  Journal of Hydroinformatics | 01.2 | 1999

# Using phase-state modelling for inferring forecasting uncertainty in nonlinear stochastic decision schemes

Ezio Todini


## ABSTRACT

The paper introduces the use of phase-state modelling as a means of estimating expected benefits or losses when dealing with decision processes under uncertainty of future events. For this reason the phase-space approach to time series, which generally aims at forecasting the expected value of a future event, is here also used to assess the forecasting uncertainty. Under the assumption of local stationarity the ensemble of generated future trajectories can be used to estimate a probability density that represents the *a priori* uncertainty of forecasts conditional on the latest measurements. This *a priori* density can then be used directly in the optimisation schemes if no additional information is available, or after deriving an *a posteriori* distribution in the Bayesian sense, by combining it with forecasts from deterministic models, here taken as noise-corrupted 'pseudo-measurements' of future events. Examples of application are given in the case of the Lake Como real-time management system as well as in the case of rainfall ensemble forecasts on the River Reno.


**Key words** | decision processes, forecasting uncertainty, phase-state modelling

**Ezio Todini**
Dipartimento di Scienze della Terra e Geologico Ambientali,
University of Bologna,
Italy

## INTRODUCTION

Management and control of flooding generally require the use of forecasting techniques to assist in the analysis of the effects of decisions under the uncertainty of future events.

In a utopian world of linear systems the expected value of the forecasts would be sufficient for taking reliable and effective decisions; unfortunately, the highly nonlinear and discontinuous behaviour of rainfall runoff and routing processes, of the control structures (such as reservoirs and dykes) and of the loss mechanisms, make the forecast, for expected value of rainfall, for flows or levels, a useless quantity if it is not associated with a description of the probability density of the forecasting errors, or at least of a measure of the uncertainty. This is essential owing to the inherently uncertain nature of forecasts, which rarely coincide with what will happen at a future time: this implies that the forecast must be taken not as a real image of future (certainty), but as a means for reducing the original *a priori* uncertainty.

Several authors have used phase-space approaches for forecasting, but have rarely provided a measure of uncertainty associated with the forecast, and in any case the use of this measure is not conceived to be operational tool. For instance, Smith (1992) provides an estimate of the confidence limits of the phase-space forecast only to show the variable performance in time.

Therefore, the novelty of this paper is in the introduction of phase-state modelling so as to construct an *a priori* probability density of future events based upon historical information and conditional on the latest measurements, to be used either directly for optimisation purposes, as in the case of the Lake Como real-time management system, or to be combined with additional information to produce an *a posteriori* distribution, as was done in the case of the real-time rainfall and flood forecasting system of the River Reno.

## THE PHASE-SPACE MODEL

Consider a deterministic system with phase space dimension $M_x$. A trajectory, $\boldsymbol{X}(t)$, of this system is reconstructed in $M$ dimensions from a time series of a single observable, $x(t)$, by the method of delays (Packard *et al.* 1980; Takens 1981). At time $t$, we consider the time series of interest that could be hourly or daily flow measured at a specific section, or rainfall averaged over a basin. The trajectory $\boldsymbol{X}(t)$ in phase-space may be represented with $M$ degrees of freedom, as:

$$X(t) = \begin{bmatrix} x(t) \\ x(t-\tau) \\ \dots \\ x(t-(M-1)\tau) \end{bmatrix}, \tag{1}$$

where $\tau$ is the sampling interval.

In general the problem is to construct a predictor (or map), $F(\boldsymbol{X}){:}\Re^M \rightarrow \Re^1$ which estimates $x$ for any $\boldsymbol{X}$. $F(\boldsymbol{X})$ is usually taken of the form:

$$F(X)=\sum_{j=1}^{n_c} \lambda_j \phi\left(\| X - X_j^c \|\right), \tag{2}$$

where:

$\boldsymbol{X}_j^c$, $j=1, 2, \dots, n_c$; $\boldsymbol{X}_j^c \in \Re^M$ represents a set of $n_c$ centres in an $M$-dimensional space; $\phi(\|\boldsymbol{X}-\boldsymbol{X}_j^c\|)$ are radial basis functions and $\lambda_j$ are constants to be determined by observations using a learning set:

$$F(X_i) \approx x_i \; . \tag{3}$$

Estimating the weights $\lambda_j$ corresponds to the solution of the linear problem:

$$\boldsymbol{b} = \boldsymbol{A}\lambda \; , \tag{4}$$

where $\lambda$ is a vector of length $n_c$ whose $j$th component is $\lambda_j$ and $\boldsymbol{A}$ and $\boldsymbol{b}$ are given by:

$$A_{ij} = \omega_i \phi\left(\| X_i - X_j^c \|\right) \tag{5}$$

and

$$b_j = \omega_i x_i \; , \tag{6}$$

where the weights $\omega_i$ reflect the confidence associated with the $i$th observation.

Casdagli (1989) solved this problem considering the special case of exact interpolation where the centres are chosen from the learning set and the number of observations $n_l$ equals the number of centres $n_c$, whereas others such as Broomhead & Lowe (1988) solved this problem in the least squares sense with $n_l > n_c$.

Because of the presence of noise and the high degree of non-stationarity in the natural systems to be analysed in the sequel, a slightly different approach has been taken in this paper, which leads to a linear locally stationary model.

In this case the learning set is based upon the latest measurements and its size equals the space dimension, namely $n_l = M$, which is smaller than the number $n_c$ of centres to be used to filter out the noise. The centres are selected using a Nearest Neighbour technique (Yakowitz 1987; Yakowitz & Karlsson 1987) based upon a Euclidean distance and assumed to be independent from each other. Under these assumptions the centres correspond to independent observations of a locally stationary process; a set of independent linear regressions can be established to relate each centre to the learning set $\boldsymbol{X}(t)$. This state vector is compared with past vectors $\boldsymbol{X}_j^c(t')$, representing the selected centres, showing the same temporal succession pattern (where $t' < t - f\tau$; $f$ being the forecasting horizon). The $j$th weight $\lambda_j$ is thus determined via a linear regression on the following model:

$$X(t) = \lambda_j X_j^c(t') + \varepsilon_j(t) \; , \tag{7}$$

where $\varepsilon_j(t)$ is taken as an $M$-dimensional Normal Independent Process.

Future values are then calculated as:

$$\hat{x}_j^c(t+k\tau) \hat{=} \hat{\lambda}_j \, x_j^c(t'+k\tau) \qquad k=1, f; \; j=1, n_c \, . \tag{8}$$

Each new estimate $\hat{x}_j^c(t+k\tau)$ is considered as a possible and equally probable event of a stochastic process, assumed to be locally stationary. The forecast at time $t+k\tau$ will be given, over a single series, in terms of the expected value and standard deviation of the conditional distribution calculated as a function of the $n_c$ generated $\hat{x}_j^c(t+k\tau)$ state vectors:

$$\mu_{\hat{x}_{t+k\tau|t}} = \frac{1}{n_c} \sum_{j=1}^{n_c} \hat{x}_j^c(t+k\tau) \tag{9}$$

$$\sigma_{\hat{x}_{t+k\tau|t}} = \left[ \frac{\sum_{j=1}^{n_c} [\hat{x}_j^c(t+k\tau)]^2 - n_c\, \mu^2_{\hat{x}_{t+k\tau|t}}}{n_c - 1} \right]^{\frac{1}{2}} . \tag{10}$$

If the spread of the generated traces $\hat{x}_j^c(t+k\tau)$ is far from being Normal, additional moments are needed to reproduce the conditional density, which can be estimated from the generated sample of $n_c$ values, with $n_c$ sufficiently large to produce reliable moment estimates.

## THE LAKE COMO REAL-TIME MANAGEMENT SYSTEM

Lake Como is in northern Italy and is mainly used for irrigation and hydro-electrical power production, and its outlet is controlled through a dam at Olginate. Management of the lake has to cope with the necessity of saving water to respond to the agricultural and hydro-electrical demands and, at the same time, must guarantee safety against the risk of flooding which has increased in the past 30 years, because of the increasing subsidence of the main square of Como.

