Multiagent Systems and Distributed Constraint Reasoning for Regulatory Mechanism Design in Water Management

Matteo Giuliani, Ph.D., M.ASCE[1]; Andrea Castelletti, Ph.D., M.ASCE[2]; Francesco Amigoni[3]; and Ximing Cai, Ph.D., M.ASCE[4]

Introduction

Water resources management in large-scale systems is often associated with multiple institutionally-independent decision makers (DMs), both within national jurisdictions and in transboundary contexts (Yoffe et al. 2003; Wolf et al. 2006; Zeitoun and Mirumachi 2008). The presence of distributed localized decision processes is challenging the traditional centralized approach to water resources management as underlying much of the literature (Wallace et al. 2003; Wu and Whittington 2006; Tilmant and Kinzelbach 2012). Classic top-down approaches might not be suitable for analyzing such problems as they neglect the principle of individual-rationality, focusing on the search of solutions that maximize the efficiency at the system-level, also known as social planner's solution or maximum social welfare (Loucks et al. 2005; Dariane and Momtahen 2009; Zoltay et al. 2010; Tilmant et al. 2010), possibly defined by multiple criteria, such as flood prevention, hydropower production, and water supply.

To actually attain such an efficient use of the available water resources, this approach requires a cooperative attitude of the involved parties, who agree on adopting a fully coordinated strategy on water allocation and distribution in time and space and full knowledge of the current system conditions. In many situations these conditions are satisfied and yield win-win solutions, where either the maximization of the individual benefits is equivalent to the maximization of the system benefit or the development of negotiation processes allows the identification of implementable compromise solutions (Lund and Palmer 1997; Lubell et al. 2002; Soncini-Sessa et al. 2007; Teasley and McKinney 2011; Bhaduri and Liebe 2013).

Increasing environmental awareness and emerging trends such as water trading, energy market, deregulation, and democratization of water-related services are reducing the general validity of the centralized approach in some contexts and, particularly, in transboundary systems, where the DMs belonging to different countries may have local internalities, which contrast the system-level goals (Young 1986). In these situations, the centralized social planner's solution, though interesting from a conceptual point of view as it allows for quantifying the best achievable performance and obtaining insights on strategies to foster cooperation (Anghileri et al. 2013; Giuliani and Castelletti 2013; Marques and Tilmant 2013), turns out to be of low practical meaning given the actual decision-making context (Waterbury 1987; Whittington et al. 2005) and the differences in the political, social, and economic status of the parties involved (Madani 2013). At the other extreme, a totally uncoordinated setting, where the DMs independently pursue their local objectives on the basis of the individual-rationality

[1]Postdoctoral Research Fellow, Dept. of Electronics, Information, and Bioengineering, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy (corresponding author). E-mail: matteo.giuliani@polimi.it

[2]Assistant Professor, Dept. of Electronics Information, and Bioengineering, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy. E-mail: andrea.castelletti@polimi.it

[3]Associate Professor, Dept. of Electronics Information, and Bioengineering, Politecnico di Milano, P.za Leonardo da Vinci, 32, 20133 Milano, Italy. E-mail: francesco.amigoni@polimi.it

[4]Professor, Dept. of Civil and Environmental Engineering, Univ. of Illinois at Urbana-Champaign, 301 N. Mathews Ave., Urbana, IL 61801.

Note. This manuscript was submitted on October 24, 2013; approved on May 21, 2014; published online on July 23, 2014. Discussion period open until December 23, 2014; separate discussions must be submitted for individual papers.

principle, represents a more realistic picture of such institutional frameworks. These individualistic behaviors generally induce undesired outcomes at the system-level, possibly leading to tragedy of the commons (Hardin 1968). The problem structure and outcome might change when other actors are involved, such as river basin authorities promoting negotiation processes to incorporate, at different levels, cooperative agreements or coordination mechanisms (Madani and Lund 2012). Suboptimal solutions may be preferred and require coupling the traditional search for efficient solutions with the assessment of their acceptability, often referred to as stability or fairness (Dinar and Howitt 1997; Cardenas and Ostrom 2004; Madani and Hipel 2011; Read et al. 2014).

The focus of this paper is on this second class of problems, where the social planner's solution might be unfeasible because the problem is characterized by multiple, originally noncooperative DMs, nontransferable utility, impossibility of economic compensation due to the presence of noncommensurable objectives, and prohibitively high transaction costs. In this context, the mismatch between system-level efficiency and individual-level acceptability can be summarized as in Fig. 1. In one extreme there is represented the fully centralized solution, where all the DMs cooperatively maximize the total benefit of the system and, in the other extreme, an uncoordinated solution where the DMs consider their local objective functions only. The figure suggests that, according to Read et al. (2014), there is a trade-off between efficiency and acceptability: 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 unacceptable or impracticable. Clear examples are the failures of many international initiatives, such as the Zambezi River Basin Water Commission (ZAMCON) Protocol for the management of the Zambezi River at the river basin scale. The government of Zambia refused to sign any international agreement because its upstream position guarantees the possibility to autonomously meet the national targets, independently from any interaction with the other countries sharing part of the Zambezi River catchment and, in particular, with Mozambique. At the other extreme, noncoordinated management practices result in low performance at the system-level. It is worth noting that the utopia point (i.e., the absolute optima of both efficiency and acceptability, represented with a white square in Fig. 1) is often not feasible, except for 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 acceptable than the centralized ones. Different

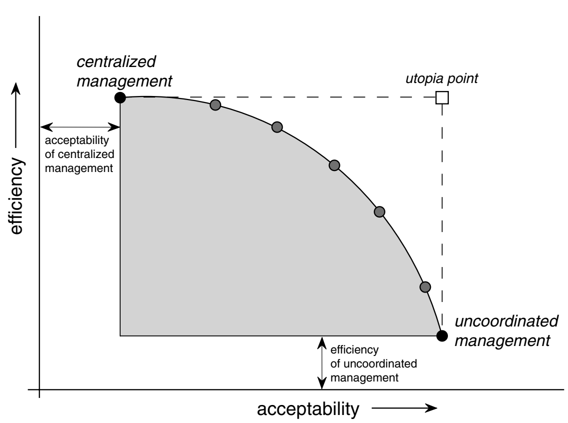

Fig. 1. Representation of the conflict between system-wide efficiency and individual-level acceptability for watershed management strategies

policy mechanisms (e.g., regulatory constraints or economic incentives) have been proposed in many research fields, such as economics and game theory, to explore this space (Pannell 2008). Yet most of these analyses provide descriptive tools based on what-if or scenario analyses, focusing on mechanisms designed on the basis of the empirical experiences of the water resources managers in charge to promote negotiated solutions. Mathematical and technological tools to identify such distributed solutions are nearly undeveloped and, therefore, the applicability of them within a decision-support systems has been so far limited.

