

## ON-LINE DESIGN OF WATER RESERVOIR POLICIES BASED ON INFLOW PREDICTION

A. Castelletti\* D. de Rigo\* R. Soncini-Sessa\*  
L. Tepsich\* E. Weber\*

\* *Dipartimento di Elettronica e Informazione, Politecnico  
di Milano, Milan, Italy*

Abstract: Stochastic Dynamic Programming (SDP) is the method most extensively adopted to design release policies for water reservoir networks. However, it suffers of the well known “curse of dimensionality”, which actually limits its applicability to small reservoir networks. In this paper we present an on-line approach to policy design that not only constitutes a viable alternative to overcome the SDP limits, but can also be used with an inflow predictor to improve the performance of SDP-based off-line policies. This latter possibility is explored and discussed through a real world case study. *Copyright*©2008 IFAC.

Keywords: on-line optimal control, reservoir management, stochastic dynamic programming, inflow predictor

### 1. INTRODUCTION

The design of release policies for water reservoir networks is a complex problem due to various aspects of the modelling involved, including the existence of multiple conflicting interests, the randomness of inflows, and the non-linearity of the release functions and the objectives. Stochastic Dynamic Programming (SDP) is by far the solution approach adopted in most of the real world applications as it can be applied under very weak conditions and can easily handle uncertainties and non-linearities. However, SDP suffers of the well known “curse of dimensionality”: the computational burden increases exponentially with the number of state variables. This actually limits its applicability to small reservoir networks. A radical, though very common, solution for overcoming the curse is to somehow manipulate the model of the system and reduce the dimensionality of its state. Many authors proposed to modify the system topology in order to reduce the number of reservoirs (storage variables) to a value that is handable by SDP. However, the effectiveness of

this approach is highly problem-dependent and it can be applied in a limited number of practical circumstances. In this paper we explore the applicability of an alternative approach (see Bertsekas (1976)) in which the state reduction is performed by eliminating the model of the uncontrolled parts of the system, i.e. the natural catchments, and considering their outputs, i.e. the outflows, among the system’s disturbances. The dynamics of the outflows is then accounted for by solving the problem on-line and updating their probability distribution functions with a dynamic outflow predictor fed with real-time information. This information may include not only the current state of the natural uncontrolled catchments but also any other variable that is useful for predicting future outflows, like, for example, precipitation or snow-cover measures. The use of real-time information makes the approach particularly interesting not only as a remedy to the curse of dimensionality but also as a way of refining the off-line policy with information that was not available when the off-line problem was formulated and solved, thus

adapting such a policy to any new and unexpected condition the system may face. This potential is explored in this paper through a real world case study of a simple reservoir network composed of a multipurpose regulated lake fed by a natural uncontrolled catchment. The Pareto efficient off-line policy, which resulted as the best compromise policy at the end of a previous planning process (see Soncini-Sessa *et al.* (2007b)), is refined on-line using different steps-ahead outflow predictions provided by an outflow predictor (e.g. persistent, perfect or real).

The paper is structured as follows. In Section 2 the general model of a water reservoir network is presented. Section 3 and 4 present the off-line and on-line control problem formulation respectively. Section 5 is entirely devoted to the case study and its results. Finally, in Section 6 we draw some conclusion on the work.

## 2. MODEL OF THE SYSTEM

We consider the general case of a water system composed of  $N$  reservoirs that are fed by  $M$  uncontrolled catchments and serve  $L$  water users, like, for example, power plants or irrigation districts.

### 2.1 Reservoirs

The model of the  $j$ -th water reservoir is based on the usual mass balance equation