The management problem is also complicated on the one hand by the relatively small reservoir control (246.5 million $m^3$, which is roughly 1/20 of the yearly inflow volume), and on the other hand by the small downstream gates that allow a release of 900–1,000 $m^3/s$ while the inflow can reach 1,800–2,000 $m^3/s$, which may lead to rapid rising and filling of the lake (3–5 days) in the case of large flood events.

The management operation must then respond to two basic requirements: the first is to optimise water resources on a yearly basis, and the second is to cope with fast reservoir responses to inflowing flood waters. Therefore the management operation was first established on a 10-day basis using a Stochastic Dynamic Programming (SDP) algorithm. This 10-day rule was then taken as the target rule for a finer optimisation based upon daily inflow forecasts.

### The Stochastic Dynamic Programming (SDP) approach

The 10-day operation rule was obtained using an SDP approach by modifying the original scheme developed at the Massachusetts Institute of Technology for the management of Lake Nasser (Alarcon & Marks 1979). The optimisation is based upon two state variables, the reservoir volume at the beginning of the 10-day period ($S_i^t$) and the last 10-day observed inflow ($Q_j^{t-1}$). The state variables discretisation has led to the identification of 97 volume values ($S_i^t$, $I = 1, \ldots, 97$) and 8 equally probable inflow values ($Q_j^{t-1}$, $j = 1, \ldots, 8$) for each 10-day period. The stochastic conditioning is performed by assuming a Markovian inflow for which the 10-day conditional probability matrices were derived as a function of the inflow discrete states.

The optimal release for each 10-day period, conditional on the previous inflow, was calculated by minimising the expected value of the following cost function:

$$f_n^t(S_i^t, Q_j^{t-1}) = \min_{R_{i,j}^t} \left\{ g_t(R_{i,j}^t) + d_t(S_i^t) + \sum_l P_{j,l}^{t-1} \cdot f_{n-1}^{t+1}(S_{i,j,l}^{t+1}, Q_l^t) \right\} \tag{11}$$

subject to the reservoir state transition equation

$$S_{i,j,l}^{t+1} = S_i^t + Q_l^t - R_{i,j}^t - E(S_i^t, Q_l^t, R_{i,j}^t) , \tag{12}$$

where:

$S_i^t$ storage at the beginning of period $t$ equal to the $i$th discrete value of that state variable,

$Q_l^t$ inflow during decade $t$ equal to the $l$th discrete value of that state variable,

$R_{ij}^t$ optimal release when at the beginning of period $t$ the storage state is $i$, the previous period inflow state is $j$,

$E(S_i^t, Q_l^t, R_{i,j}^t)$ the reservoir evaporation when the storage is equal to the $i$th discrete value for period $t$, the previous period inflow equals the $j$th discrete value and the release decision is $R$,

$P_{j,l}^t$ probability of occurrence of inflow $j$ in decade $t$, given that inflow $l$ was realised in decade $t-1$,

$d_t(S_i^t)$ cost in period $t$ associated with initial volume $S_i^t$ (namely due to flooding or low flow),

$g_t(R)$ cost in period $t$ when $R$ is not the optimal release,

$f_n^t(S_i^t, Q_j^{t-1})$ minimum expected cost for the release policy from the present period until the end of the process.

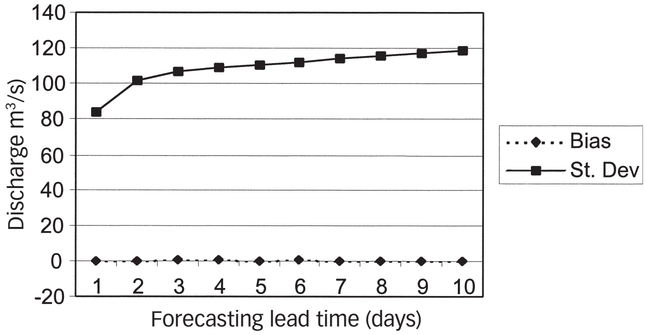


**Figure 1** | Expected value of error statistics for the phase-space model lake inflow forecasts.

## The daily rule

As previously mentioned, owing to the small size of the operational volume and the limited outflow from the lake due to the downstream gates, the 10-day rule, derived on the assumption of a stochastic cyclo-stationary process, although reflecting the long-term objectives of expected minimum losses, cannot be considered the most efficient rule for coping with short transient events such as high inflows that can fill the lake in 3–5 days.