In the present paper, the adoption of an agent-based approach and the combination of descriptive and prescriptive methods are proposed in order to provide informative tools capable of representing the actual decision-making context (e.g., removing the simplifying assumption of fully cooperative DMs), as well as decision support procedures, which recommend proper coordination mechanisms between the originally self-interested decision-making actors. The selection of a framework based on multiagent systems, or MAS (Wooldridge 2009) naturally allows the representation of multiple DMs or stakeholders (agents), which act in the same environment, thus influencing each other, and need to coordinate to maximize the system-wide efficiency. The adoption of this bottom-up, agent-based perspective aims to move beyond the traditional top-down, centralized, social planner's approach to water resources management, and it allows the analysis of different levels of cooperation (Watkins 2006). 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 watershed. Despite the illusory simplicity of this nondynamic, numerical 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 determining the releases from the dam), and secondary environmental concerns. In this paper, the focus is supporting the design of a regulatory-based coordination mechanism where a watershed authority is in charge of imposing soft (normative) constraints on the originally self-interested agents' decisions. Since the problem has nontransferable utilities, alternative approaches based on water market mechanisms (Matthews 2004; Ambec and Ehlers 2008; Zhao et al. 2013) are not viable options. The imposition of normative constraints yields a distributed, constrained problem, which can be effectively managed through distributed constraint reasoning, comprising algorithms developed for the resolution of distributed constraint satisfaction problems, or DCSP (Yokoo and Hirayama 2000), and distributed constraint optimization problems, or DCOP (Modi et al. 2005). In contrast with the approach proposed by Yang et al. (2009), these formulations allow the separation of soft (normative) and hard (physical) constraints, thus guaranteeing a priori the feasibility of the solutions. Conversely, in Yang et al. (2009), the feasibility of the solutions is checked a posteriori, allowing some candidate alternatives that violate physical constraints. Moreover, the use of the *Adopt* algorithm (Modi et al. 2005) for the resolution of DCSP and DCOP ensures global solution quality operating efficiently in a distributed process that attempts to narrow the number of exchanged messages.

In summary, this paper provides three main contributions: (1) the adoption of an MAS approach based on formal methods and tools, derived from the distributed constraint reasoning theory and related algorithms, to support the design of regulatory mechanisms improving the system-level efficiency in problems characterized by multiple, noncooperative DMs, (2) the discussion about the trade-off between efficiency-acceptability is coupled with methods allowing the development of numerical analysis and evidence supporting the arguments, and (3) for the first time *Adopt* algorithm, developed in the Autonomous Agents and MultiAgent Systems (AAMAS) domain, is applied in water resources problems, where the asynchronous and distributed nature of the algorithm (and of its extensions) has a great potential for solving large-scale distributed optimization problems.

The rest of the paper is organized as follows: the next section introduces the methodology, followed by the description of the case study. Results and discussion are then reported, while final remarks, along with issues for further research, are presented in the last section.

Methods and Tools

Multiagent Systems

The theory of multiagent systems (MAS) has been developed first in the distributed artificial intelligence (DAI) domain, a subfield of artificial intelligence, and has become a rather independent research field referred to as autonomous agents and multiagent systems (AAMAS). Nowadays, 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 work MAS is considered more as a mindset than a technology (Bonabeau 2002) that relies on the general definitions given in Wooldridge (2009) and Shoham and Leyton-Brown (2009) as follows: "an agent is defined as a computer system situated in some environment and capable of autonomous actions to meet its design objectives; multiagent systems are those systems that include multiple autonomous entities with either diverging information or diverging interests, or both."

Depending on the specific application domain, the agent design process might differently balance the following features (Wooldridge and Jennings 1995): (1) *reactivity*, namely agents being able to perceive the environment and timely respond to its changes, (2) *proactiveness*, namely agents being able to exhibit goal-directed behaviors, and (3) *social ability*, namely agents being 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 according to a stimulus-response type of behavior and respond to the present state of the environment (Sycara 1998). They usually do not look at history or plan their strategy over the future. This characteristic allows the design of simple *if-then* agent behaviors. Yet they make decisions based on local information only and do not predict the effect of their decisions on global behavior. These 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). 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 in which the agents are placed to a real number. Fully 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). 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).

From the early works developed in DAI, the applications of MAS dealing with distributed reasoning problems have rapidly covered a variety of domains [for a review, see Yeoh and Yokoo (2012) and references therein], such as distributed allocation of resources in disaster evacuation scenarios (Lass et al. 2008), management of power distribution networks (Kumar et al. 2009), or distributed coordination of logistics operation (Léauté and Faltings 2011). MAS approaches have become a widely used tool in several environmental modeling contexts (Athanasiadis 2005). The primary goal of most of these studies, also referred to as multiagent simulations [for a review, see Bousquet and Le Page (2004), An (2012), and references therein], is to simulate complex systems in order to evaluate macrolevel 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): (1) it provides a more natural description of a system, especially when it is composed of multiple, distributed, and autonomous agents, (2) it relaxes the hypothesis of homogeneity in a population of actually heterogeneous individuals, (3) it allows an explicit representation of spatial variability, and (4) it captures emergent global behaviors resulting from local interactions.

Purely reactive agents are largely adopted as a modeling approach to define behavioral rules that react to environmental changes (Le et al. 2012; Shafiee and Zechman 2013; Kanta and Zechman 2014). However, the prescriptive use of MAS models in decision support systems remains a challenge due to the mathematical complexity of the models, which requires a 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). In the water resources literature, the first contribution adopting proactive MAS for a nondynamic optimization problem was presented in Yang et al. (2009) and further developed by Giuliani et al. (2012). A similar approach was then adopted to optimize preseason farmers' decisions (Ng et al. 2011) and to simulate an optimization-driven water market (Huskova and Harou 2012) and emission trading (Nguyen et al. 2013). Giuliani and Castelletti (2013) proposed an agent-based optimization framework to assess the value of cooperation in large-scale, transboundary water resources systems.

