Hydrology and Earth System Sciences, 2(3), *283–298* (1998) © EGS

Hydrology & Earth System Sciences

A grid-based distributed flood forecasting model for use with weather radar data: Part 2. Case studies

V.A. Bell and R.J. Moore

Institute of Hydrology, Wallingford, Oxfordshire, OX10 8BB, UK

Abstract

A simple distributed rainfall-runoff model, configured on a square grid to make best use of weather radar data, was developed in Part 1 (Bell and Moore, 1998). The simple form of the basic model, referred to as the Simple Grid Model or SGM, allows a number of model variants to be introduced, including probability-distributed storage and topographic index representations of runoff production and formulations which use soil survey and land use data. These models are evaluated here on three catchments in the UK: the Rhondda in south Wales, the Wyre in north-west England and the Mole in the Thames Basin near London. Assessment is initially carried out in simulation mode to focus on the conversion of rainfall to runoff as influenced by (i) use of radar or raingauge input, (ii) choice of model variant, and (iii) use of a lumped or distributed model formulation. Weather radar data, in grid square and catchment average form, and raingauge data are used as alternative estimates of rainfall input to the model. Results show that when radar data are of good quality, significant model improvement may be obtained by replacing data from a single raingauge by 2 km grid square radar data. The performance of the Simple Grid Model with optimised isochrones is only marginally improved through the use of different model variants and is generally preferred on account of its simplicity. A more traditional lumped rainfall-runoff model, the Probability-Distributed Moisture model or PDM, is used as a benchmark against which to assess the performance of the distributed models. This proves hard to better, although the distributed formulation of the Grid model proves more reliable for some storm and catchment combinations where spatial effects on runoff response are evident. Assessment is then carried out in updating mode to emulate a real-time forecasting environment. First, a state updating form of the Grid Model is developed and then assessed against an ARMA error-prediction technique. Both state updating and error prediction give much improved model performance when compared with simulation mode results. No one updating technique is superior, with the simulation model formulation having greatest impact on forecast accuracy. However, when the results from the different catchments are considered together it is apparent that in the rapidly responding Rhondda catchment state updating gives slightly better results, while in the slower responding Wyre and Mole catchments, error prediction is slightly superior. This is attributed to the greater difficulty of reliably adjusting states when there are significant time delays associated with the catchment response. In general, the influence of rainfall input type, model variant and distributed versus lumped model reflect the results obtained in simulation mode. Updating doesn't fully compensate for a poor rainfall input or a deficient rainfall-runoff model formulation, especially for longer forecast lead times.

Introduction

A grid-based distributed rainfall-runoff model designed for use with weather radar data, and called the Grid Model, was presented in Part 1 (Bell and Moore, 1998). Part 2 aims to assess the basic Grid Model, called the Simple Grid Model or SGM, and its variants when applied to three catchments in the UK. The key science issues addressed are whether the use of distributed measurements of rainfall can provide significant forecast improvement and whether a distributed model can outperform a lumped model, given the benefit of distributed weather radar data. Selection of a distributed model appropriate for real-time use from the set of model variants considered forms an important science and operational issue. Assessments of distributed models employing data from comparatively sparse raingauge networks have generally concluded that model performance is input limited and the potential of distributed models awaits improvement in the spatial measurement of rainfall.

Evaluation of the models is carried out in both simulation and forecast mode, the latter emulating a real-time environment in which recent flow observations are used to update the model in order to improve forecast accuracy. Both grid square and catchment average radar data are used in the assessment along with raingauge data as a baseline reference. Performance of the distributed models is compared with that of a lumped conceptual rainfall-runoff model used operationally in the UK for flood forecasting and warning. Perfect foreknowledge of measured rainfall is assumed so as not to introduce errors associated with rainfall forecasts which would only serve to confound the assessment.

The model variants and the three catchments involved in the evaluation are outlined first along with the data sources used and the procedures adopted for model calibration and assessment. The evaluation of model performance is presented in two stages, the first in 'simulation mode' and the second in 'updating mode'. In simulation mode the model acts as a transformation of rainfall (and potential evaporation) into basin runoff without reference to observed runoff, other than as a basis for assessment. This allows the evaluation to focus on (i) the influence of radar and raingauge inputs in grid and catchment average form on flow simulation performance, (ii) the deterministic model formulation and the relative merits of the different model variants and (iii) the relative merits of lumped versus distributed models. In the latter case, a lumped conceptual model, the Probability-Distributed Moisture or PDM model, is used as a baseline reference. Results are presented and discussed for each catchment in turn followed by a summary of the overall findings.

In updating mode a form of the model suitable for use in real-time is developed. This employs a method of state adjustment, in which the internal water store contents of the model are adjusted using observed runoff, to achieve better accordance between observed and forecast flows. The results obtained are compared with an error-predictor form of updating which exploits the dependence seen in past model errors to predict future ones, which are then used to improve the initial form of the forecast. Assessment in forecast (or updating) mode is restricted here to the use of the 'best' Grid Model for a catchment and the lumped PDM model. More extensive results presented in Moore *et al.* (1994) indicate that the effect of type of rainfall input and choice of model variant reflect the results presented in simulation mode. Finally, a summary of the main findings is presented along with some suggestions for further research.

Framework for model assessment

MODEL VARIANTS FOR EVALUATION

The Grid Model and its variants formulated in Part 1(Bell and Moore, 1998) are reviewed briefly here so as to identify clearly the set of models to be considered in the assessment. The basic Grid Model employs a simple water balance within each grid square, in which storage capacity is related to the mean slope gradient in the square as calculated from a Digital Terrain Model (DTM). Overflow and drainage from the grid square storage are translated separately to the basin outlet using two discrete kinematic routing models in parallel, representing fast and slow response pathways respectively, and with each reach coincident with an isochrone band. The two routing models share the same isochrone-based discretisation into reaches and differ only in the kinematic wave speed parameter used to characterise their response. Isochrones are derived using path lengths over land and river inferred from a DTM together with an estimate of the velocity of flow along land and river paths. In the basic Grid Model initial estimates of 0.1 and 0.5 m s^{-1} for land and river paths are assumed.

The first two model variants relate to the translation components of the Grid Model. In the first model variant the velocities used to calculate the isochrones are optimised as part of the model calibration process. This formulation is used as the basis of the model variants that follow. In the second variant the kinematic routing model representing the slow response pathways is allowed to have its own isochrones. These are obtained using Darcy velocities inferred from the gradient and length of the paths as calculated from the DTM.

The formulations of the third and fourth variants introduce representation of spatial variability in runoff production *within* a model grid square. In the third variant the simple storage water balance for a grid square is replaced with a probability-distributed storage representation, which is linked to the distribution of slope gradient within the square as measured from the DTM. The fourth variant replaces the storage water balance with a topographic index representation of soil saturation within a model grid square.

Soil survey and land use information form the basis of the remaining two model variants which again influence runoff production. The first replaces the storage capacity linkage with slope with a linkage to Integrated Air Capacity data and average slope in the square. Land use data in the form of Landsat-classified urban areas are used to introduce a fraction of impervious area within a grid square for the simple storage water balance.

CATCHMENTS USED FOR MODEL EVALUATION

Three catchments in the UK are used to assess the performance of the basic Grid Model and its variants. These are the Mole to Kinnersley Manor in the Thames basin near London, the Wyre to St Michaels in north-west England and the Rhondda to Trehafod in south Wales. Radar rainfall coverage is provided by the C-band weather radars at Chenies, Hameldon Hill and Dyfed respectively. These radars form part of the UK operational network and provide data on a 2 km grid out to 75 km and 5km out to a range of 210 km. Details of each catchment are given below and Fig. 1 provides maps of the catchments including locations of the raingauges used.

The Mole to Kinnersley Manor drains an area of 142 km^2 to the south of London with relatively low relief, elevation ranging from 48 to 178 m. The catchment is largely impervious, with Weald Clay lithology being predominant.

