24 PUB as Robust State Estimation

JOHN P. NORTON

INTRODUCTION

This chapter takes a look at prediction in ungauged basins (PUB) from the perspective of state estimation. There are several potential advantages in stepping back and viewing PUB in the first instance as a state-estimation problem. State estimation is based on versatile state-space models and, in its original form as Kalman filtering (Jazwinski, 1970; Gelb, 1974; Grewal & Andrews, 1993), its aim, to make the estimation errors orthogonal to (i.e. uncorrelated with) the estimates, makes sense in most circumstances[1]. The versatility of the model form and the broad appeal of the aim make the Kalman filter widely applicable, and it has succeeded in fields very different from aerospace engineering, its original home. In 40 years of diverse applications, a good understanding has grown up of the difficulties facing it and much effort has gone into techniques for dealing with them. Over the past decade, the potential scope of state estimation has been greatly extended by progress in two areas: optimization-based robust state estimation for use when the model is subject to uncertainties, clearly relevant to PUB, and Monte Carlo techniques to deal with nonlinear systems and arbitrary error distributions. The exercise of trying to put PUB into the framework of state estimation should throw light on its particular difficulties. More specifically, comparison with established state-estimation problem formulations should help focus on the choice of model structure and performance criterion for prediction.

The connection between PUB and state estimation is that they both concern the internal variables of a system. Flow is predicted from rainfall and other climatic inputs by use of some sort of model of the basin dynamics. The model describes, implicitly or explicitly, how the inputs affect variables (storages and flows) within the basin, including the flow to be predicted. The state-estimation problem is precisely to estimate the current values of the essential internal variables of a system, and thence the outputs. The test of whether the internal variables comprising the state are essential is that the future state, and hence outputs, can be predicted solely from knowledge of the present state and the future forcing inputs. In other words, at any point in time the

[1] The Kalman filter can also be interpreted in several less general ways, sometimes requiring restrictive assumptions. For instance, it is often quoted as producing the maximum-likelihood estimate, assuming that the initial state error, unknown forcing and observation error are all Gaussian. As well as being unnecessarily specialized, this interpretation has the unappealing aspect that maximum-likelihood estimates are not generally unbiased in finite samples. Kalman's original derivation, making the errors orthogonal to the estimates, makes no such assumptions. Nor does the interpretation that the filter produces the minimum-covariance estimate among all linear, unbiased estimates. This last interpretation implies that the estimates minimize (among linear, unbiased estimates) the variance of any linear function of state, including any individual state variable.

state variables are the means by which the history of the system affects its future behaviour. Starting from the current state, prediction is no more than substitution of the future inputs into the state-space model.

Of course, prediction could make use of an input–output model (e.g. a unit hydrograph) without reference to state variables. The advantages of bringing in state variables are that:

(a) differences between the gauged and ungauged basins may be easier to assess in terms of their state variables than their outputs alone. The reason for thinking so is that state-space models disaggregate input–output relations into contributions which have some physical significance, and which should be easier than the overall input–output relations to relate to known physical differences between basins;
(b) when ungauged flow is to be predicted from an actual climatic sequence, estimates of the state of a gauged basin, using flow observations, may be combined with known differences in climatic inputs and state behaviour between the basins to give estimates of the ungauged basin's state and hence output. How the state of the ungauged basin might be linked to state estimates in a gauged basin is discussed here.

State estimators assume that the model configuration is as in Fig. 24.1. They produce an updated state estimate each time a new observation of one or more outputs, subject to error, is received. The updated estimate is the result of correcting the predicted state value according to how well it predicted the newly observed output. The size of the correction depends on the relative uncertainties in the predicted state and the observation. State estimators require:

– a model of the dynamics, giving state from inputs: the state equation;
– a model of the instantaneous state-to-output relation: the observation equation;
– an estimate of the initial state and its uncertainty;
– known inputs, if any;
– specifications of the uncertainties in the imprecisely known forcing inputs (process noise) and observed outputs (observation noise); and
– an optimality objective.