$$s_{t+1}^j = s_t^j + q_{t+1}^j - r_{t+1}^j \quad (1a)$$

where  $s_t^j$  is the storage in the  $j$ -th reservoir at time  $t$ ,  $q_{t+1}^j$  is the inflow volume in the time interval  $[t, t+1)$  and  $r_{t+1}^j$  is the release in the same interval<sup>1</sup>. Other terms like direct precipitation on the reservoir, infiltration and evaporation have been neglected but they can easily be added to the mass balance when necessary. The inflow  $q_{t+1}^j$  is the outflow of a drainage network fed by the releases  $r_{t+1}^i$  ( $i = 1, \dots, i \neq j$ ) of the upstream reservoirs (if any) and by the outflows  $a_{t+1}^k$  ( $k = 1, \dots$ ) from the natural uncontrolled catchments. The release  $r_{t+1}^j$  is a function of the control variable  $u_t^j$  (which is the release decision made at time  $t$  for reservoir  $j$ ), of the storage  $s_t^j$  and of the inflow  $q_{t+1}^j$

$$r_{t+1}^j = R_t^j(s_t^j, u_t^j, q_{t+1}^j)$$

The function  $R_t^j(\cdot)$  is called release function and it is a non-linear periodic function by mean of which any potential deviation of the actual release  $r_{t+1}^j$  from the release decision  $u_t^j$  (e.g. when the

available water is not sufficient to realize the decision or when spill takes place) can be described appropriately (see Soncini-Sessa *et al.* (2007a) for a detailed description).

### 2.2 Uncontrolled catchments

Simple statistical models are usually adopted to describe the outflows from natural uncontrolled catchments. For instance, the outflow  $a_{t+1}^k$  from the  $k$ -th uncontrolled catchment can be assumed to be a cyclostationary, lognormal, stochastic process with periodic mean  $\mu_t^k$  and standard deviation  $\sigma_t^k$ , and its dynamics be described as

$$a_{t+1}^k = \exp(y_{t+1}^k \cdot \sigma_t^k + \mu_t^k) \quad (2a)$$

$$A^k(z^{-1})y_{t+1}^k = \varepsilon_{t+1}^k \quad (2b)$$

where  $A^k$  is a polynomial in the backward shift operator  $z^{-1}$  and  $\varepsilon_{t+1}^k$  is a zero mean Gaussian white noise with constant variance.

### 2.3 Water users

The presence of the  $L$  water users can be formalized by defining for each one of them a step-cost function associated to the system's transitions.

### 2.4 Global model

The global model of the water system is obtained by suitably aggregating the models of the reservoirs, catchments and water users, that compose it. The result is a discrete-time, periodic, non-linear, stochastic system of the form

$$\mathbf{x}_{t+1} = f_t(\mathbf{x}_t, \mathbf{u}_t, \varepsilon_{t+1}) \quad (3)$$

where  $\mathbf{x}_t \in \mathbb{R}^{n_x}$ ,  $\mathbf{u}_t \in \mathbb{R}^{n_u}$  and  $\varepsilon_t \in \mathbb{R}^{n_\varepsilon}$  are the state, control and disturbance vectors. The state is composed of the state variables of the  $N$  reservoirs, i.e. their storages, the state variables of the  $M$  catchments, and, when the case, the state of water users

$$\mathbf{x}_t = [s_t^1, \dots, s_t^N; y_t^1, \dots, y_{t-p_1}^1; \dots; y_t^M, \dots, y_{t-p_M}^M, \dots]^T \quad (4)$$

where  $p_k$  is the order of polynomial  $A^k(z^{-1})$  in equation (2b). The control vector is composed of the  $N$  release decisions for the  $N$  reservoirs

$$\mathbf{u}_t = [u_t^1, \dots, u_t^N]^T$$

The disturbance vector is composed of the  $M$  random disturbances that appear in the models of the uncontrolled catchments and any other random variable that could be used to describe random terms in the reservoir mass balance equation (e.g. evaporation, infiltration, etc.) or in the model of

<sup>1</sup> The time subscript of each variable denotes the time instant at which it assume a deterministic value.