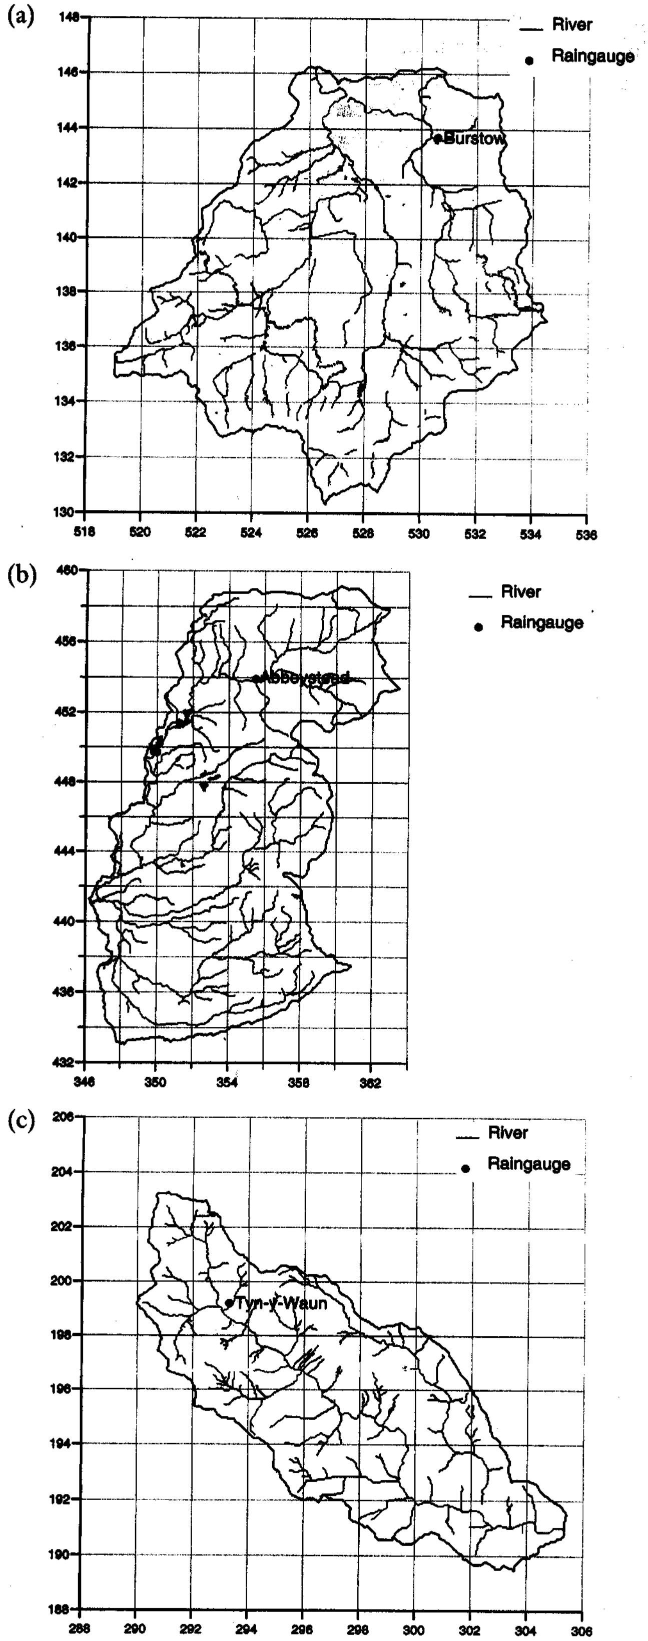

Fig. 1. *Maps of the catchments used for model evaluation showing the raingauge locations. (a) Mole, (b) Wyre, (c) Rhondda*

Land use is very mixed ranging from rural tracts to urban centres (notably Crawley) and the airport at Gatwick. The two other catchments are in areas of higher relief. The Wyre catchment to St Michaels drains an area of 275 km^2 with an elevation range of 4 to 560 m. Land use is almost wholly rural with moorland on limestone, Millstone Grit and Bunter Sandstone. The gauging station is slightly affected by tides at low flows. The smaller Rhondda catchment drains an area of 100 km^2 to the gauging station at Trehafod, situated below the confluence of twin valleys: the Rhondda Fach and the Rhondda Fawr. Here, the area upstream of Castell Nos reservoir is discounted, the reservoir being fully impounding and contributing only when spilling, reducing the modelled catchment to 94 km^2. The elevation range is from 68 to 600 m. Urban areas occupy 13% of the catchment and these are concentrated in the valley bottoms. The Rhondda is the wettest of the three catchments with an annual average rainfall (for the standard period 1941–70) of 2200 mm, compared to 1211 and 811 mm for the Wyre and Mole catchments respectively.

AVAILABILITY OF DATA FOR MODEL EVALUATION

Radar coverage for the Mole catchment is provided by the London Weather Radar at Chenies, north-west of London. The '75 km' range circle for which 2 km data are available just encompasses the catchment; 53 grid squares are associated with the catchment. Raingauge data are available for a site at Burstow, located within the catchment. Data for the period 22 December 1990 to 22 January 1991 are used for model evaluation.

For the Wyre catchment, weather radar coverage is available from the Hameldon Hill radar and 95 grid squares are associated with the catchment. The radar data required pre-processing to interpolate across areas of anomaly due to blockages in the beam (notably a television mast and hills) using the correction method described in Moore *et al.* (1992). Raingauge data are available from the Abbeystead site located within the catchment. Data used for model evaluation are for the period 23 October to 22 November 1986.

Radar coverage for the Rhondda catchment is not ideal with the Dyfed radar at Crug-y-Gorllwyn in part being beyond the 75 km range. For about three-quarters of the catchment 5 km radar data are mapped onto the 2km model grid; 39 model grid squares are associated with the catchment. Raingauge data are available for the Tyn-y-Waun site located within the catchment. Data for the two periods 10 June to 5 July 1991 and 1 January to 1 February 1993 are used for model evaluation.

A simple sine curve over the year with a mean value of 1.4 mm day^{-1} is used for daily potential evaporation (PE) in water balance calculations for the two higher relief catchments, the Wyre and the Rhondda. This value is based on information supplied by the UK Met. Office; note that the value is intermediate between the constant value of 1.37 and the sine curve climatological mean value of 1.5 given by Calder *et al.* (1983) as typical for the UK. For the Mole catchment, historical monthly Penman PE values for the 'Upper Mole' sub-area of the Thames are disaggregated to form a daily series using linear interpolation between the 16th day of successive months. A constant value within a day provides an adequate approximation for water balance purposes. The data and model time-step employed throughout is 15 minutes. Flow data are available at this time-step for each catchment and are used here for model evaluation and in forecast updating.

MODEL CALIBRATION AND EVALUATION PROCEDURES

Model calibration followed a procedure of interactive visual inspection of hydrographs and manual adjustment of model parameters followed by automatic optimisation to minimise a chosen objective function. The objective function used is the sum of the square of the difference between the observed flow, Q_t, and modelled flow, q_t, at each time step, t; that is

$$J = \sum \varepsilon_t^2 \qquad (1)$$

where the model error $\varepsilon_t = Q_t - q_t$ and the summation is computed over n values. A robust and straightforward simplex (polytope) minimisation procedure (Nelder and Mead, 1965), modified to incorporate ideas suggested by Gill, Murray and Wright (1981), is used for optimisation of the model parameters. It is the authors' experience that a model which is successful at representing the hydrological response of a catchment is likely to have parameters that are interdependent. This applies even when a parsimonious parameterisation is used. The parallel routing formulation of the Grid Model is a good example where the representation provides a realistic hydrological response but where interdependence of the velocity parameters may arise. This problem in part is alleviated by using information from the digital terrain model to support configuration of the routing model. However, such parallel formulations are inherently prone to give rise to parameter interdependence. Interactive visualisation of the model hydrographs as part of a manual calibration process is used, in combination with automatic optimisation, for this reason.

The main criteria used for assessment are the root mean square error, defined as

$$rmse = \sqrt{\left(n^{-1}\sum \varepsilon_t^2\right)} \qquad (2)$$

and the R^2 statistic

$$R^2 = 1 - \frac{\sum \varepsilon_t^2}{\sum (Q_t - \bar{Q})^2}. \qquad (3)$$

where $\bar{Q}$ is the mean of the n observed flow values. The R^2 statistic indicates the proportion of the variance in the original observations accounted for by the model.

Evaluation of the models is carried out using the same period of data used for calibration rather than employing a formal 'split sample testing' procedure incorporating an independent dataset. This was primarily due to the difficulty of putting together more extensive datasets incorporating weather radar data. Use of a long period of data, typically one month in duration and containing a number of flood events, along with models containing relatively few parameters (usually between 5 and 8 active parameters out of the set of 11) helps alleviate the effects of overfitting. However, an extension of this work would aim to include split sample testing and employ further datasets encompassing winter and summer periods, and convective summer storms displaying significant spatial variability.

Assessment of models in simulation mode

The models are assessed first in simulation mode for each of the three catchments in turn, and the results used to compile a synthesis of the main conclusions. The focus is on the deterministic model formulation and to establish which of the model variants offers the best performance. Models are calibrated and assessed using grid square radar data, catchment average radar data and raingauge data in order to assess any benefit of a distributed model formulation and the use of radar compared with raingauge data. The assessment in forecast mode that is presented later incorporates the effects of updating and provides an indication of the forecast performance at different lead times that can be obtained in a real-time context.

MODEL EVALUATION FOR THE MOLE CATCHMENT

The results obtained from the Grid Model and its variants are shown in Table 1 in terms of the R^2 and *rmse* performance statistics. Model simulated flows using radar data as input are seen to be more accurate than those using raingauge data throughout, with R^2 values of circa 0.86 compared to 0.67. Figure 2 shows a comparison of the simulated hydrographs obtained from the optimised isochrone models using raingauge and radar data as rainfall input, demonstrating the significantly improved performance obtained when using radar data. The distributed nature of the radar data clearly has a beneficial effect on model performance.

