Bayesian networks and participatory modelling

A. Castelletti and R. Soncini-Sessa

Dipartimento di Elettronica e Informazione, Politecnico di Milano
Piazza L. da Vinci, 32, I-20133 Milano, Italy
{*castelle,soncini*}*@elet.polimi.it*

Abstract

Bayesian networks (Bns) are emerging as a widely-adopted approach for modelling and supporting decision making in the water resource field. Being based on the coupling of an interaction graph and a probabilistic model they have the potential to improve participation and allow integration with other models. The wide availability of ready-to-use software through which Bn models can be easily designed and implemented on a PC is further contributing to their diffusion. Although a number of papers is available, where the application of Bns to water related problems is investigated, the majority of the works use the Bn semantic to model the whole water system, thus avoiding their integration with other types of model. In this paper pros and cons of adopting Bns in water resource planning and management are analyzed, framing them within the wider context of a participatory and integrated planning procedure, and exploring their integrability with other types of model.

Key words: Bayesian networks, participatory modelling, model integration, water resources planning, decision making

1 Introduction

The last decade has witnessed a growing interest to the application of graphical models, namely Bayesian networks (Bns), in environmental and resource modelling. This new trend is closely linked to the recognition of participation and uncertainty as having a key role in integrated natural resource management and to the need of tools and methodologies that facilitate dealing with them.

Starting from Varis (1985), who first generalized and rearranged the mathematical framework provided by Pearl (1988) to fit the features of environmental resource systems, many other authors contributed to illustrate the use of

Bns in the environmental field. Applications range from ecological issues, such as fisheries and related problems (Kuikka *et al.*, 1999; Borsuk *et al.*, 2002; Little *at. al*, 2004), to the assessment of climate changes on crop production (Gu *et al.*, 1996). In the water resource context they have been used by Batchelor and Cain (1999) in irrigated and rainfed farming system modelling, by Varis and Kuikka (1997) to investigate the effect of climate change on surface waters and by Borsuk *et al.* (2001, 2004) in studying the eutrophication of river estuaries. Particularly focusing on Bns as tools to support and improve participation are the works of Baran and Jantume (2004) and Bromley *et al.* (2005).

Bns may be essentially used with two different aims:

(1) *modelling*, when they are used to describe the system being studied;
(2) *supporting decision making*, when they includes decision and utility nodes, and are employed as a decision support system (DSS).

The DSS use of the Bns implies the description of the whole system model through the Bn semantic, and this in fact is what is generally found in many of the above mentioned works. However, this approach could be limiting in some circumstances and is what we argue in the following of the paper. To reach this target we will first consider the problem of decision-making in the water field by a more general perspective, regardless the modelling approach adopted. Then we will go through the model construction process comparing Bns with other types of models. Thus the Bns' role will clearly appear.

Precisely the paper is organized as follow. The structure and the alternative ways of using Bayesian networks are described in the first section. As their usefulness can be evaluated only by looking at the problem of decision making in water resource planning from a general perspective, the second section introduces a Participatory and Integrate Planning (PIP) procedure, from which the key role of participatory modelling emerges. To this latter topic is entirely devoted the third section, where we will go through the construction of a water system model, showing where participation comes in, how the model is an aggregation of models that describes the subsystems that constitute the water system and how the type of each one of these models can be selected among four different types, one of which is the Bn. At that point we will have all the necessary ingredients to focus the role of Bn in water management in the last section.

2 Bayesian networks

Bayesian networks (also known as Belief networks or Bayesian Belief Networks) are a powerful modelling technique that replicates the essential features of *plausible reasoning* (reasoning under conditions of uncertainty) in a consistent, efficient and mathematically sound way (Charniak, 1991). They have been firstly developed by the artificial intelligence and machine learning community (Pearl, 1986, 1988 and Jensen, 1996) and successfully applied in the fields of medical diagnosis (Andreassen *et al.*, 1991 and Hamilton *et al.*, 1994) and system maintenance and reliability (see for instance Heckerman *et al.* 1995 and Yu *et al.* 1999). The mathematical background of the approach is extensively covered by the above mentioned references and by Neapolitan (1990), Lauritzen (1996), Pearl (2000) and Jensen (2001). However it is worthwhile outlining here some essential notions on the structure and use of Bayesian networks that can be useful in the following of the paper and of the Special Issue.

Bns provide a framework for graphically representing the logical relationship between variables and quantitatively defining the strength of this relationship using conditional probabilities. Strictly, Bns are directed acyclic graphs, in which nodes represent random variables and the lack of arcs between two nodes represents the conditional probabilistic independence of the two unlinked variables. This structure is nothing, but a factorization of a joint probability distribution (Lauritzen and Spieghelarter, 1988). Practically, given two nodes A and B, a directional arc from A to B can be informally regarded as indicating that A *causes* B, thereby A and B are usually said to be *parent* and *child* respectively. A node which does not have any parents is called a root node and represents an input variable for the network. A node without children is a leaf node and constitutes an output variable. All the internal nodes can be either internal or state variables. Each node in the network is associated with a finite set of discrete, mutually exclusive values (in jargon *states*, not to be confused with the state variables), which represent all the possible conditions of the node, and can be either quantitative or qualitative. For each node (except the root nodes) a conditional probability table (CPT) is specified, which lists the probability that the variable associated to the node assumes a particular value, given the values taken on by the variables associated to its parents. Root nodes, being uncaused, are associated with an unconditional probability.

To be more concrete, consider a very simple example concerning the water quality of a lake. Agricultural and civil practices are responsible respectively for the production of *nitrogenous* and *phosphorous loads* that reach the lake, through superficial runoff and drainage flow. The result is an increase of the *trophic level* that leads to an alteration of the *water quality* of the lake. These

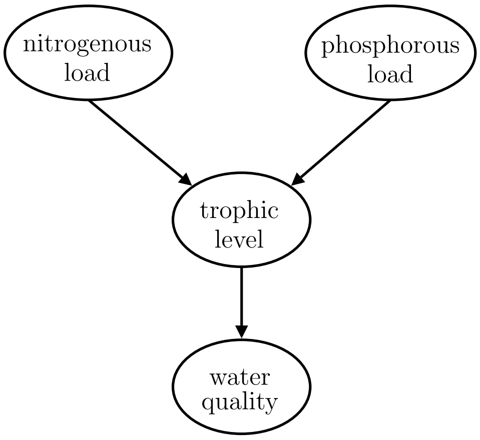

Fig. 1. The Bn of the lake in the example.

a)

nitr. load		L		H	
phosph. load		L	H	L	H
trophic lev.	L	1.0	0.3	0.5	0.0
	H	0.0	0.7	0.5	1.0

b)

trophic lev.		L	H
water quality	L	0.0	0.8
	H	1.0	0.2

Fig. 2. The conditional probability tables of a) the *trophic level* and b) the *water quality*; where L=low and H=high.

cause-effect relationships are described by the network in Fig. 1. To further simplify the example each variable is assumed to take value in a binary set, whose values are *high* and *low*.

Once the topology of the Bn is given, its structure is completely defined, but, in order to be used as a quantitative model, its CPTs have to be *populated*, i.e., filled in with probability values. These may be derived from data, simulation of other models or elicited from experts. For example the CPTs associated to the *trophic level* and *water quality* nodes are depicted in Fig. 2. Populating CPTs with numbers is conceptually the same of calibrating any other model; though an important difference exists: the conditional probability values in the CPT of each node are independent upon the values in the CPTs of the others and as a consequence they can be *locally* updated. This property allows to use the best information currently available for each variable to populate the relative CPT, to keep it updated as more data or knowledge become available and to operate interactively and on-line.

a)	phosph. load	*L*	0.3	nitr. load	*L*	0.1
		H	0.7		*H*	0.9
b)	trophic lev.	*L*	0.1	water quality	*L*	0.7
		H	0.9		*H*	0.3
c)	trophic lev.	*L*	0.0	water quality	*L*	0.8
		H	1.0		*H*	0.2

Fig. 3. a) The a-priori beliefs of *phosphorous* and *nitrogenous loads*; b) the consequent a-priori belief of *trophic level* and *water quality*; c) the a-posteriori belief of the same variables when there is evidence that *phosphorous load* is high; L=low and H=high.

