Towards environmental models of everywhere: advances in modelling and data assimilation

Keith Beven

Lancaster University, UK

INTRODUCTION

Computers are getting ever more powerful. Local parallel PC systems are ever more affordable. Distributed parallel computers linked by GRID-type networks will be a widely available resource in the future (see: www.gridcomputing.com; www.gridforum.org; www.globus.org). It will be increasingly possible to access these resources from portable machines in the field; while increasingly intelligent measurement devices, including satellites, will feed back information in real time about an expanding range of environmental variables.

How then should such powerful data resources be used to make real advances in hydrological and environmental modelling? There are some obvious answers. We can add more detail and complexity of processes in our models. We can use finer time steps and spatial discretisations in our models. We can develop increasingly sophisticated data assimilation methodologies. We can avoid the limitations of linearisation assumptions in applying non-linear models. We can make many more model runs than previously possible to assess the uncertainty in model predictions. We can take advantage of a wider range of data sources in trying to constrain those uncertainties.

But these types of advantages will be mostly straightforward to implement. There are also some other interesting issues that will arise. These issues are being driven not only by hydrological theory and increases in computer power and data availability but also by the requirements of users to meet new statutory requirements. In particular, in Europe, the Water Framework Directive will be a major driving force in the implementation of integrating environmental modelling systems at the large catchment and national scales. The Water Framework Directive implies whole basin and interbasin management of water resources, including quality and habitat considerations, to allow sustainable use of water bodies of 'good status'. What do these factors imply about how such environmental models of everywhere should be implemented?

THE CURRENT STATUS OF ENVIRONMENTAL MODELLING AND DATA ASSIMILATION

Current environmental models range in scale from molecular to global scale. Molecular models might be used to predict the detail of geochemical interactions between different chemical species; global models are used to predict the weather, pollutant dispersion in the atmosphere and oceans, and climate change. In both cases, more and more sophisticated models are possible as a result of the use of increasingly large parallel computer systems, including the massively parallel *Earth Simulator* in Japan (Earth Simulator Centre, 2003).

Let us assume that numerical analysts and professional programmers will ensure that any of these models will be internally consistent with the equations on which they are based (which is not always so easy in non-linear dynamical equations). Consistency does not, however, imply accurate predictions, even when the model equations are a correct description of the processes (also not so easy in some of the coupled systems in complex geometries that we are interested in). In any application of such models there are two (related) issues that arise in applications and necessarily limit the accuracy of predictions. The first is scale (Beven, 1995; Blöschl and Sivapalan, 1995); the second is uniqueness of place (Beven, 2000, 2002). The scale problem has many aspects, but most result from a mismatch of measurement and modelling scales. This can limit how well we understand the systems we are trying to model and how well we can represent them in practical model applications. Environmental systems are invariably heterogeneous. Measurements provide us with only an index of that heterogeneity, integrated over some (perhaps variable) effective measurement scale. The measurement scale may be too large to properly reveal the controls on a set of processes; it may be too small to provide an adequate characterisation of the processes in a heterogeneous domain. Only very rarely is there a direct match between the scale of a measured variable and that same variable as it appears in a model.

Hydrology: Science & Practice for the 21st Century. Volume I

For example, a rainfall-runoff model requires an input of rainfall. There may be a few point measurements of rainfall available in and close to a catchment area of interest but these represent a very small scale sample of the changing rainfall fields. Radar data can help to reveal spatial patterns of rainfall, but also do not always provide accurate estimates of the inputs required by a model at the scale of discretisation used in the model. Every rainfall-runoff modeller knows of examples where it is suspected that the rainfall inputs to a model were poorly estimated. One of the very first hydrological models, the Stanford Watershed Model (Crawford and Linsley, 1966; Fleming, 1975), allowed a rainfall adjustment coefficient as one of its parameters to be calibrated for this very reason. Recent work has suggested allowing rainfall adjustment factors in every storm in calibration (Kavetski *et al.*, 2002). There is no general solution here, however. Any adjustment factors will be space and time specific, with no guarantee that they will apply outside of the calibration period.