The state-estimation approach is dictated by the form of the uncertainty specification. In the Kalman filter, it is in the form of error covariances, assuming that the process noise, observation noise and initial state error are zero-mean and white (serially uncorrelated). A well established but less widely used alternative is to specify bounds

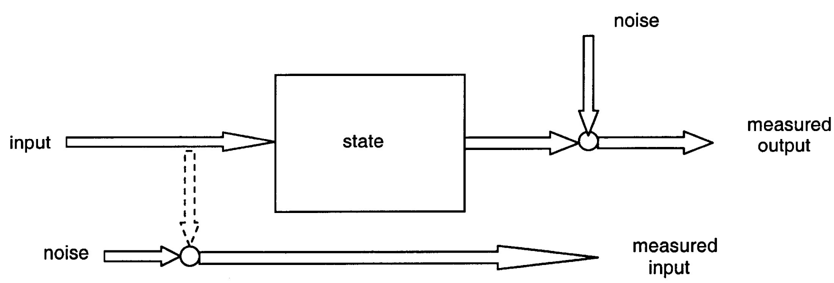

Fig. 24.1 *Configuration of state-space model.*

on the instantaneous values of the noises and initial state error, and compute bounds on the state (Schweppe, 1973; Kurzhanski & Valyi, 1994; Norton; 1996; Maksarov & Norton, 1996). A third possibility is to specify the probability density functions (PDFs) of the noises and initial state error and find the PDF of state by Bayesian estimation (Gordon *et al.*, 1993; Doucet *et al*, 2001); clearly this is much more ambitious.

All the items listed above need examination in the context of PUB, but there is one glaring discrepancy: standard state estimators assume that the model itself is accurate, while PUB faces substantial model uncertainty. The problem posed by PUB is one of *robust* state estimation. Robust state estimation takes account of error in the model and in the uncertainty specification. It aims to provide state estimates which remain usable over a given range of model uncertainty. The approaches to robust state estimation differ in how this aim is formalized and how the uncertainty is prescribed; some of the possibilities will be discussed in the following section. Whatever approach is adopted, robust state estimation is computationally demanding and only in the past few years has it begun to be a practical proposition for realistic problems. Techniques for robust state estimation, especially their assumptions and limitations, are reviewed. Then, the next section tries to fit the PUB problem into the robust-state-estimation framework, and suggests some state-estimation strategies for PUB. Conclusions follow in the final section.

APPROACHES TO ROBUST STATE ESTIMATION

Before getting into details, it is worth noting that although some model uncertainty is always present in state-estimation problems, it can often be taken into account without explicitly representing the model parameters as uncertain. The simplest tactic is just to inflate the process and observation noise covariances to cover error originating in the state equation and observation equation, respectively. Common practice is to tune the covariances by trial and error to optimize estimation performance. The main effect is to alter the balance between the predicted state value from the previous update and the new information in the latest observation. Raising the process noise covariance reduces trust in the predicted value, while raising the observation noise covariance reduces trust in the latest observation. Covariance tuning is hard to defend, since it implies, usually wrongly, that the effects of model errors are indistinguishable from those of unstructured (white, zero-mean) process and observation noise. Even so, it often works quite well so long as the model accounts for the dominant behaviour of the system.

A more refined strategy allows for structure in the modelling error. The error is represented as being produced by passing white noise through a linear dynamical submodel, adjoined to the main model. The submodel parameters are added to the unknowns being estimated, or suitable values are guessed. Again trial and error may help. Application of this idea to PUB can be imagined, developing an auxiliary model for prediction-error structure to fit the observed errors in gauged basins, then using it, perhaps modified on the basis of basin attributes, for the ungauged basin.

If neither of these tricks to allow for model uncertainty is enough, a robust state estimator is needed. Robust state estimation has grown up in the control-engineering community so, not surprisingly, the main approaches reflect popular control-design paradigms. Four options will be discussed: H_∞, bounding, guaranteed-cost and Monte Carlo Bayesian.