2.1 How Bns can be used

Once the unconditional probability of the root nodes have been specified, it is possible to calculate a-priori unconditional probabilities (*a-priori beliefs*) for all the nodes in the network (*belief propagation*). This can be done by employing basic probability calculus and Bayes' theorem. Depending on the features of the network, the calculus can be either exact, if the network is relatively small (Pearl, 1988; Peot and Shachter, 1991), or approximated in the case of large networks (Huang and Darwiche, 1996). The a-priori beliefs is modified as new knowledge on the system is obtained, in the form of observation of the values (*evidences*) taken on by one or more variables. Hence the evidences are substituted to the a-priori belief of these variables, and the beliefs of the others are updated through belief propagation.

Coming back to the example, if the a-priori beliefs of the *nitrogenous* and *phosphorous loads* are as in Fig. 3a, the a-priori beliefs of *trophic level* and *water quality* are the ones in Fig. 3b. On the contrary if an high *phosphorous load* is observed the a-posteriori beliefs of *trophic level* and *water quality* turn into the ones in Fig. 3c.

Belief propagation provides the computational tool for using Bn models in probabilistic *inference*, which is essentially the purpose for which Bns have been conceived. They can be used either in "bottom-up" reasoning to address diagnostic tasks or "top-down" for descriptive/explanatory purposes. In the first case, the evidence of an effect is given and the most likely cause is inferred. In the second, the probability of an effect is computed once the evidence of one or more of its causes is provided.

a)

nitr. load		
	L	0.2
	H	0.8

phosph. load		
	L	0.0
	H	1.0

b)

water quality		
	L	0.8
	H	0.2

Fig. 4. a) The a-posteriori beliefs of *phosphorous* and *nitrogenous loads* when there is evidence that *water quality* is low; b) the a-posteriori belief of the *water quality* when there is evidence that both the loads are high; L=low and H=high.

These two types of inference can be easily illustrated on the lake example. The bottom-up reasoning occurs, for instance, when a low value of *water quality* is observed, thereby the associated probability being set to 1. The possible causes for this evidence are: either the *nitrogenous load* is high or the *phosphorous load* is high or both are high. To find out the more likely cause, the belief has to be propagated back from the outcome to compute the a-posteriori probability of each explanation (see Fig. 4a). Hence the Bn model is used as a diagnostic tool.

The top-down reasoning is necessary if one wonder what could be the *water quality* if both the pollutant loads were high. In this event the probability of the outcome (*water quality*) can be computed by propagating the evidence of the inputs (results are given in Fig. 4b). The network is now used in an explanatory way, which is definitively the most interesting in planning and managing water resources. Note in fact that, depending on whether the evidence on the inputs is directly observed or a-priori assumed, the network can be used either to forecast the outcomes as new evidence become available over time, or to perform a "what if" analysis by entering the network with different sets of evidences and then assessing their effects on the outcomes. In this latter case, when the input variables are uncertain and not directly decidable (i.e. uncontrollable), as it is in the example, every set of evidences defines a different evaluation scenario.

A Bn model can also include *decision* variables among its inputs. These can be interpreted as special types of root nodes, whose states are not associated with a marginal probability, but deterministically known (as they are provided by the Decision Maker) and mutually esclusive. For instance, one might want to evaluate whether the *water quality* could be enhanced by setting a norm that limits the usage of nitrogenous manure in farming. To analyze this case a new node (*norm*) has to be included at the top of the network (Fig. 5), which may assume two states: low and high manure. Consequently, the *phosphorous load* has to be described through a CPT. By putting the evidence on one of the *norm*'s states, the corresponding effects on the probability of *water quality*

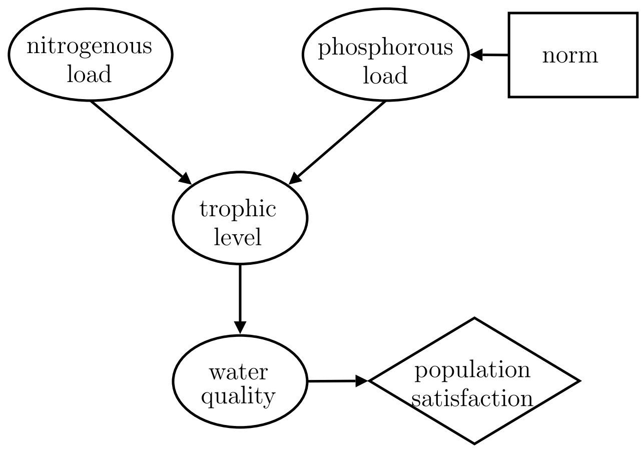

Fig. 5. The Bn of the lake in presence of decision and utility nodes.

might be assessed.

Very often decision nodes are coupled with *utility* nodes linked to the outcomes and representing the utility (or the value), which results from a given decision. In this way the probabilistic information provided by the network is synthesized in a deterministic value reflecting the stakeholder (or Decision Maker) satisfaction. In the example the utility node is given by the *satisfaction of the population* living on the lake shores. So the effect of the norm in terms of population satisfaction can be assessed. A Bn that contains decision and utility nodes is called an Influence Diagram (Oliver and Smith, 1990) and it is usually represented with the graphical convention adopted in Fig. 5, where rectangles denote decisions, rhombs utilities and ovals all the other variables (not controllable inputs, internal and output variables).

3 Decision making in water management

Participation and integration are increasingly accepted as central principles for decision making in the environmental field by international organizations and by national authorities in many countries. For instance, in the water resource sphere, they have been included by the European Union in the Water Framework Directive (WFD, Directive/2000/60/EC), which defines the rules for planning and managing water resources in the EU countries. The promotion and application of such concepts in practical decision making can be effectively enhanced through the adoption of a procedural approach that sets out the general methodological framework within which specific methods, models and tools can be easily integrated.

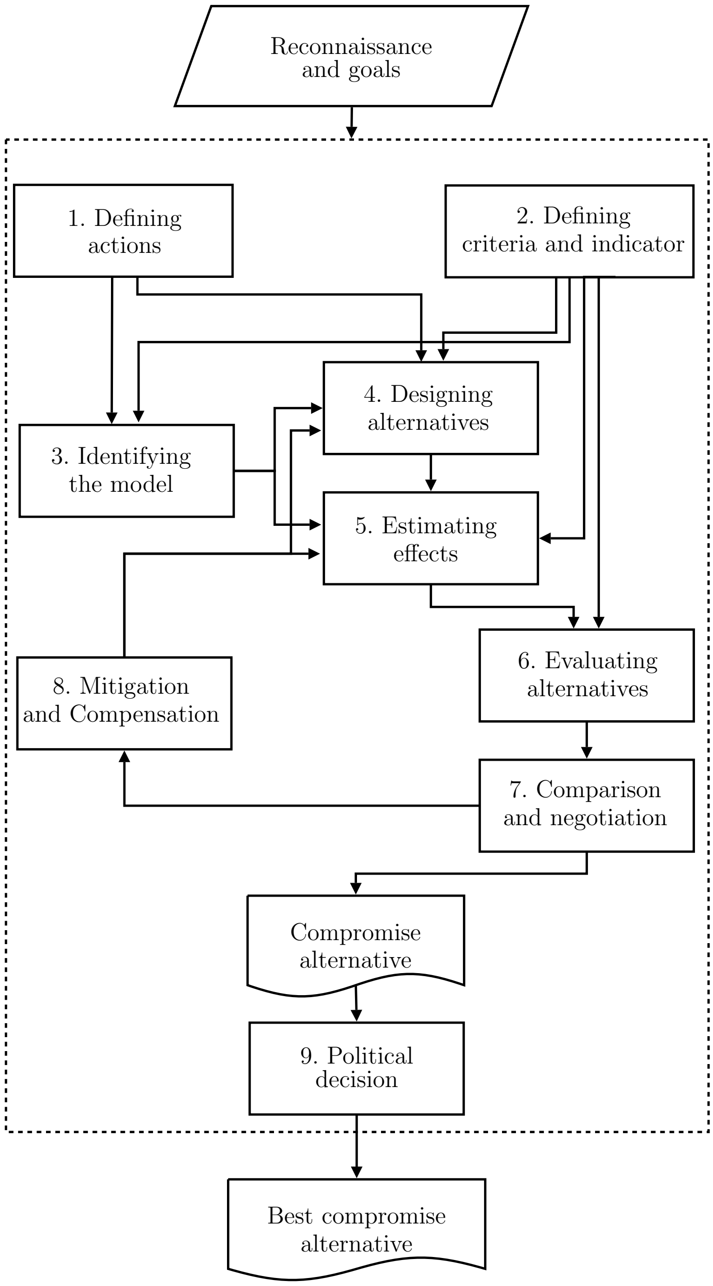

Fig. 6. The PIP procedure for water resource projects.