Similar considerations apply to the integration of sub-grid process representations in heterogeneous domains to larger model grid scales. Typical examples are the representation of convective processes in atmospheric models that are important in cloud and rain formation and hence on the wider atmospheric circulation; the representation of sub-grid scale turbulence in fluid dynamics models of all types; and the integration of small scale heterogeneities in subsurface flow in soils and aquifers to the model grid scale. Where there are important nonlinearities at the small scale, then the mathematics suggests that relationships developed from small scale measurements may not be applicable at the larger scale. The mathematical form of an integral representation at the grid scale may then be scale dependent. This has been recognised in some sub-grid scale representations in atmospheric models but not generally in hydrological models (see the discussion of the use of Darcy's law at different scales in Beven, 1989; 1995).

Discharge estimates at the catchment outlet are, however, one example of where the scale of a measurement does match that of a predicted variable. There may be inaccuracies in the measurement (discharge is not often measured *directly*) and in the model predictions but the variables are commensurable. Such measurements may not give a good indication, however, of how to disaggregate process representations to smaller scales.

The second problem is uniqueness of place. The application of any model requires the specification of model parameter values in the characterisation of the particular model domain for that application. Since it is generally impossible to measure values of parameters directly at the scales required by the model, some extrapolation from other applications will be required, or the parameters will need to be 'calibrated' for the new application. It has generally proven surprisingly difficult to extrapolate model parameters from one application to another, because each place of application has its own unique characteristics and each model implementation may require different effective values of the same parameter (parameters of the same name are not necessarily commensurate in different models). In part, this may itself be due to the scale problem, in that if process representations in heterogeneous flow domains are expected to be scale dependent, so will the effective parameter values for those representations. This means that the value of a parameter for an application at one scale and in one location may not have the same meaning for another scale or another location or another model. It might also, of course, be a result of the fact that the characteristics of the new location are simply different, despite having apparently similar characteristics of geology, soil, land use, etc. There is then a fundamental uncertainty in characterisation that should be taken into account in model applications.

These scale and uniqueness problems will arise even if the model formulation that is being used is correct. They are compounded in real environmental modelling applications by the fact that our model structures are not correct. It is worth repeating again that they are wrong, and we know them to be wrong, and treating them *as if* they were correct is not necessary the best management strategy (Morton, 1993; Beven, 2002, 2004). Our aim is, of course, to improve them over time as we learn more about the controlling processes and can incorporate more of that understanding into the model formulation. This can, itself, however, lead to difficulties in that it is commonly found that incorporating more understanding requires the inclusion of more parameter values that might still be scale dependent and not directly measureable.

Similarly, the use of finer spatial and temporal discretisations (as increasing computer power allows) might also not always help in improving model predictions. It is true that finer discretisations of variability will allow small scales of heterogeneity and flow structures to be resolved in the model — but this might then require finer information on boundary conditions and parameters of that heterogeneity that might be difficult to make available.

The result is that uncertainty is generic to the application of environmental models and that it might be difficult to decide between competing model representations (both model structures and sets of effective parameter values) of environmental systems (the equifinality problem of Beven, 1993). Beven (2002) has discussed how the constraint of uncertainty and the problem of deciding between different model representations might be handled in a scientific way, effectively through data assimilation that constrains the mapping of the real system (the landscape space) into a model space. This is essentially a learning process that provides a key to how environmental models might be used in the future.

An essential part of learning is data assimilation. Data assimilation / model updating is a highly recommended strategy for real-time forecasting or nowcasting applications. It has been the subject of extensive work in atmospheric weather prediction (mostly using variational techniques, Swinbank *et al*., 1993) and flood forecasting (mostly using variants on recursive state space approaches, e.g. Young, 2002; or Bayesian methods e.g. Krzysztofowicz, 2002; or combinations of both, e.g. Romanowicz and Beven, 1998). Some recent hydrological studies have looked at Monte Carlo based ensemble Kalman filtering for non-linear models (e.g. Margulis *et al.*, 2002) and assimilation of spatially distributed datasets from remote sensing (Houser *et al.*, 1998; Hoeben and Troch, 2000). The important point here is that a variety of methodologies is available to help in formulating the learning process within a context of using new data to constrain predictive uncertainties. These techniques are not, however, without limitations and we will return to this subject later.

