INVITED COMMENTARY

HYDROLOGICAL PROCESSES
Hydrol. Process. **17**, 0–0 (2003)
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/hyp.5090

On environmental models of everywhere on the GRID

Keith J. Beven*

Department of Environmental Science, Lancaster University

*Correspondence to:
Keith J. Beven, Department of Environmental Science, IENS, Lancaster University, Lancaster LAI 4 YQ, UK.
E-mail: k.Beven@lancaster.ac.uk

New computer technologies, such as the GRID, seem likely to change the way that environmental models are constructed and used. The GRID is a new hardware and software initiative based on distributed high-performance parallel computers, linked by fast network connections that, to the user, should appear as a single machine. The concept is that the user should not have to worry about where the data necessary for a project are stored, nor where any computational tasks are run. To the user, the software (or 'middleware') should make the GRID appear as a desktop machine. The possibility of using GRID-scale computer networking to link together distributed database and computational engines means that it will become possible to couple together models of many more different environmental systems across disciplinary boundaries and across national administrative boundaries. In fact, this is already possible and is already happening on a limited basis, as demonstrated, for example, in the regional water resources models under construction in Denmark and in the national environmental management models being used in the Netherlands.

There is, however, a real question raised about how these types of interdisciplinary model might be best implemented. In the past, comprehensive modelling systems have been constructed as large complex computer programs. These programs were intended to be general, but have proven to be expensive to develop, difficult to maintain and difficult to apply because of their data demands and needs for parameter identification. With GRID computing technology it will be possible to continue in the same vein, but with more coupled processes and finer spatial and temporal resolutions for the predictions. It is not clear, however, whether this will result in a real improvement in model accuracy and use, because the problems inherent in the current generation of distributed environmental model do not necessarily easily go away with improvements in space and time resolutions of the component models.

There may be another approach, one that will be explored in this commentary. One of the features of having the possibility of these large-scale models is that everywhere is represented. We will have environmental models of everywhere. Once all places are represented within the flexible GRID-based system outlined in what follows below, the data may assume a greater importance than model structures as a means to refine the representation of each place within a learning framework. The result may be a new way of looking at environmental modelling, one that transcends the

Received 30 October 2002
Accepted 30 October 2002

traditional goal of incorporating all our understanding of the complexity of coupled environmental systems into a single mathematical framework with a multitude of parameters that cannot easily be identified for any particular place (Beven, 2002).

We can distinguish two different (albeit overlapping) types of model here in terms of constraints. In the first, the accuracy of solutions is still constrained by computational resources; in the second model, this accuracy is primarily a result of lack of knowledge of appropriate process representations and boundary conditions. In the first types of model, real advances may still be possible as computational constraints are relaxed. In atmospheric modelling, for example, there is still scope for improvements of the representation of local convection and rainfall forecasts, in the representation of sub-grid spatial variability of energy fluxes, and in the representation of topography by finer grid scales. Ultimately, however, this type of model will be constrained by the need to know increasingly finer detail of boundary conditions and parameter values, as is already the case in the second type of model. An example of this second type is classical distributed hydrological models.

New developments in environmental modelling philosophy allow a new approach to be taken to this problem based on matching 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*. It will be possible, in fact, to have models of *all* places of interest. However, as argued by Beven (2000, 2001, 2002), as a result of scale, nonlinearity and incommensurability issues, the representation of place will be inherently uncertain so that this learning process should be implemented *within an uncertainty estimation framework*.

Sites of interest for a particular prediction 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, based perhaps on only geographical information system databases and limited local information, may be relatively uncertain, but experience in monitoring and auditing of predictions will gradually improve the representation of sites and boundary conditions. It is this learning process that will be critical in the development of a new generation of environmental models that are geared towards the management of specific places, rather than general process representations.

That is not to say that models of places will not require process representations, but there is a real research question about how detailed a process representation is necessary to be useful in predicting the dominant modes of response of a system, given the uncertainties inherent in representing the processes in places that are all unique. This appropriate complexity issue has become obscured by the desire to build more and more scientific understanding into model, including physical, chemical and biological components. This desire is perfectly understandable, it is a way of demonstrating that we do understand the science of the environment, but it results in models that have lots of parameter values that cannot be easily measured or estimated in applications to real places. There is always a certain underlying principle in science, that as we add more understanding and eliminate empiricisms, then the application of scientific principles should become simpler and more robust. This does not seem to have been the experience in the practical application of environmental models.