Centralized and Distributed Problem Formulations

As described in the introduction, this work proposes an agent-based approach to support the definition of regulatory mechanisms favoring water allocation alternatives that are able to balance system-level efficiency and individual-level acceptability in the context of watersheds modeled as MAS. The presence of multiple agents (DMs), nontransferable utility, impossibility of economic compensation, and high transaction costs exacerbate the gap between the fully cooperative solution and the totally uncoordinated one, while distributed solutions may represent a sort of compromise between these two extreme situations. To introduce the differences between the considered approaches, the mathematical model of a centralized k-objective optimization problem is first defined

$$\max_{\mathbf{x}} \mathbf{f}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_k(\mathbf{x})] \tag{1a}$$

subject to

$$c_1(\mathbf{x}), c_2(\mathbf{x}), \ldots, c_r(\mathbf{x}) \leq 0 \tag{1b}$$

$$\mathbf{x} \in \mathcal{D} \tag{1c}$$

where $\mathbf{x} \in \mathbb{R}^n$ = decision vector; $\mathbf{f}(\mathbf{x}) = k$-dimensional objective vector; $f_i(\mathbf{x})(i = 1, \ldots, k) = i$th objective to be maximized; and Eq. (1*b*) defines r constraints on the values of $\mathbf{x}$, whose domain is $\mathcal{D}$. The solution of Eq. (1) does not yield a unique optimal solution but a set of Pareto-optimal solutions. Yet this set of Pareto-efficient solutions is obtained by assuming the presence of a social planner and a fully-cooperative agents' attitude. According to this hypothesis, some agents (e.g., the upstream water users) agree on decreasing their local benefit to improve the benefit of other users, a condition that is hardly pursuable in most of real-world institutional contexts. In such cases, it is likely that the DMs act according to the individual-rationality principle and look at their local objectives only, without considering the potentially negative externalities that their decisions produce at the system-level. Consequently, Eq. (1) can be reformulated as a sequence of k independent problems (assuming one-to-one correspondence between objectives and DMs), each one defined as

$$\max_{\mathbf{x}_i} f_i(\mathbf{x}_i, \mathbf{x}_{-i}) \tag{2a}$$

subject to

$$c_{i,1}(\mathbf{x}_i, \mathbf{x}_{-i}), c_{i,2}(\mathbf{x}_i, \mathbf{x}_{-i}), \ldots, c_{i,r}(\mathbf{x}_i, \mathbf{x}_{-i}) \leq 0 \tag{2b}$$

$$\mathbf{x}_i \in \mathcal{D}_i \tag{2c}$$

where $i = 1, \ldots, k$, $\mathbf{x}_i$ = decision variables of the ith agent; and $\mathbf{x}_{-i}$ = decision variables of all the other agents except i. The ith agent's decisions are optimized with respect to his local objective function f_i, which also depends on the decisions of the other agents. It is also worth noting that the constraints in Eq. (2*b*) depend on both the values of $\mathbf{x}_i$ and $\mathbf{x}_{-i}$. Note that the assumption that each DM solves a single-objective optimization problem can be easily removed to consider DMs dealing with multiobjective optimization problems.

Problem can be modeled by adopting an MAS framework. Each DM is represented by an agent, which acts in order to maximize its local objective function while at the same time satisfying all the constraints. This MAS model is inspired by distributed reasoning principles and is not assuming a strategic context (as in game theory-based approaches). For example, it does not require that the rationality of the agents is common knowledge between the agents themselves and that the agents know the whole system. In addition, this MAS framework allows also for considering different levels of interactions between the agents and for developing mechanism design strategies through the imposition of *normative constraints* to be added to those in Eq. (2*b*). The aim of these new constraints is to condition the individualistic, independent decisions of the agents, in order to increase the efficiency of the uncoordinated solution according to the ideal reference of the fully cooperative, centralized strategy. 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). Agent-based methods can be adopted to solve the resulting distributed constrained problems as discussed in the next section.

MAS Methods for Distributed Constraint Reasoning

Within the MAS field, a number of distributed constrained problems have been addressed by methods relying on distributed constraint reasoning, which comprise algorithms developed for distributed constraint satisfaction problems, or DCSPs (Yokoo and Hirayama 2000), and distributed constraint optimization problems, or DCOPs (Modi et al. 2005).

Formally, a DCSP consists of n decision variables, $\mathbf{x} = [x_1, x_2, \ldots, x_n]$, each assigned to a different agent, where the values of the variables are taken from finite, discrete domains $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_n$, and of a number of Boolean constraints over the values of these variables $c_1(\mathbf{x}), \ldots, c_r(\mathbf{x})$, where $c_j(\mathbf{x}) \in \{\text{true}, \text{false}\}\forall j$. 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 constraint satisfaction problems (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.

A DCOP instead consists of n decision variables, $\mathbf{x} = [x_1, x_2, \ldots, x_n]$, each assigned to a different agent, where the values of the variables are taken from finite, discrete domains $\mathcal{D}_1, \mathcal{D}_2, \ldots, \mathcal{D}_n$, and of a number of valued constraints over the values of these variables $c_1(\mathbf{x}), \ldots, c_r(\mathbf{x})$, where $c_j(\mathbf{x}) \in \mathbb{R}\forall j$. 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 $\min_{\mathbf{x}} g = \min_{\mathbf{x}} \sum_{j=1}^{r} w_j \cdot c_j(\mathbf{x})$, where w_j is the weight of $c_j(\mathbf{x})$.

Probably, the most known (though not the most efficient for all problems) distributed algorithm for solving DCOPs in the AAMAS literature is *Adopt* [Asynchronous Distributed OPTimization, see Modi et al. (2005)], which performs a global search of the solution space and provides theoretical guarantees on the optimality of the identified solutions. Also in the case of *Adopt*, agents are responsible for choosing the values of their variables. 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. *Adopt* assumes that constraints are at most binary (i.e., they involve one or two variables/agents), while no restrictions are placed on the topology of the constraints, so loops are allowed. *Adopt* operates asynchronously using only local communication, meaning that agents do not broadcast messages to every agent but only to neighboring agents in the attempt to limit the number of exchanged messages. This has the nice consequence that although agents might not know directly all the other agents, they are able to find the optimal solution of the DCOP. Extensions of the basic formulations of DCSP and DCOP allow each agent to control multiple variables as well as to consider agents solving multiobjective 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, consider the simple problem represented in Fig. 2, where each variable $x_i (i = 1, 2, 3)$ is assigned to an agent and has the same domain $\mathcal{D}_i = \{0, 1\}$. The objective function of each agent is $f_i = 10$ if $x_i = 1$ and $f_i = 0$ if $x_i = 0$. The two links in the figure represent two inequality constraints defined as $x_1 \neq x_3$ and $x_2 \neq x_3$. In the case of the DCSP, the constraints are Boolean.

The problem has two feasible solutions $\{x_1 = 1, x_2 = 1, x_3 = 0\}$ and $\{x_1 = 0, x_2 = 0, x_3 = 1\}$, 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 Fig. 2(b), the DCOP has one optimal solution only, namely $\{x_1 = 1, x_2 = 1, x_3 = 0\}$, which corresponds to a cost for constraint violations $g(x_1, x_2, x_3) = c(x_1, x_3) + c(x_2, x_3) = 0$ and to the maximum of system-level benefit (i.e., $f_1 + f_2 + f_3 = 20$). Note that, the DCOP problem has another feasible solution, $\{x_1 = 0, x_2 = 0, x_3 = 1\}$, with a higher cost for constraint violations $g(x_1, x_2, x_3) = 2$ and a lower system-level performance (i.e., $f_1 + f_2 + f_3 = 10$). Finally, the DCOP has other two solutions, which are the ones considered unfeasible in the DCSP formulation, that assume an infinite cost in the DCOP formulation and, consequently, will never be considered. The DCOP formulation is therefore more flexible as it allows for discarding the same unfeasible solutions as in the DCSP problem and also for distinguishing between the feasible solutions depending on their system-level performance.