MOVING TOWARDS ENVIRONMENTAL MODELS OF EVERYWHERE: REQUIREMENTS

With the advent of GRID distributed computing facilities and tools, it is becoming feasible to construct flexible modelling systems for environmental management in a way that has not been possible hitherto. GRID middleware will mean that models developed by different institutions and running on different machines will be able to exchange information, including that for real-time forecasting applications involving data assimilation from distributed monitoring equipment. Users will be able to obtain visualisations of the results and test different scenarios. In addition, with computing power continuing to develop at a rapid rate, the user will be able to take advantage of new uncertainty estimation tools within decision support frameworks.

For models of everywhere, the (sometimes subtle) coupling between atmospheric forcing, catchment characteristics and response, river runoff and coastal interaction with tidally dominated sea level requires the dynamic coupling of many processes and components to capture the integrated response of the system. Components would be a representation of the coastal seas, the regional atmosphere and the terrestrial surface and subsurface hydrology that would interact through different boundary conditions. Built on the fluxes within those models, air and water pollutant transport models and geochemical models could be implemented locally within the basin, regional or national scale domains. Each component would be able to assimilate data transmitted from field sites and assess the uncertainty in the predictions.

So what would be the requirements of such a system?

- It would have to deal with the need for predictions across a wide range of scales from small ecological niches to the regional atmosphere. This implies some form of nested structure for spatial objects representing domains of interest at different scales.
- It would have to deal with the way in which process parameterisations and effective values of parameters might change with scale.
- It would have to deal with the difficulty of defining effective parameter values for everywhere that predictions are required, taking account of uniqueness of place.
- It would have to allow for the various sources of uncertainty in representing particular places.
- It would have to allow for the use of monitoring data of different types to evaluate model predictions and refine the representation of particular places, or at least constrain uncertainties.
- It would have to allow for the evaluation of model predictions showing that a model representation appears to be wrong.
- It would have to present the results of predictions and the associated uncertainties in an effective interactive graphical interface.

MOVING TOWARDS ENVIRONMENTAL MODELS OF EVERYWHERE: IMPLEMENTATION

This integrated approach to multi-scale environmental modelling is only made possible by advances in computer power that have happened and that are planned, particularly the GRID. It requires interdisciplinary ways of thinking that have been only partially considered in the past. It brings together existing databases, existing models within the active spatial object structure, developing and making use of GRID middleware tools. The overall objective is to develop a system that links environmental modelling objects, spatial objects, databases, monitoring equipment and interactive visualisation tools over the GRID in a way that is flexible in design and implementation but robust in use.

In the past, comprehensive modelling systems have been constructed as large complex computer programmes. These programmes were intended to be general, but are expensive to develop and difficult to maintain. GRID computing technology and new developments in environmental modelling philosophy allow a new approach to be taken to this problem based on an approach that will match scale dependent model objects, databases and spatial objects in applications within the areas of interest. One of the most exciting benefits of the possibilities provided by the GRID in environmental modelling is the potential to implement models available from different institutions as a process of *learning about specific places within an uncertainty estimation framework.* Sites of interest can be implemented as *active objects*, seeking the information across the GRID to achieve a specified purpose and using the power of parallel computing resources to estimate the uncertainty associated with the predictions as constrained by *site-specific observations*, including those accessed over the GRID in real time. Initially, model results may be relatively uncertain but experience in monitoring and auditing of simulations will gradually improve the representation of sites and boundary conditions.

Such an integrated system should operate both in real time, assimilating data and boundary conditions from larger scale models, and displaying the “current state of the environment”. It should also provide the potential to update model predictions into the future under different scenarios. Requirements such as the Water Framework Directive are increasing demands for predictions of this type about the *responses of specific locations to change* in a way that integrates hydrological and ecological considerations in management. The system should be powerful enough to be used for assessing *uncertainties in model predictions and resulting risks.* It should also be able to be used off-line for “what-if” management purposes or decision support including developing strategies for risk-based sustainable management in the context of climate and other changes. This will include the evaluation licensing of air-borne emissions and effluents to water courses; strategies for remediation of contaminated land, rivers and estuaries, etc. in managing various subsystem components.

Implementation of this type of system on the GRID gives the opportunity to think about the nature of a modelling system of this type. In the past, it was relatively difficult to integrate large-scale databases and modelling systems as single monolithic program packages (such as the Water Information System (WIS) that was trialed by the UK Water Authorities in the 1990s). Now, the problem can be addressed in a way in which the different components can be treated as program objects that can be running on different machines and that can be upgraded individually.