Therefore it was necessary to modify the long-term policy in real time on the basis of daily inflow forecasts. Data of daily inflows to the lake over 35 years were used to develop the forecasting model based upon the phase-space approach presented above, with $n_l = M = f = 10$ (days) and $n_c = 20$ under the assumption of a Normal distribution of forecasting errors.

The first 20 years were used as *a priori* knowledge, whereas the additional 15 years were used as a test set to analyse the quality of the forecasting mechanism. Figure 1 shows the results in terms of the expected value of the error statistics. One can observe that the forecasting mechanism is virtually unbiased with a variance, expressing the uncertainty that increases with the forecasting lead time.

## The Stochastic Conditional Optimisation (SCOOP)

Using the new information provided by the phase-space model in terms of a probability density of future inflows to the lake conditional on the latest measurements, it is possible to set up a new stochastic optimisation scheme aimed at coping with emergency situations.

For this purpose a daily objective function is formulated by penalising on the one hand the departure from the 10-day rule here taken as a target, and on the other hand the expected value of losses due to low flows or to flooding, conditional upon the inflow forecast. This gives rise to the following objective function:

$$
\begin{aligned}
E\{J(x)\} &= c_t(R_t - x) + c_m(R_m - x) + c_v(R_v - x) \\
&+ \sum_{i=1}^{n_{max}} \int_{V_i}^{V_{i+1}} \left[\alpha_i + \beta_i(\xi - V_i) + \gamma_i(\xi - V_i)^2\right] \\
&\times \left(\exp\left\{-\tfrac{1}{2}\left[\frac{\xi - u - x}{\sigma}\right]^2\right\} \Big/ \sigma\sqrt{2\pi}\right) d\xi \qquad (13) \\
&+ \sum_{j=1}^{n_{min}} \int_{V_j}^{V_{j+1}} \left[\alpha_j + \beta_j(V_j - \xi) + \gamma_j(V_j - \xi)^2\right] \\
&\times \left(\exp\left\{-\tfrac{1}{2}\left[\frac{\xi - u - x}{\sigma}\right]^2\right\} \Big/ \sigma\sqrt{2\pi}\right) d\xi
\end{aligned}
$$

where:

- $R_t$ 10-day target release,
- $R_m$ minimum allowed release,
- $R_v$ river vital minimum flow,
- $c_t$ cost for failing the target release ($c_t \neq 0 \forall\ x < R_T$),
- $c_m$ cost for failing the minimum release ($c_m \neq 0 \forall\ x < R_M$),
- $c_v$ cost for failing the minimum river flow ($c_v \neq 0 \forall\ x < R_V$),
- $\alpha_i, \beta_i, \gamma_i$ coefficients of the stepwise quadratic objective function when the lake volume exceeds its allowed maximum; they differ from zero only for $V_i \leq x \leq V_{i+1}$,
- $\alpha_j, \beta_j, \gamma_j$ coefficients of the stepwise quadratic objective function when the lake volume drops below its allowed minimum; they differ from zero only for $V_j \leq x \leq V_{j+1}$.

As can be seen from Equation (13), in the SCOOP formulation the chance constraints used by the Extended Linear Quadratic Gaussian Control (ELQG) (Georgakakos & Marks 1987; Georgakakos 1989) to limit the lake surface elevation are introduced in the objective function in terms of expected losses from the penalty function given in Figure 2.

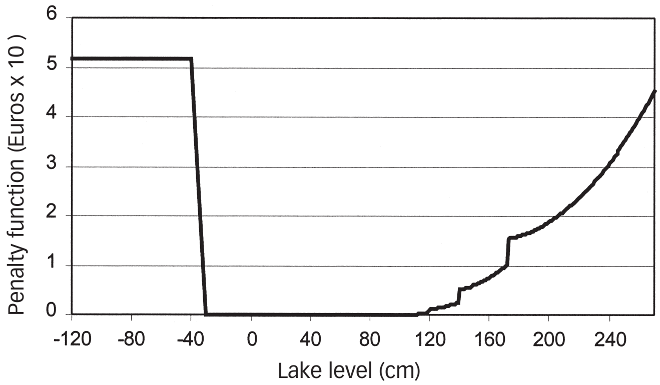


**Figure 2** | The penalty function applied to the lake levels from which an expected cost objective function is computed.