Case Study Description

The proposed MAS-based mechanism design strategy is tested on a numerical case study first introduced by Yang et al. (2009). The system is composed of a Y-shaped river with one mainstream and one main tributary (Fig. 3). The mainstream one 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 water reservoir (just downstream with respect to the city water supply diversion) in order to produce hydropower energy. A second agricultural district diverts water from the tributary for irrigation purposes. 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; and
- A_6: ecological preservation in the mainstream.

A quadratic concave objective function is associated with each agent, which preserves the nonlinear characteristics of a real objectives: $f_i = a_i x_i^2 + b_i x_i + c_i$ [the values of the parameters, defined as in Yang et al. (2009), are reported in Table 1]. The quadratic formulation of the objective functions allows for obtaining 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. It is worth noting that the objective functions are different and, therefore, the same volume of water has a different effect on the benefit of different agents and, consequently, on the system benefit. Moreover, the six agents have different natures: a first group, i.e., A_1-A_2-A_3-A_4, includes *active agents*, 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 $f_i(x_i)$ (with $i = 1, \ldots, 4$); on the other side, agents A_5-A_6 are defined as *passive agents*, who do not make decisions but represent the ecological interests through the functions $f_5(x_5)$ and $f_6(x_6)$, which are explicitly optimized only in the centralized case.

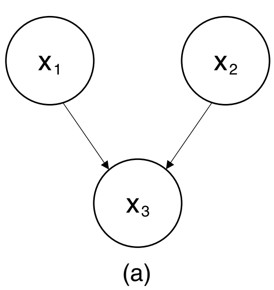

x_i	x_j	$c(x_i, x_j)$
0	0	∞
1	1	∞
0	1	1
1	0	0

(b)

Fig. 2. 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

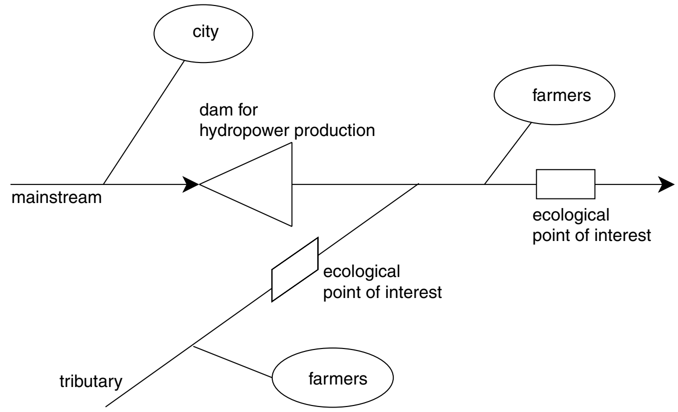

Fig. 3. Schematic map of the system [adapted from Yang et al. (2009)]

Table 1. Values of Parameters Defining Agents Objective Functions

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

Note: Coefficients a_i, b_i and c_i are dimensionless.

Assuming for simplicity a nondynamic 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

$$\max_{x_1} f_1(x_1) \text{ s.t. } x_1 \leq Q_1 \tag{3a}$$

$$\max_{x_2} f_2(x_2) \text{ s.t. } x_2 \leq S + Q_1 - x_1 \tag{3b}$$

$$\max_{x_3} f_3(x_3) \text{ s.t. } x_3 \leq Q_2 \tag{3c}$$

$$\max_{x_4} f_4(x_4) \text{ s.t. } x_4 \leq x_2 + Q_2 - x_3 \tag{3d}$$

$$f_5(x_5) \text{ s.t. } x_5 = Q_2 - x_3 \tag{3e}$$

$$f_6(x_6) \text{ s.t. } x_6 = x_2 + x_5 \tag{3f}$$

where Q_1 = mainstream inflow; Q_2 = tributary inflow; and S = 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 (Table 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.

Comparison among the Various Methods

In this paper the ideal, fully cooperative centralized solution is compared to three different distributed alternatives:

- An uncoordinated solution, where each active agent acts independently considering his 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 functions 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 physically feasible and has to satisfy the soft (normative) constraints. 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. In this case, there is no difference between hard (physical) and soft (normative) constraints; and
- 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.

Table 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

Results

The proposed MAS approach requires first computing the two extreme solutions of Eq. (3), namely the centralized and the uncoordinated one, for each hydrological scenario. We assume 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, Eq. (3) 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 Fig. 4. Not surprisingly, the centralized solution (black bars) produces higher benefits than the uncoordinated one (white 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 regulatory mechanism, the watershed authority imposes the following set of normative constraints in order to protect the environment, especially in the medium and low flow

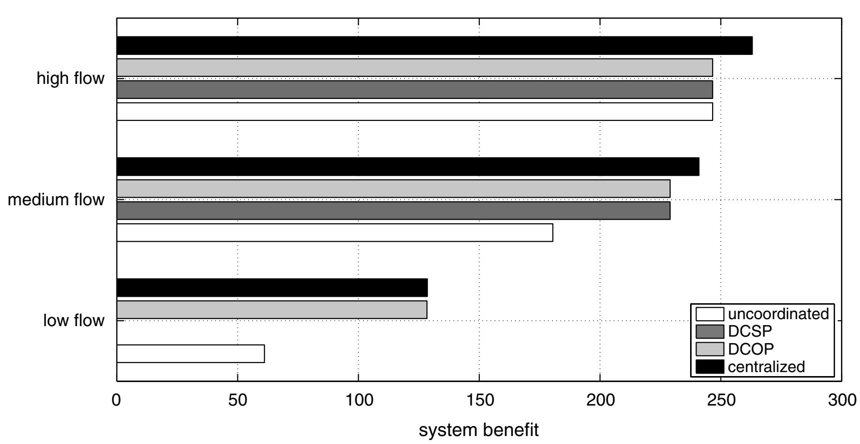

Fig. 4. Comparison of system benefits for centralized, uncoordinated, and DCSP/DCOP solutions

scenarios in which upstream agents overuse the available water producing externalities over the downstream agents suffering water shortage:

$$\alpha_1 - x_1 \leq 0 \quad \alpha_2 - Q_1 + x_1 \leq 0 \quad \alpha_3 - x_3 \leq 0 \quad \alpha_4 - x_4 \leq 0$$
$$\alpha_5 - Q_2 + x_3 \leq 0 \quad \alpha_6 - x_2 - (Q_2 - x_3) + x_4 \leq 0 \qquad (4)$$

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; and $\alpha_5 = 6$ 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 Eq. (3) are obviously nonviolable, while the normative constraints just defined [Eq. (4)] 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, i.e., $x_3 \geq 8$, being $x_3 = 2$).