The major technical challenges in implementing the proposed system will be to devise ways of handling the interactions between these spatial objects and the modelling components required for prediction across hardware and scales in a way that is as far as possible future-proof. This will be based on the concept of arbitrary active spatial objects that, once defined for a particular project, are informed autonomously by existing databases, on-line monitoring data and dynamic results from other model components (including real-time data assimilation). A hierarchy of modelling tools is envisaged that are both scale and purpose dependent but which will be themselves replaceable objects within the system.

The idea of places as active objects within the system will require some novel middleware development to support active, interacting, distributed spatial objects. This may involve developing intermediate gateway objects between the active objects and the Globus framework which manages data transformation and activity coordination on the GRID. One of the most interesting areas of development in this respect will be how to develop interfaces between the available databases and the model components at different scales of application.

A further issue in implementation will be dynamic resource allocation and reconfiguration. These facilities allow new computational resources to be deployed as soon as they become available and provide mechanisms for fault-tolerance as active objects are moved from failed nodes to other active nodes in the system. In addition, there will be times when maximum demand would be seen, for example, during short periods of flood alert conditions when joint atmosphere, ocean and hydrological models might be needed in real time, coupled to telemetry of radar and other measurement devices.

ENVIRONMENTAL MODELS OF EVERYWHERE AND PUB

The interaction of places as active objects and databases is a problem that is at the very heart of the IAHS Prediction of Ungauged Basins (PUB) initiative. Even in gauged basins, the vast majority of places *within* a basin will be ungauged. There may be maps available of geology, soils, land use and topography. These may be linked to relational databases for hydrological and hydraulic characteristics. There may be remote sensing images available for different types of sensors and at different resolutions. But none of this information gives the parameter values required by a model representation of places within the catchment directly. If they did, there would be no need for the PUB initiative.

Model representations must be based on such data, of course. Hence the current interest in distributed modelling strategies (see, or example, www.nws.noaa.gov/oh/hrl/dmip). Treating places as active objects will require that places will be able to identify the relevant databases over the GRID. To be useful, the process of model application will require the definition of a self-coding system attached to places to record and retrieve the methods that have been applied to (or by) that place in the past so that they can be easily reviewed and evaluated by the user. There is then a further interesting question that arises as to how far the place, once defined for a problem, can learn about itself from the data and model predictions available. Tools such as fuzzy classification or genetic programming might be used to extrapolate data from that and other sites to develop predictive methods of appropriate complexity, fit for the purpose of an application, within the limitations of the uncertainties implied by the data available. This approach has been advocated, for example, by the proponents of 'hydroinformatics' (e.g. Abbott, 1991, 1992). In hydrology, it is still a research question as to how to best use different types of data in the characterisation of places. It is likely to work something like this:

- First estimates of the parameters required by a model structure (or multiple model structures) are derived from GIS overlays and linked databases of parameter values. These first estimates may be of a range of potential values, or possible distribution of potential values rather than single values for different parameters. There is a difficulty in doing this, in that the *effective* values required in a particular model structure may be different for different model implementations (e.g. Beven, 2000, 2001, 2004).
- First estimates (or ranges or distributions) are made of the initial and boundary conditions for the place of interest. These estimates might also vary with model implementation.
- Predictions are made of observable responses for the application, including the uncertainty arising form the uncertainty in the effective parameter values and initial and boundary conditions. These predictions are compared with observations and the difference is used to refine the representation of the system (effective parameter values, boundary conditions) and uncertainty in the predictions. This may involve a model rejection step if it is found that the model predictions are incompatible with the observations.
- The refined information on feasible model structures, effective parameter values and boundary conditions are fed back to augment the databases.

