jury, W., and H. Vaux (2007), The emerging global water crisis: Managing scarcity and conflict between water users, in *Advances in Agronomy*, *Advances in Agronomy*, vol. 95, edited by Sparks, DL, pp. 1–76, Elsevier Academic Press INC, doi:10.1016/S0065-2113(07)95001-4.

Kang, Y., S. Khan, and X. Ma (2009), Climate change impacts on crop yield, crop water productivity and food security - A review, *Progress* *in Natural Science-Materials InternationalPROGRESS IN NATURAL SCIENCE-MATERIALS INTERNATIONAL*, *19*(12), 1665–1674, doi: 10.1016/j.pnsc.2009.08.001.

Kantanantha, N., N. Serban, and P. Griffin (2010), Yield and price forecasting for stochastic crop decision planning, *Journal of agricultural, biological, and environmental statistics*, *15*(3), 362–380.

Kasprzyk, J., P. Reed, G. Characklis, and B. Kirsch (2012), Many-objective de Novo water supply portfolio planning under deep uncertainty , *Environmental Modelling & Software*, *34*(0), 87–104, doi: http://dx.doi.org/10.1016/j.envsoft.2011.04.003.

Keim, D., F. Mansmann, J. Schneidewind, and H. Ziegler (2006), Challenges in visual data analysis, in *Information Visualization, 2006. IV 2006. Tenth International Conference on*, pp. 9–16, IEEE, London, UK.

Kibaroglu, A., and I. Unver (2000), An Institutional Framework for Facilitating Cooperation in the Euphrates-Tigris River Basin, *International Negotiation*, *5*(2), 311–330.

Klemeš, V. (1977), Value of information in reservoir optimization, *Water Resources Research*, *13*(5), 837–850.

Knowles, J., and D. Corne (2002), On metrics for comparing non–dominated sets, in *Proceedings of the 2002 World Congress on Computational Intelligence (WCCI)*, pp. 711–716, IEEE Computer Society.

Kollat, J., and P. Reed (2007), A framework for Visually Interactive Decision-making and Design using Evolutionary Multi-objective Optimization (VIDEO), *Environmental Modelling & Software*, *22*(12), 1691–1704.

Kollat, J., P. Reed, and R. Maxwell (2011), Many–objective groundwater monitoring network design using bias–aware ensemble kalman filtering, evolutionary optimization, and visual analytics, *Water Resources Research*, *47*(W02529), –.

Korf, R. (1985), Depth–first iterative–deepening: an optimal admissible tree search, *Artificial Intelligence*, *27*(1), 97–109.

Korhonen, P., and J. Wallenius (1988), A pareto race, *Naval Research Logistics (NRL)*, *35*(6), 615–623.

Kormushev, P., and D. Caldwell (2012), Direct policy search reinforcement learning based on particle filtering, in *Proceedings of the 10th European Workshop on Reinforcement Learning.*

Koutsoyiannis, D., and A. Economou (2003), Evaluation of the parameterization-simulation-optimization approach for the control of reservoir systems, *Water Resources Research*, *39*(6), 1170–1187.

Kurková, V., and M. Sanguineti (2001), Bounds on rates of variable-basis and neural-network approximation, *IEEE Transactions on Information Theory*, *47*(6), 2659–2665.

Labadie, J. (2004), Optimal operation of multireservoir systems: State-of-the-art review, *Journal of Water Resources Planning and Management*, *130*(2), 93–111.

Lall, U., and A. Sharma (1996), A nearest neighbor bootstrap for resampling hydrologic time series, *Water Resources Research*, *32*(3), 679–693, doi:10.1029/95WR02966.

Laplace, P. (1951), *Essai Philosophique sur les probabilités (1820). English translation: Philosophical Essays on Probabilities*, New York: Dover.

Larichev, O. (1992), Cognitive validity in design of decision–aiding techniques, *Journal of Multi–Criteria Decision Analysis*, *1*(3), 127–128.