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 versus actuators) as well as from the mathematical formulation (four single-objective problems versus 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 in order to obtain distributed solutions in between the two extreme situations so far described. Again the problem is solved as a sequence of optimization problems, where each active agent is considering his local objective function only from upstream to downstream or, possibly, according to a tree defined by agents-ordering algorithms such as the *depth first search* (Korf 1985; Collin and Dolev 1994). In both DCSP and DCOP formulations, the distributed constraint reasoning problem is subject to a new set of constraints, comprising the hard (physical) ones already defined in Eq. (3) along with the soft (normative) constraints introduced in Eq. (4). 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. Conversely, 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 $x_1 - \alpha_1 \leq 0$ is 0, when actually $x_1 - \alpha_1 \leq 0$, and $x_1 - \alpha_1$, 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 weight is assumed for each soft constraint (namely, $w_j = 1 \forall j$).

A graphical comparison of the system benefit for the centralized and the uncoordinated solutions with respect to the results obtained solving the new distributed constraint reasoning problem, adopting the proposed DCSP- and DCOP-based approaches, is represented in Fig. 4 by dark and light gray bars, respectively. In the high flow scenario, the uncoordinated, DCSP, and DCOP solutions are all equivalent because in this scenario, characterized 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 (f_5 and f_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 Fig. 5, where the benefit of every 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

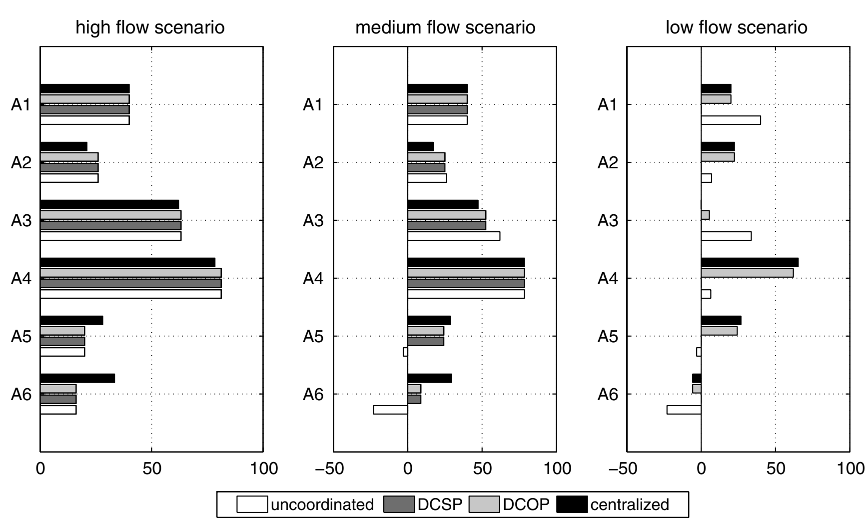

Fig. 5. Agents benefits in the three considered scenarios

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 for partially taking 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].

Discussion

The results presented in the previous section 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 to commit a *rational crime*, namely a violation of the imposed regulation (i.e., the set of normative constraints) since the agents believe the potential benefits of the violation outweight the consequences of the violation (Cooter and Ulen 1988; Souza Filho et al. 2008). The effects of these individualistic behaviors then involve the entire system and, usually, tend to decrease the benefits of the other agents. In particular, we assess the consequences of the individualistic behaviors in the low flow scenario 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 (Charness and Rabin 2002; Jager and Janssen 2003; Yang et al. 2009; Poteete et al. 2010), the consequences at a system-level produced by individualistic behaviors of the agents can be assessed by modifying Eq. (3), introducing a parameter β_i that multiplies the objective function of the ith agent ($\beta_i = [0, 10]$). These coefficients represent the *selfishness* of the ith 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 $\beta_i = 1 \forall i$. On the other hand, it is possible to represent an individualistic behavior of the ith agent by increasing the value of β_i. Results are represented in Fig. 6, 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 from looking at Fig. 6 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 nonunique preference representation. Moreover, looking at the trade-offs existing between the Pareto efficient alternatives, the watershed authority might estimate the marginal costs of individualistic 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 Fig. 6(b), 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 Fig. 6(d) 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 on the agents in order to better explore the space between the two extreme centralized and uncoordinated solutions, looking for another compromise between efficiency and acceptability. In particular, we concentrated on the medium flow scenario and we focus on the DCOP-based approach. We tested different sets of normative constraints α^i (with $i = 1, 2, 3$), reported in Table 3. The system benefits obtained solving Eq. (3) with these sets of normative constraints are represented in Fig. 7; it can be observed that more strict (weak) constraints yield solutions that are closer to the centralized (uncoordinated) one. In particular,

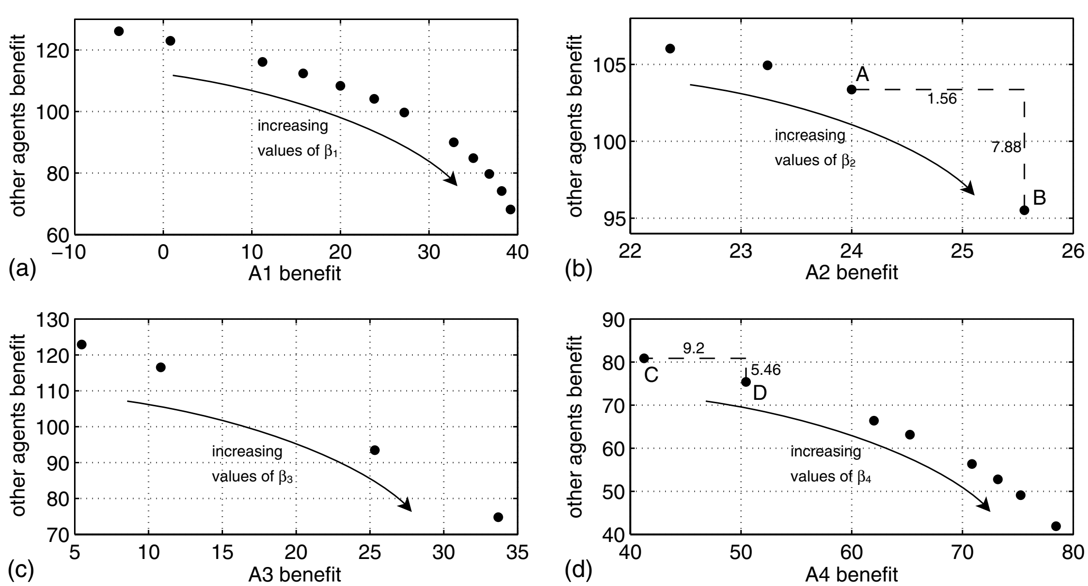

Fig. 6. Individualistic behaviors of the active agents for different values of β_i