H_∞ robust state estimation

Historically, much control-system design has been based on frequency responses. The two main requirements of control are accurate following of demanded output changes and good regulation of the output in the face of disturbances. Close following of the demanded output requires the closed-loop gain from reference input to controlled output to be near unity over a sufficient bandwidth to ensure rapid enough response. At the same time, disturbances and measurement noise have to be rejected by making the gain from where they enter to the output small over as much as possible of their bandwidth. In H_∞ robust control (Green & Limebeer, 1995; Zhou *et al.*, 1996), this is done by minimizing, or keeping below a prescribed value, the H_∞ norm from noise to output (the highest value, over the frequency range, of the largest singular value of the transfer function matrix), over a given range of uncertainty in the plant transfer function. The plant uncertainty may be unstructured (e.g. as a maximum deviation of the frequency response from its nominal value) or structured (specified ranges for a selected set of model parameters). There is a trade-off between performance, poor when the control is required to compromise between all sorts of unlikely behaviour within the unstructured uncertainty bounds, and computing load, which tends to be high if highly structured uncertainty is specified.

H_∞ robust state estimation is a close cousin of H_∞ robust control (Nagpal & Khargonekar, 1991; Xie *et al.*, 1991; Li & Fu, 1997; Petersen & Savkin, 1999). The aim is to keep the maximum gain from model uncertainty to state error below a given level. The gain is the ratio of the "noise" and error energies, so the model uncertainties are prescribed by bounds on equivalent process and observation noise energy, rather than covariances or any other probabilistic properties. The resulting estimator resembles a Kalman filter and is easy to implement. However, it may well perform poorly, as it has to provide robustness against modelling errors which are allowed by the noise-energy bounds but do not arise in practice. Moreover, filter divergence is possible and a condition for the filter to remain stable has to be checked at each update.

Robust state bounding

State bounding (Schweppe, 1973; Kurzhanski & Valyi, 1994; Norton; 1996; Maksarov & Norton, 1996), sometimes called set-membership estimation, is a natural basis for worst-case control design as it gives the extremes of state consistent with the observations to date and specified bounds on the forcing and observation noise. It too has been extended to include model uncertainties (Savkin & Petersen, 1995, 1998; Petersen & Savkin, 1999; El Ghaoui & Calafiore, 2001; Polak *et al*., 2004), described by bounds on, e.g. the sum of a weighted quadratic function of the initial state error and the sum, over time, of weighted quadratic functions of the process and observation noises (Savkin & Petersen, 1995, 1998).

State bounding has some advantages. Bounds are easy and direct to interpret, and noise bounds are easier to supply, as tentative values, when too little is known for a probabilistic specification to be supplied. Bounds on state are clearly appropriate when worst-case ranges of state or future output values are of interest. Worst-case predictions of the output are readily computed, in principle, from the bounds on state, forcing and model uncertainty. Sensitivity assessment is also straightforward, simply loosening or tightening bounds on the uncertainty or on the process or observation noise to see the effects on the state bounds or output-prediction bounds.

Robust state bounding has drawbacks as well. A limitation is that the bounds on state at any instant are usually approximated by an ellipsoid, which is analytically convenient and economical to describe but, being convex, may not fit non-convex bounds well; non-convex bounds do arise in robust state bounding. Another weakness is that approximation of the state bounds at every update may cumulatively make the bounds over-conservative, too loose to be useful. Finally, there are technical conditions to be satisfied if the filter updates are not to break down.

Guaranteed-cost/minimum-covariance state estimation

Guaranteed-cost state estimation (Bolzern *et al.*, 1995; Theodor & Shaked, 1996; Sayed, 2001) aims to produce estimates with error variance upper-bounded over a range of model uncertainty specified by bounds on parameters and noise covariances. Filter stability is again an issue, with design parameters to be tuned to ensure stability. An appealing procedure to avoid this difficulty, interpretable as regularization, is proposed in Sayed (2001). It stems from an interpretation of the Kalman as minimizing a cost function which sums quadratic terms, each weighted by the inverse of the corresponding covariance (Bryson & Ho, 1975). The terms respectively represent the discrepancy between the present state and its value predicted from the previous estimate, the output prediction error and the inferred value of the forcing. Robust state estimation also can be treated as minimization of the same quadratic function, but now of its maximum value over the specified model uncertainty. The technique involves optimizing a scalar design parameter, but in practice the value found by a fixed simple rule seems satisfactory.