Le, Q., R. Seidl, and R. Scholz (2012), Feedback loops and types of adaptation in the modelling of land-use decisions in an agent-based simulation, *Environmental Modelling & Software*, *27-28*, 83–96.

Léauté, T., B. Ottens, and R. Szymanek (2009), FRODO 2.0: An Open-Source Framework for Distributed Constraint Optimization, in *Proceedings of the IJCAI'09 Distributed Constraint Reasoning Workshop (DCR'09)*, pp. 160–164, Pasadena, California, USA, url: http://frodo2.sourceforge.net.

Leavesley, G. (1994), Modeling the effects of climate–change on water resources - a review, *Climatic Change*, *28*(1-2), 159–177, doi: 10.1007/BF01094105.

Lenderink, G., B. van den Hurk, E. van Meijgaard, A. van Ulden, and H. Cuijpers (2003), *Simulation of present-day climate in RACMO2: first results and model developments*, Koninklijk Nederlands Meteorologisch Instituut.

Liu, J., et al. (2007), Complexity of coupled human and natural systems., *Science*, *317*, 1513–1516, doi:10.1126/science.1144004.

Lotov, A., and K. Miettinen (2008), Visualizing the Pareto frontier, in *Multiobjective optimization*, edited by J. Branke, K. Deb, K. Miettinen, and R. Slowinski, pp. 213–243, Springer.

Lotov, A., O. Chernykh, V. Bushenkov, H. Wallenius, and J. Wallenius (1998), Interactive decision maps, with an example illustrating ocean waste management decisions, *Lecture notes in economics and mathematical systems*, pp. 313–323.

Lotov, A., V. Bushenkov, and G. Kamenev (2004), *Interactive Decision Maps Approximation and Visualization of Pareto Frontier*, Springer–Verlag, Heidelberg, D.

Loucks, D., and O. Sigvaldason (1982), Multiple-reservoir operation in North America., in *The operation of multiple reservoir systems*, edited by Z. Kaczmarck and J. Kindler, IIASA Collab. Proc. Ser., pp. 1–103.

Loucks, D., J. Stedinger, and D. Haith (1980), *Water resources systems, planning and analysis*, Prentice-Hall, Englewood Cliffs, NJ.

Loucks, D., E. van Beek, J. Stedinger, J. Dijkman, and M. Villars (2005), *Water Resources Systems Planning and Management: An Introduction to Methods, Models and Applications*, UNESCO, Paris, France.

Louie, P., W. Yeh, and N. Hsu (1984), Multiobjective water resources management planning, *Journal of Water Resources Planning and Management*, *110*(1), 39–56.

Lund, J., and J. Guzman (1999), Derived operating rules for reservoirs in series or in parallel, *Journal of Water Resources Planning and Management*, *125*(3), 143–153.

Maass, A., M. Hufschmidt, R. Dorfman, H. Thomas Jr, S. Marglin, and G. Fair (1962), *Design of water–resource systems*, Harvard University Press Cambridge, Mass.

Madani, K. (2010), Game theory and water resources, *Journal of Hydrology*, *381*(3-4), 225–238.

Madani, K., and J. Lund (2012), California's Sacramento–San Joaquin Delta Conflict: From Cooperation to Chicken, *Journal of Water Resources Planning and Management*, pp. 90–99, doi: 10.1061/(ASCE)WR.1943-5452.0000164.

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, Wisa, Poland.

Maier, H., A. Jain, G. Dandy, and K. Sudheer (2010), Methods used for the development of neural networks for the prediction of water resource variables in river systems: Current status and future directions, *Environmental Modelling & Software*, *25*(8), 891–909.

Mall, R., R. Singh, A. Gupta, G. Srinivasan, and L. Rathore (2006), Impact of climate change on Indian agriculture: A review, *Climatic Change*, *78*(2–4), 445–478, doi:10.1007/s10584-005-9042-x.