Table 1 shows that for both radar and raingauge data, optimisation of the velocities involved in the calculation of isochrones improves model performance appreciably (R^2 increasing from 0.73 to 0.86 for radar). No significant improvement is gained through the use of further model

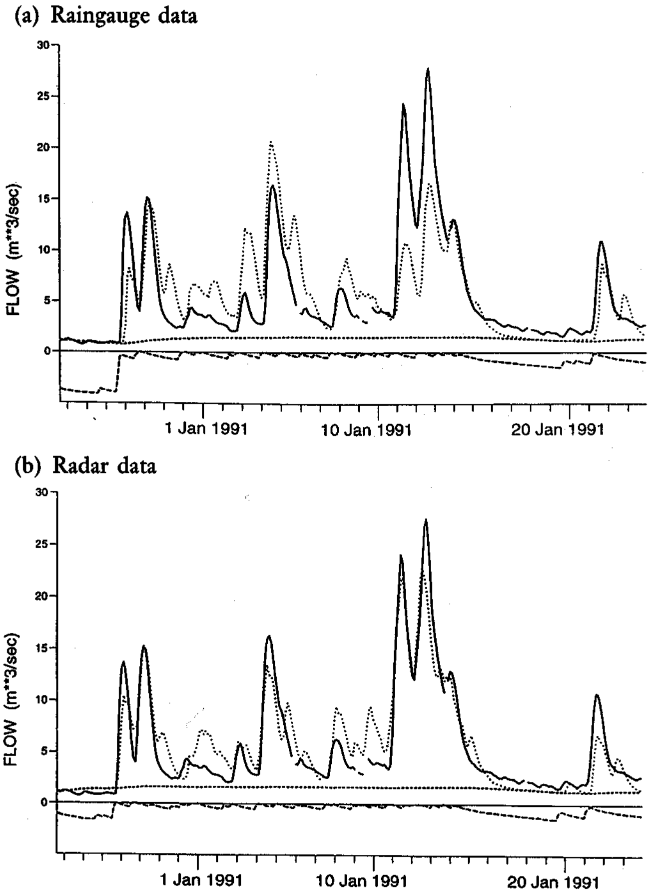

Fig. 2. *Simulations from the Simple Grid Model using raingauge and radar data: Mole catchment. 22 December 1990 to 22 January 1991 (solid line: observed flow; dotted line: simulated total flow; dashed line: simulated 'baseflow'; lower dashed line: storage deficit).*

variants. The best model employs separate Darcy slow response isochrones and uses radar data as input. For raingauge data used as input, the best model performance is obtained using the probability-distributed storage model formulation; the full set of model variants have not been tested on account of the poorer performance obtained using raingauge data for this catchment. Typical values of the model parameters obtained for the Mole catchment are shown in Table 2; Table 2 of Part 1 (Bell and Moore, 1998) provides a brief description of the parameters and their units. Note that the regional upper limit of storage capacity, $c_{\max}$, is allowed to vary, but 75 mm was found to be appropriate for the Mole across different model types.

As expected, optimising the isochrone velocities on land and river increases accuracy. Automatic calibration has the effect of increasing the river velocity from 0.50 to 0.59 m s^{-1}, and decreasing the land velocity from 0.10 to about 0.02 m s^{-1}. The overall influence on the unit hydrograph base is an increase from 22 to 70 hours, after which the response becomes negligible (the optimisation was halted to force the velocities to remain at acceptable values). A comparison of the hourly unit hydrographs obtained using the standard and optimised velocities reveals that with a change to the latter, although the time-to-peak increases only slightly from 10 to 13 hours, a more pronounced secondary peak develops at 36 hours and the base is significantly lengthened (Fig. 3). This arises mainly on account of the very small land velocity which gives small closely spaced isochrones in regions away from the river, and a unit hydrograph with a long narrow tail. The double peak effect in the flow hydrograph after isochrone optimisation follows from the similarly shaped unit hydrograph. Note that the unit hydrographs displayed in Fig. 3 characterise the advective response of the basin to uniform rainfall and do not incorporate the additional diffusion effect of the discrete kinematic wave scheme used for flow routing.

MODEL EVALUATION FOR THE WYRE CATCHMENT

For model evaluation purposes the Wyre catchment provides an interesting contrast to the Mole, being located in an upland, hilly region of north-west England. Table 3 summarises the performance obtained by the different model variants. Results for the urban area impervious fraction are not included because the urban fraction is so small for this rural catchment. Figures 4 (a) and (b) show the observed and simulated hydrographs for the case of

Table 1 **Model assessment for the Mole catchment**

Model variant	Raingauge		Radar	
	R^2	*rmse* $m^3\ s^{-1}$	R^2	*rmse* $m^3\ s^{-1}$
Isochrones with standard velocities	0.436	3.969	0.733	2.730
Isochrones with optimised velocities	0.578	3.432	0.864	1.950
Darcy slow response isochrones			**0.869**	**1.913**
Probability-distributed storage model	**0.667**	**3.048**	0.822	2.229
Topographic index model			0.862	1.960
Integrated Air Capacity model			0.862	1.963
Urban area impervious fraction			0.869	1.914

Table 2. Model parameter estimates for the Mole catchment

Data and model type	S_0 —	c_{max} mm	f_r —	$k_d.10^7$ $h^{-1}mm^{-2}$	θ_s —	θ_b —	v_L $m\ s^{-1}$	v_R $m\ s^{-1}$	g_{max} —
Radar Optimal isochrones	0.9	75	1.350	2.780	0.880	0.490	0.015	0.587	8.25
Raingauge Optimal isochrones	0.635	75	1.036	2.517	0.831	0.519	0.015	0.605	8.25
Raingauge Probability-distribution storage model	0.135	75	1.139	3.666	0.910	0.722	0.017	0.584	10.25

(a) Before optimisation: land velocity = 0.1 m s^{-1}, river velocity = 0.5 m s^{-1}

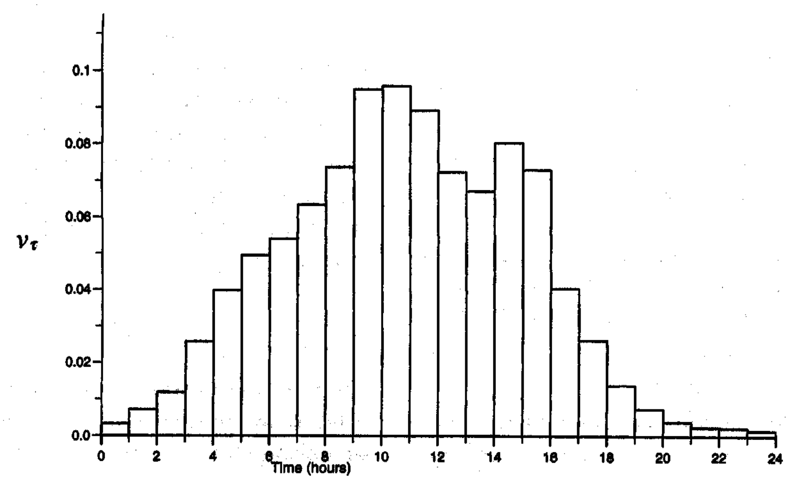

(b) After optimisation: land velocity = 0.015 m s^{-1}, river velocity = 0.587 m s^{-1}

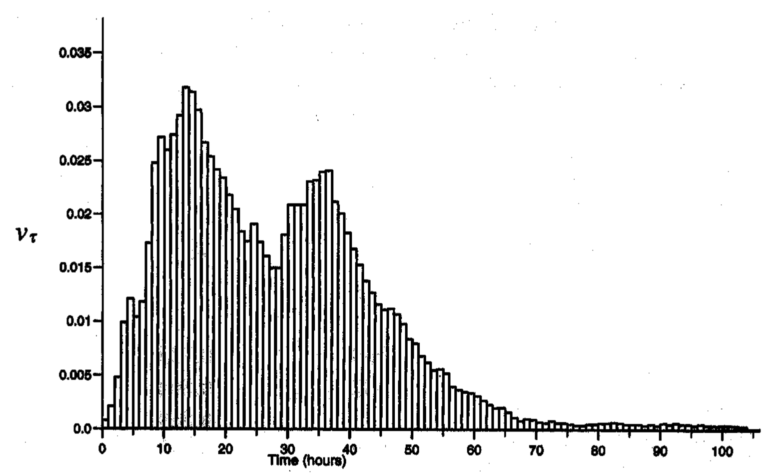

Fig. 3. *DTM-derived 'unit hydrographs' for the Mole catchment before and after optimisation.*

Table 3. Model assessment for the Wyre catchment.

Model variant	Raingauge		Radar	
	R^2	*rmse* $m^3 s^{-1}$	R^2	*rmse* $m^3 s^{-1}$
Isochrones with standard velocities	0.552	12.41	0.624	11.370
Isochrones with optimised velocities	0.679	10.51	0.738	9.502
Darcy slow response isochrones	0.722	9.784	0.731	9.626
Probability-distributed storage model	0.732	9.608	0.705	10.082
Topographic index model	0.718	9.862	0.710	9.991
Integrated Air Capacity model	**0.760**	**9.101**	**0.759**	**9.103**
Urban area impervious fraction	—	—	—	—

optimised isochrones, first using raingauge and then weather radar data as rainfall input to the Grid Model.