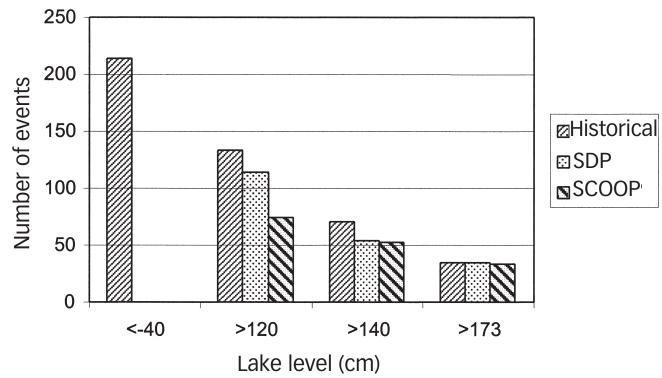


**Figure 3** | Improvement in lake management in terms of exceeding lake critical levels.

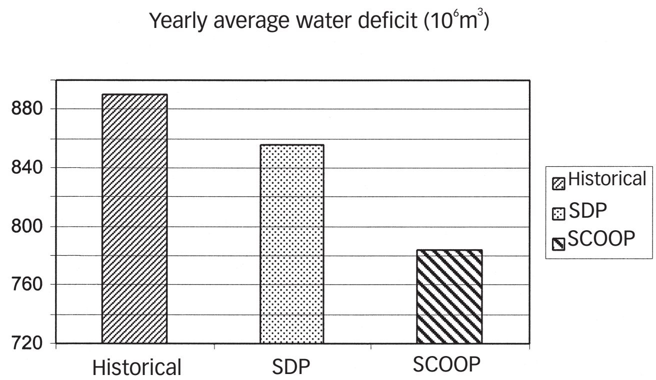


**Figure 4** | Average water deficit in 15 years for historical management, SDP and SCOOP.

It should also be noted that in the SCOOP formulation, the conditional inflow forecasts, as well as the conditional variance of forecasting errors, are expressed in terms of cumulative values (total inflow volume in increasing time intervals) because, as shown by Todini (1990), it is possible to avoid the use of stochastic dynamic programming if the cost coefficients do not vary during the optimisation period by solving $f$ independent problems to estimate the optimal cumulative release. The optimal release value for each time step (1 day) is then calculated backward as the difference between two optimal consecutive solutions.

The optimal solution is found by setting the first derivative of the objective function, for which an analytical solution is available, equal to zero and solving with a dichotomic method between zero and a large release value. The final daily release is then decided on the basis of a weighted average between the 10-day long-term cyclo-stationary rule and the 1-day non-stationary one on the basis of a weighting parameter representing, in a Bayesian sense (Berger 1980) the relative uncertainty of the daily forecast and of the 10-day conditional probabilities.

The results of the different simulations using the 15 test years (1981–95) are synthesised in Figures 3 and 4.

Figure 3 indicates the number of times that four critical levels in the lake are exceeded. The first level, $-40$ cm, represents the minimum allowed level for the lake; at 120 cm the main square of the city of Como starts flooding; at 140 cm traffic must be stopped; and the legal status of flooding is finally associated with 173 cm (i.e. a level of 53 cm above the main square in Como). The results of the two levels of optimisation, SDP and SCOOP, clearly show the reduction in the frequency of undesired events when compared with what has happened in reality: the lake is never below $-40$ cm and the frequency of the intermediate floods is noticeably reduced. Nothing can be done to mitigate the very large floods owing to the small size of the lake operational volume.

What is interesting is that at the same time as improving flood control, the water deficit drops by more than $100 \times 10^6$ $m^3$ per year on average (see Figure 4), passing from historical management to that resulting from SCOOP. Even using an extremely conservative estimate such as 0.1 Euro per $m^3$, one can estimate an average gain of 10 million Euros per year.