3.1 The PIP procedure

In this paper we will consider as reference the Participatory and Integrated Planning (PIP) procedure described in Soncini-Sessa (2004a and 2004b) and in Castelletti and Soncini-Sessa (2005), developed specifically for the water resource systems. It is a 9 phases procedure, as shown in Fig. 6, that starts from the reconnaissance of the system and the identification of the goals of the Project (Phase 0) and ends with a set of compromise alternatives to be submitted to the Decision Maker(s) for the the final political decision (Phase 9). The actions, a mixture of which is presumed to be apt to achieve the goals, are first identified (Phase 1). Each mixture is called an alternative. Then the goal is translated into a hierarchy of operational criteria (Phase 2), reflecting the point of view of the stakeholders in evaluating the alternatives (e.g., farmers are typically interested in maximizing the income from the agricultural production). Each one of the lower level criteria is associated to a quantitative indicator through which the criterion can be verified. These indicators are functionals of the trajectories of the variables that describe the system conditions. Their calculations in correspondence of a given alternative requires to identify a model of the system (Phase 3), that describes the dynamics of the system variable as a consequence of the actions that constitute the alternative. The set of the alternatives to be considered is designed in Phase 4, either by experts or by solving a Mathematical Programming (or an Optimal Control) problem, which selects, among the alternative that may be obtained by combining in all the possible ways the actions specified in Phase 1, the alternatives that are efficient (in the Pareto's sense). Then the effects produced by the efficient alternatives are assessed by simulating (Phase 5) the system behaviour over a properly specified time horizon. For each alternative in Phase 6 each stakeholder transforms the values of the indicators in a dimensionless index that expresses her/his satisfaction for the effects of such alternative. Through these indexes the alternatives can be compared and negotiated among the stakeholders (Phase 7). The process ends with the identification of a set of alternatives (attractive alternatives) that enjoy the consensus of the majority of the stakeholders. For the unsatisfied minority of them mitigation and/or compensation measures (Phase 8) can be analyzed in order to broaden further the consensus. Eventually, a set of compromise alternatives, on which the consensus can not be further enlarged, is identified. It is now up to the Decision Maker(s) (Phase 9) to select, among these, the best compromise alternative.

3.2 The key role of modelling

To carry out successfully the PIP procedure, as well as any other decisional procedure, it is essential that stakeholders be fully involved in all the phases and in particular that they actively participate the negotiation process. In this way the democratic legitimization of the final political decision is ensured by the co-deciding role (Hare *et al.*, 2003) they play. However, being the alternatives negotiated on the basis of indexes obtained from indicators (which are in turn computed, via simulation, from the model of the system), the inescapable condition for the active and profitable participation of the stakeholders to the negotiation is that they completely agree on the model of the system and trust its outcomes. This emphasizes the central role of participatory modelling in the effective implementation of the procedure.

4 Participatory modelling

Participatory modelling requires the stakeholders not only to be aware of assumptions, limitations and uses of a particular model, but also to ideally go through the same thinking process and be exposed to the same information and arguments as the modeller, along the entire model identification process. Only in this way they can share a quantitative understanding of the system and be ensured that what is being modelled actually reflects their feeling on the system behaviour.

Looking at the PIP procedure (Fig. 6) one would naturally think that the modelling activity starts and ends within the Phase 3. However, the object to be modelled and the goals the model has to serve are defined in Phase 0. Moreover, among the inputs of the models must appear the variables that quantify the actions defined in Phase 1 and among its outputs must appear all the variables required to compute the criteria that are defined in Phase 2. Hence the reader will understand why in describing how to make a model we have to start from the description of the Project.

4.1 Structuring the model

4.1.1 From the system to the domains

To exemplify the exposition we refer to a realistic, though didactic, Project, while a real world one can be found in Castelletti and Soncini-Sessa (*this issue*). Consider a lake shared between two States, U(pstream) and D(ownstream). The lake is regulated by a dam located in D, which supplies an irrigation district located in the same State. The seasonal changes of lake levels, induced by the regulation, are responsible for the frequent flooding of town C, which is located on the lake shores within U. The inflow to the lake is modulated by the activity of a complex of hydropower reservoirs. The Government of U thinks that by changing the way the dam is daily operated, i.e. its regulation policy, the flooding can be considerably reduced. As any modification of the policy has to be negotiated with D, before starting the negotiation, the U's Government wish to ascertain if a regulation policy really exists that may reduce the floods (goal of the Project). The lake shore inhabitants are accustomed to evaluate the degree of flooding through the following indicator: the number i of days (t) in which the lake level h_t in the town C exceeded the flooding threshold $\bar{h}$ along a given horizon H (e.g. the last twenty years), i.e.

$$i = \sum_{t \in H} g_t(h_t) \tag{1a}$$

where g_t is a binary variable, defined as

$$g_t(h_t) = \begin{cases} 1 \text{ if } h_t > \bar{h} \\ 0 \text{ otherwise} \end{cases} \tag{1b}$$

The value $i(A0)$ taken on by the indicator in correspondence to the current regulation policy (alternative A0) is easy to obtain as the historical trajectory of the lake levels were recorded during the horizon H. To estimate the value $i(Aj)$ that the indicator would assume in correspondence to any other policy Aj, requires to have the trajectory of the lake level that would have been obtained if the policy Aj had been used over H, and this, in turn, requires a model of the lake.

The exact framing of the Project goals, for which the model is being identified, is the starting point of any modelling exercise. The model identification must then proceed with the definition of the boundaries of the system being modelled and with the choice of the level of detail to use in its description. The contribution of the stakeholders is here of great importance to complete, and in some cases to substitute, analyses of physical and socio-economical nature. In the example they might suggest not to limit the attention to the lake itself, but to enlarge the system to include both the upstream catchment, where melting snow, rainfall, and releases from hydropower reservoirs contribute to generate the inflow to the lake, and the downstream users, as the release decisions depend on their water demand. The lake can be therefore considered as a component of a larger system. More in general any water system can be taught as an aggregate of individual components. Each of these components serves a precise function and the decomposition of the system must be based on the identification of these functions, which should be those that are most relevant to the modelling scopes. The decomposition of the system into interconnected components and the individuation of the topology of their interconnections is the first level of formalizing reality. It is not yet a model, but just a logical organization of the knowledge available (forethought, theories, information, data relative to the system) into sets of information (*domains*), each of which contains the information that concerns a component, regardless the type of model that will be adopted to describe it in the following. The domain is the first level of abstraction from reality, which does not yet require to assume hypotheses about the mathematical relationships that there exist among the variables, but simply to define the data which will be used and how they are represented. As a consequence many different models can be associated with the same domain.

The lake for instance is a complex system, in which a number of physical and chemical processes take place. The lake domain is the set of all the quantities and information that pertains to it (inflow, release, level, chemical characteristics of the water, biota, algae, topography, the water authority, etc.), including their sources. The model of the lake is instead a simplified representation of reality that should reproduce the characteristics that are relevant from the point of view of the problem for which it is created. In the example the aim is the reduction of its flooding and therefore only the quantities that contribute either directly or indirectly to form the levels have to be selected among those available within the domain (i.e. inflow, release, level). These are the candidate variables for the model and have to be carefully defined, by specifying how they are measured, at what time, etc. For instance the inflow to the lake can be considered either as the entire volume of water feeding the lake or as the separate contribution of each individual tributary. These definitions must be communicated to the stakeholders, and understood and shared by them in such a way that in successive phases their meaning is always clear.