Monte Carlo Bayes state estimation: "particle filtering"

Kalman filtering relies on linearity of the model to make updating of the state-error covariance possible. A nonlinear model may sometimes be successfully linearized for state estimation, but the linearization error may cause the computed state-error covariance to diverge badly from the actual covariance. As well as giving a misleading indication of the estimates' accuracy, the diverging covariance may even cause the estimator to break down. Another limitation of covariance-based state estimators is that they do not register important factors such as severe skew in the errors, ambiguity in the estimates, or heavy- or light-tailed error distributions, which would be apparent in the probability density function (PDF) of the state. These weaknesses have led to development of Bayesian estimators. They compute the entire state PDF at each update, and derive any desired state or state-dependent estimate from it. The best-known example is the conditional mean, which minimizes the mean square error in any linear function of the state variables. The state PDF is updated from $p(\mathbf{x}_{t-1} \mid \mathbf{y}_{t-1},\mathbf{y}_{t-2},\ldots,\mathbf{y}_1)$ to $p(\mathbf{x}_t \mid \mathbf{y}_t, \mathbf{y}_{t-1},\ldots,\mathbf{y}_1)$ when observation $\mathbf{y}_t$ is processed, by Bayes' rule:

$$p(\mathbf{x}_t \mid \mathbf{y}_t, \mathbf{y}_{t-1},\ldots,\mathbf{y}_1) = \frac{p(\mathbf{y}_t \mid \mathbf{x}_t)\, p(\mathbf{x}_t \mid \mathbf{y}_{t-1},\ldots,\mathbf{y}_1)}{p(\mathbf{y}_t \mid \mathbf{y}_{t-1},\ldots,\mathbf{y}_1)} \quad (1)$$

where:

$$p(\mathbf{x}_t \mid \mathbf{y}_{t-1},\ldots,\mathbf{y}_1) = \int_{\mathbf{x}_{t-1}} p(\mathbf{x}_t \mid \mathbf{x}_{t-1})\, p(\mathbf{x}_{t-1} \mid \mathbf{y}_{t-1},\ldots,\mathbf{y}_1)\,\mathrm{d}\mathbf{x}_{t-1} \quad (2)$$

Here $p(\mathbf{x}_t \mid \mathbf{x}_{t-1})$ is found by using the state equation to write the PDF of the process noise in terms of $\mathbf{x}_t$ and $\mathbf{x}_{t-1}$, and $p(\mathbf{y}_t \mid \mathbf{x}_t)$ by using the observation equation to write the observation noise PDF in terms of $\mathbf{y}_t$ and $\mathbf{x}_t$. The denominator on the right-hand side of equation (1) does not depend on $\mathbf{x}_t$ and just serves to scale the state PDF to unity area. The big attraction of equations (1) and (2) is that they do not care whether the state and observation equation are linear, and are not confined to any particular form for the PDFs. Algebraic derivation of the posterior PDF of state from equation (1) is hardly ever possible, so numerical solution is necessary. However, the integration over all possible $\mathbf{x}_{t-1}$ in equation (2) imposes a heavy load. Recursive Bayesian state estimation was considered long ago (Jazwinski, 1970) in the form of the Chapman-Kolmogorov equation, but because of the computing load remained impracticable for realistic problems.

Over the past decade, the increase in computing power has made Monte Carlo evaluation of the posterior PDF of state feasible for some non-trivial problems (Gordon *et al.*, 1993; Doucet *et al.*, 2001). The idea is simple. First take samples of $\mathbf{x}_{t-1}$ according to $p(\mathbf{x}_{t-1} \mid \mathbf{y}_{t-1},\mathbf{y}_{t-2},\ldots,\mathbf{y}_1)$. For each sample, compute $\mathbf{x}_t$ by inserting $\mathbf{x}_{t-1}$ and a sample of the process noise into the state equation. Next weight each sample by its likelihood $p(\mathbf{y}_t \mid \mathbf{x}_t)$, by substituting the observation $\mathbf{y}_t$ and $\mathbf{x}_t$ into the expression for the observation noise PDF obtained by rewriting the observation equation in terms of them. The weighted samples are effectively samples of the numerator of equation (1), and thus make up an unnormalized description of the posterior PDF of the state. This completes one update cycle, and the next starts by sampling of the state $\mathbf{x}_t$ according to that PDF. The scheme, often given the foolish name “particle filtering”, has been widely used in target tracking problems, where the state variables are the position and velocity components of the target and where either the state equation or the observation equation is nonlinear. It has had considerable refinement (Doucet *et al.*, 2001) to maximize computational economy and to avoid problems which arise in naïve implementations, notably over-concentration of the re-sampled state onto a small number of values. An immediate limitation of the Monte Carlo Bayesian approach for *robust* state estimation is the need to sample, or include in a likelihood calculation, every uncertain item. If there are a number of uncertain model parameters, it may well not be feasible to propagate a sample set large enough to be representative of all the uncertainties simultaneously.