water users. For example, if uncontrolled catchments are described with models of the form (2b) and no other disturbance affects the water system, the disturbance vector is given by

$$\varepsilon_{t+1} = [\varepsilon_{t+1}^1, \dots, \varepsilon_{t+1}^M]^T$$

The disturbance vector  $\varepsilon_{t+1}$  is described in terms of a pdf  $\phi_t(\cdot)$ , which at each time  $t$  may be function of the state and control at the same time

$$\varepsilon_{t+1} \sim \phi_t(\cdot | \mathbf{x}_t, \mathbf{u}_t) \quad (5)$$

### 3. OFF-LINE CONTROL PROBLEM

For each of the  $L$  users present in the system an objective function  $J^i$  (with  $i = 1, \dots, L$ ) can be defined to express the cost (s)he pays over an infinite time horizon,

$$J^i = \lim_{h \rightarrow \infty} \frac{1}{h} E_{\varepsilon_1, \dots, \varepsilon_h} \sum_{t=0}^{h-1} g_t^i(\mathbf{x}_t, \mathbf{u}_t, \varepsilon_{t+1}) \quad (6)$$

where  $g_t^i(\cdot)$  for  $t = 1, \dots, h-1$  are the step-cost functions associated to the transitions from  $t$  to  $t+1$ . At each time step, the release decision for each reservoir is given by the control law

$$\mathbf{u}_t = m_t(\mathbf{x}_t) \quad (7)$$

The scope of the off-line control problem (OffCP) is to define the sequence of control laws  $m_t(\cdot)$  over the horizon  $[0, h-1]$ , i.e. the release policy

$$p = [m_0(\cdot), \dots, m_{h-1}(\cdot)] \quad (8)$$

Therefore the multi-objective OffCP is formulated as

$$\min_p [J^1, J^2, \dots, J^L] \quad (9)$$

subject to the constraints (3), (5), (7), (8) and given  $\mathbf{x}_0$ . The solution is not a unique optimal policy but is constituted by the set  $\mathcal{P}$  of Pareto efficient policies and it is usually obtained by applying SDP.

### 4. ON-LINE APPROACH

The idea of the on-line approach is as follows: the models of the uncontrolled catchments are eliminated and their outflows  $a_{t+1}^k$ ,  $k = 1, \dots, M$ , are included among the disturbances of the reduced water system model. This is possible because these subsystems are not influenced by the control  $\mathbf{u}_t$ . By doing so, the number of components in the state vector (4) is reduced, since the components  $y_{t-i}^k$  do not appear. We denote with  $\tilde{\mathbf{x}}_t$  the reduced state. At each time  $t$ , an on-line optimal control problem (OnCP) over a finite horizon  $[t, t+h]$  is formulated and solved. For each time  $\tau$  in the finite horizon  $[t, t+h]$ , the pdf  $\phi_\tau(\cdot)$  of the disturbance is provided by a dynamic predictor that uses all information  $I_t$

available at time  $t$  that are significant for the prediction of the catchments outflow. Once the on-line problem has been solved, only the control for the first time step  $[t, t+1]$  is actually applied and, at time  $t+1$ , a new problem is formulated over the horizon  $[t+1, t+1+h]$  with pdfs for the disturbances based on  $I_{t+1}$  (receding horizon principle). In other words, the on-line updating of the outflow pdfs can be based on a model more sophisticated than model (2). In most of the cases, in fact, the description of the uncontrolled catchment provided by model (2) is a rough approximation but it can not be improved due to the need of limiting the state dimension in the off-line solution with SDP. The multi-objective OnCP can be formulated as a stochastic closed-loop control problem, usually known as POLFC (Partial Open-Loop Feedback Control) problem

$$\min_p [J^1, J^2, \dots, J^L] \quad (10a)$$

with

$$J^i = E_{\varepsilon_{t+1}, \dots, \varepsilon_{t+h}} \left[ \sum_{\tau=t}^{t+h-1} g_\tau^i(\tilde{\mathbf{x}}_\tau, \mathbf{u}_\tau, \varepsilon_{\tau+1}) + g_{t+h}^i(\tilde{\mathbf{x}}_{t+h}) \right]$$

subject to

$$\tilde{\mathbf{x}}_{\tau+1} = \tilde{f}_\tau(\tilde{\mathbf{x}}_\tau, \mathbf{u}_\tau, \varepsilon_{\tau+1}) \quad (10b)$$

$$\varepsilon_{\tau+1} \sim \phi_\tau(\cdot | I_t) \quad \tau = t, \dots, t+h-1 \quad (10c)$$

$$\tilde{\mathbf{x}}_t \text{ given} \quad (10d)$$

$$\mathbf{u}_\tau = m_\tau(\tilde{\mathbf{x}}_\tau), \quad \tau = t, \dots, t+h-1 \quad (10e)$$

$$p = [m_t(\cdot), \dots, m_{t+h-1}(\cdot)] \quad (10f)$$

and can be solved by means of SDP.