When all the components, the topology that maps how they are interlinked, their domains and the variables within each domain have been well defined, one can proceed to the construction of the model of each component.

4.1.2 From the domain to its model: the causal network

In this step the cause-effect relationships that link the variables must be identified. Once again the process develops by trial and error, working with the stakeholders interested in the component being described. When the system is not well-known (for instructional purpose we will pretend that the lake is not) the best way to proceed is to construct, by trials, a *causal network*. This type of representation guarantees transparency in the construction of the model and aids in creating it with and explaining it to the stakeholders, since it reflects an intuitive cognitive procedure which is inductive and which everyone understands. When the stakeholders are not used to the formal language of graphics, the network can be constructed manually, together with the stakeholders, by writing the names of the variables on pieces of paper and asking them to arrange the papers so that the causes precede the effects (Hodgson, 1992).

The level[1] h_{t+1} at the beginning of day $t+1$ depends on the net volume of water that enters the lake during day t (the interval $[t, t+1)$), i.e. it depends on the difference between the inflow a_{t+1} and the release r_{t+1} in the same interval. This latter is produced by the release decision u_t taken on by the Regulator at time t, while the former is the sum of the natural inflow ε_{t+1} and the volume w_t that is going to be released from the hydroelectric reservoirs during day t. These cause-effect relationships are represented with the causal network in Fig. 7.

[1] In the symbol of a variable the subscript denotes the time at which it assumes a deterministic value.

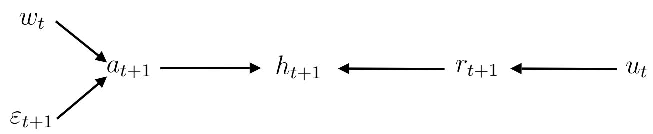

Fig. 7. The causal network of the lake at the first attempt.

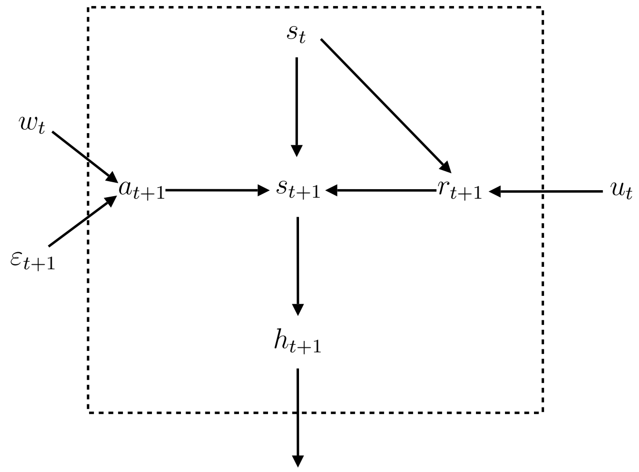

Fig. 8. The causal network of the lake in the definitive version.

However it is easy to show that this network does not constitute yet a good description of the reality. As a matter of fact if the regulator of the lake were to decide to release a very large volume u_t, the volume r_{t+1} actually released would differ according to whether the lake were full or empty (if for example the lake was empty r_{t+1} would be null). Therefore, r_{t+1} is influenced not only by the decision u_t, but also by the volume s_t that is stored in the lake at the time of the decision. Analogously, h_{t+1} is not completely defined once a_{t+1} and r_{t+1} are known, because it also depends on s_{t+1}. The causal network must therefore be extended as in Fig. 8.

Note that from the storage s_t at time t one might calculate the storage at time $t+1$ and from this derive the level h_{t+1} that is needed to calculate the indicator i defined by (1)).

The causal network contains all the kind of variables that could be included in a model. The storage s_t is the state variable and in fact it influences itself at time $t{+}1$. The level h_{t+1} is the output for which the model is being constructed. All the uncaused variables, but the storage s_t, which indeed is the effect of the level h_t at the previous time, are input variables. These can be either controllable variable, such as the release decision u_t, or disturbances that can not be decided, as it is the case of w_t and ε_{t+1}. In their turn disturbances can be deterministic, such as the volume w_t, or uncertain, such as the natural inflow ε_{t+1}, depending on whether their value is known or not at the time t when the release decision is taken. There are finally other variables, such as the total inflow a_{t+1} and the release r_{t+1}, that do not fall into any of these categories. These are called internal variables, because they serve only to retain, during the calculations, a partial result to be used afterwards. For example, the release r_{t+1} is calculated only as a step towards determining the state s_{t+1}.

It is worthwhile noting that, with the exception of the state variables, the other variables can be classified in different way according with the modelling scopes. For example, if the scope were to be to study the effects of the regulation on the downstream users, the output variables would be the release r_{t+1}, while the level h_t would not even appear in the causal network. If the aim was instead to evaluate the effects of the regulation of the hydroelectric reservoir, the control variable would be w_t, while u_t would assume the role of a deterministic disturbance.

Generally, the control u_t cannot assume any value, only those that are feasible in that moment, given the state of the system. For example, in the case of the lake, it is not possible to decide to deliver something if the storage is null. It is usual to represent all of this information with two equations and a condition. The first equation expresses the future state as function of the present state and the inputs

$$s_{t+1} = f\left(s_t, u_t, w_t, \varepsilon_{t+1}\right) \tag{2a}$$

and the function $f(\cdot)$ is said, for obvious reasons, *state transition function.* To it is associated the condition

$$u_t \in \mathcal{U}(s_t) \tag{2b}$$

that defines the feasible controls as a function of the storage. The second equation expresses the output as a function of the state

$$h_t = h(s_t) \tag{2c}$$

and the function $h(\cdot)$ is called *output transformation function.* It is important to note that the forms of the functions $f(\cdot)$ and $h(\cdot)$ are not yet known. Their individuation, either explicit or implicit, is the aim of the next step in the modelling process.

What this simple example points out has a general validity. Any model, regardless of the mathematical formulation with which it is expressed (i.e. its *type*), is described by variables that can be classified as state ($\mathbf{x}_t$), output, ($\mathbf{y}_t$) and input variables. These last ones in turn are subdivided into control variables (actual controls $\mathbf{u}_t$, planning decisions $\mathbf{u}^p$ that are controls whose values do not change in time) and disturbances, which are subdivided in deterministic $\mathbf{w}_t$ and uncertain ε_{t+1}. Their dynamics is described by the two equations already encountered (Kalman, 1969):

(1) The *state equation*

$$\mathbf{x}_{t+1} = f_t\left(\mathbf{x}_t, \mathbf{u}^p, \mathbf{u}_t, \mathbf{w}_t, \varepsilon_{t+1}\right) \tag{3a}$$

defined by the *state transition function*, to which the following conditions are associated

$$\mathbf{u}^p \in \mathcal{U}^p \tag{3b}$$

$$\mathbf{u}_t \in \mathcal{U}_t\left(\mathbf{x}_t, \mathbf{u}^p\right) \tag{3c}$$

which define the feasible planning decisions and controls at time t.

(2) The *output equation*, defined by the *output transformation function*, that assumes the following form

$$\mathbf{y}_t = h_t\left(\mathbf{x}_t, \mathbf{u}^p\right) \tag{3d}$$

With the identification of the causal network the phase of qualitative modelling ends. The causal network has now to be translated into a model, by quantitatively expressing the qualitative cause-effect relationships that it defines. There are a number of ways to address this task that give rise to different types of models. We will present them in section 4.2. Before, in the next section, we will assume that the models of the components have already been identified, in order to show how the model of the whole system can be finally obtained.

4.1.3 From the component models to the system model

Once a model of the form (3) has been specified for each one of the components, the model of the whole water system is obtained by interlinking all the component models on the basis of the topology that expresses the interconnections between the components. Technically, the output of the model of a component, which is logically upstream of at least one component, becomes an input to the models of the components that are downstream of it. Thus the model of the whole water system has again the form (3).

4.2 Type of models

We come back now to the different types of models that can be used to describe each individual component. They can be classified on the basis of the degree of the a-priori information that is involved in the quantification of the cause-effect relationships appearing in the causal network of the component. This information can be available either in the form of theories provided by scientific disciplines or as empirical knowledge obtained directly from the stakeholders and their experts. Obviously, when dealing with a lake, one falls in the first category, since hydrology and hydraulics provide a well structured and established system of knowledge to describe its behaviour. However, for didactic reasons, we will assume for a while that no a-priori information is available.