Flows simulated using radar data are slightly better than those obtained from raingauge data (R^2 of 0.74 compared to 0.68 for the SGM). The simulated hydrographs can be seen to reflect the differences between radar and raingauge measurements in individual events, some rainfall events being well measured by one and not the other.

Best overall model performance (R^2 of 0.76) is obtained from the models employing Integrated Air Capacity, with similar results for raingauge and radar data. Using raingauge data, the probability-distributed storage Grid Model is next best (R^2 of 0.73). Overall the results show that similar performance is obtained using raingauge and radar data, except for the simplest models when radar gives slightly better flow simulations. The model variants improve performance only marginally; the basic model with optimised isochrones might be judged best, being simple and fast to calibrate whilst still providing reasonable performance (R^2 of 0.74).

MODEL EVALUATION FOR THE RHONDDA CATCHMENT

Calibration and assessment of the Grid Model and its variants for the Rhondda catchment are carried out for both a summer and winter period: 10 June to 5 July 1991 and 1 January to 1 February 1993. For the summer period, radar failed to measure rainfall well and for this reason tables and figures of results are omitted. Good model simulations are obtained using raingauge data with an R^2 of 0.72 for the optimised isochrone Grid Model. A particular problem is that the radar fails to detect the storm on 25 and 26 June. Radar data from higher elevation scans were examined in an attempt to account for the inability of the radar to 'see' the rain on this occasion. Examination of vertical sections through the rainfall field, derived from the multiple scan data, provides some evidence for low-level enhancement below the radar beam leading to underestimation of rainfall at the ground. Support also comes from the work of Hill *et al.* (1981) which shows that in winter 80% of the orographic enhancement can occur in the lowest 1.5 km over the south Wales hills, and is particularly associated with strong winds and high relatively humidity below 2 km. A seeder-feeder mechanism appears to operate over hills where raindrops, falling from pre-existing higher-level clouds, wash out smaller cloud droplets within lower-level clouds whose liquid water content is replenished by low-level air flows ascending over the hills.

The second period of data for January 1993 is analysed here as a response to the poor performance of the radar during the June/July 1991 period. The results are summarised in Table 4. Underestimation of rainfall by radar again badly affects model performance with the best model using radar data giving an R^2 of only 0.60. The storms of 20 and 21 January are barely detected by the radar with rainfall formation apparently occurring below the radar beam. Using raingauge data most of the model variants result in only a slight improvement in performance when compared with that from the basic model. Best overall performance is obtained using the probability-distributed storage model with optimised isochrones and raingage data (R^2 of 0.95). Figure 5 shows the simulated hydrograph from this model which provides an excellent correspondence with the observed flows. Optimisation of isochrones has the effect of increasing the land velocity from 0.08 to 0.95 m s^{-1} and river velocity from 0.5 to 1.7 m s^{-1}. The latter compares well with flow velocities measured in the field which indicate that at average flows of around 30 $m^3 s^{-1}$ velocities are in the region of 1.75 m s^{-1}.

SUMMARY OF MODEL SIMULATION RESULTS

Model variants

The model variants—probability-distributed stores, topographic index, integrated air capacity and urban area impervious fraction—all result in some improvement in some circumstances, but no variant can be relied upon to improve model performance. Using the performance of the

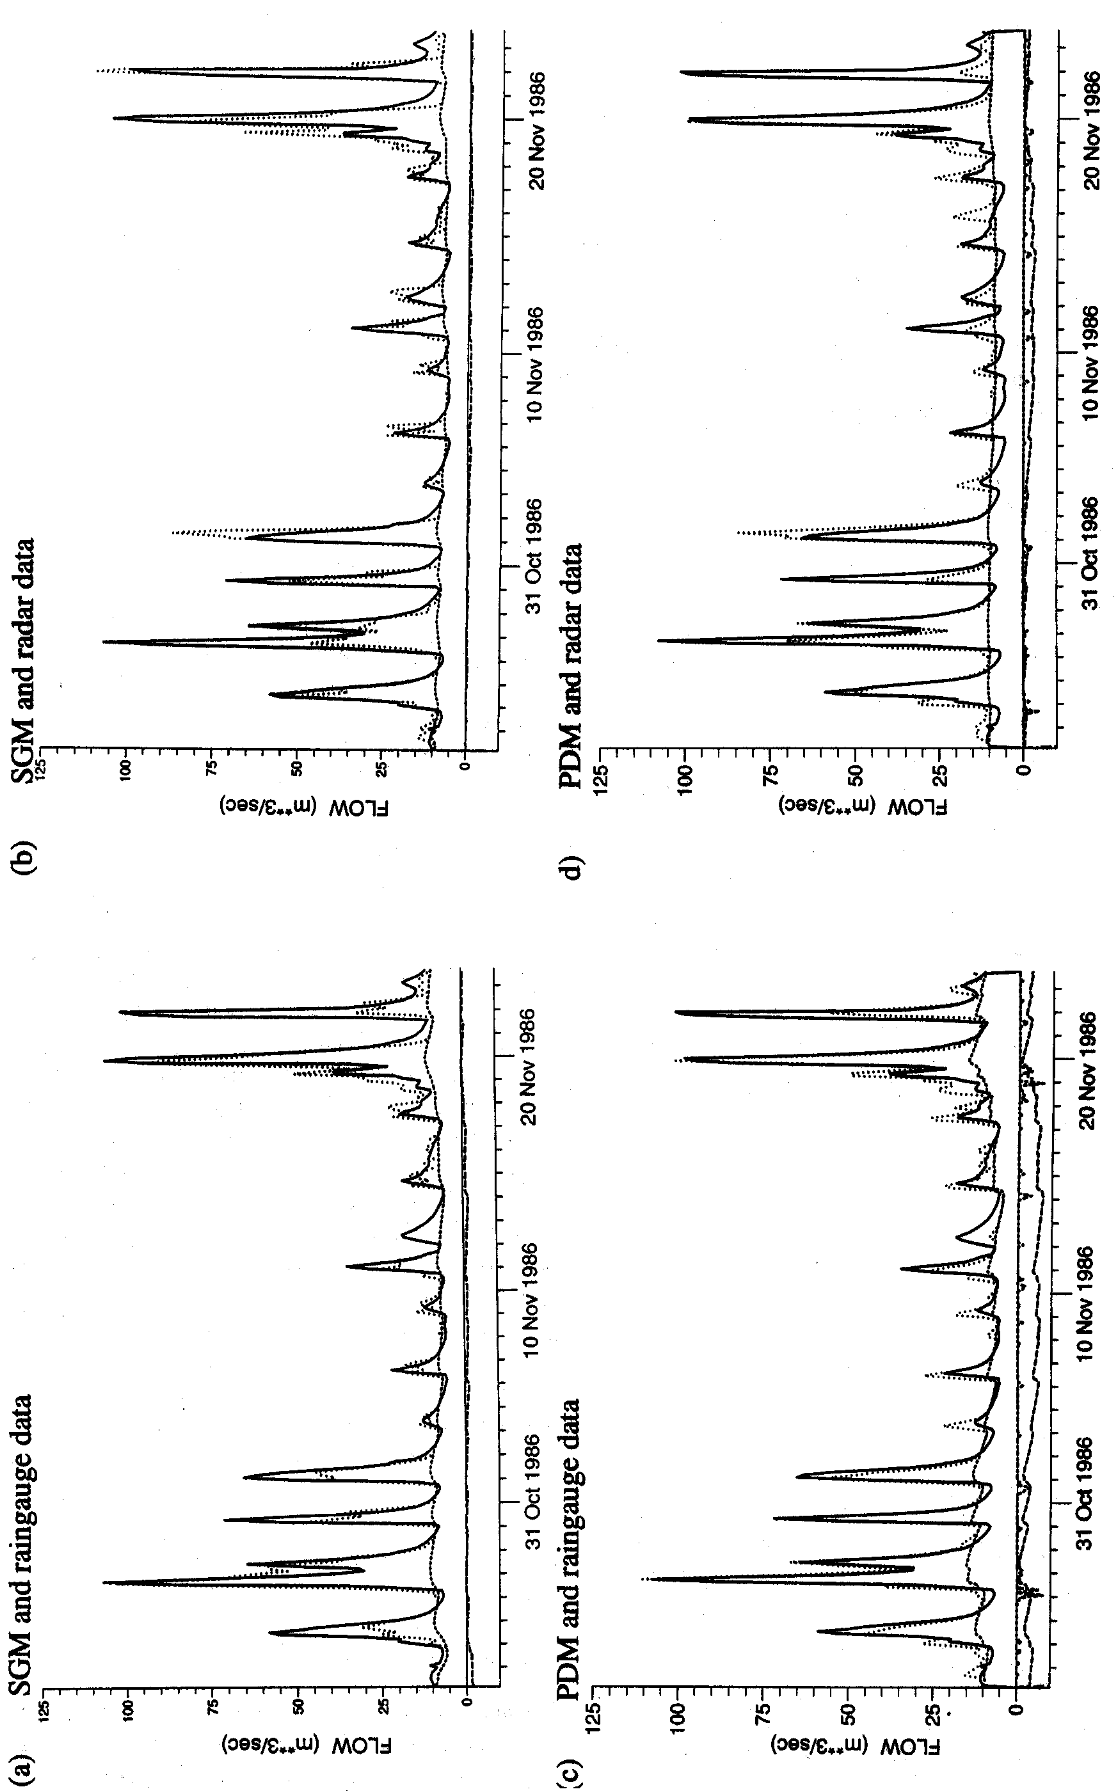

Fig. 4. Simulations from the distributed Simple Grid Model and the lumped PDM model using raingauge and radar data: Wyre catchment, 23 October to 22 November 1996. (Above origin:– solid line: observed flow; dotted line: simulated total flow; dashed line: simulated 'baseflow'. Below origin (on a proportional scale):– lower dashed line: simulated storage deficit; spiky dashed line ((c) and (d) only): rainfall).