The choice of the penalty function  $g_h(\cdot)$ , which influences either the performances of the closed loop scheme and its stability, is particularly difficult. One possibility (Nardini *et al.*, 1994) is to let  $g_h(\cdot)$  be equal to the optimal cost-to-go  $H_h(\cdot)$  obtained by solving an off-line infinite horizon problem with the reduced model and a trivial predictor, i.e. with a priori pdf for the description of the disturbance. However, since the solution of the latter problem requires using SDP, this approach can be followed only if the reservoir network is composed of few reservoirs, as in the case case study discussed next.

### 5. APPLICATION TO THE DAILY MANAGEMENT OF LAKE MAGGIORE

Lake Maggiore is a regulated lake located south of the Alps between Italy and Switzerland. It is the most important water system of the sub-alpine area on account of its multiple and conflicting socio-economic uses (irrigation, hydropower generation, navigation, flooding reduction, etc.). At the end of 1999 an EU-Interreg project was funded with the purpose of exploring whether any planning alternative (i.e., combination of structural and

normative interventions and a release policy) exists that the parties can agree on to resolve, or at least mitigate, the existing conflict (see Soncini *et al.* (2007b) for a detailed description of the project and its outcomes). Although the project actually ended with the identification of a set of reasonable alternatives, i.e. the alternatives gathering a large consent among the parties, here we assume that the Italian and Swiss governments have agreed on the choice of the best compromise alternative from that set. More precisely, we assume that they have chosen alternative A34 which foresees the excavation of the lake outlet (with an increase of 600 m<sup>3</sup>/s of the outflow capacity), the modification of the regulation range (with the setting of its upper extreme to 1.5 m all through the year), and, finally, an efficient, off-line, release policy  $p^{A34}$ , designed (a) by taking into consideration, as objective of the OffCP, a convex linear combination of flooding reduction around the lake and satisfaction of the downstream irrigation users; (b) by assuming the a priori probability distribution for the inflow  $\varepsilon_{t+1}$ .

In autumn 2000, three flood waves occur in succession (dot-dashed line in Figure 1(a)) and under the historic regulation produced two flooding events (dotted line in Figure 1(b)). The level and release trajectories that would have been obtained with A34 are reported in Figure 1(a) and (b) (continuous line). Notice how the second flooding event (B) would have been completely avoided and the third (C) significantly reduced: this marked improvement is due to the effect both of the excavation and of a more efficient regulation. The first allows, the levels being equal, the release of a larger flow and thereby a faster decrease of the level. Consider in fact what occurs on 22nd September (point D in the figure), when the historic level and the one produced by A34 coincide: the historic release is 153 m<sup>3</sup>/s (the maximum releasable given the lake level), while the release with the A34 is 709 m<sup>3</sup>/s. It is this higher release that allows A34 to significantly reduce the lake level. As far as the second reason is concerned, the high expected costs associated with the high levels induce A34 to maintain, as far as possible, the reservoir level around zero, which is the level to which the minimum expected costs are associated in this season. That is the reason why, after 22nd September, A34 completely opens the dam gates, although the reservoir level is still within the regulation range. The historical regulation behaves differently: first it releases the maximum releasable volume, but later, in the following days, it maintains the release at the same value, in spite of the increase of the level. Finally, let us observe what happens at the end of the flood event. After 15th September, A34 encourages the lake manager to gradually increase the level, be-