Mareels, I., E. Weyer, S. Ooi, M. Cantoni, Y. Li, and G. Nair (2005), Systems engineering for irrigation systems: Successes and challenges, *Annual reviews in control*, *29*(2), 191–204, doi: 10.1016/j.arcontrol.2005.08.001.

Marques, G., J. Lund, and R. Howitt (2005), Modeling irrigated agricultural production and water use decisions under water supply uncertainty, *Water Resources Research*, *41*(8), doi:10.1029/2005WR004048.

Marris, E. (2008), More Crop per Drop, *Nature*, *452*, 273–277.

Maskin, E. (2008), Mechanism Design : How to Implement Social Goals, *The American Economic Review*, *98*(3), 567–576.

Mayne, D., J. Rawlings, C. Rao, and P. Scokaert (2000), Constrained model predictive control: stability and optimality, *Automatica*, *36*(2), 789–814.

Mazvimavi, D., and P. Wolski (2006), Long–term variations of annual flows of the Okavango and Zambezi Rivers, *Physics and Chemistry of the Earth*, *31*(15–16), 944–951, doi:10.1016/j.pce.2006.08.016.

McCann, L., and K. Easter (1999), Transaction costs of policies to reduce agricultural phosphorous pollution in the Minnesota river, *Land Economics*, *75*, 402–414.

McCann, L., B. Colby, K. Easter, A. Kasterine, and K. Kuperan (2005), Transaction cost measurement for evaluating environmental policies, *Ecological Economics*, *52*, 527–542.

McCarl, B. (2006), U.S. Agriculture in the Climate Change Squeeze: Part 1: Sectoral Sensitivity and Vulnerability, *Tech. rep.*, National Environmental Trust.

McDonald, R., P. Green, D. Balk, B. Fekete, C. Revenga, M. Todd, and M. Montgomery (2011), Urban growth, climate change, and freshwater availability, *Proceedings of the National Academy of Sciences*, *108*(15), 6312–6317.

Menache, I., S. Mannor, and N. Shimkin (2005), Basis function adaptation in temporal difference reinforcement learning, *Annals of operations research*, *134*(1), 215–238, doi:10.1007/s10479-005-5732-z.

Miller, B., and D. Goldberg (1996), Optimal sampling for genetic algorithms, in *Proceedings of the Artificial Neural Networks in Engineering (ANNIE'96) conference*, vol. 6, pp. 291–297.

Milly, P., J. Betancourt, M. Falkenmark, R. Hirsch, Z. Kundzewicz, D. Lettenmaier, and R. Stouffer (2008), Stationarity is dead, *Science*, *319*, 573–574.

Modi, P., H. Jung, M. Tambe, W. Shen, and S. Kulkarni (2001), A dynamic distributed constraint satisfaction approach to resource allocation, in *Principles and Practice of Constraint Programming – CP 2001*, pp. 685–700, Springer.

Modi, P., W. Shen, M. Tambe, and M. Yokoo (2005), Adopt: asynchronous distributed constraint optimization with quality guarantees, *Artificial Intelligence*, *161*(1-2), 149–180.

Momtahen, S., and A. Dariane (2007), Direct search approaches using genetic algorithms for optimization of water reservoir operating policies, *Journal of Water Resources Planning and Management*, *133*(3), 202–209, doi:10.1061/(ASCE)0733-9496(2007)133:3(202).

Momtahen, S., and A. Dariane (2009), Optimization of Multireservoir Systems Operation Using Modified Direct Search Genetic Algorithm, *Jurnal of Water Resources Planning and Management*, *135*(3), 141–148, doi:10.1061/(ASCE)0733-9496(2009)135:3(141).

Naivinit, W., C. Le Page, G. Trébuil, and N. Gajaseni (2010), Participatory agent-based modeling and simulation of rice production and labor migrations in northeast thailand, *Environmental Modelling & Software*, *25*(11), 1345–1358.