HOW DOES ROBUST STATE ESTIMATION FIT PREDICTION IN UNGAUGED BASINS?

For PUB, the known forcing consists of rainfall and temperature and the unknown forcing consists of groundwater flows, unmeasured climatic variables such as humidity and wind speed, and the errors in the assumed values of the known inputs. The outputs of interest are flows at specified locations, and the state variables are storages. In principle there are an infinite number of state and forcing variables, since the values of distributed variables are required at an infinite number of points to describe them fully. In practice, it is well known that very low-order models suffice for purposes where a good deal of spatial and temporal averaging is permissible. For example, an IHACRES-type model (Jakeman *et al.*, 1990; Jakeman & Hornberger, 1993) relating spatially and temporally aggregated (e.g. daily) rainfall to flow at one location, typically has only three state variables indicating soil moisture and the storages of the quick and slow flow components. Higher spatial or temporal resolution would necessitate more state variables. The fact that experience shows that useful hydrological models may have very few state variables is highly significant for robust state estimation, where computing load typically rises with the cube, or even higher power, of the number of state variables.

The most obvious non-standard feature of PUB, viewed as a state-estimation problem, is that it asks for prediction of an output not directly observed. This could be attempted in several ways, discussed in the three following sections. All of them assume that an uncertain model for the ungauged basin has been identified by reference to gauged basins with topographical, geological, land-use and climatic similarities. The model parameter estimates are accompanied by either probabilistic or deterministic measures of uncertainty (covariance, joint PDF or bounds).

PUB by running a state-space model for the ungauged basin alone

As there are no observations directly related to the ungauged output, the model is run in simulation mode, driven by estimated climatic and possibly groundwater inputs. Even this apparently simple situation raises some questions:

– Which of the forcing inputs should be treated as uncertain? If they are, are their estimates and uncertainties to be provided by a subsidiary state estimator?
– Should the uncertain model parameters be treated as constant? The answer is far from clear. Consider first probabilistically estimated parameters (and their uncertainties). Their mean values are valid for average behaviour over the records on which they were based, so there is no problem in using the mean values to give "best" predictions judged on their average quality over a similar range of conditions. However, the best parameter values of an approximate model in the short term have to vary as conditions change, since the best local approximation by a simple model (e.g. linear, low-order) of a less simple system (e.g. nonlinear or higher-order) depends on the operating point. This is one source of parameter uncertainty. The conclusion is that fixed parameter values cannot give the best predictions as measured in the short term. One can conceive of an auxiliary model describing just how the parameter values vary with the conditions, but that is the same thing as making the model nonlinear and more elaborate. Condition-dependent parameter variation could be *conservatively* allowed for by computing the range of predicted output values covered by all parameter histories allowed by the computed parameter uncertainty (i.e. within the computed parameter bounds, or in a confidence region based on the estimated parameter covariance). Such histories will include a wide range of unrealistic parameter variations and will thus give very pessimistic results. Accepting that the parameters have to be regarded as constant, there remains the problem of computing the predicted-output uncertainty due to the parameter (as well as input) uncertainty. Even in a model linear in its parameters and variables, the presence of products of uncertain parameters and uncertain variables makes the prediction of uncertain outputs by any means other than Monte Carlo simulation difficult. This applies equally to uncertainties described by covariances and bounds on vectors of parameters or variables.
– Can the quality of prediction be improved by indirect information about the flow output? Examples might be the condition of vegetative cover or farm dams. The most obvious is flow elsewhere; this is considered further in the next two sections.

Robust state estimation is relevant here only for the time-update part, where standard state estimators do not consider model-parameter uncertainty but robust estimators do.

As outlined so far, this scheme could employ an input–output model rather than a state–space model. The latter has the advantage that monitoring of the state estimates may allow unrealistic behaviour to be detected, or the model to be adjusted, e.g. as soil moisture gets very high or very low.