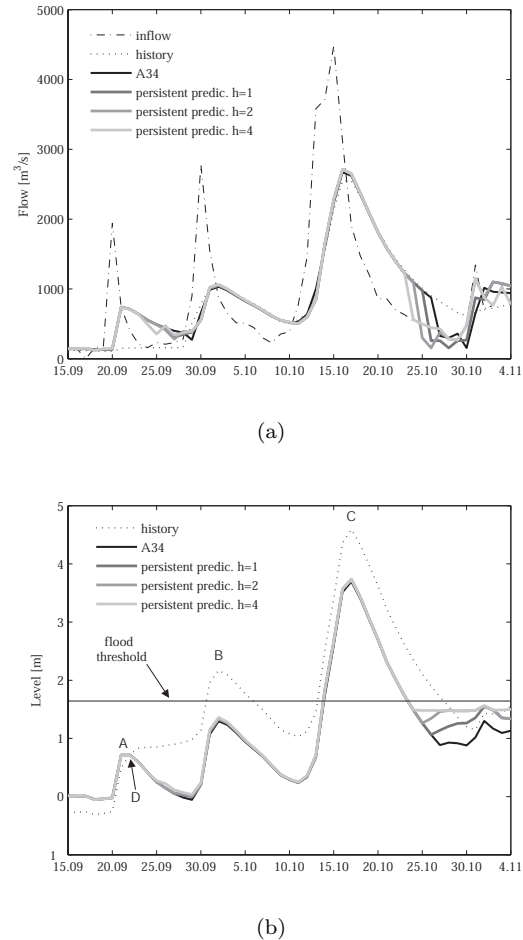


Fig. 1. The flood event of autumn 2000. (a) Lake inflow, historical release and releases generated by A34 and by a POLFC scheme fed by a persistent predictor with  $h = 1, 2, 4$ . (b) Historical level and levels that would have been obtained with A34 and with a POLFC scheme fed by a persistent predictor with different forecasting horizon ( $h = 1, 2, 4$ ).

cause in this period the thermal zero elevation generally rises, the flood probability lessens and as a consequence the interest in storing water for the spring irrigations prevails. That is the reason why  $p^{A34}$  tends to fill the lake in this period.

Since in the hydrometeorological sub-Alpine regime of the Lake Maggiore area the autumn is characterized by heavy rain and therefore by floods, it could be productive to use an inflow predictor combined with a POLFC scheme to improve the performance of policy  $p^{A34}$ . Then, at every time step, a problem of the form shown in (10) has to be solved, on a horizon of  $h$  steps, assuming as penalty  $g_{t+h}(\cdot)$  the optimal cost-to-go obtained by solving the OffCP with SDP, and using for the inflow the values  $\hat{\varepsilon}_\tau$  supplied by a predictor.

### 5.1 Persistent predictor

The simplest  $h$ -step-ahead predictor that can be imagined is the *persistent predictor*, which at every instant  $t$  forecasts a value  $\hat{\varepsilon}_{t+\tau|t}$  for the inflow in the interval  $[t + \tau - 1, t + \tau)$  equal to the inflow  $\varepsilon_t$  which has occurred in the last time interval  $[t - 1, t)$ , i.e.

$$\hat{\varepsilon}_{t+\tau|t} = \varepsilon_t \quad \tau = 1, \dots, h \quad (11)$$

In other words, the last inflow recorded at instant  $t$  is the only information which we exploit to predict the next  $h$  inflows; we limit ourselves to assuming that its value does not change through time. We shall also assume that the variance of  $\hat{\varepsilon}_{t+\tau|t}$  is zero. The level trajectories obtained with the on-line policy, with the persistent predictor and for different forecasting horizons ( $h = 1, 2, 4$ ), are reported in Figure 1(b) (lines with different tones of grey). For  $h = 1$ , the trajectory practically overlaps that of the A34 while for higher values of  $h$  the peak levels (A) and (B) are higher than those produced by A34: the longer the forecasting horizon  $h$ , the higher the level. It is thus evident that it is dangerous to use the POLFC scheme when the predictor is not a good one. It is therefore opportune to look for a better predictor.