4.2.1 Bayesian networks

When no theory exists from which the modeller can drawn on support in quantitatively formulating the model, the only information he may hope to obtain are observations of statistical nature, such as: "thirty per cent of the time that the inflow and the storage are high, and the release scarce, there is a flood the following day". These observations can be either subjective evaluations provided by the stakeholders or quantitative evaluations derived from conditional frequency counts on time series of recorded data. They are, at least in subjective form, always available. The natural and simplest way of mathematically organizing such kind of information is through a Bayesian network that associates to every variable z of the causal network, except the input variables, a CPT $\phi(z|w_1, \ldots, w_r)$ expressing the probability of z conditioned to the values taken on by the variables $w_1, \ldots, w_r$ influencing it. For example the causal network in Fig. 8 can be quantitatively expressed through the Bayesian network depicted in Fig. 9.

It is easy to observe that a Bn implicitly includes the two functions (3a) and (3d) that characterize the general structure of a model, which, in the particular case of the lake, take the form (2a) and (2c). The output transformation function (2c) is the CPT $\phi(h_t|s_t)$ that provides the probability of the level h_t as a function of the storage s_t, while the state transition function (2a) is the

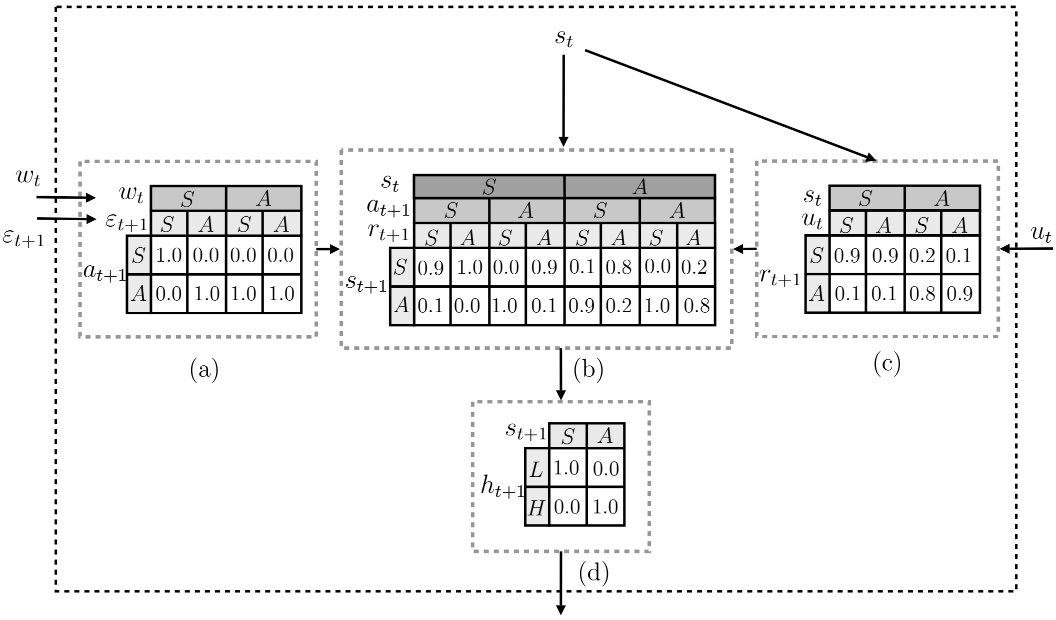

Fig. 9. The Bn that describes the lake, where S =*scarse*, A =*abundant*, L =*low* and H =*high*.

s_t		*S*								*A*							
w_t		*S*				*A*				*S*				*A*			
ε_{t+1}		*S*		*A*		*S*		*A*		*S*		*A*		*S*		*A*	
u_t		*S*	*A*	*S*	*A*	*S*	*A*	*S*	*A*	*S*	*A*	*S*	*A*	*S*	*A*	*S*	*A*
s_{t+1}	*S*	0.9	0.9	0.1	0.1	0.1	0.1	0.1	0.1	0.7	0.7	0.2	0.2	0.2	0.2	0.2	0.2
	A	0.1	0.1	0.9	0.9	0.9	0.9	0.9	0.9	0.3	0.3	0.8	0.8	0.8	0.8	0.8	0.8

Fig. 10. The CPT that expresses the transition function of the lake, where S =*scarse*, A =*abundant*, L =*low* and H =*high*.

CPT $\phi_t(s_{t+1}|s_t, u_t, w_t, \varepsilon_{t+1})$ depicted in Fig. 10 that can be easily obtained by properly concatenating tables (a), (b) and (c) in Fig. 9. Note that in general the function $\phi_t(s_{t+1}|s_t, u_t, w_t, \varepsilon_{t+1})$ may describe a non-deterministic relationship, where the uncertainty may derive not only from the input uncertainty, but may also be due to the poor and not certain knowledge on the process that it describes. One of the peculiarities of the Bns is actually that they allow to implicitly account for the structural uncertainty of the model (usually called process noise). Indeed, the probability in a CPT may assume values different from zero or one only when the conditioned variable (in the example s_{t+1}) is affected by disturbances that do not explicitly appear among the entry variables (e.g. evaporation). An analogous remark can be made for the output transformation function. In the example this function defines a deterministic relationship between the storage and the level. Nevertheless, the CPT of Fig. 9d could have been filled in also with non unitary values if one would have to reflect the presence of a measurement error (measurement noise) of the level records.

Bns are very useful for representing systems, such as social systems, for which quantitative theories are not available and, as a consequence, the knowledge is unstructured and/or limited (see Castelletti and Soncini-Sessa, *this issue*). In such cases the values within the CPTs can be easily obtained through interviews to the stakeholders. Note in fact that, as already mentioned in Sec. 2, the CPTs are locally updatable and this makes easy for the modeller to properly formulate the questions and for the stakeholders to understand them, and consequently provide the answer. They are however not very well suited to represent systems, such as the lake, in which there are known deterministic relationships and/or the number of values that each variable can assume can be very high (high cardinality of the set of states of each node). An example of the first case is the not easily understandable representation of a simple arithmetic operation such as the sum of w_t and ε_{t+1}, that is expressed with the CPT (a) in Fig. 9. To understand the difficulties posed by the second, think of the model of a real lake such as the lake Maggiore in Northern Italy (Castelletti *et al.*, 2004). The cardinality of the sets of the feasible values of the variables s_t, u_t, w_t and ε_{t+1} is respectively 120, 150, 25 and 25; it follows that if the model of the lake were represented in the form of a Bn the CPT that defines the state transition function would have 1 350 000 000 elements

(120×120×150×25×25). In both the cases the graphical nature of the Bns, which generally is of great utility in the interaction with the stakeholders, turns out to be a limit in the communication with them. This problem also emerges when the Bn is used as a dynamic model, as in the case of the lake of the example. Note in fact that the Bn in Fig. 9 represents the lake at a given time t. To simulate the dynamic model of the lake over the time horizon H, one has to create a cascade of networks, one for each time instant (see Fig. 11). Moreover when the system is time-variant, the CPTs have to be properly filled in for each time instant, thus resulting in a huge number of questions to be posed to the stakeholders.

When the knowledge on the system behaviour is highly structured Bns can anyway be useful to elicit the stakeholders that are not experienced in using quantitative models. In such cases CPTs of the network can be filled in using the output of more complex models (e.g. mechanistic models) and the resulting Bn used for instructional purpose, to enhance for instance the social learning on the system. Obviously in this case it is not convenient to use the Bn as a computational model, as a more flexible model is available.