Table 3. Different Sets of Normative Constraints α^i Expressed in L^3

Parameter	α^1 (original)	α^2	α^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

the original set of constraints α^1 allows for obtaining a solution almost equivalent to the centralized one. Yet the rigidity of these constraints might induce some individualistic behaviors. More flexible constraints, instead, slightly decrease the system benefit and the solutions move towards the uncoordinated one, which is the most preferable for the active agents. A trade-off characterizes the mechanism design process because the definition of strict constraints, which allows for approaching the centralized solution with good performances for the passive ecological agents and with respect to the system-level 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 practicable solutions with lower benefits for the passive ecological agents and, consequently, lower system-level efficiency. To analyze the trade-off, two indices measuring system-level efficiency and individual-level acceptability are necessary. Efficiency is measured in terms of the resulting total system benefit (i.e., the sum of the benefits of the six agents), and acceptability is defined as the percentage of available water that is not constrained by the normative constraints [e.g., in the medium flow scenario and soft constraints α^1, the total available water is $Q_1 + Q_2 + S = 68$, and the constrained water is equal to $\alpha_1 + \alpha_2 + \alpha_3 + \alpha_4 + \alpha_5 + \alpha_6 = 61$, hence acceptability is $(68 - 61)/\ 68 = 7/68 = 0.1029$]. This definition relies on the idea that the agents are more willing to accept solutions that allow them to freely make their decisions, while the imposition of normative constraints somewhat reduces the probability that the agents will accept and implement the solution. As a consequence, a solution that requires constraining a large portion of the available water will be impracticable and, conversely, a solution with a low level of regulation will be more acceptable. 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 acceptability is equal to 0. In the uncoordinated management, there are no soft constraints and, consequently, acceptability is 100%. Fig. 8 represents the values of efficiency and acceptability for

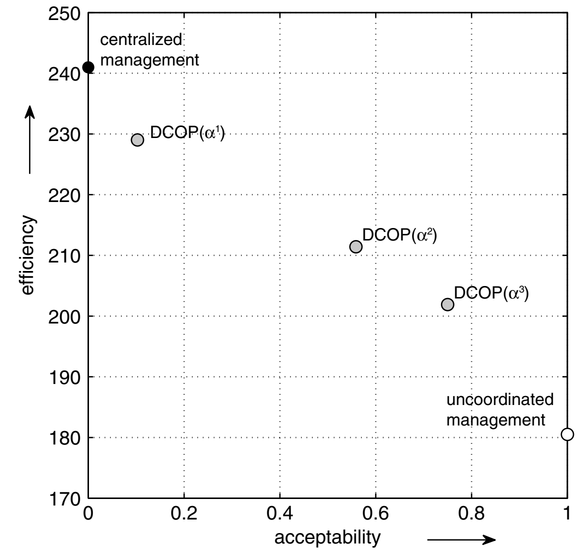

Fig. 8. Representation of the conflict between system-wide efficiency and individual-level acceptability for solutions of Problem (3) obtained with different sets of normative constraints

different solutions obtained by varying the set of normative constraints in Eq. (3). It is evident that it is not possible to simultaneously maximize both efficiency and acceptability. However, the analysis of this trade-off curve, coupled with the assessment of the individualistic behaviors' effects, can support the watershed authority in designing effective coordination mechanisms. Finally, note that the two extremes of the trade-off curve in Fig. 8 correspond to the centralized and the uncoordinated solutions (as in Fig. 1). On the contrary, the best compromise solution is generally set in the middle of the trade-off curve, and these extreme solutions have therefore low practical value.

Conclusions

In this paper a decision analytic framework based on multiagent systems (MAS) was introduced to model multiple and self-interested decision makers acting in a distributed decision-making context. In particular, regulatory-based mechanism design strategies were developed to drive the inefficient uncoordinated practices toward solutions that are balanced with respect to efficiency (social welfare) and acceptability (stability). The approach is tested on a hypothetical nondynamic problem, characterized by the presence

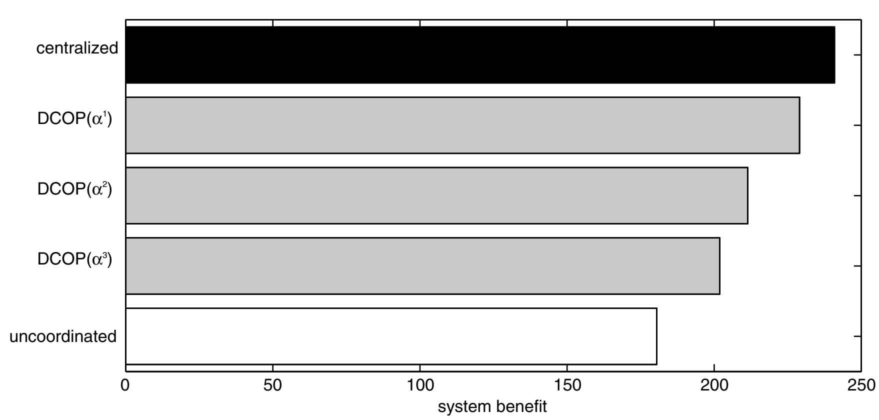

Fig. 7. Effects of different set of normative constraints on system benefit

of several human and ecological agents with nontransferable utilities.

Results show that the proposed distributed constraint reasoning approach successfully identified regulatory-based mechanism design alternatives, which are more efficient than the uncoordinated ones and more acceptable, and hence politically practicable in real decision-making processes, than the social planner's solution. The latter remains the most efficient alternative (i.e., it produces the highest system benefit), but the improvement of the DCSP/DCOP solutions with respect to the uncoordinated management is substantial.

The considered case study allows also the assessment of the consequences of individualistic behaviors (i.e., rational crimes) by the agents. Especially in situations of water scarcity, it is possible that some agents might consider not complying 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 identifying critical situations, which have to be carefully considered in designing the normative constraints.

Finally, different sets of normative constraints are considered to quantitatively estimate the trade-off curve between efficiency (i.e., the total system benefit) and acceptability (i.e., the percentage of available water that is not constrained by the normative constraints). Results confirm that the proposed DCSP- and DCOP-based approaches allow the identification of regulatory mechanisms able to produce solutions in the space in between the two extremes.