Table 4. Model assessment for the Rhondda catchment using January 1993 data.

Model variant	Raingauge		Radar	
	R^2	*rmse* $m^3 s^{-1}$	R^2	*rmse* $m^3 s^{-1}$
Isochrones with standard velocities	0.640	8.223	0.340	11.133
Isochrones with optimised velocities	0.904	4.251	0.506	9.635
Darcy slow response isochrones	0.914	4.012	—	—
Probability-distributed storage model	**0.953**	**2.987**	—	—
Topographic index model	0.947	3.151	—	—
Integrated Air Capacity model	0.917	3.959	**0.600**	**8.672**
Urban area impervious fraction	0.896	4.427	0.486	9.822

optimal isochrone model as a benchmark against which to compare the performance of other model forms, it is apparent that for the higher relief catchments (the Rhondda and the Wyre) the Integrated Air Capacity model and the probability-distributed storage model often, but not always, result in improved flow simulations. This might be expected for steeper catchments where slope effects on storage and in turn runoff production are likely to be more dominant. These formulations appear to be less useful for modelling flows for the Mole catchment where low relief reduces the influence of slope on storage and runoff production. Flow simulations for the Mole are improved slightly through the use of the model variant which incorporates separate Darcy slow response isochrones. Model simulations using raingauge data, however, are improved considerably by the use of probability-distributed stores, with R^2 increasing from 0.58 for optimal isochrones to 0.67. The Grid Model variant which estimates soil saturation using a topographic index improves model performance for the Rhondda catchment

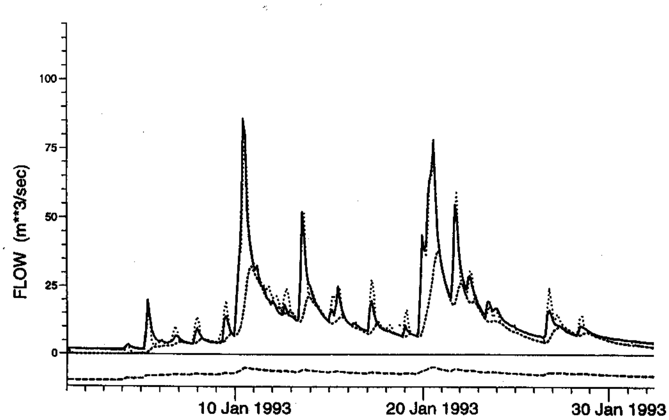

Fig. 5. *Simulations from the probability-distributed storage Grid Model using raingauge data: Rhondda catchment, January 1993 (solid line: observed flow; dotted line: simulated total flow; dashed line: simulated 'baseflow'; lower dashed line: simulated storage deficit).*

when radar rainfall data are used, but the probability-distributed storage model performs slightly better. However, for the Mole and using radar data the topographic index model performs slightly better than the probability-distributed storage model, whilst both are outperformed by the variant using Darcy isochrones. Use of Landsat-classified urban areas for delineating impervious areas is usually helpful in improving model performance and in the latter case gives essentially equal performance to the Darcy isochrone variant.

Radar versus raingauge data

A substantial improvement in performance is gained when radar data are used in place of data from a single raingauge in the 126 km² Mole catchment, with R^2 increasing from 0.58 to 0.86 for the basic model. Problems with blockages (notably a television mast and hills) in the radar field for Hameldon make the model results obtained for the Wyre catchment difficult to assess. Within the month-long event considered, the raingauge provides the better rainfall estimate for one storm and the radar for another, with both performing about equally overall, giving an R^2 of circa 0.72 (Fig. 4 (a) and (b)). Results obtained using the Dyfed radar data for the Rhondda catchment are disappointing. Here, radar underestimates or even fails to detect rainfall on a number of occasions. This is thought to be associated with low level enhancement of rainfall below the radar beam in this hilly region.

Optimised isochrones

In general the greatest improvement in model performance, for the three catchments considered, is provided by optimisation of model isochrones via calibration of land and river velocities. This procedure improves performance in all cases and usually by a considerable amount, for example increasing R^2 from 0.73 to 0.86 in the case of the Mole catchment using radar data.

Distributed versus catchment average radar data

To assess whether the improvement in model performance for the Mole when radar data are used is due to the distributed nature of the data, or whether the radar simply measures rainfall better than the raingauge, a comparison is made using grid square and catchment average radar data as input to the Grid Model. The results are presented in Table 5(a) where it can be seen that a slight improvement in performance is obtained using grid square rather than catchment average radar data.

Table 5. Comparison of model performance obtained using grid square and catchment average radar data

(a) Mole catchment		
Rainfall data type	R^2	*rmse* ($m^3 s^{-1}$)
Grid square radar	0.859	1.966
Catchment average radar	0.827	2.179
(b) Wyre catchment		
Rainfall data type	R^2	*rmse* ($m^3 s^{-1}$)
Grid square radar	0.737	9.510
Catchment average radar	0.580	12.020

A similar comparison is made for the Wyre catchment and the results summarised in Table 5(b). A substantial improvement in model performance is obtained using grid square rather than catchment average radar rainfall as input. Examination of the simulated hydrographs reveals an underestimation in the simulated peaks for 30 October, 9 and 20 November and the spurious appearance of a flow peak around 14 November when catchment average radar data are used.

Whilst only slightly improved performance using grid square radar data is obtained for the Mole catchment, for the larger more variable Wyre catchment the value of such distributed data is more apparent. Simulated peak flows seem particularly sensitive to the spatial variability of rainfall used as input to the Grid Model for the Wyre.

Lumped versus distributed models

To determine whether there is any benefit in using a distributed grid-based model over a more conventional lumped rainfall-runoff model, the Institute of Hydrology's Probability Distributed Moisture model or PDM (Moore, 1985, 1993; Institute of Hydrology, 1992, 1996) is applied to the three catchments and same periods of record. In an evaluation of three lumped conceptual models used operationally for flood forecasting in the UK (Moore *et al.*, 1993; Austin and Moore, 1996), the PDM was found to provide the most resilient forecasts. It therefore provides a good benchmark for comparison.

The PDM is a conceptual rainfall-runoff model which uses a probability distribution to describe the spatial variation of water storage capacity across a catchment. Saturation excess runoff generated at any point in the catchment is integrated over the catchment to give the total direct runoff entering fast response pathways to the basin outlet. Drainage from the soil enters slow response pathways. Storage routing representations of the fast and slow response pathways yields a fast and 'baseflow' response at the basin outlet which, when summed, gives the total basin flow.

The performance of the PDM for each of the three catchments is summarised in Table 6 along with the performance of the best Grid-Model variant for each catchment. The assessment of the (distributed) Grid Model against the (lumped) PDM model indicates that the latter is in general marginally superior and for the Wyre using raingauge data it performs particularly well; the flow simulations are shown in Fig. 4 (c). However, when radar data are used the PDM simulates spurious flow peaks whilst the Simple Grid Model is able to simulate the observed flows more reliably (Fig. 4 (d) and (b)).

Table 6. Comparison of the 'best' Grid Model for a catchment with the lumped PDM model using the R^2 statistic performance measure

Catchment	Data type	'Best' Grid Model	PDM Model
Mole	Radar	0.869	**0.876**
	Raingauge	0.667	0.646
Wyre	Radar	0.759	0.679
	Raingauge	0.760	**0.853**
Rhondda	Radar	0.600	0.623
	Raingauge	0.953	**0.979**

A general conjecture is that a distributed approach to flood forecasting may prove worthwhile in situations where the scale of the storm system is less than that of the basin for which flood forecasts are required, or where there is significant internal heterogeneity in hydrological response. However, a lumped model may prove more resilient for flood forecasting for smaller, more homogeneous catchments, particularly at times of widespread rain associated with low pressure systems. The lumped model has the further advantages of simplicity and ease of initialisation.

Assessment of models in forecast mode

The evaluation of the Grid Model and its variants has been confined so far to judging performance in simulation mode, where a model acts to transform rainfall (and potential evaporation) to runoff without reference to observed flow. This has allowed attention to focus on the choice of an appropriate deterministic model structure. For real-time applications, model simulations can be further improved through incorporating recent measurements of flow, usually received via telemetry. The technique of incorporating such measurements to improve forecast performance is called 'updating'. The terms 'forecast mode' and 'updating mode' are used synonymously to refer to forecasts obtained using some form of updating technique.