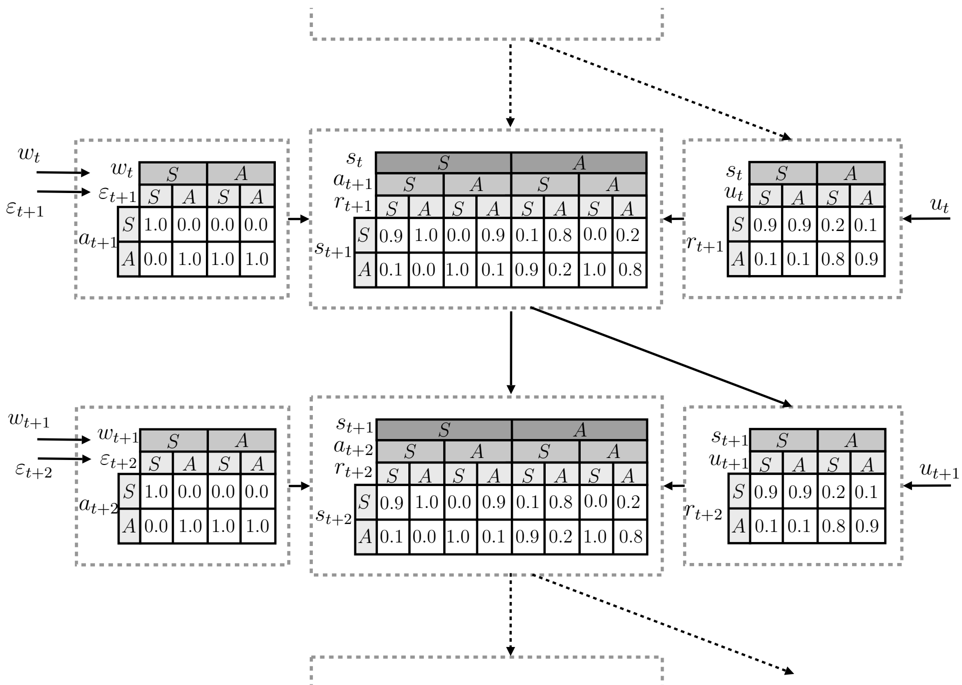

Fig. 11. Concatenation of Bns of the type in Fig. 9, in order to represent the time evolution of the lake.

4.2.2 Mechanistic models

When the knowledge on the system to be modelled is well established, i.e. there are theories explaining its internal processes in a quantitative way, it is rational to use them to built up a quantitative model that describes the mechanism operating within the system. Given the way it is derived, this type of model is said *mechanistic*.

For example, in the case of the lake, Physics suggests that a simple equation of mass conservation fairly accurately describes the storage dynamics

$$s_{t+1} = s_t + w_t + \varepsilon_{t+1} - r_{t+1} \tag{4}$$

Hydraulics provides the functional relationship (stage-discharge relation) linking the release r_{t+1} to the storage s_t when the dam is completely open

$$r_{t+1} = N(s_t) \tag{5}$$

where $N(s_t)$ generally assumes the form

$$N(s_t) = \alpha s_t^{\beta} \tag{6}$$

being α and β two positive parameters. When the dam is operated the release r_{t+1} also depends on the release decision u_t: in general r_{t+1} is equal to u_t whenever it is physically possible and the legal constraints that must be obeyed permit. By saying $\bar{s}$ the storage level, above which for precautionary reasons the dam must be completely opened, the release r_{t+1} can be expressed as follows:

$$r_{t+1} = \begin{cases} N(s_t) & \text{if } s_t > \bar{s} \\ \begin{cases} N(s_t) & \text{if } u_t > N(s_t) \\ u_t & \text{otherwise} \end{cases} & \text{otherwise} \end{cases} \tag{7}$$

Finally the output transformation function linking the levels and the storage is easily derivable once the lake bathymetry and its shores' altimetry are known

$$h_t = h(s_t) \tag{8}$$

Observe that, once the disturbances w_t and ε_{t+1} are known, the above equations define a deterministic mechanistic model. In order for such a model to be equivalent to the Bn presented in the previous section, the process and output noises, which are implicitly included in the Bn, have to be somehow included also in this model. These noises can both be modelled as random disturbances and added to the corresponding equation. The output noise has always the obvious meaning of a measurement error and adds to equation (8); the process noise in turn accounts for the simplifications introduced in the modelling of the state transition function (4), such as the adoption of a stage-discharge relation to describe the release as a function of the storage: as it is well known from Hydraulics this one-to-one relationship only holds in a permanent flow regime, when the storage does not vary in time, while in the practice the storage is highly non stationary.

Mechanistic models are commonly used in water resource modelling and plausibly some of the stakeholders (e.g., hydropower managers) have already experienced or are currently using a model of this class. This should suggest the modeller not to force them to move to a new type of model, e.g. a Bn. However, when the construction of the model has to be started from scratch, one should consider that mechanistic models may contain an high number of parameters that could make their identification not easy at all. Estimating such parameters might require a wide variety of data, which have in some case to be collected through expensive measurement campaigns, and then estimated with ad-hoc and highly time consuming techniques.

4.2.3 Empirical models

In many cases the construction of a Bn, or of a mechanistic model, proves to be a too costly operation compared with the goals of the model. Maybe the most emblematic case is the catchment. The processes that take place within it and that transform precipitation into inflow are very complex and, therefore, Bns and the mechanistic models that represent them are equally complex. However, when the alternative under evaluation does not directly involve the catchment, as it is in the lake example, the stakeholders may be mainly interested in a model that produces precise and reliable outputs (namely inflows) in response to given inputs, rather than an accurate and realistic description of the system's internal mechanisms.

In these cases it can be a good idea to simply identify a model that describes the relationship between the input and the output of the system (*input-output relationship*) such as the following

$$\begin{aligned}\mathbf{y}_{t+1} = \mathsf{y}_t \Big(&\mathbf{y}_t, \ldots, \mathbf{y}_{t-(p-1)}, \mathbf{u}_t, \ldots, \mathbf{u}_{t-(r'-1)}, \mathbf{w}_t, \ldots \\ &\ldots, \mathbf{w}_{t-(r''-1)}, \varepsilon_{t+1}, \ldots, \varepsilon_{t-(q-1)}\Big)\end{aligned} \tag{9}$$

that explains the output at time $t+1$ on the basis of the values that the input and output variables assumed in an suitable number of preceding time steps. Equation (9) is called *external form* (or *representation*) while in contrast, the pair of equations (3a)-(3d) is called *internal form* or *representation.*

For example, it is easy to see that in the case of the lake there is a representation of the form

$$h_{t+1} = \mathsf{y}\left(h_t, u_t, w_t, \varepsilon_{t+1}\right) \tag{10}$$

that can be obtained by manipulating (4), (7) and (8).

Since one doesn't have any information about the structure of the system, the external form has to be fixed *a priori* in a class of functions such that each one of its functions can be identified by specifying a finite (and small) number of parameters. Both linear (PARMAX models) and non linear (e.g. Neural Networks) relationships may be adopted, depending on the component that is to be modelled and on the data available.

Note that also this model is deterministic, once the disturbances w_t and ε_{t+1} are known. In order to transform it in a stochastic model a noise term must be included in its expression. However, because the internal structure of the model is not available, it is not possible distinguishing between process and measurement noise. The noise term will be therefore a combination of the two. The general form of the model (9) should be then rewritten as

$$\begin{aligned}\mathbf{y}_{t+1} = \mathsf{y}_t \Big(&\mathbf{y}_t, \ldots, \mathbf{y}_{t-(p-1)}, \mathbf{u}_t, \ldots, \mathbf{u}_{t-(r'-1)}, \mathbf{w}_t, \ldots \\ &\ldots, \mathbf{w}_{t-(r''-1)}, \varepsilon_{t+1}, \ldots, \varepsilon_{t-(q'-1)}, \mathbf{e}_t, \ldots, \mathbf{e}_{t-(q''-1)}\Big) + \mathbf{e}_{t+1}\end{aligned} \tag{11}$$

where $\mathbf{e}_t$ is the above mentioned noise.

The external form can be identified only if time series of the input and output variables are both available and long enough to allow the estimation of the parameters. When these two conditions are verified a number of very efficient algorithms exists that allows to estimate the parameter in a very fast and accurate way, making also possible the comparison in real time of different models (e.g. different class for the external relationship or different combination of the inputs). Empirical models are usually simpler than mechanistic and as a consequence more suitable to be used in designing the alternative.

However, because the parameters are estimated from historical data, the empirical models cannot and should not ever be used when the alternatives under evaluation could modify the system and therefore the input-output relationship. For example, in the case of the lake, this happens when a modification of the lake's outlet is included among the considered actions.