**Table 1** | Chosen flood events at Casalecchio.

| Event | Date | Year |
|---|---|---|
| 1 | 18–25 September | 1994 |
| 2 | 10–17 June | 1994 |
| 3 | 2–11 December | 1992 |
| 4 | 27–2 October/November | 1992 |
| 5 | 29–9 September/October | 1993 |
| 6 | 4–11 November | 1994 |
| 7 | 30–3 December/January | 1993–94 |
| 8 | 13–18 April | 1994 |
| 9 | 18–22 October | 1993 |

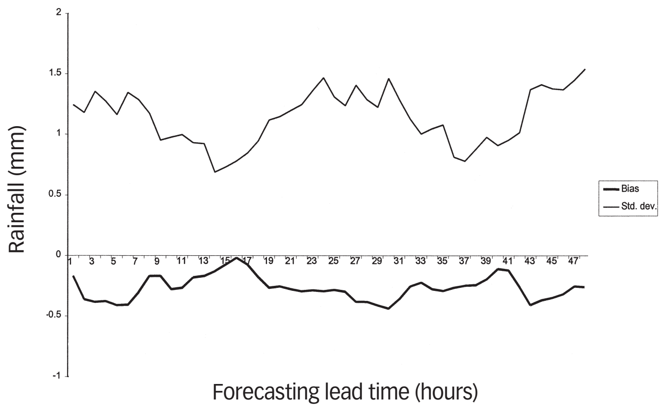


**Figure 5** | Bias and standard deviation of differences between LAM forecasts up to 48 hours in advance, and raingauge-measured hourly areal rainfall.

## THE RIVER RENO REAL-TIME FLOOD FORECASTING SYSTEM

The second use of the phase-space forecast relates to the problem of extending the lead time of rainfall and flood forecasts, which has become of interest to meteorologists dealing with the recently developed Limited Area Models of the atmosphere (LAMs). Unfortunately, for several reasons these models have not yet reached the accuracy needed for issuing reliable flood forecasts. It was therefore necessary to find a technique that allowed for the use of the LAM rainfall forecasts in a statistical framework in order to take into account their uncertainty. The following approach, set up under the EU-funded TELFLOOD Project (Todini & Cerlini 1999), uses a Bayesian combination of a set of 20 precipitation traces, conditional on the latest ground measurements generated with a phase-space approach similar to the one described above, with the LAM forecasts, which are taken as biased and noise-corrupted measurements of future precipitation.

The procedure was tested on nine real flood events at Casalecchio, listed in Table 1, which were chosen to represent various rain and flood typologies.

The LAM forecasts from 1 to 48 hours in advance were compared for hourly areal rainfall over the River Reno catchment at Casalecchio (1,051 $km^2$), with the corresponding rainfall measured by a dense raingauge network. The results in terms of bias and standard deviation are shown in Figure 5.

Because the forecast is issued at midnight, from Figure 5 a diurnal error pattern can be detected showing problems in the LAM parametrisation. The figure also shows that there is no substantial increase of errors with the lead time. Nevertheless, both the bias and the standard deviation are quite large, varying from 0.7 to 1.5 mm. Given the size of the catchment, 1 mm per hour corresponds to approximately 300 $m^3/s$, which means that $\pm 2\sigma$ may approximately lead to errors of the order of magnitude of 600 $m^3/s$ in a catchment where the one in 100 years flood is being estimated as 1,900 $m^3/s$.

To improve the forecast a phase-space model of rain was developed on the same lines described above with $n_l = M = f = 12$ (hours) and $n_c = 20$ under the assumption of a Normal distribution of forecasting errors. The results of this model, as bias and standard deviation of 1 hour of areal precipitation forecast compared with actual rain-gauge measurements, are given in Figure 6. It is noticeable that, although of the same order of magnitude of the LAM errors, the standard deviation and the bias of the phase-space model errors increase with the lead time.