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 solving 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 the Borg multiobjective evolutionary algorithm (Hadka and Reed 2013; Reed et al. 2013), are able to find optimal solutions for challenging problems characterized by many-objective formulations, multimodality, nonlinearity, discreteness, severe constraints, and stochastic objectives. It is worth noting that these algorithms have been so far adopted only in traditional, centralized optimization problems, while their application in a distributed, bottom-up approach is still an open research topic. For DCSP/DCOP the dimensionality of the problem can limit the applicability of *Adopt* 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. Alternative algorithms have been proposed in the AAMAS literature trying to improve the scalability of the *Adopt* algorithm when the number of agents increases, such as the distributed breakout algorithm (Hirayama and Yokoo 2005). Although they can be computationally more efficient than *Adopt*, these algorithms do not guarantee the optimality of the identified solution. An extensive analysis of the performance of different algorithms in solving DCOPs is, however, beyond the purpose of this paper.

Further research will concentrate on the extension of the present study to a dynamic water system and its application on a real-world case study. To realistically represent the actual institutional or political dimensions of a real-world problem, the proposed framework should be further extended in order to take into account the socio-political status of the parties involved. Moreover, the proposed approach will require a validation with the real DMs; for example, through the organization of games testing whether the decisions that the real DMs make are reproduced by the modeled agents. Finally, interesting results might be achieved by removing the assumption that each agent has a single objective function through the adoption of multiobjective DCOP algorithms (Delle Fave et al. 2011), which allow for the exploration of potential trade-offs between the active agents benefits.

Acknowledgments

The authors would like to thank the editor, the associate editor, and the five reviewers for their very useful suggestions that contributed to improving the manuscript. The authors would like to also thank Dr. Mashor Housh at University of Illinois at Urbana-Champaign, who provides useful comments on the manuscript. Matteo Giuliani was partially supported by *Fondazione Fratelli Confalonieri*.

References

Ambec, S., and Ehlers, L. (2008). "Cooperation and equity in the river sharing problem." *Game theory and policymaking in natural resources and the environment*, Routledge, London, 112.

An, L. (2012). "Modeling human decisions in coupled human and natural systems: Review of agent-based models." *Ecol. Modell.*, 229, 25–36.

Anghileri, D., Castelletti, A., Pianosi, F., Soncini-Sessa, R., and Weber, E. (2013). "Optimizing watershed management by coordinated operation of storing facilities." *J. Water Resour. Plann. Manage.*, 139(5), 492–500.

Athanasiadis, I. (2005). "A review of agent-based systems applied in environmental informatics." *MODSIM 2005 Int. Congress on Modelling and Simulation*, Univ. of Western Australia, Nedlands, WA, Australia.

Bhaduri, A., and Liebe, J. (2013). "Cooperation in transboundary water sharing with issue linkage: Game-theoretical case study in the Volta Basin." *J. Water Resour. Plann. Manage.*, 10.1061/(ASCE)WR .1943-5452.0000252, 235–245.

Bonabeau, E. (2002). "Agent-based modeling: Methods and techniques for simulating human systems." *Proc. Nat. Acad. Sci. U.S.A.*, 99(3), 7280–7287.

Bousquet, F., and Le Page, C. (2004). "Multi-agent simulations and ecosystem management: A review." *Ecol. Modell.*, 176(3–4), 313–332.

Cardenas, J., and Ostrom, E. (2004). "What do people bring into the game? Experiments in the field about cooperation in the commons." *Agric. Syst.*, 82(3), 307–326.

Charness, G., and Rabin, M. (2002). "Understanding social preferences with simple tests." *Q. J. Econ.*, 117(3), 817–869.

Collin, Z., and Dolev, S. (1994). "Self-stabilizing depth first search." *Inf. Process. Lett.*, 49(6), 297–301.

Cooter, R., and Ulen, T. (1988). *Law and economics*, 3rd Ed., Scott Foresman, Glenview, IL.

Dariane, A., and Momtahen, S. (2009). "Optimization of multireservoir systems operation using modified direct search genetic algorithm." *J. Water Resour. Plann. Manage.*, 10.1061/(ASCE)0733-9496(2009)135:3(141), 141–148.

Delle Fave, F., Stranders, R., Rogers, A., and Jennings, N. (2011). "Bounded decentralised coordination over multiple objectives." *Proc., 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Vol. 1, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 371–378.

Dinar, A., and Howitt, R. (1997). "Mechanisms for allocation of environmental control cost: Empirical tests of acceptability and stability." *J. Environ. Manage.*, 49(2), 183–203.

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

Giuliani, M., and Castelletti, A. (2013). "Assessing the value of cooperation and information exchange in large water resources systems by agent-based optimization." *Water Resour. Res.*, 49(7), 3912–3926.

Giuliani, M., Castelletti, A., Amigoni, F., and Cai, X. (2012). "Multi-agent systems optimization for distributed watershed management." *Proc., Int. Congress on Environmental Modeling and Software (iEMSs 2012)*, Leipzig, Germany.

Hadka, D., and Reed, P. (2013). "Borg: An auto-adaptive many-objective evolutionary computing framework." *Evol. Comput.*, 21(2), 231–259.

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

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

Huskova, I., and Harou, J. (2012). "An agent model to simulate water markets." *Proc., Int. Congress on Environmental Modeling and Software (iEMSs 2012)*, Leipzig, Germany.

Jager, W., and Janssen, M. (2003). "The need for and development of behaviourally realistic agents." *Multi-agent-based simulation II*, J. Sichman, F. Bousquet, and P. Davidsson, eds., Vol. 2581, Springer, Berlin, 36–49.

Kanta, L., and Zechman, E. (2014). "Complex adaptive systems framework to assess supply-side and demand-side management for urban water resources." *J. Water Resour. Plann. Manage.*, 10.1061/(ASCE)WR .1943-5452.0000301, 75–85.

Korf, R. (1985). "Depth-first iterative-deepening: An optimal admissible tree search." *Artif. Intell.*, 27(1), 97–109.

Kumar, A., Faltings, B., and Petcu, A. (2009). "Distributed constraint optimization with structured resource constraints." *Proc., 8th Int. Conf. on Autonomous Agents and Multiagent Systems*, Vol. 2, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 923–930.

Lass, R., Sultanik, E., and Regli, W. (2008). "Dynamic distributed constraint reasoning." *Proc., 23rd AAAI Conf. on Artifical Intelligence*, 1469–1466.

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

Léauté, T., and Faltings, B. (2011). "Coordinating logistics operations with privacy guarantees." *Proc., Twenty-Second Int. Joint Conf. on Artificial Intelligence*, Vol. 3, AAAI Press, 2482–2487.

Loucks, D., Van Beek, E., Stedinger, J., Dijkman, J., and Villars, M. (2005). *Water resources systems planning and management: An introduction to methods, models, and applications*, UNESCO, Paris.

Lubell, M., Schneider, M., Scholz, J., and Mete, M. (2002). "Watershed partnerships and the emergence of collective action institutions." *Am. J. Polit. Sci.*, 46(1), 148–163.