PUB by running state-space models for gauged and ungauged basins in parallel

The idea here is to benefit from comparison between predicted and observed flow in a gauged basin. As each flow observation in the gauged basin is processed, the estimate is corrected by the state estimator according to the output-prediction error, and is thereby influenced by the effects of model deficiencies and unmeasured inputs as well as the error in the previous state estimate and observation noise. Some of the correction might reasonably be expected to apply to the ungauged basin too. For example, if the basins are adjacent the effects of an untypical rainfall event might be similar, and if they have similar topography, the modelling errors due to omitting nonlinearity might be partly in common. It has recently been pointed out (Reichert & Borsuk, 2005) that the uncertainty in a difference between predictions that share some sources of error may be much less than the sum of the uncertainties in the individual predictions.

This scheme relies on some science-based and/or empirical modelling of the shared characteristics of the output-prediction errors in the gauged and ungauged basins. At its crudest, it might be no more than choice of a factor to scale the state corrections in the gauged basin. Apart from the state corrections borrowed from the gauged basin, the ungauged basin model is run in simulation mode, as in the previous scheme.

PUB by running a combined model for a gauged and an ungauged basin

A more coordinated scheme is to regard the gauged and ungauged basins as together forming a single system, with a single state-space model. One of the system outputs is observed but a different one is to be predicted. In the absence of direct interaction between the two subsets of state variables, one per basin, the amount of information about the ungauged output conveyed by the gauged output depends on the correlation between the estimates of the two subsets. For a linear model, a single observed output y_t at sampling instant t is related to the state $\mathbf{x}_t$ by:

$$y_t = \mathbf{h}_t^T \mathbf{x}_t + v_t \tag{3}$$

where superscript T denotes transpose, v_t is the zero-mean observation noise, of variance σ_v^2 , and uncertainty in $\mathbf{h}_t$ is ignored for simplicity. The Kalman filter correction equation is:

$$\hat{\mathbf{x}}_{t|t} = \hat{\mathbf{x}}_{t|t-1} + \frac{\mathbf{P}_{t|t}\mathbf{h}_t}{\sigma_v^2}(y_t - \mathbf{h}_t^T \hat{\mathbf{x}}_{t|t-1}) \tag{4}$$

where $\mathbf{P}_{t|t}$ is the estimated covariance of the state estimate $\hat{\mathbf{x}}_{t|t}$ and $\hat{\mathbf{x}}_{t|t-1}$ is the state estimate predicted from instant t – 1 using the state equation. The correction to state variable i induced by the prediction error in y_t is proportional to (row i of $\mathbf{P}_{t|t}$)$\mathbf{h}_t$. Thus, if a flow which is also a state variable is observed, the correction to state variable i depends only on that state variable, say state variable j, and the correlation p_{ij} between them. With no direct physical interaction (e.g. through groundwater or because they are at points on the same stream), they are mutually correlated as a result of being affected by correlated inputs (true for adjacent basins) and persistence of any initial correlation. One mechanism for ensuring that mutual correlation is properly registered by the model would be to constrain suitable elements of the state-estimate covariance matrix $\mathbf{P}_{\cdot|\cdot}$, perhaps even keeping it constant. This would have to be done on the basis of judgement of the strength of similarity between the conditions in the two basins. Alternatively, or additionally, the climatic inputs might be represented as having a common component and a difference component.

Input prediction

Another unusual feature is that PUB may involve long-term prediction of flow in unknown future climatic conditions, i.e. with the need to predict uncertain input. By contrast, in most state-estimation applications either the current state is to be estimated (as in state-feedback control systems) or the past state (as in fault and change detection). Where prediction is required, it is often as part of a real-time system, with longer-term predictions given less weight and soon superseded (as in model predictive control).

The model almost universally employed for unknown input in probabilistic state estimation is vector white noise, i.e. a collection of wide-sense stationary sequences with no autocorrelation, usually mutually uncorrelated. Where the sequences are in fact structured, the structure might be modelled by passing the white sequences through linear, constant-coefficient filters. This type of model is clearly inadequate for climatic variables, so a more refined and specialized stochastic simulation model would have to be used. Some state-estimation applications have generated more flexible models; for instance, Markov transition matrix models have been used for tracking manoeuvring targets, and models switching between modes are common in that application area.