Nalbantis, I., and D. Koutsoyiannis (1997), A parametric rule for planning and management of multiple-reservoir systems, *Water Resources Research*, *33*(9), 2165–2177.

Nash, J. (1951), Non-cooperative games, *The Annals of Mathematics*, *54*(2), 286–295.

Neitsch, S., J. Arnold, J. Kiniry, and J. Williams (2011), Soil and Water Assessment Tool Theoretical Documentation Version 2009, *Tech. Rep. 406*, Grassland, Soil and Water Research Laboratory–Agricultural Research Service Blackland Research Center–Texas AgriLife Research, College Station, Texas.

Neumann, K., E. Stehfest, P. Verburg, S. Siebert, C. Muller, and T. Veldkamp (2011), Exploring global irrigation patterns: A multilevel modelling approach, *Agricultural Systems*, *104*(9), 703–713, doi:10.1016/j.agsy.2011.08.004.

Ng, T., J. Eheart, X. Cai, and J. Braden (2011), An agent-based model of farmer decision-making and water quality impacts at the watershed scale under markets for carbon allowances and a second-generation biofuel crop, *Water Resources Research*, *47*, doi:10.1029/2011WR010399.

Nguyen, N., J. Shortle, P. Reed, and T. Nguyen (2013), Water quality trading with asymmetric information, uncertainty and transaction costs: A stochastic agent-based simulation, *Resource and Energy Economics*, *35*(1), 60–90, doi:10.1016/j.reseneeco.2012.09.002.

Nicklow, J., et al. (2010), State of the Art for Genetic Algorithms and Beyond in Water Resources Planning and Management, *Journal of Water Resources Planning and Management*, *136*(4), 412–432, doi:10.1061/(ASCE)WR.1943-5452.0000053.

Niewiadomska-Szynkiewicz, E., K. Malinowski, and A. Karbowski (1996), Predictive methods for real-time control of flood operation of a multireservoir system: Methodology and comparative study, *Water Resources Research*, *32*(9), 2885–2895.

Nowak, K., J. Prairie, B. Rajagopalan, and U. Lall (2010), A nonparametric stochastic approach for multisite disaggregation of annual to daily streamflow, *Water Resources Research*, *46*, doi:10.1029/2009WR008530.

O'Hare, G., and N. Jennings (1996), *Foundations of distributed artificial intelligence*, vol. 9, Wiley.

Oliveira, R., and D. P. Loucks (1997), Operating rules for multi reservoir systems, *Water Resources Research*, *33*, 839–852.

Ostadrahimi, L., M. A. Mariño, and A. Afshar (2012), Multi-reservoir operation rules: multi-swarm pso-based optimization approach, *Water resources management*, *26* (2), 407–427.

Pannell, D. (2008), Public benefits, private benefits, and policy intervention for land-use change for environmental benefits, *Land Economics*, *84*, 225–240.

Pareto, V. (1964), *Cour d'economie politique*, Geneve: Librarie Droz (the first edition in 1896).

Paudel, K., and L. Hatch (2012), Global warming, impact on agriculture and adaptation strategy, *Natural Resource Modeling*, *25* (3), 456–481, doi:10.1111/j.1939-7445.2012.00127.x.

Peng, H., F. Long, and C. Ding (2005), Feature selection based on mutual information: criteria of max-dependency, max-relevance, and minimum redundancy, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *27* (8), 1226–1238.

Phalan, B., M. Bertzky, S. Butchart, P. Donald, J. Scharlemann, and A. Stattersfield, A.J. Balmford (2013), Crop expansion and conservation priorities in tropical countries, *Plos one*, *8* (1), doi: 10.1371/journal.pone.0051759.

Pianosi, F., and R. Soncini-Sessa (2009), Real-time management of a multipurpose water reservoir with a heteroscedastic inflow model, *Water Resources Research*, *45* (10), W10,430.