4.2.4 Markov chains

The utility of Markov chains emerges only when one examines how the model can be used to evaluate the effects that the alternative produces. As already pointed out this means the system must be simulated. When a Bn or a mechanistic model is fed by a stochastic white noise the state $\mathbf{x}_t$ of the system is at every instant (except for the initial instant) a stochastic variable that, to be described, requires that its probability distribution π_t be specified. This can be determined in a purely numerical way through a Monte Carlo simulation approach. As one can imagine, this method is however computationally expensive in terms of time. To avoid this difficulty, instead of using a description of the system based on the state $\mathbf{x}_t$, one can utilize a description that assumes that the state is the probability distribution π_t of $\mathbf{x}_t$ and describes the dynamics of this latter through a *Markov chain*

$$\pi_{t+1}^T = \pi_t^T \mathbf{B}_t \tag{12}$$

where the superscript T denotes the transposition of the vector to which it is applied. The matrix $\mathbf{B}_t$ has a simple meaning: its element b^{ij} is the probability that the state of the system passes from its i-th value at time t to its j-th value at time $t+1$, when the disturbance that acts on the system is the white noise considered and the system is controlled by a given policy (since the alternative is being evaluated the policy it includes is certainly defined).

To clarify, consider again the lake example under the following assumptions:

- the system is described by the Markov chain that corresponds to the Bn in Fig. 9,
- the two disturbances w_t and ε_{t+1} that appear in it are both random and white and their values are equiprobable,
- the adopted policy establishes that the supply decision is *scarse* (S) when the storage is *scarce* (S), and *abundant* (M) when the storage is *abundant* (A).

Equation (12) is therefore the following

$$\begin{vmatrix} \pi_{t+1}^S & \pi_{t+1}^E \end{vmatrix} = \begin{vmatrix} \pi_t^S & \pi_t^E \end{vmatrix} \begin{vmatrix} 0.29 & 0.91 \\ 0.32 & 0.68 \end{vmatrix}$$

where π_t^S and π_t^E represent the probability that the storage will be *scarce* and the one that it will be *abundant* at time t respectively. It is easy to see that the probability π_t tends, as the time goes on, to the following condition of equilibrium

$$\pi = \begin{vmatrix} 0.31 \\ 0.69 \end{vmatrix}$$

that shows how, with the adopted policy, 31% of the time the storage will be *scarce* and 69% *abundant*; a result with a great informative value that would have been only very laboriously obtained operating directly on the Bn in Fig. 9.

As the example reveals, the matrix $\mathbf{B}_t$ can be easily derived once the following are known:

- the state transition equation of the model (Bn or mechanistic),
- the models of the disturbances,
- the adopted regulation policy.

This is the same information that is required for the simulation using the Monte Carlo method.

By adopting a Markov chain, the mechanistic system with stochastic state $\mathbf{x}_t$ is described by a model with a very simple structure, that has the appearance of an empirical model, whose state π_t is deterministic. As this system is an autonomous system its simulation is very easy. The analysis of the characteristics of the matrix $\mathbf{B}_t$ allows also to recognize several important properties of the system controlled by the regulation policy adopted (for example if it tends towards an equilibrium probability and what that would be, as we have seen in the example).

Through the Markov chain the simulation is notably simplified, but at the price of an increase in the dimension of the state of the system to be simulated, inasmuch as the components of the vector π_t are as many as the values that the state $\mathbf{x}_t$ can assume. The number of elements that describe the matrix $\mathbf{B}_t$ is therefore very high, but, in contrast to what happens for the CPT of the Bn, the elements that compose it do not have to be directly provided by the stakeholders, but can be calculated, given a mechanistic model (or a Bn). For this reason hardly a model is directly described in the form of a Markov chain, but rather it is first given in one of the three forms that have been previously presented (Bns, mechanistic models, and empirical models) and then derived from these.

5 The Bns' role in water management

The previous sections have focused the key role of participatory modelling in water resource decision-making and introduced four types of models, among which the modeller has to conveniently select the one to be adopted for each one of the components of the water system considered. Translating such a convenience into few and precise selection criteria is not easy at all (see Jakeman *et al.*, 2005)**.

The social acceptance of a model, i.e. its being widely and regularly applied in the same or similar contexts, might undoubtedly make it more appealing and encourage its adoption, thus simplifying the work of the modeller. However, stakeholders might have already experienced the use of a different model and be reluctant in moving from it. On the other hand, the model accuracy, in terms of quantitative performance indexes (e.g. explained variance), which would be the selection criterion naturally adopted by any modeller, might convince only the stakeholders that have a strong technical expertise. In short, the selection has to be made weighting social acceptance and accuracy with other issues of great relevance, that in the experience of the authors include:

Easiness of identification. To foster the stakeholders participation to the model building process it is of primal importance that the model be easily identifiable. This requires that its parametric expression (i.e. the meta-model) be legible and quickly understandable by the stakeholders and, besides, that empirical and/or mathematical methods (algorithms) to estimate its parameter be available.

Integrability. Water systems are generally composed of many components (catchment, reservoir, farms, fish, etc.), which may also be different in nature (physical, social, economical, ecological, etc.). As the model of the whole system is obtained by integrating the models of the components, the type of these latter models must be such to make technically feasible such integration, without any loss of information or accuracy. Obviously this condition may be satisfied by adopting the same type of model for all the components, but this restriction is unnecessarily heavy, as it we will be apparent in a while.

Dynamics and parsimoniousness. Within a water system, components with a dynamic behaviour are not the exception, but the rule. These components must be described by means of dynamic models and, when their management has to be considered, in the Phase of Designing Alternatives (Phase 4 in Fig. 6), an Optimal Control problem has to be solved. The computational burden for its solution exponentially increases with the dimensionality (i.e. the number of state variables) of the model of the whole water system. Therefore, the type of model selected should have the capacity of capturing in a synthetic and essential way all the complexity of the component, without loosing transparency, being too mathematically abstruse, or trustability for the stakeholders, because the many simplifications introduced. When this is not possible a good practice would be first to built up a model (*evaluation model*) for estimating the effects in Phase 5. Then to obtain a *parsimonious* version of it (*screening model*) for designing the alternatives in Phase 4 (an example is in (Chappell et al., 2001)).

If we now reconsider the Bns' features presented in Section 4.2.1 in the light of the these issues, it clearly emerges that Bns do not always meet the above requirements. As a consequence forcing their use as a general modelling tool, regardless the features of the individual component being modelled, may result not only in poor performance, in terms of the accuracy of the system description, but also in the lost of sense of ownership and confidence of the stakeholders toward the model. Bns are a type of models that provides a simplified semantic that could be useful when the knowledge on the system to be modelled is poor or not structured, and mainly empirical in nature. In this case Bns are definitively a well-suited tool to quantitatively transform such a kind of a-priori information. Their identification might be sometimes facilitated by their structural properties, but could become critical when the number of values required to describe correctly the system variables is fairly high. Moreover, Bns partially loose their great potential when the considered system is dynamic, and includes recursive decisions (e.g. regulation policy). Therefore we suggest to limit the use of Bns to the water system components that do meet the above three characteristics (unstructured knowledge, the system variables assume a low number of values, absence of dynamics), while to adopt other types of models for the other components. This is possible because the Bns can easily be integrated with other types of models, by expressing them in the general form of equation (3), as is shown in Fig. 10. For instance, in the lake example it is convenient to describe the upstream catchment by means of an empirical model, the lake by means of a mechanistic one and the irrigation district through a Bn. The resulting model of the whole water system can therefore be used first to derive from it a Markov chain for designing the alternatives, and then to simulate the system behaviour in estimating the effects of the alternatives. An example of the integration of different types of models could be found in Castelletti and Soncini-Sessa (*this issue*).