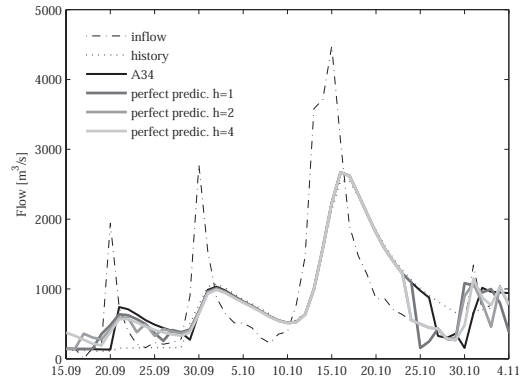
### 5.2 Perfect predictor

Intuitively, the best  $h$ -step-ahead inflow predictor is the *perfect predictor*: for every instant  $\tau$  between  $t + 1$  and  $t + h$ , it supplies the value  $\varepsilon_\tau$ , which effectively will occur, i.e.

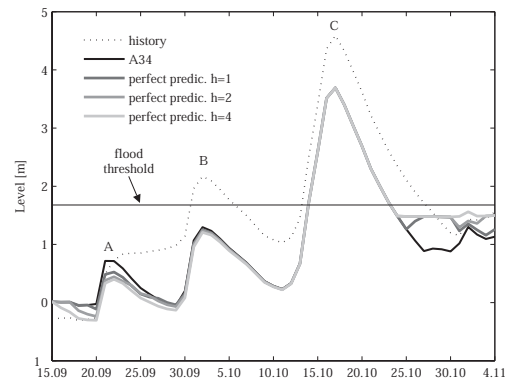
$$\hat{\varepsilon}_{t+\tau|t} = \varepsilon_{t+\tau} \quad \tau = 1, \dots, h \quad (12)$$

Clearly, such a predictor is not realizable in practice.

The trajectories produced with the perfect predictor on different forecasting horizons ( $h = 1, 2, 4$ ) are reported in Figure 2. There is now an improvement compared with A34; the longer the horizon, the more marked it is. Observing in particular the first event (A) (see Figure 3), one notes that the bringing forward of the reservoir's spilling in order to buffer the first inflow peak corresponds precisely to the forecasting horizon used: for example, with the 4-step-ahead predictor (lighter grey line) the policy begins to release more than the other policies on 15th September (4 days before the increase of the inflow) and this allows it to reduce the peak level on 21st by 0.40 m compared with A34. With all the predictors, the improvement is marked on the first peak, more contained on the second (0.06 m with a 4-step-ahead predictor and 0.03 m with a 1-step-ahead one) and zero on the third. This is easily explained by observing that from 22nd September the lake is in free regime with all the policies (also with the POLFC, which is driven by the penalty that it inherits from A34)



(a)



(b)

Fig. 2. Same as Figure 1 one but with a perfect predictor.

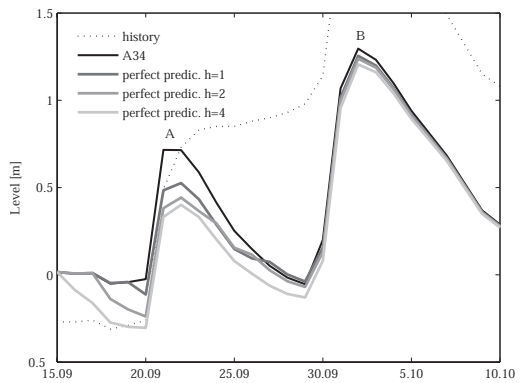


Fig. 3. An enlargement of the first two flooding events reported in Figure 2(b).

and as a consequence of the asymptotic stability of the system all the trajectories tend to overlap as time goes on. Of course, they do not overlap the historic trajectory, because the historic system was not excavated and thus the behaviours of the two systems in free regime are different.

### 5.3 Real predictor

As we have already emphasized a perfect predictor actually is not realizable: the performance obtained with it are therefore only useful to have an idea of the upper bound of the performance that one expects from the use of an on-line policy. With a real predictor, plausibly the improvement with respect to policy A34 will be less marked. And in fact with the best predictor that we were able to create (1-step-ahead<sup>2</sup>) we earned only 0.02 m on peak B.