Pianosi, F., X. Quach, and R. Soncini-Sessa (2011), Artificial Neural Networks and Multi Objective Genetic Algorithms for water resources management: an application to the Hoabinh reservoir in Vietnam, in *Proceedings of the 18th IFAC World Congress*, Milan, Italy.

Piccardi, C., and R. Soncini-Sessa (1991), Stochastic dynamic programming for reservoir optimal control: Dense discretization and inflow correlation assumption made possible by parallel computing, *Water Resources Research*, *27* (5), 729–741.

Pope, V., M. Gallani, P. Rowntree, and R. Stratton (2000), The impact of new physical parametrizations in the Hadley Centre climate model: HadAM3, *Climate Dynamics*, *16* (2-3), 123–146.

Portmann, F., S. Siebert, and P. Döll (2010), Mirca2000–global monthly irrigated and rainfed crop areas around the year 2000: A new high–resolution data set for agricultural and hydrological modeling, *Global Biogeochemical Cycles*, *24*(1), doi:10.1029/2008GB003435.

Poteete, M. J., A.M. (2010), *Working Together: Collective Action, the Commons and Multiple Methods in Practice*, Princeton University Press.

Powell, W. (2007), *Approximate Dynamic Programming: Solving the curses of dimensionality*, Wiley, NJ.

Raes, D., P. Steduto, T. Hsiao, and E. Fereres (2009), AquaCrop–The FAO Crop Model to Simulate Yield Response to Water: II. Main Algorithms and Software Description, *101*(3), 438–447, doi:10.2134/agronj2008.0140s.

Rajagopalan, B., J. Salas, and U. Lall (2010), Stochastic methods for modeling precipitation and streamflow, in *Advances in Data–based Approaches for Hydrologic Modeling and Forecasting*, edited by B. Sivakumar and R. Berndtsson, World Scientific, pp. 17–52, Singapore.

Ralha, C., C. Abreu, C. Coelho, A. Zaghetto, B. Macchiavello, and R. Machado (2013), A multi-agent model system for land-use change simulation, *Environmental Modelling & Software.*

Randall, D., L. Cleland, C. Kuehne, G. Link, and D. Sheer (1997), Water Supply Planning Simulation Model Using Mixed–Integer Linear Programming "Engine", *Journal of Water Resources Planning and Management*, *123*(2), 116–124, doi:10.1061/(ASCE)0733-9496(1997)123:2(116).

Reed, P., D. Hadka, J. Herman, J. Kasprzyk, and J. Kollat (2013), Evolutionary Multiobjective Optimization in Water Resources: The Past, Present, and Future, *Advances in Water Resources*, *51*, 438–456.

Rocha, A., and I. Simmonds (1997), Interannual variability of South–Eastern African summer rainfall. Part1: relationships with air–sea interaction processes, *International Journal of Climatology*, *17*, 235–265.

Rosenstein, M., and A. Barto (2001), Robot weightlifting by direct policy search, in *International Joint Conference on Artificial Intelligence*, vol. 17, pp. 839–846, Citeseer.

Rosenzweig, C., and F. Tubiello (2007), Adaptation and mitigation strategies in agriculture: an analysis of potential synergies, *Mitigation and Adaptation Strategies for Global Change*, *12*, 855–873, doi: 10.1007/s11027-007-9103-8.

Rückstiess, T., F. Sehnke, T. Schaul, D. Wierstra, Y. Sun, and J. Schmidhuber (2010), Exploring parameter space in reinforcement learning, *Paladyn, Journal of Behavioral Robotics*, *1*(1), 14–24.

Russell, S., and P. Norvig (1995), *Artificial Intelligence: a Modern Approach*, Prentice-Hall, NY.

Salas, J., J. Delleur, V. Yevjevich, and W. Lane (1980), *Applied Modelling of Hydrologic Time Series*, Water Resources Publications, Littleton, CO, USA.

Saltelli, A., M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and T. S. (2008), *Global sensitivity analysis: the primer*, Wiley.