Events, such as the river flooding in the UK in 2000 and 2002 and the consequences of the 2001 fires in the USA, have demonstrated the need for a new generation of systems for environmental forecasting. The subtle (and sometimes not so subtle) coupling between atmospheric forcing, catchment response, river runoff and coastal interaction with tidally dominated sea level requires the dynamical coupling of many processes and components to capture these subtleties. 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 biogeochemical models could be implemented locally within the regional-scale domain. Each component would be able to assimilate data transmitted from field sites and assess the uncertainty in the predictions. The components would share 4-D/5-D visualization tools with appropriate interactive user interfaces. Users will be able to access the current data, visualize predictions for particular locations and play what-if scenario games over different time scales. The structure of the system would be such as to facilitate and even stimulate improvements to the representation of different components and the constraint of predictive uncertainty by field data collection. The potential capabilities of the GRID underlie all these components, though much could already be achieved using the Web technology of today. Examples of steps toward this type of integrated system (albeit essentially raster based) include the US Inter-Agency Object Modelling System (OMS).

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', as well as provide the potential to update model predictions into the future under different scenarios. Initiatives such as the European Union 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 would need to be powerful enough to be used for assessing *uncertainties in model predictions and the consequent risks of potential outcomes*. 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 of management of subsystems, including for licensing of air-borne emissions and effluents to water courses; strategies for remediation of contaminated land, rivers and estuaries, etc.

An essential element of this strategy will be the need, as far as possible, to 'future proof' the model and database systems used; avoiding, for example, a strict raster based approach or a commitment to one particular modelling framework. The key will need to be flexibility. Raster databases will continue to be driven by remote-sensing imaging inputs to the modelling process, and, in some cases, by convenient numerical solution schemes for partial differential equations. However, it is often inappropriate to force an environmental problem into a raster straightjacket. Treating places as flexible active objects might be one way around this future-proofing problem. Defining the spatial domain of a prediction problem would allow that place, as an active object, to search on the GRID for appropriate methods and data for resolving that problem, and also for appropriate methods and data for providing the boundary conditions for the problem (which might then involve other modelling or data extrapolation techniques).

There are some interesting implications of such an approach. One is that the variety of modelling methods available across the GRID to solve a prediction problem might be able to be compared more readily, leading to better understanding of issues of appropriate model complexity for different modelling problems. This will especially be the case if, as part of the learning process, simulations are saved to be compared with later observations of the real outcome. This use of 'post-audit' analysis has been rarely used in environmental modelling, but it has been instructive in the field of groundwater modelling (Konikow and Bredehoeft, 1992; Anderson and Woessner, 1992) and is routine in atmospheric modelling in the evaluation of forecast skill (although the evaluation of global climate model predictions still requires an element of compromise at the regional level (Shackley *et al.*, 1997).

To be useful, however, 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; using methods such as fuzzy classification or genetic programming as tools to extrapolate data from that and other sites to develop predictive methods of appropriate complexity to the problem at hand 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).

The learning framework that underlies this framework is best suited to systems that are not changing. In that way new data should allow a refinement of the feasible model representations and reduction in the predictive uncertainty. Many of the predictions required of environmental systems, however, involve questions of current or future change under different scenarios. Such predictions will be even more uncertain than the simulation of current conditions, but 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 of a changing conditions changes. Data assimilation, in this framework then becomes a tool for following drift in system response (within the limitations of data uncertainties).

Perhaps a theme can be identified in running through this discussion of environmental models of everywhere: the focus on data. Data will be required to characterize places, to drive model predictions, to evaluate the results of model predictions 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 *places*, and taking account of the inherent uncertainty in doing so, will become more important than using particular model structures.

References

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

Abbott MB. 1992. The theory of the hydrological model, or: the struggle for the soul of hydrology. In *Advances in Theoretical Hydrology*, O'Kane JP (ed.). Elsevier: 237–254.

Anderson MP, Woessner WW. 1992. The role of the post-audit in model validation. *Advances in Water Resources* 15: 167–174.

Beven KJ. 2000. Uniqueness of place and the representation of hydrological processes. *Hydrology and Earth System Sciences* 4(2): 203–213.

Beven KJ. 2001. How far can we go in distributed hydrological modelling? *Hydrology and Earth System Sciences* 5(1): 1–12.

Beven KJ. 2002. Towards a coherent philosophy of environmental modelling. *Proceedings of the Royal Society of London, Series A* 458: 2465–2484.

Konikow LF, Bredehoeft JSD. 1992. Groundwater models cannot be validated. *Advances in Water Resources* 15: 75–83.

Shackley S, Young PC, Parkinson S, Wynne B. 1997. Uncertainty, complexity and concepts of good science in climate change modelling: are GCMs the best tools? *Climate Change* 38(2): 159–235.