MOVING TOWARDS ENVIRONMENTAL MODELS OF EVERYWHERE: DATA ASSIMILATION AND MODEL EVALUATION

These steps summarise a framework for learning about places — or as formulated here for places as active objects learning about themselves by a process of data assimilation (McLaughlin, 1995). This is data assimilation used in a way that is different to many applications in the past. The classic use of data assimilation, as for example is still the case in atmospheric prediction, is to use observable variables to refine the estimates of the suite of model state variables with a view to improving future forecasts. Methods such as 4D variational assimilation are used to re-initialise atmospheric models in weather forecasting (see, for example, Swinbank *et al*., 2003). 4D methods have also been used to update model estimates of surface soil moisture based on passive microwave observations (Houser *et al*., 1998). In other situations, methods such as Tikhonov regularisation (van Loon and Troch, 2002) and the extended Kalman filter (EKF) are used to update both state variables and parameter values, for example, in real-time flood forecasting (Young, 2002). The EKF, and the related ensemble Kalman filter (Margulis *et al*., 2002), can also be considered as a form of recursive Bayesian analysis. Bayesian methods have been used to refine parameter estimates of rainfall-runoff models (Thiemann *et al*., 2001; Vrugt *et al*., 2002); to refine estimates of rainfall inputs to rainfall-runoff models (Kavetski *et al*., 2002) and in real-time flood forecasting (Krzysztofowicz, 2002). Kalman filter and Bayesian methods can also provide estimates of the variances and covariances associated with both predicted variables and parameter values, at least under standard Gaussian assumptions about the nature of that uncertainty.

The possibility of the routine application of data assimilation within a learning framework raises some interesting questions about the nature of modelling and model evaluation. Effectively, repeated updating and correction of model predictions will allow the data assimilation process to compensate for errors in model inputs and model structure. Model evaluation therefore becomes more difficult. In real-time forecasting this may not be such a problem. In fact, we would *wish* for the data assimilation process to compensate for errors in model inputs and model structures if this results in improved forecasts with maximised accuracy and minimised uncertainty.

This may not be the case, however, in simulation where such compensation may not be desirable if it leads to a model structure being accepted when it should be rejected. There is, of course, a fundamental difficulty in deconstructing the causes of model error and isolating the model structural error alone (see discussions in Kavetski *et al*., 2002; Beven and Young, 2003; Beven, 2004). While we do not want to accept a model structure because of the compensation allowed by data assimilation, equally we would not want to make the error of rejecting a perfectly good model because of errors in the input data. Differentiating between these sources of error may be very difficult, if only because the *nature* of both input errors and model structural errors may be non-stationary in time.

Thus there is a question as to whether the use of data assimilation can, in a simulation context, reveal deficiencies in either model structures or input data. There is an analogy here with the State Dependent Parameter (SDP) estimation methodology used by Young (2003; Romanowicz *et al*., 2004). In the SDP approach, an initial estimate of a linear transfer function model is used within a recursive data assimilation framework to examine how the best estimates of the parameters change over time or with respect to some other variable. This can lead to the identification of the dominant non-linear modes of behaviour of the system, based directly on observations rather than prior conceptual assumptions about the system response.

Generalised Likelihood Uncertainty Estimation (GLUE) can be used in a similar recursive way, updating the weights associated with Monte Carlo model realisations as more data periods are simulated (e.g. Beven and Freer, 2001; Freer *et al*., 2002). This has been formalised in the Dynamic Identifiability Analysis (DYNIA) of Wagener *et al*. (2003). However, one important step in the GLUE methodology is the possibility of allowing a model rejection step if the predictions go beyond some limits of acceptable error (Freer *et al*., 2002; Beven, 2004). This is an important point since we will not learn very much about the nature of the model from using data assimilation to compensate for model error. We might learn more from model rejection, and particularly where there is a reason for rejecting all the models tried.

However, these types of parameter and uncertainty estimation methodologies are best suited to systems that are not changing and where conditioning data are available. In that way new data should allow a refinement of the feasible model representations and reduction in the predictive uncertainty. Many applications of simulation to environmental systems, however, involve questions of past, current or future change under different scenarios. Data assimilation may be valuable in the analysis of past change. Recursive analysis may reveal the impact of past change in terms of changing effective parameter values or model rejections.

Data assimilation may also be useful in following the nature of current change and refining predictions into the future under different scenarios. There has been very limited work on estimating the uncertainties of potential outcomes in future scenario simulations, and less on the conditioning of those predictions as monitoring reveals changing conditions. Data assimilation, in this framework, then becomes a tool for following drift in system response (within the limitations of data uncertainties).