Finally note that the direct use of Bns, in the form of Influence Diagrams, as a Decision Support System (see Section 2.1) implies that all the components of the water system be described as Bns. It should now be clear why this approach is weak: it can only be applied when the system is non-dynamic or, if dynamic, it is stationary and the time horizon over which the alternatives are evaluated not too long. However, another issue exists that contributes to this weakness: Bn-based DSSs are able to support the decision-making procedure only when the alternatives have to be compared, but not when they have to be negotiated. With a Bn (i.e. an Influence Diagram), in fact, alternatives can be only compared by running a sequence of "what if": a first alternative is specified by putting evidence on a given combination of values of the decision nodes, and the corresponding utilities calculated through belief propagation; a second alternative is then specified by changing the evidence, and the process repeated until all the alternatives have been evaluated. This process, however, can not be negotiated. In order for it to be, after the evaluation of the first alternative, each stakeholder should say if (s)he is satisfied or not by it. This could easily be done even with a Bn (i.e. an Influence Diagram) under the condition that it contains as many utility nodes as are the stakeholders' points of view. The critical point comes afterward, when, in order to enlarge the consensus, an alternative has to be found that increase the utilities of the unsatisfied stakeholders without lowering the utilities of the satisfied ones. However this type of search can not be performed by an Influence diagram, but requires a more complex and flexible tool.

6 Conclusions

In the last decade Bns have captured the interest of the environmental modelling community thanks to their friendly semantic and the graphical support they provide to the interaction with the stakeholders. The wide availability of ready-to-use software that allows to easily design and implement Bn models on a PC further contributed to their diffusion. In this paper we have shown pros and cons of their use in water resource planning and management, framing it within a general Participatory and Integrated Planning procedure (PIP). We argued that Bns are well suited for modelling system components on which the knowledge is poor or not structured, and mainly empirical in nature. Moreover their use as stand alone DSS is limiting, since it imposes all the components to be modelled by means of Bns, and definitively unsuitable when the compromise alternative has to be determined via negotiation.

References

Andreassen, S., Jensen, F., Olesen, K., 1991. Medical expert systems based on causal probabilistic networks. International Journal of Biomedical Computing 28, 1–30.

Baran, E., Jantunen, T., 2004. Stakeholder consultation for bayesian decision suport systems in environmental management. In: proceedings of ECOMOD, September 15-16. Penang, Malaysia.

Batchelor, C., Cain, J., 1991. Application of belief networks to water management studies. Agricultural Water Management 4 (1), 51–57.

Borsuk, M., P.Reichert, Burkhardt-Holm, P., 2002. A bayesian belief network for modelling brown trout (salmo trutta) populations in switzerland. In: Integrated Assessment and Decision Support, proceedings of 1st biennial meeting of IEMSS, June 24-27. Lugano, CH.

Borsuk, M., Stow, C., Higdon, D., Reckhow, K., 2001. A bayesian hierarchical model to predict benthic oxygen demand from organic matter loading in estuaries and coastal zones. Ecological Modelling 143, 165–181.

Borsuk, M., Stow, C., Reckhow, K., 2004. A bayesian network of eutrophication models for synthesis, prediction, and uncertainty analysis. Ecological Modelling 173, 219–239.

Bromley, J., Jackson, N., Clymer, O., Giacomello, A., Jensen, F., 2005. The use of hugin to develop bayesian network as an aid to integrated water resource planning. Environmental Modelling and Software 20, 231–242.

Castelletti, A., Cellina, F., Soncini-Sessa, R., 2004. The PIP procedure in the planning of lake verbano. In: IFAC Workshop on Modelling and Control for Participatory Planning and Managing Water Systems, September 28 - October 1. Venezia, I.

Castelletti, A., Soncini-Sessa, R., 2005. A procedural approach to strengthening integration and participation in water resource planning. Environmental Modelling & Software , submitted to.

Castelletti, A., Soncini-Sessa, R., this issue. Coupling real time control and socio-economic issues in river basin planning. Environmental Modelling & Software .

Chappell, N., McKenna, P., Bidin, K., Douglas, I., Walsh, R., 2001. Parsimonious modelling of water and suspended-sediment flux from nested-catchments affected by selective tropical forestry. Imperial College Press, London, UK, Ch. In Changes and disturbance in tropical rainforest in South-East Asia, pp. 107–122.

Charniak, E., 1991. Bayesian networks without tears. AI Magazine 12 (4), 50–63.

European Commission, 2000. Directive 2000/60/EC of the European Parliament and of the Council establishing a framework for community action in the field of water policy. Official Journal OJ L 327, european Commission, Brussels, B.

Gu, Y., McNicol, J., Peiris, D., Crawford, J., Marshall, B., Jefferies, R., 1996. A belief network-based system for predicting future crop production. AI Applications 10 (1), 13–24.

Hamilton, P., Anderson, N., Bartels, P., Thompson, D., 1994. Expert system support using bayesian belief networks in the diagnosis of fine needle aspiration biopsy specimens of the breast. Journal of Clinical Pathology 47, 329–336.

Hare, M., Letcher, R., Jakeman, A., 2003. Participatory modelling in natural resource management: a comparison of four case studies. Integrated Assessment 4, 62–72.

Heckerman, D., Mamdani, A., Wellman, M., 1995. Real world applications of bayesian networks. Communicarions of the ACM 38 (3), 24–26.

Hodgson, A., 1992. Hexagons for systems thinking. European Journal of Operational Research 59, 220–230.

Huang, C., Darwiche, A., 1996. Inference in belief networks: A procedural guide. International Journal of Approximate Reasoning 15 (3), 225–263.

Jakeman, A., Letcher, R., Norton, J., 2005. Ten interactive steps in model development and evaluation. Environmental Modelling & Software , submitted to.

Jensen, F., 1996. An Introduction to Bayesian Networks. Springer-Verlag, Heidelberg, D.

Jensen, F., 2001. Bayesian Networks and Decision Graphs. Springer-Verlag,

Heidelberg, D.
Kalman, R., Falb, P., Arbib, M., 1969. Topics in Mathematical System Theory. McGraw-Hill, New York, NY.
Kuikka, S., Hilden, M., Gislason, H., Hansson, S., Sparholt, H., Varis, O., 1999. Modeling environmentally driven uncertainties in baltic cod (gadus morhua) management by bayesian influence diagrams. Can. J. Fish. Aquat. Sci. 56, 629–641.
Lauritzen, S., 1996. Graphical Models. Clarendon Press, Oxford, UK.
Lauritzen, S., Spiegelhalter, D., 1988. Local computations with probabilities on graphical structures and their application to expert systems (with discussion). Journal of the Royal Statistical Society 50, 157–224.
Little, L., Kuikka, K., Punt, A., Pantus, F., Davies, C., Mapstone, B., 2004. Information flow among fishing vessels modelled using a bayesian network. Environmental Modelling and Software 19, 27–34.
Neapolitan, R., 1990. Probabilistic reasoning in Expert Systems: Theory and Algorithms. John Wiley, New York, NY.
Oliver, R., Smith, J., 1990. Influence Diagrams, Beliefe Nets and Decision Analysis. John Wiley, New York, NY.
Pearl, J., 1986. Fusion, propagation, and structuring in belief networks. Artificial Intelligence 29, 241–288.
Pearl, J., 1988. Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, San Mateo, CA.
Pearl, J., 2000. Causality: Models, Reasoning, and Inference. Cambridge University Press, Cambridge, UK.
Peot, M., Shachter, R., 1991. Fusion and propagation with multiple observations in belief networks. Artificial Intelligence 48 (3), 299–318.
Soncini-Sessa, R., 2004a. Planning and managing water resources: integrated modelling and participatory decision making in practice: the Verbano Project (in Italian). McGrawHill, Milano, I. (English version forthcoming with Elsevier).
Soncini-Sessa, R., 2004b. Planning and managing water resources: MODSS for integrated and participatory decision making (in Italian). McGrawHill, Milano, I. (English version forthcoming with Elsevier).
Varis, O., 1997. Bayesian approach to expert judgment elicitation with case studies on climatic change impact assessment on surface waters. Climate Change 37 (3), 539–563.
Varis, O., Kuikka, S., 1997. Bayesian approach to expert judgment elicitation with case studies on climatic change impact assessment on surface waters. Climate Change 37 (3), 539–563.
Yu, D., Nguyen, T., Haddawy, P., 1999. Bayesian network model for reliability assessment of power systems. IEE transactions on Power Systems 14, 426–432.