For state bounding, the commonest vector input model is a collection of sequences bounded instantaneously and either independently (i.e. with box bounds on the vector) or jointly (usually by an ellipsoid). For technical convenience, additively combined bounds on forcing and the state transition matrix, and on observation noise and the observation matrix, have been used (Polak *et al.*, 2004). These models have great flexibility as they do not prescribe any distributional properties (apart from the support) or time structure. Specifically they do not require whiteness or even stationarity, except that the bounds are usually taken to be constant, for simplicity but not essentially. Conversely, they embrace a vast range of behaviour, much of it unrealistic. The same comments go for bounded observation noise models. A crucial open question is therefore: What additional constraints should be applied to models of bounded noise to enforce realism in the context of PUB?

Estimator performance criterion for PUB

Robust state estimators use the state error energy or covariance, the tightness of state bounds, or the accuracy, resolution and coverage of the state PDF (in Monte Carlo Bayesian estimation) as criteria. To see how well PUB matches a robust state estimation technique, its performance criterion must be named. Prediction for management purposes might take the accuracy of some integrated single quantity, such as total water availability above a nominated flow level over a given period, as its criterion. An integrated output quantity poses no difficulty, being handled simply by making it an extra state variable with, e.g. *rate of change of measure = flow above level* as the corresponding extra scalar state equation. Similarly, a criterion such as the proportion of the time for which a threshold is exceeded could be handled by an extra state variable that increments by one at the end of each sampling interval in which the threshold is exceeded. The rate of change switches between zero and one as the threshold is passed. However, a sharp nonlinearity as imposed by a flow threshold makes it difficult to relate, analytically, uncertainty in the nonlinearity input to uncertainty in the nonlinearity output.

Any more synoptic criterion might or might not make analysis difficult. For instance, if the primary aim of management is to avoid unacceptable extremes of flow behaviour, the worst-case-optimal nature of most robust state-estimation schemes fits well. As another example, the joint acceptability of the behaviour of a number of state variables (e.g. flows at a number of locations) might be expressed as a quadratic function of their deviations from desirable values. A case where choice of an estimation performance criterion could be difficult is where multiple considerations affect the choice between two or more alternative flow scenarios.

HOW DO EXISTING STATE ESTIMATORS MATCH UP TO THE NEEDS OF PUB?

Some broad conclusions can be drawn from the discussion above. Existing robust state estimators can handle:

- model error lumped with output error and/or uncertain forcing (by any robust estimation method); or
- structured model uncertainty (by bounding but conservatively, by H_∞ robust filtering or particle filters at high computational cost);
- uncertain uncertainty specifications (by guaranteed-cost estimators, particle filters or traditional extended simulation trials).

That said, the prescriptions of uncertainty in existing robust state estimators are largely dictated by computational convenience and particularly the need to make optimization convex (Polak *et al.*, 2004). A positive development is the widespread use of norm bounds as an alternative to hard-to-supply statistical information on uncertainties. The flexibility of Monte Carlo Bayesian state estimators is still exploitable only for models with low state dimension, but for at least some PUB purposes, low-dimension models are enough.

Further hydrological modelling, including modelling of uncertainties, and augmentation or modification of existing estimators are needed to handle:

- output-poor identification: construction of an ungauged basin model according to behaviour of its estimated input, the inputs and outputs of gauged basins, topography, land use, geology and perhaps some ungauged-basin variables not directly related to flow (such as vegetation and dam conditions);
- updating of state estimates in an ungauged basin on the basis of observations elsewhere;

and more generally:

- loosely coupled state-space models for gauged and ungauged basins, sharing some forcing and characteristics but with little or no direct interaction;

– systematic error in both the model and any indirectly flow-related measurements in the ungauged basin;
– long-term forcing by climatic variables, requiring models more specialized than the mainly white-noise-driven and/or Markov transition probability models conventional in state estimation;
– awkward estimation performance criteria, related to the use of the predictions.

References

Bolzern, P., Colaneri, P. & De Nicolao, G. (1995) Optimal design of robust predictors for linear discrete-time systems. *Systems & Control Letters* **26**, 25–31.