Lund, J., and Palmer, R. (1997). "Water resource system modeling for conflict resolution." *Water Resour. Update*, 3(108), 70–82.

Madani, K. (2013). "Modeling international climate change negotiations more responsibly: Can highly simplified game theory models provide reliable policy insights?" *Ecol. Econ.*, 90, 68–76.

Madani, K., and Hipel, K. (2011). "Non-cooperative stability definitions for strategic analysis of generic water resources conflicts." *Water Resour. Manage.*, 25(8), 1949–1977.

Madani, K., and Lund, J. (2012). "California's Sacramento-San Joaquin delta conflict: From cooperation to chicken." *J. Water Resour. Plann. Manage.*, 10.1061/(ASCE)WR.1943-5452.0000164, 90–99.

Marques, G., and Tilmant, A. (2013). "The economic value of coordination in large-scale multireservoir systems: The Parana River case." *Water Resour. Res.*, 49(11), 7546–7557.

Maskin, E. (2008). "Mechanism design: How to implement social goals." *Am. Econ. Rev.*, 98(3), 567–576.

Matthews, O. (2004). "Fundamental questions about water rights and market reallocation." *Water Resour. Res.*, 40(9).

Modi, P., Shen, W., Tambe, M., and Yokoo, M. (2005). "Adopt: Asynchronous distributed constraint optimization with quality guarantees." *Artif. Intell.*, 161(1–2), 149–180.

Ng, T., Eheart, J., Cai, X., and Braden, J. (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 Resour. Res.*, 47(9).

Nguyen, N., Shortle, J., Reed, P., and Nguyen, T. (2013). "Water quality trading with asymmetric information, uncertainty, and transaction costs: A stochastic agent-based simulation." *Resour. Energy Econ.*, 35(1), 60–90.

Pannell, D. (2008). "Public benefits, private benefits, and policy mechanism choice for land-use change for environmental benefits." *Land Econ.*, 84(2), 225–240.

Poteete, A. R., Janssen, M. A., and Ostrom, E. (2010). *Working together: Collective action, the commons, and multiple methods in practice*, Princeton University Press, Princeton, NJ.

Read, L., Madani, K., and Inanloo, B. (2014). "Optimality versus stability in water resource allocation." *J. Environ. Manage.*, 133, 343–354.

Reed, P., Hadka, D., Herman, J., Kasprzyk, J., and Kollat, J. (2013). "Evolutionary multiobjective optimization in water resources: The past, present, and future." *Adv. Water Resour.*, 51, 438–456.

Schreinemachers, P., and Berger, T. (2011). "An agent-based simulation model of human-environment interactions in agricultural systems." *Environ. Modell. Softw.*, 26(7), 845–859.

Shafiee, M., and Zechman, E. (2013). "An agent-based modeling framework for sociotechnical simulation of water distribution contamination events." *J. Hydroinf.*, 15(3), 862–880.

Shoham, Y., and Leyton-Brown, K. (2009). *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*, Cambridge University Press, Cambridge, U.K.

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

Souza Filho, F., Lall, U., and Porto, R. (2008). "Role of price and enforcement in water allocation: Insights from game theory." *Water Resour. Res.*, 44(12).

Sycara, K. (1998). "Multiagent systems." *AI Mag.*, 19(2), 79–92.

Teasley, R., and McKinney, D. (2011). "Calculating the benefits of transboundary River Basin cooperation: Syr Darya basin." *J. Water Resour. Plann. Manage.*, 10.1061/(ASCE)WR.1943-5452.0000141, 481–490.

Thomas, J., and Sycara, K. (1998). "Heterogeneity, stability, and efficiency in distributed systems." *Proc., IEEE Int. Conf. on Multi Agent Systems*, IIEEE, 293–300.

Tilmant, A., Beevers, L., and Muyunda, B. (2010). "Restoring a flow regime through the coordinated operation of a multireservoir system: The case of the Zambezi River basin." *Water Resour. Res.*, 46(7), 1–11.

Tilmant, A., and Kinzelbach, W. (2012). "The cost of noncooperation in international river basins." *Water Resour. Res.*, 48(1).

Wallace, J., Acreman, M., and Sullivan, C. (2003). "The sharing of water between society and ecosystems: From conflict to catchment-based co-management." *Philos. Trans. R. Soc. London Ser. B*, 358(1440), 2011–2026.

Waterbury, J. (1987). "Legal and institutional arrangements for managing water resources in the Nile basin." *Int. J. Water Resour. Dev.*, 3(2), 92–104.

Watkins, K. (2006). *Human Development Rep. Beyond Scarcity: Power, Poverty, and Global Water Crisis*, United Nations Development Programme, New York.

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

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

Wooldridge, M. (2009). *An introduction to multi agent systems*, 2nd Ed., Wiley, New York.

Wooldridge, M., and Jennings, N. (1995). "Intelligent agents: Theory and practice." *Knowl. Eng. Rev.*, 10(2), 115–152.

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

Yang, Y. C. E., Cai, X., and Stipanović, D. M. (2009). "A decentralized optimization algorithm for multiagent system-based watershed management." *Water Resour. Res.*, 45(8).

Yeoh, W., and Yokoo, M. (2012). “Distributed problem solving.” *AI Mag.*, 33(3), 53–65.

Yoffe, S., Wolf, A., and Giordano, M. (2003). “Conflict and cooperation over international freshwater resources: Indicators of basins at risk.” *J. Am. Water Resour. Assoc.*, 39(5), 1109–1126.

Yokoo, M., Durfee, E., Ishida, T., and Kuwabara, K. (1992). “Distributed constraint satisfaction for formalizing distributed problem solving.” *Proc., Twelfth IEEE Int. Conf. on Distributed Computing Systems*, IEEE, Los Alamitos.

Yokoo, M., and Hirayama, K. (2000). “Algorithms for distributed constraint satisfaction: A review.” *Auton. Agents Multi-Agent Syst.*, 3(2), 185–207.

Young, R. (1986). “Why are there so few transactions among water users?” *Am. J. Agric. Econ.*, 68(5), 1143–1151.

Zeitoun, M., and Mirumachi, N. (2008). “Transboundary water interaction I: Reconsidering conflict and cooperation.” *Int. Environ. Agreements Politics Law Econ.*, 8(4), 297–316.

Zhao, J., Cai, X., and Wang, Z. (2013). “Comparing administered and market-based water allocation systems through a consistent agent-based modeling framework.” *J. Environ. Manage.*, 123, 120–130.

Zoltay, V., Vogel, R., Kirshen, P., and Westphal, K. (2010). “Integrated watershed management modeling: Generic optimization model applied to the Ipswich River basin.” *J. Water Resour. Plann. Manage.*, 10.1061/(ASCE)WR.1943-5452.0000083, 566–575.