From this analysis we can conclude that the improvement of the off-line policy A34 by an on-line policy is, all in all, modest. The reason lies in the fact that the period considered (autumn) is a usual flood season and thus the off-line policy already takes due account of this. The advantage of the POLFC becomes more significant when an unexpected event occurs. In 1987, for example, surprisingly a flood took place in the month of July (a unique case in the hydrological series). Figure 4 shows that with a 1-step-ahead real predictor the peak reduction compared with A34 is of 0.07 m with respect to 0.08 m of the 1-step-ahead perfect predictor. The reduction may appear very modest, but notice that a reduction of 0.01 m corresponds to a reduction of 1 ha of the flooded area!

## 6. CONCLUSIONS AND REMARKS

An on-line approach to policy design for water reservoir networks has been presented in this paper. It has been proposed both as a viable way to overcome the well know SDP's curse of dimensionality in large reservoir networks and as a valuable approach to exploit real-time information to improve off-line, SDP based policies. This latter aspect has been explored through a real world case study. Results clearly show that the improvements of the off-line policy is, all in all, modest when floods actually occur in the flood season while become significant with unexpected events. Our research reveals that an increase in the forecasting horizon causes an increase in performance, but that increase is less than proportional to the increase of the forecasting horizon. On the other hand, the identification of a good inflow predictor entails costs and these increase as the forecasting horizon increases. In the choice of predictor it is thus necessary to find a trade-off between such costs and the resulting improvement in the performance. In this sense the information given by the perfect predictor can constitute important information.

<sup>2</sup> Notice that the concentration time of the catchment is less than 24 hours.

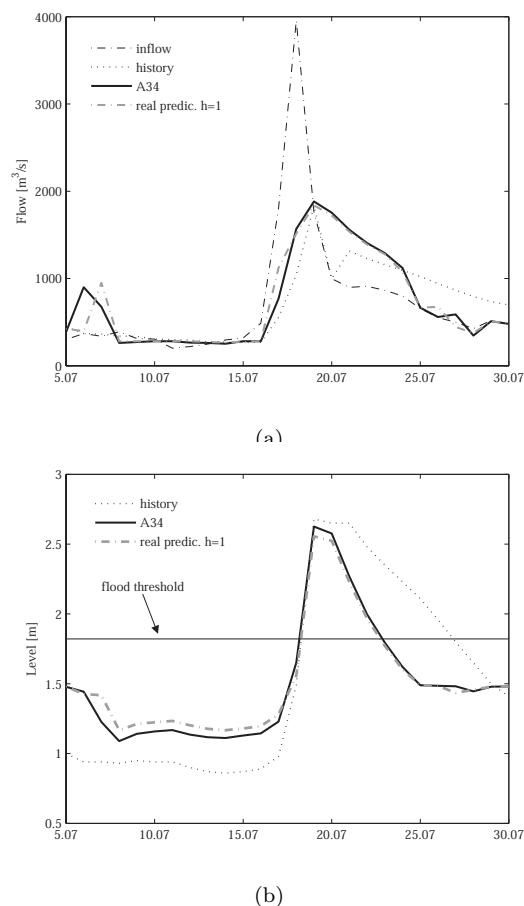


Fig. 4. The flood event of summer 1987. (a) Inflow, historical release and releases generated by A34 and by a POLFC scheme fed by a real predictor with  $h = 1$ . (b) Historical level and levels that would have been obtained with A34 and with a POLFC scheme fed by a real predictor with  $h = 1$ .

## 7. ACKNOWLEDGMENT

Partially supported by FONDAZIONE CARIPLO TWOLE-2004.

## REFERENCES

- Bertsekas, D.P. (1976). *Dynamic Programming and Stochastic Control*. Academic Press. New York.
- Nardini, A., C. Piccardi and R. Soncini-Sessa (1994). A decomposition approach to sub-optimal control of discrete-time systems. *Optimal Control Applications and Methods* **15**(1), 1–12.
- Soncini-Sessa, R., A. Castelletti and E. Weber (2007a). *Integrated and participatory water resources management. Theory*. Elsevier. Amsterdam, The Netherlands.
- Soncini-Sessa, R., F. Cellina, F. Pianosi and E. Weber (2007b). *Integrated and participatory water resources management. Practice*. Elsevier. Amsterdam, The Netherlands.