Following the Bayesian approach described in Berger (1980) under the assumption of Gaussian errors, the

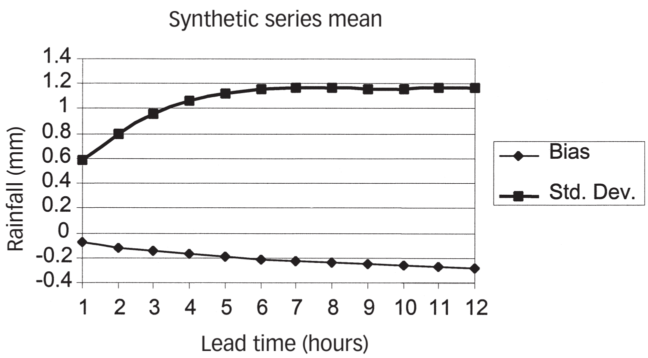


**Figure 6** | Bias and standard deviation of the phase-space areal precipitation forecast as compared with actual rain-gauge measurements.

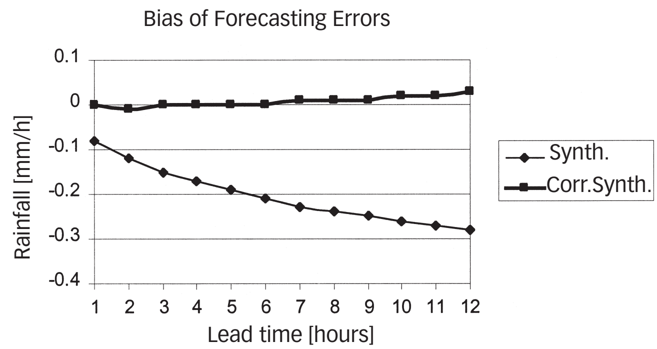


**Figure 7** | Bias of forecasting errors for the synthetic series and the ones corrected using the LAM.

following correction scheme typical of the Kalman filter can be found (Gelb 1974). If at a future time $t+k\tau$, a measure $z_{t+k\tau}$ of the state vector $x_{t+k\tau}$ is available, although affected by an error $\varepsilon_{t+k\tau}$, with bias $\mu_{\varepsilon_{t+k\tau}}$ and variance $\sigma^2_{\varepsilon_{t+k\tau}}$, which can be expressed in scalar form as:

$$z_{t+k\tau} = x_{t+k\tau} + \varepsilon_{t+k\tau} \tag{14}$$

and an *a priori* estimate of $x_{t+k\tau}$, indicated by $\hat{x}_{t+k\tau\,|\,t}$, is also available at time $t+k\tau$ ($\hat{x}_{t+k\tau\,|\,t}$ represents each individual *a priori* series conditionally forecasted by means of the phase-space technique), one can obtain an unbiased and minimum variance *a posteriori* estimate of the state variable, namely $\hat{x}_{t+k\tau\,|\,t+k\tau}$, using the following scalar relationship:

$$\hat{x}_{t+k\tau|t+k\tau} = \hat{x}_{t+k\tau|t} + g_{t+k\tau} v_{t+k\tau}\ , \tag{15}$$

where $v_{t+k\tau}$ represents the innovation defined as:

$$v_{t+k\tau} = z_{t+k\tau} - \hat{x}_{t+k\tau|t} - \mu_{\varepsilon_{t+k\tau}} \tag{16}$$

and $g_{t+k\tau}$ is the gain given by the following expression:

$$g_{t+k\tau} = \frac{\sigma^2_{\hat{x}_{t+k\tau|t}}}{\sigma^2_{\hat{x}_{t+k\tau|t}} + \sigma^2_{\varepsilon_{t+k\tau}}}\ . \tag{17}$$

Replacing in Equation (15) expressions given by Equations (16) and (17) one can obtain the *a posteriori* estimate $\hat{x}_{t+k\tau\,|\,t+k\tau}$ as a function of the *a priori* forecast $\hat{x}_{t+k\tau\,|\,t}$ and of the meteorological forecast $z_{t+k\tau}$.