Savage, L. (1951), The theory of statistical decision, *Journal of the American Statistical association*, *46*(253), 55–67.

Scattolini, R. (2009), Architectures for distributed and hierarchical Model Predictive Control - A review, *Journal of Process Control*, *19*(5), 723–731.

Schaefli, B., B. Hingray, and A. Musy (2007), Climate change and hydropower production in the Swiss Alps: quantification of potential impacts and related modelling uncertainties, *Hydrology and Earth System Sciences*, *11*(3), 1191–1205.

Schmidhuber, J. (2001), Sequential Decision Making Based on Direct Search, in *Sequence Learning*, *Lecture Notes in Computer Science*, vol. 1828, edited by R. Sun and C. Giles, pp. 213–240, Springer Berlin Heidelberg.

Schreinemachers, P., and T. Berger (2011), An agent–based simulation model of human–environment interactions in agricultural systems, *Environmental Modelling & Software*, *26*(7), 845–859.

Seo, S., R. Mendelsohn, and M. Munasinghe (2005), Climate change and agriculture in sri lanka: A ricardian valuation, *Environment and Development Economics*, *10*(Part 5), 581–596, doi: 10.1017/S1355770X05002044.

Sheer, D. (2010), Dysfunctional Water Management: Causes and Solutions, *Journal of Water Resources Planning and Management*, *136*(1), 1–4.

Sheer, D., and A. Dehoff (2009), Science–Based Collaboration: Finding Better Ways to Operate the Conowingo Pond, *Journal of the American Water Works Association*, *101*(6), 20–22.

Sheikhmohammady, M., K. Madani, M. Moradi, and I. Behmanesh (2011), Finding the Best Legal Governance Regime for the Caspian Sea through Multi-Criteria Decision-Making Methods, in *Proceedings of the IEEE International Conference on Systems Man and Cybernetics Conference*, pp. 3029–3034.

Shen, W., C. Chuong, and P. Will (2002), Simulating self-organization for multi-robot systems, in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3, pp. 2776–2781, IEEE.

Shoham, Y., and K. Leyton-Brown (2009), *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Princeton University Press.

Siebert, S., and P. Döll (2010), Quantifying blue and green virtual water contents in global crop production as well as potential production losses without irrigation, *Journal of Hydrology*, *384*(3-4, SI), 198–217, doi:10.1016/j.jhydro1.2009.07.031.

Sobol, I. (2001), Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Mathematics and computers in simulation*, *55*(1-3), 271–280.

Soncini-Sessa, R., A. Castelletti, and E. Weber (2007a), *Integrated and participatory water resources management: Theory*, Elsevier, Amsterdam, NL.

Soncini-Sessa, R., , F. Cellina, F. Pianosi, and E. Weber (2007b), *Integrated and participatory water resources management: Practice*, Elsevier, Amsterdam, NL.

Steduto, P., T. Hsiao, D. Raes, and E. Fereres (2009), AquaCrop–The FAO Crop Model to Simulate Yield Response to Water: I. Concepts and Underlying Principles, *Agronomy Journal*, *101*(3), 426–437, doi: 10.2134/agronj2008.0139s.

Swartz, P. (2006), Conowingo pond management plan, *Tech. Rep. 242*, Susquehanna River Basin Commission.

Sycara, K. (1998), Multiagent systems, *AI magazine*, *19*(2), 79.

Sycara, K., S. Roth, N. Sadeh, and M. Fox (1991), Distributed constrained heuristic search, *IEEE Transactions on Systems, Man and Cybernetics*, *21*(6), 1446–1461.

Teegavarapu, R., and S. Simonovic (2001), Optimal operation of water resource systems: trade–offs between modeling and practical solutions, in *Integrated water resources management*, vol. 272, edited by S. S. Mariño M.A., pp. 257–263.