CONCLUSIONS

Perhaps a theme can be identified running through this discussion of environmental models of everywhere: the focus on data. Data will be required to characterise places, to drive model predictions, to evaluate and update the results of model predictions in data assimilation and constrain predictive uncertainty, to reject some models previously considered feasible, and to monitor changes in system response. The role of models has always been, albeit sometimes rather implicitly, to extrapolate data in both time and space. This role will now become more explicit in extrapolating from those sites where data are available to the more numerous sites without data and where the characteristics are poorly known. There will still be an argument for using models based on understanding to do that extrapolation (particularly for predicting the impact of changes into the future) but, given the demonstrated limitations and uncertainties of current models based on understanding, there will also be the opportunity to reconsider the extrapolation problem for particular places. In essence, it would appear that learning about the specifics of *places*, and taking account of the inherent uncertainty in doing so, will become more important than using particular model structures, until a model is rejected. Rejection may then provide an important driving force for model improvement (where an application justifies the costs of further site specific studies).

A further challenge will be to embed the system within an uncertainty analysis framework for both forecasting and scenario modelling uses. This also raises some complex issues. It is worth noting that a statistical approach to uncertainty, treated as an additive error model, can also serve to compensate for model deficiencies, albeit with suitably large error bounds (see the example in Thiemann *et al*., 2001, where the error bounds for discharge predictions do not contain the observations at successive hydrograph peaks, and discussions of Beven and Young, 2003; Gupta *et al*., 2003; and Beven, 2004). There has been very little discussion in the literature about how much error is allowable before a model is rejected, presumably because in most cases the uncertainty is evaluated with respect to the ‘optimal’ model found in calibration, even if the predictions might not be satisfactory.

Thus, the potential for implementing environmental models of everywhere on GRID scale computing resources provides some particularly interesting questions for the design of a flexible system that will allow continued learning about the nature of places across a range of scales and applications while allowing for continuing uncertainties and their constraint by the assimilation of different types of data. In this process lies the resolution of the ungauged basin problem, by the very fact that such models necessarily require the representation of everywhere. The challenges posed are numerous: of places as active objects, of uniqueness of place, of the value of different types of data; of model rejection and the potential for compensation of model and data deficiencies by data assimilation and uncertainty estimation. Regardless of these scientific challenges, however, environmental models of everywhere will happen, and in the not-too-distant future. It is time to prepare.

REFERENCES

Abbott, M.B. 1991. *Hydroinformatics, Information Technology and the Aquatic Environment*. Avebury Technical, Aldershot, UK and Brookfield, USA.

Abbott, M.B. 1992. The theory of the hydrological model, or: the struggle for the soul of hydrology. In: *Advances in Theoretical Hydrology*, J P O'Kane (Ed.), Elsevier. 237–254.

Beven, K.J. 1989. Changing ideas in hydrology: the case of physically based models. *J. Hydrology*, **105**, 157–172.

Beven, K.J. 1993. Prophecy, reality and uncertainty in distributed hydrological modelling. *Adv. in Water Resourc.*, **16**, 41–51.

Beven, K.J. 1995. Linking parameters across scales: sub-grid parameterisations and scale dependent hydrological models, *Hydrolog. Process.*, **9**, 507–526.

Beven, K.J. 2000. Uniqueness of place and the representation of hydrological processes, *Hydrol. Earth Syst. Sci.*, **4**, 203–213.

Beven, K.J. 2001. How far can we go in distributed hydrological modelling? *Hydrol. Earth Syst. Sci.*, **5**, 1–12.

Beven, K.J. 2002. Towards a coherent philosophy of environmental modelling. *Proc. Roy. Soc. Lond. A*, **458**, 2465–2484.

Beven, K.J. 2004. A manifesto for the equifinality thesis, submitted to *Water Resources Research*.

Beven, K.J. and Freer, J. 2001. Equifinality, data assimilation, and uncertainty estimation in mechanistic modelling of complex environmental systems. *J. Hydrol.,* **249**, 11–29.

Beven, K.J. and Young, P.C, 2003. Comment on Bayesian Recursive Parameter Estimation for Hydrologic Models by M Thiemann, M Trosset, H Gupta and S Sorooshian. *Water Resour. Res.,* **9**, doi: 10.1029/2001WR001183

Blöschl, G. and Sivapalan, M., 1995, Scale issues in hydrological modelling – a review. *Hydrolog. Process.*, **9**, 251–290.