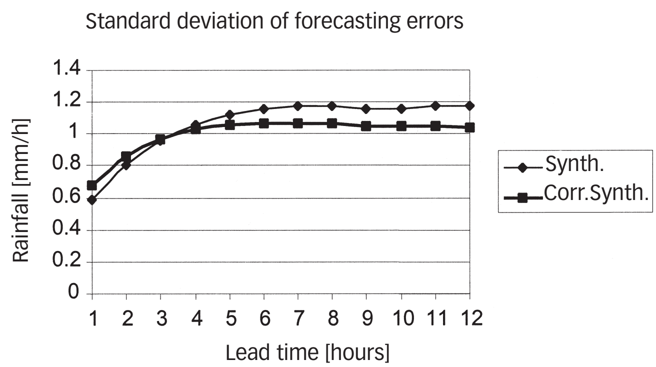


**Figure 8** | Standard deviation of forecasting errors for the synthetic series and the ones corrected using the LAM.

The results of the experiment on the River Reno data are quite encouraging. Figure 7 shows that the bias that was present in the phase-space generated precipitation traces has now been totally eliminated, whereas the standard deviation of errors is reduced at lags longer than 3 hours (see Figure 8). No improvement is obtained at lags up to 3 hours owing to the relatively better quality of the *a priori* phase-space forecasts on a short-term basis.

## CONCLUSIONS

It is shown that phase-space modelling can be a useful tool not only for providing actual forecasts, but also for generating ensembles of conditional scenarios describing the *a priori* uncertainty. This type of information is essential in any decision-making process for the estimation of the expected value of some objective function.

The same description of the *a priori* probability distribution can be combined with 'pseudo measurements' of future events independently generated by means of available physically based models, to produce *a posteriori* distributions of forecasting errors which will be unbiased and of smaller variance.

From a philosophical standpoint the interpretation of the physically based model as a noise-corrupted measurement may seem arbitrary, but if one thinks carefully the analogy is pertinent and useful in order to account for additional uncertain information.

It is intended to proceed by improving the phase-space model forecasts at the same time as meteorologists improve the rainfall parametrisation of LAMs, so as to increase the reliability of the *a posteriori* scenarios and consequently improve the decision-making process.

## REFERENCES

Alarcon, L. & Marks D. H. 1979 A stochastic dynamic programming model for the operation of the High Aswan Dam. *Rep. TR. 246.* Ralph M. Pearson Laboratory For Water Resources And Hydrodynamics, Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Berger, J. O. 1980 *Statistical Decision Theory. Foundations, Concepts, and Methods*. Springer-Verlag, New York.

Broomhead, D. S. & Lowe, D. 1988 Multivariable functional interpolation and adaptive networks. *J. Complete Systems* **2**, 321–355.

Casdagli, M. 1989 Non-linear prediction of chaotic time series. *Physica* D **35**, 335–356.

Gelb, A. (editor) 1974 *Applied Optimal Estimation*. MIT Press, Cambridge, Massachusetts.

Georgakakos, A. P. & Marks, D. H. 1987 A new method for the real-time operation of reservoir systems. *Wat. Resour. Res.* **23**(7), 1376–1390.

Georgakakos, A. P. 1989 Extended linear quadratic Gaussian control: further extensions. *Wat. Resour. Res.* **25** (2), 191–201.

Packard, N. H., Crutchfield, J. P., Farmer, J. D. & Shaw, R. S. 1980 Geometry from a time series. *Phys. Rev. Lett.* **45** (9), 712.

Smith, L. A. 1992 Identification and prediction of low dimensional dynamics. *Physica* D **58**, 50–76.

Takens, F. 1981 Detecting strange attractors in turbulence. *Lect. Notes Math* no. 898. (390 pages.) Springer Verlag, Berlin.

Todini, E. 1990 La valutazione dei benefici derivanti dalla previsione operativa dei deflussi futuri: uno schema semplificato applicato al caso del Nilo. In *Proceedings of the XXII Conference in Hydraulics and Hydraulic Construction, Cosenza*. (In Italian.)

Todini, E. & Bongioannini Cerlini, P. 1999 *TELFLOOD*. Technical Report, DISTGA, Università di Bologna.

Yakowitz, S. 1987 Nearest neighbour methods for time series analysis. *J. Time Series Analysis* **8** (2), 235–247.

Yakowitz, S. & Karlsson, M. 1987 Nearest neighbour methods for time series with application to rainfall-runoff prediction. In *Stochastic Hydrology* (ed. J. B. MacNeil & G. H. Humphries), pp. 149–160. Reidel, Hingham, Massachusetts.