Tejada-Guibert, J., S. Johnson, and J. Stedinger (1995), The value of hydrologic information in stochastic dynamic programming models of a multireservoir system, *Water Resources Research*, *31*(10), 2571–2579.

Thiele, L., K. Miettinen, P. Korhonen, and J. Molina (2009), A preference–based evolutionary algorithm for multi-objective optimization, *Evolutionary Computation*, *17*(3), 411–436.

Thomas, J., and K. Cook (2005), *Illuminating the path: The research and development agenda for visual analytics*, IEEE Computer Society Press.

Thomas, J., and K. Sycara (1998), Heterogeneity, stability, and efficiency in distributed systems, in *Proceedings of the IEEE International Conference on Multi Agent Systems*, pp. 293–300, IEEE.

Tilmant, A., and W. Kinzelbach (2012), The cost of noncooperation in international river basins, *Water Resources Research*, *48*(W01503).

Tilmant, A., L. Beevers, and B. Muyunda (2010), Restoring a flow regime through the coordinated operation of a multireservoir system: The case of the Zambezi River basin, *Water Resources Research*, *46*(W07533), 1–11.

Tilmant, A., W. Kinzelbach, D. Juizo, L. Beevers, D. Senn, and C. Casarotto (2012), Economic valuation of benefits and costs associated with the coordinated development and management of the Zambezi river basin, *Water Policy*, *14*(3), 490–508.

Timberlake, J. (2000), *Biodiversity of the Zambezi Basin*, Occasional Publications in Biodiversity No. 9. Biodiversity Foundation for Africa, Famona, Bulawayo, Zimbabwe.

Timmerman, J., and S. Langaas (2005), Water information: what is it good for? The use of information in transboundary water management, *Regional Environmental Change*, *5*(4), 177–187.

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

Tubiello, F., J. Soussana, and S. Howden (2007), Crop and pasture response to climate change, *Proceedings of the National Academy of Sciences of the United States of America*, *104*(50), 19,686–19,690, doi: 10.1073/pnas.0701728104.

UNEP (2002), *Atlas of international freshwater agreements*, United Nations Environment Programme, Nairobi, Kenya.

U.S. Army Corps of Engineers (1977), *Reservoir System Analysis for Conservation, Hydrologic Engineering Methods for Water Resources Development*, Hydrologic Engineering Center, Davis, CA.

USDA-SCS (1972), *National engineering handbook.*, vol. 630 hydrology, section 4, Department of Agricolture - Soil Conservation Service, U.S. Government Printing Office, Washington, DC, USA.

van Oel, P., M. Krol, A. Hoekstra, and R. Taddei (2010), Feedback mechanisms between water availability and water use in a semi–arid river basin: A spatially explicit multi–agent simulation approach, *Environmental Modeling & Software*, *25*(4), 433–443, doi: 10.1016/j.envsoft.2009.10.018.

Vidal, J. (2009), *Fundamentals of Multiagent Systems*, http://multiagent.com/2010/02/multiagent-systems-textbook.html.

Washington, W., and C. Parkinson (2005), *An introduction to three-dimensional climate modeling*, University science books.

Washington, W., L. Buja, and A. Craig (2009), The computational future for climate and earth system models: on the path to petaflop and beyond, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *367*(1890), 833–846.

Waterbury, J. (1987), Legal and institutional arrangements for managing water resources in the Nile Basin, *International Journal of Water Resources Development*, *3*(2), 92–104.

Waterbury, J. (1997), Between Unilateralism and Comprehensive Accords: Modest Steps toward Cooperation in International River Basins, *International Journal of Water Resources Development*, *13*(3), 279–290.

Watkins, K. (2006), *Human Development Report. Beyond scarcity: power, poverty and global water crisis*, United Nations Development Programme, New York, USA.

Wehenkel, L. (1998), *Automatic Learning Techniques in Power Systems*, Kluwer Academic, Boston, MA.