Two forms of forecast updating—state updating and error prediction—will be considered here. State updating is based on using the simulation-mode model errors to adjust the amount of water in store (treated as state variables) in various parts of the model to make the model forecasts better agree with observed flows. Error prediction attempts to predict future simulation-mode model errors from present and past ones, and to add the predicted error to the model simulation to obtain an updated forecast. Whilst error prediction operates independently of the deterministic model structure, state updating schemes must be tailored to a particular model structure. First, a new state updating scheme is developed specifically for the Grid Model. The framework for model assessment in forecast mode is then outlined and the results of the assessment, carried out on each of the three catchments, are summarised.

METHODS FOR FORECAST UPDATING

State updating

The term 'state' is used to describe a variable of a model which mediates between inputs to the model and the model output (Szollosi-Nagy, 1976). In the case of the Grid Model the main input is rainfall and basin flow is the model output. Typical state variables are the water contents of the conceptual storage elements, typically representing soil/vegetation, channel and groundwater systems. The flow rates out of the conceptual stores can also be regarded as state variables: examples are q_s, the flow out of the fast response ('surface' or channel) system, and q_b, the flow out of the slow response ('baseflow' or groundwater) system.

When an error, $\varepsilon = Q - q = Q - (q_s + q_b)$, occurs between the model simulation, q, and the observed value of basin runoff, Q, it would seem sensible to 'attribute the blame' to mis-specification of the state variables and attempt to 'correct' the state values to achieve concordance between observed and model simulated flow. Mis-specification may, for example, have arisen through errors in rainfall measurement which, as a result of the model water accounting procedure, manifest themselves through the values of the store water contents, or equivalently the flow rates out of the stores. A formal approach to 'state correction' is provided by the Kalman filter algorithm (Jazwinski, 1970; Gelb, 1974; Moore and Weiss, 1980). For nonlinear dynamic models, such as the Grid and PDM models, an extended form of Kalman filter based on a linearisation approximation is required which is no longer optimal in the adjustment it provides. The implication of this is that simpler, intuitive adjustment schemes can be devised which potentially provide better adjustments than the more complex and formal extensions of the Kalman filter which accommodate nonlinear dynamics through approximations. We will call such schemes which make physically sensible adjustments 'empirical state adjustment schemes'.

A simple example is the apportioning of the error, ε, between the surface and groundwater flow routing components of the Grid Model in proportion to their contribution to the total flow. Mathematically this may be expressed as

$$q_b^* = q_b + \alpha g_b \varepsilon \tag{4}$$

$$q_s^* = q_s + (1 - \alpha) g_s \varepsilon \tag{5}$$

where the baseflow proportion

$$\alpha = q_b/(q_s + q_b) \tag{6}$$

and the superscript * indicates the value after adjustment. The 'gain' coefficients, g_b and g_s, when equal to unity yield the result that $q_s^* + q_b^*$ equals the observed flow, Q, thus achieving exact correction of the model flow to equal the observed value. Values of the coefficients other than unity allow for different adjustments to be made, and g_b and g_s can be regarded as model parameters whose values are established through optimisation to achieve the 'best' fit between state-adjusted forecasts and observed flows.

A generalisation of the above is to define α to be

$$\alpha = \frac{q_b}{\beta_1 q_s + \beta_2 q_b} \tag{7}$$

and to choose the incidental parameters β_1 and β_2 to weight the apportionment towards or away from one of the flow components; in practice β_1 and β_2 are assigned values of 10 and 1.1 to apportion more of the error adjustment to the surface store. Note that the adjustment is carried out at every time-step and the time subscripts have been omitted for notational simplicity. The scheme with α defined by (6) is referred to as the proportional adjustment scheme and that defined by (7) is the super-proportional adjustment scheme. Replacing α and $(1 - \alpha)$ in (4) and (5) by unity yields the simplest non-proportional adjustment scheme.

Application of state updating to the Grid Model

The Grid Model has two separate routing components, one representing fast translation typically along channel paths and the other slow translation associated with sub-surface paths. Since the routing procedure is similar for both, and based on a cascade of discrete kinematic reaches, the same form of state updating scheme can be used. For simplicity of presentation a single routing path is assumed below.

Consider the one-step ahead forecast for time t+1 made from a time origin t. For the kinematic reach model (Bell and Moore, 1998; Moore *et al.*,1994), the flow out of the j th reach at time t+1 is given by

$$q_{t+1}^j = (1 - \theta) q_t^j + \theta q_t^{j+1} + \theta r_t^j \tag{8}$$

for $j = 1,2,...,N$. Here, N is the number of reaches, with each reach chosen to be coincident with an isochrone band,

and r_t^j is the lateral inflow (runoff or drainage) from the jth band; θ is a dimensionless wave speed parameter. It is possible to update this simulation forecast of flow using the observed simulation error at the time origin t, $\varepsilon = Q_t - q_t^1$, where Q_t is the measured flow at the catchment outlet and q_t^1 the model simulation. The form of adjustment is to modify the flows out of each reach so that the adjusted model outflow equals the measured outflow; that is

$$q_1^* \equiv (q_t^1)^* = Q_t = q_t^1 + \varepsilon. \quad (9)$$

Also, the adjustments to upstream reach outflows are chosen to decrease smoothly as a power function to zero at the topmost (Nth) reach, such that at time t

$$q_j^* = q_j + f_p(j)\varepsilon \quad j = 1,2,...,N \quad (10)$$

where

$$f_p(j) = \left(\frac{N-j}{N-1}\right)^p \quad (11)$$

and the exponent p is a constant parameter.

The updated forecast corresponding to equation (8) is then given by

$$(q_{t+1}^j)^* = q_{t+1}^j + \{(1-\theta)f_p(j) + \theta f_p(j+1)\}\varepsilon \quad (12)$$

where the last term is the correction that is applied to the forecast obtained using (8). This completes the development for a single channel path and for 'total' adjustment to match the observed flow.

In practice two parallel channels, representing 'surface' runoff q_s and 'baseflow' q_b, are used in the normal form of Grid Model. The adjustment follows equations (4) to (7) allowing for partial adjustment, proportional adjustment or super-proportional adjustment, but with g_b (and similarly for g_s) replaced by

$$g_b' = \{(1-\theta_b)f_p(j) + \theta_b f_p(j+1)\}g_b \quad (13)$$

where g_b (and g_s) are gain coefficients estimated by optimisation; here, θ_b denotes the dimensionless wave speed of the 'groundwater' channel path.

It is also possible to formulate an adjustment to the water contents of the soil/vegetation stores in the Grid Model. The form of adjustment investigated, for a given grid square with capacity $S_{\max}$ and water content S, is

$$S^* = S + \alpha\frac{S}{S_{\max}}g_s\varepsilon \quad (14)$$

with g_s a regional storage gain parameter and α allowing for proportional, super-proportional and direct (equal to unity) adjustments as before. Initial trials indicated that adjustment of the soil/vegetation store of the Grid Model provided little improvement and this approach is not included in the assessment reported later. This lack of success may be attributed to the time delays in the routing components of the Grid Model making allocation of errors to the soil/vegetation stores problematic.

It should be noted that all the above empirical state updating techniques utilise the same basic form of adjustment employed by the Kalman filter in which an updated state estimate is formed from the sum of the current state value and the model error multiplied by a gain coefficient. However, instead of defining the gain statistically, as the ratio of the uncertainty in the observation to that of the current state value, it is first related to a physical apportionment rule multiplied by a gain factor. This gain factor acts as a relaxation coefficient which is estimated through an off-line optimisation using past flow records.

Error prediction

As an alternative to empirical state updating, the Grid Model has been coupled with an ARMA error prediction scheme which exploits the dependence structure of model simulation errors in order to forecast future errors. A forecast of the error is added to the deterministic model simulation to obtain the updated model forecast. In contrast to the state adjustment scheme, which internally adjusts values within the model, the error prediction scheme is wholly external to the deterministic Grid Model operation.

Error prediction is now a well established technique for forecast updating in real-time (Box and Jenkins, 1970; Moore, 1982) which may be used in combination with any model. However, whilst error prediction provides a general technique which is easy to apply, its performance in providing improved forecasts will depend on the degree of persistence in the model errors. Unfortunately in the vicinity of the rising limb and peak of the flood hydrograph this persistence is least and errors show a tendency to oscillate rapidly and most widely. Dependence is at its strongest for errors on the falling limb, where improved forecast performance matters least for flood warning applications. In addition, timing errors in the model forecast may lead to erroneous error predictions being made, a problem which is also shared by the technique of state updating. The general applicability and popularity of error prediction as an updating tool commends its use as an 'off-the-shelf' technique, but empirical state adjustment schemes should also be considered as viable alternatives. The two techniques are therefore both included in the assessment that follows.

FORECAST ASSESSMENT FRAMEWORK

An assessment of model performance in forecast mode, incorporating real-time updating using both error prediction and state updating methods, is again carried out for the three catchments the Mole, the Wyre and the Rhondda. The assessment is restricted to a comparison of the basic form of Grid Model with optimised isochrones, denoted 'Simple Grid Model' (SGM), and the lumped conceptual model, PDM. The SGM is assessed using both state updating and ARMA error prediction whilst results for the PDM are presented using only state updating.