Bryson, A. E. & Ho, Y. -C. (1975) *Applied Optimal Control: Optimization, Estimation, and Control*. Taylor & Francis, London, UK.

Doucet, A., de Freitas, J. F. G. & Gordon, N. J. (eds) (2001) *Sequential Monte Carlo Methods in Practice.* Springer-Verlag, New York, USA.

El Ghaoui, L. & Calafiore, G. (2001) Robust filtering for discrete-time systems with bounded noise and parametric uncertainty. *IEEE Trans. on Autom. Control* **46**(7),1084–1089.

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

Green, M. & Limebeer, D. J. N. (1995) *Linear Robust Control*. Prentice-Hall, Englewood Cliffs, New Jersey, USA.

Grewal, M. S. & Andrews, A. P. (1993) *Kalman Filtering: Theory and Practice*. Prentice-Hall, Englewood Cliffs, New Jersey, USA.

Gordon, N. J., Salmond, D. J. & Smith, A. F. M. (1993) Novel approach to non-linear/non-Gaussian Bayesian state estimation. *IEEE Proceedings-F* **140**, 107–113.

Jakeman, A. J. & Hornberger, G. M. (1993) How much complexity is warranted in a rainfall-runoff model? *Water Resour. Res.* **29**, 2637–2649.

Jakeman, A. J., Littlewood, I. G. & Whitehead, P. G. (1990) Computation of the instantaneous unit hydrograph and identifiable component flows with application to two small upland catchments *J. Hydrol.* **11**(7), 275–300.

Jazwinski, A. (1970) *Stochastic Processes and Filtering Theory*. Academic Press, New York, USA.

Kurzhanski, A. & Valyi, I. (1994) *Ellipsoidal Calculus for Estimation and Control.* Birkhäuser, Boston, Massachusetts, USA.

Li, H. & Fu, M. (1997) A linear matrix inequality approach to robust filtering. *IEEE Trans. on Signal Processing* **45**(9), 2338–2350.

Maksarov, D. & Norton, J. P. (1996) State bounding with ellipsoidal set description of the uncertainty. *Int. J. Control* **65**(5), 847–866.

Nagpal, K. M. & Khargonekar, P. P. (1991) Filtering and smoothing in an H_∞ setting. *IEEE Trans. on Autom. Control* **36**, 152–166.

Norton, J. P. (1996) Roles for deterministic bounding in environmental modelling. *Ecol. Modelling* **6**, 157–161.

Petersen, I. R. & Savkin, A. V. (1999) *Robust Kalman Filtering for Signals and Systems with Large Uncertainties.* Birkhäuser, Boston, Massachusetts, USA.

Polak, B. T., Nazin, S. A., Durieu, C. & Walter, E. (2004) Ellipsoidal parameter or state estimation under uncertainty. *Automatica* **40**, 1171–1179.

Reichert, P. & Borsuk, M. E. (2005) Does high forecast uncertainty preclude effective decision support? *Environ. Modelling & Software* **20**(8), 991–1001

Savkin, A. V. & Petersen, I. R. (1995) Recursive state estimation for uncertain systems with an integral quadratic constraint. *IEEE Trans. on Autom. Control* **40**(6), 1080–1083.

Savkin, A. V. & Petersen, I. R. (1998) Robust state estimation and model validation for discrete-time uncertain systems with a deterministic description of noise and uncertainty. *Automatica* **34**(2), 271–274.

Sayed, A. H. (2001) A framework for state-space estimation with uncertain models. *IEEE Trans. on Autom. Control* **46**(7), 998–1013.

Schweppe, F. C. (1973) *Uncertain Dynamic Systems*. Prentice-Hall, Englewood Cliffs, New Jersey, USA.

Xie, L., de Souza, C. & Fu, M. (1991) H_∞ estimation for discrete-time linear uncertain systems. *Int. J. Robust & Nonlinear Control* **1**, 111–123.

Theodor, Y. & Shaked, U. (1996) Robust discrete-time minimum-variance filtering. *IEEE Trans. on Signal Processing* **44**(2), 181–189.

Zhou, K., Doyle, J. C. & Glover, K. (1996) *Robust and Optimal Control*. Prentice-Hall, Englewood Cliffs, New Jersey, USA.