Whittington, D., X. Wu, and C. Sadoff (2005), Water resources management in the Nile basin: the economic value of cooperation, *Water Policy*, *7*, 227–252.

Wierzbicki, A. (1980), *The use of reference objectives in multiobjective optimization*, Springer.

Wolf, A. (1998), Conflict and Cooperation Along International Waterways, *Water Policy*, *1*(2), 251–265.

Wolf, A., A. Kramer, A. Carius, and G. Dabelko (2006), *Water can be a pathway to peace, not war*, Woodrow Wilson International Center for Scholars.

Woodruff, M., P. Reed, and T. Simpson (2013), Many objective visual analytics: rethinking the design of complex engineered systems, *Structural and Multidisciplinary Optimization*, pp. 1–19, doi: 10.1007/s00158-013-0891-z.

Woodward, R., and Y. Wui (2001), The economic value of wetlands services: A meta-analysis, *Ecological Economics*, *37*, 257–270.

Wooldridge, M. (2009), *An Introduction to MultiAgent Systems*, second ed., Wiley, NY.

Wooldridge, M., and N. Jennings (1995), Intelligent agents: Theory and practice, *Knowledge engineering review*, *10*(2), 115–152.

Wu, W., A. Simpson, and H. Maier (2010), Accounting for Greenhouse Gas Emissions in Multiobjective Genetic Algorithm Optimization of Water Distribution Systems, *Journal of Water Resources Planning and Management*, *136*(2), 146–155.

Wu, X., and D. Whittington (2006), Incentive compatibility and conflict resolution in international river basins: A case study of the Nile Basin, *Water Resources Research*, *42*(2).

Yang, Y., X. Cai, and D. Stipanović (2009), A decentralized optimization algorithm for multiagent system-based watershed management, *Water Resources Research*, *45*(8), 1–18.

Yates, D., S. Gangopadhyay, B. Rajagopalan, and K. Strzepek (2003), A technique for generating regional climate scenarios using a nearest–neighbor algorithm, *Water Resources Research*, *39*(7), doi: 10.1029/2002WR001769.

Yeh, W. (1985), Reservoir management and operations models: a state of the art review, *Water Resources Research*, *21 (12)*, 1797–1818.

Yoffe, S., A. Wolf, and M. Giordano (2003), Conflict and cooperation over international freshwater resources: indicators of basins at risk, *Journal of the American Water Resources Association*, *39*(5), 1109–1126, doi:10.1111/j.1752-1688.2003.tb03696.x.

Yokoo, M., and K. Hirayama (2000), Algorithms for Distributed Constraint Satisfaction: A Review, *Autonomous Agents and Multi-Agent Systems*, *3*(2), 185–207.

Yokoo, M., E. Durfee, T. Ishida, and K. Kuwabara (1992), Distributed constraint satisfaction for formalizing distributed problem solving, in *Proceedings of the Twelfth IEEE International Conference on Distributed Computing Systems*.

Zeitoun, M., and N. Mirumachi (2008), Transboundary water interaction I: Reconsidering conflict and cooperation, *International Environmental Agreements: Politics, Law and Economics*, *8*(4), 297–316.

Zeitoun, M., and J. Warner (2006), Hydro-hegemony: A framework for analysis of transboundary water conflicts, *Water Policy*, *8*(5), 435–460.

Zhang, H., and T. Oweis (1999), Water–yield relations and optimal irrigation scheduling of wheat in the Mediterranean region, *Agricultural Water Management*, *38*, 195–211.

Zitzler, E., L. Thiele, M. Laumanns, C. Fonseca, and V. da Fonseca (2003), Performance assessment of multiobjective optimizers: an analysis and review, *IEEE Transactions on Evolutionary Computation*, *7*(2), 117–132.

Zoppoli, R., M. Sanguineti, and T. Parisini (2002), Approximating networks and extended ritz method for the solution of functional optimization problems, *Journal of Optimization Theory and Applications*, *112*(2), 403–440.