A transfer function noise (TFN) modelling package provides time series analysis tools used here to identify the appropriate order of the ARMA error predictor (number of autoregressive and moving average terms to include) and to provide initial estimates of its parameters. In all cases a fourth order autoregressive model, with dependence on the last four simulation-mode model errors, proved to be most appropriate. A preliminary trial on the relative merits of using proportional and super-proportional methods of state updating was carried out for the Simple Grid Model. The results suggest a marginal improvement can be obtained using the super-proportional form. This has been adopted in the following assessment when state updating is applied. The gain parameters are first found from optimisation by minimising the sum of the squares of the one-step ahead forecast errors. Generally, the power parameter p, involved in adjusting the flows in the kinematic routing model of the SGM, takes a value between 1 and 3.

Both fixed-origin and fixed lead-time forecasts are made. The fixed lead-time forecasts are made for 1,2,3,4,6,9 and 12 hours ahead. For fixed origin forecasts the selection of origin is based on first identifying flow peaks: initial forecasts are made with an origin at the start of the hydrograph rising limb and then using a forecast origin halfway up the peak. Forecasting for the Mole and the Wyre catchments is carried out using radar rainfall data, whilst for the Rhondda raingauge data are used because of the poor quality of the radar data. Results are presented below for each catchment in turn.

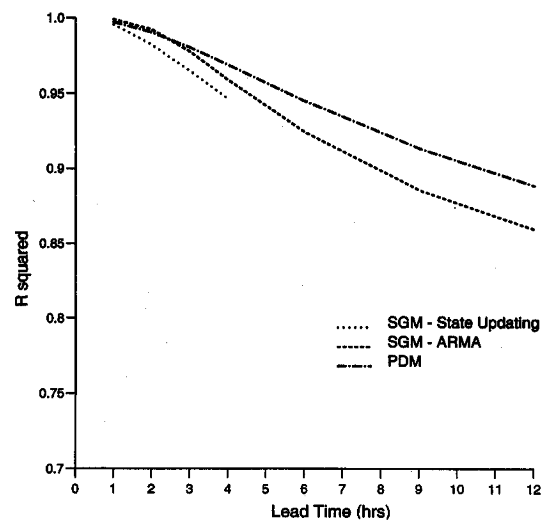

Fig. 6. *Forecast model performance at different lead times: Mole catchment.*

RESULTS OF FORECAST ASSESSMENT

The Mole

The large number of isochrones needed by the Grid Model for the Mole catchment coupled with the extra processing required by the state-updating procedure to calculate fixed lead-time forecasts, results in excessively long model run times. Calculation of fixed lead-time forecasts is therefore limited to a maximum of 4 hours for the SGM with state updating. The results are displayed graphically in Fig. 6. For very short lead times of one or two hours it can be seen that the SGM with error prediction gives the best performance, whilst for longer lead times the PDM with state-updating proves superior. Figure 7 compares the 2-hour ahead forecasts obtained using the the SGM with error prediction and the PDM, illustrating the better performance from the SGM for this lead time. A 2-hour forecast is chosen as operationally useful for flood warning purposes, considering both reliability and timeliness. Figure 8 shows the detail of the fixed origin forecasts covering the period 9–11 January 1991, obtained using the SGM with error prediction.

The Wyre

For the Wyre Fig. 9 displays the graph of R^2 forecast performance against lead-time. This shows that the SGM is better than the PDM using either form of updating, with error prediction proving best. The poorer performance of the PDM is consistent with its performance in simulation mode. Use of the distributed radar data in the Grid Model appears to be particularly beneficial for the larger Wyre catchment, which the lumped model PDM is unable to benefit from. Examination of the fixed origin forecasts from the SGM with state updating reveals a slight time lag in the forecast flows of one to two hours. This appears to account for the better performance of the SGM using error-prediction which is unaffected by any time lag.

The Rhondda

The PDM provides the best updated forecasts for the Rhondda catchment where raingauge data are used because of the poor quality of the radar data. The fixed origin forecasts are seen to be excellent (Fig. 10) with R^2 values above 0.97 for all lead times. Forecasts are also very good from the SGM using either updating method with R^2 values above 0.9 for all lead times; error prediction proves better than state updating for shorter lead times below 2½ hours.

OVERVIEW OF FORECAST ASSESSMENT

An overview of the above results across catchments suggests that significant improvement in forecast performance is obtained when an updating technique is used, at least for shorter lead times. For the more slowly responding Wyre and Mole catchments, error prediction is the marginally better updating method, whilst for the more rapidly responding Rhondda catchment state updating provides the slightly better performance. It is more likely that apportionment of errors to the routing reach storages can be done more reliably for catchments with smaller time delays; also any errors in assignment will dissipate more rapidly.

(a) SGM with error prediction

(b) PDM

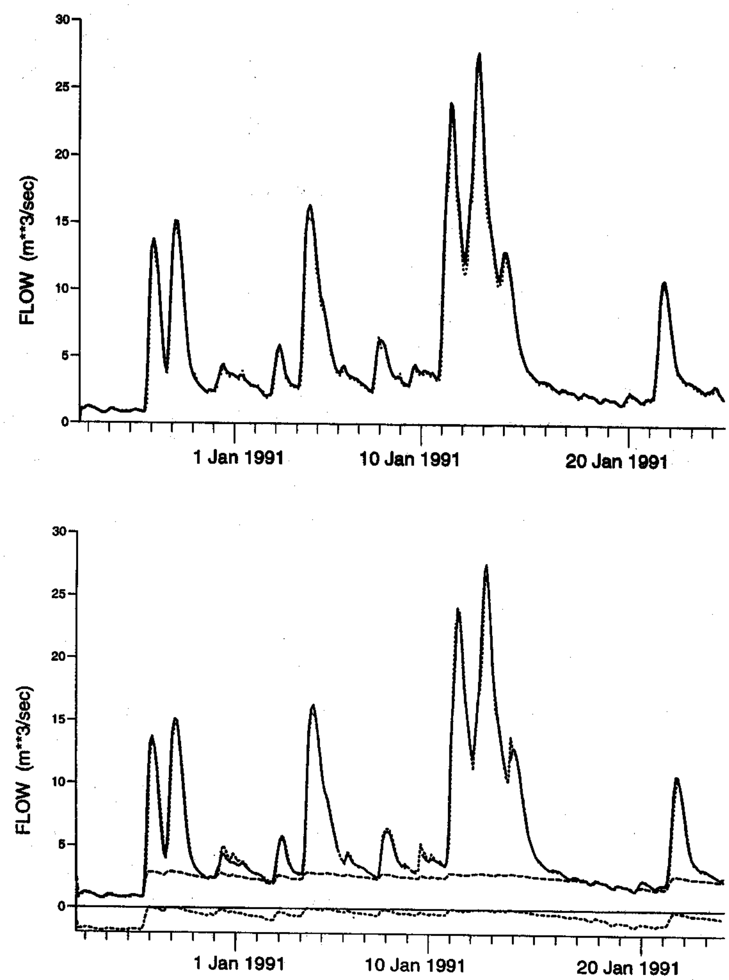

Fig. 7. *Comparison of the 2-hour ahead forecasts obtained using the SGM with error prediction and the PDM: Mole catchment. (Solid line: observed flow; dotted line: forecast flow. For (b) only:– dashed line: baseflow; dashed line below origin: storage deficit on a proportional scale.)*

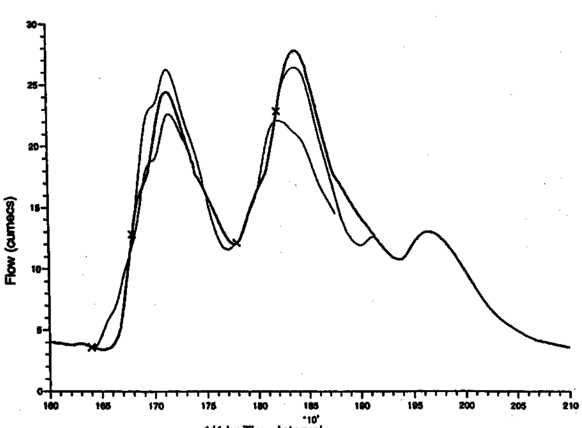

Fig. 8. *Fixed origin forecasts, for the event of 9–11 January 1991 on the Mole catchment, obtained using the SGM with error prediction.*

When the different forecast models are compared, performance in updating mode reflects that obtained in simulation mode, highlighting the fact that updating never fully compensates for a deficient deterministic model structure or unreliable radar data. Thus the lumped PDM model with state updating and using raingauge data performs best overall for the small, rapidly-responding Rhondda catchment. In contrast, for the larger, more variable and slower-responding Wyre catchment the Simple Grid Model with error prediction and radar data provides the best performance. The Mole is somewhat intermediate with the PDM performing best overall but with the SGM with error prediction being marginally better at very short lead times.