Crawford, N.H. and Linsley, R.K. 1966. Digital simulation in hydrology: Stanford Watershed Model IV. *Technical Report 3*9, Dept. of Civil Engineering, Stanford University, CA.

Earth Simulator Center 2003. Annual Report of the Earth Simulator Centre, April 2002-March 2003, The Earth Simulator Center, Yokohama, Japan.

Fleming, G. 1975. *Computer Simulation Techniques in Hydrology*. Elsevier, Amsterdam.

Freer, J.E., Beven, K.J. and Peters, N.E.. 2002 Multivariate seasonal period model rejection within the generalised likelihood uncertainty estimation procedure. In: *Calibration of Watershed Models*, Q. Duan, H. Gupta, S. Sorooshian, A.N. Rousseau and R. Turcotte (eds.), AGU Books, Washington, 69–87.

Hoeben, R. and Troch, P. 2000, Assimilation of active microwave observation data for soil moisture profile estimation. *Water Resour. Res.*, **36,** 2805–2819.

Houser, P.R., Shuttleworth, W.J., Gupta, H.V., Famiglietti, J.S., Syed, K.H. and Goodrich, D.C. 1998, Integration of soil moisture remote sensing and hydrologic modelling using data assimilation. *Water Resour. Res.*, **34**, 3405–3420.

Kavetski, D., Franks, S.W. and Kuczera, G., 2002, Confronting input uncertainty in environmental modelling. In: *Calibration of Watershed Models*, Q. Duan, H. Gupta, S. Sorooshian, A.N. Rousseau, and R. Turcotte (eds.), AGU Books, Washington, 49–68.

Krzysztofowicz, R. 2002. Bayesian system for probabilistic river stage forecasting. *J. Hydrol.,* **268**, 16–40.

Margulis, S.A., McLaughlin, D., Entekhabi, D. and Dunne, S. 2002. Land data assimilation and estimation of soil moisture using measurements from the Southern Great Plains 1997 Field Experiment. *Water Resour. Res.*, **38**, DOI 10.1029/2001-WR001114.

McLaughlin, D. 1995. Recent developments in hydrologic data assimilation. *Reviews Geophysics (suppl.)*, 977–984.

Morton, A. 1993, Mathematical models: questions of trust-worthiness. *Brit. J. Phil. Sci*., **44**, 659–674.

Romanowicz, R. and Beven, K.J. 1998. Dynamic real-time prediction of flood inundation probabilities. *Hydrol. Sci. J.,* **43**., 181–196.

Romanowicz, R., Young, P.C. and Beven, K.J. 2004. Data assimilation in the identification of flood inundation models: derivation of on-line multi-step ahead predictions of water levels (*this volume*)

Swinbank, R., Shutyaev, V. and Lahoz, W.A. 2003. *Data assimilation for the Earth System*, NATA Science Series. IV. Earth and Environmental Sciences, Vol. 26, Kluwer Academic, Dordrecht.

Thiemann, M., Trosset, M., Gupta, H. and Sorooshian, S. 2001. Bayesian recursive parameter estimation for hydrologic models. *Water Resourc. Res.,* **37**, 2521–2535.

Van Loon, E.E, and Troch, P. 2002. Tikhonov regularisation as a tool for assimilating soil moisture data in distributed hydrological models. *Hydrolog.Process.*, **16**., 531–556.

Vrugt, J.A., Bouten, W., Gupta, H.V. and Sorooshian, S. 2002. Toward improved identifiability of hydrologic model parameters: The information content of experimental data, *Water Resour. Res.*, **38**, DOI 10.1029/2001WR001312.

Wagener, T., McIntyre, N., Lees, M.J., Wheater, H.S. and Gupta, H.V., 2003. Towards reduced uncertainty in conceptual rainfall-runoff modeling: Dynamic identifiability analysis. *Hydrolog. Process.,* **17**, 455–476.

Young, P.C. 2002. Advances in Real-Time Flood Forecasting. *Phil. Trans. R. Soc. Lond.*, **360**, 1433–1450.

Young, P.C. 2003. Top-down and data-based mechanistic modelling of rainfall-flow dynamics at the catchment scale. *Hydrolog. Process.*, **17**, 2195–2217.