Summary and conclusions

The Grid Model is a grid-based rainfall-runoff model developed to exploit the distributed nature of weather

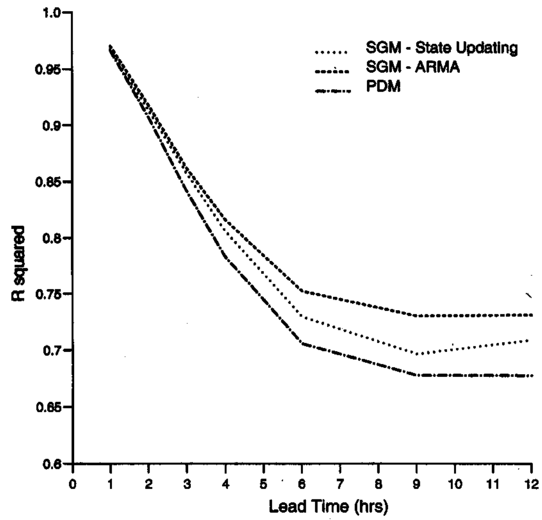

Fig. 9. *Forecast model performance at different lead times: Wyre catchment.*

radar data. The model formulation and its variants are presented in Part 1 (Bell and Moore, 1998) whilst here an assessment of performance is carried out, in simulation and forecast mode, on three catchments in the UK. The catchments are the Mole in the Thames basin near London, the Wyre in north-west England and the Rhondda in South Wales, the latter two catchments being in hillier areas. The results show that when radar data are of good quality, significant model improvement is obtained by replacing data from a single raingauge by grid square radar data. However, the possibility of low level enhancement of rainfall below the beam and blockage effects can adversely affect radar rainfall estimates for the hillier catchments, reducing the accuracy of flow forecasts.

The performance of the different model variants have been assessed using radar data, in distributed and catchment average form, and using raingauge data. It is apparent that for the two hillier catchments the Integrated Air Capacity model and the model with probability-distributed storages often, but not always, result in improved flow forecasts. These formulations are less useful in modelling flows for the Mole catchment, characterised by low relief, perhaps because they are more influenced by slope effects on storage and its impact on runoff production. Forecasts for the Mole are improved slightly through the use of model variants which incorporate separate Darcy-based baseflow isochrones. Results obtained using raingauge data, however, are improved considerably through the use of probability-distributed stores with an increase in R^2 from 0.58 (using optimal isochrones) to 0.67. Use of the topographic index model variant can improve model performance but it never provides the best model for a given catchment and overall gives similar performance to the probability-distributed storage model variant. It performs slightly worse than the probability-distributed storage variant on the Wyre and the Rhondda when raingauge data are used, and slightly better on the Mole and the Wyre when using radar data. The use of Landsat-classified urban areas for defining the impervious fraction is usually slightly helpful in improving model performance. A more traditional lumped rainfall-runoff model, the PDM, is included in the assessment as a benchmark of performance. The PDM proves hard to better, although the Grid Model can provide more reliable simulations for some flood events.

For real-time use a state updating form of the Grid Model is developed and assessed against an error prediction technique for forecast updating. Both state updating and error prediction clearly give improved forecast performance when compared with simulation mode forecasts; however, forecast accuracy is much reduced with increased lead time. An evaluation of the relative benefits of the use of error prediction and state updating suggests that there is little to choose between them: model type has greater impact on forecast accuracy than updating method. However, when the forecast performances obtained for the different catchments are considered together it becomes apparent that in the rapidly responding Rhondda catchment state updating gives slightly better results, whilst in the slower Wyre and Mole catchments, error prediction is slightly superior. A possible reason for this is the difficulty of assignment of errors to routing stores, because of the longer time lags involved in slower responding catchments, and the rapidity with which errors dissipate in faster responding catchments. For slower responding catchments this might favour error prediction, which exploits the dependence structure in previous forecast errors to estimate future errors and so yield a better updated forecast.

The model assessment described here has been limited to three catchments, the use of one or two months of data

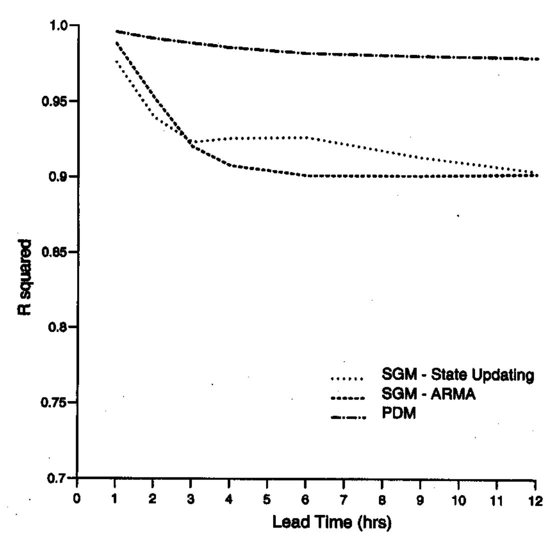

Fig. 10. *Forecast model performance at different lead times: Rhondda catchment.*

from each, and their use for both model calibration and assessment. Conclusions are regarded as indicative whilst firmer ones would involve the use of additional catchments and more extensive data, including the use of separate data for model calibration and assessment. The opportunity for a distributed model to outperform a lumped rainfall-runoff model clearly increases with spatial variability in both rainfall and catchment response. This suggests that further storm events and catchments be selected bearing in mind the synoptic and topographic conditions of each. Firmer guidelines for the operational use of distributed versus lumped models (when and where) might emerge from further investigations of this kind. A well constructed lumped rainfall-runoff model, such as the PDM, may provide the more reliable flood forecasts in widespread stratiform rain and for smaller, more homogeneous catchments. In situations where the scale of the storm is smaller than that of the catchment, typically for convective storms and larger catchments, then a distributed model such as the Grid Model may prove of benefit. This will, in part, depend on the availability of distributed radar data providing reliable rainfall estimates, which will be more likely in catchments of lower relief than hillier ones experiencing blockage effects and precipitation formation below the height of the radar beam. The results obtained here only partly support these conjectures and this highlights the need for more extensive testing on further catchment datasets.

Overall, the results here suggest that a well constructed lumped rainfall-runoff model is to be preferred for routine operational flood forecasting purposes. However, the results also provide evidence that in certain situations a distributed model may prove superior. This suggests that there may be merit in operating a distributed model alongside a lumped model as part of a decision support system approach to flood forecasting and warning. The use of weather radar data may be viewed in a similar context as part of a scenario approach to flood forecasting.

Acknowledgements

This research has been supported by the National Rivers Authority (now part of the UK Environment Agency), the Ministry of Agriculture, Fisheries and Food under its Flood and Coastal Defence Commission, and the Commission of the European Communities under the EPOCH Project 'Weather Radar and Storm and Flood Hazard' and the Environment Programme Project 'Storms, Floods and Radar Hydrology'. Particular thanks are due to George Merrick of the Environment Agency for his role as project co-ordinator. The referees are thanked for ideas leading to improvement in the clarity of presentation.

References

Austin, R.M. and Moore, R.J., 1996 Evaluation of radar rainfall forecasts in real-time flood forecasting models. *Quaderni Di Idronomia Montana*, **16**, (Special Issue, Proc. Workshop on 'Integrating Radar Estimates of Rainfall in Real-Time Flood Forecasting'), 19–28.

Bell, V.A. and Moore, R.J., 1998 A grid-based distributed flood forecasting model for use with weather radar data: Part 1. Formulation, *Hydrol. Earth System Sci.*, **2**, 265–281.

Box, G.E.P. and Jenkins, G.M,. 1970 *Time series analysis forecasting and control.* Holden-Day, 553pp.

Calder, I.R., Harding, R.J. and Rosier, P.T.W., 1983 An objective assessment of soil-moisture deficit models, *J. Hydrol.*, **60**, 329–355.

Gelb, A. ed., 1974 *Applied optimal estimation.* MIT Press, 374 pp.

Gill, P.E., Murray, W. and Wright, M.H., 1981 *Practical Optimisation*, Academic Press.

Hill, F.F., Browning, K.A. and Bader, M.J., 1981 Radar and raingauge observations of orographic rain over south Wales, *Quart. J. Roy. Meteorol. Soc.*, **107**, 643–670.

Institute of Hydrology, 1992 *PDM: A generalised rainfall-runoff model for real-time use.* Developers Training Course, National Rivers Authority River Flow Forecasting System, Version 1.0, Institute of Hydrology, Wallingford, Oxon, UK, 26pp.

Institute of Hydrology, 1996 *A guide to the PDM.* Version 1.0, January 1996, 45pp.

Jazwinski, A.H., 1970 *Stochastic processes and filtering theory*, Academic Press, 376 pp.

Moore, R.J., 1982 Advances in real-time flood forecasting practice, Invited paper, *Symposium on Flood Warning Systems*, Winter Meeting of the River Engineering Section, The Institution of Water Engineers and Scientists, 23pp.

Moore, R.J., 1985 The probability-distributed principle and runoff prediction at point and basin scales, *Hydrol. Sci. J.*, **30**, 273–297.

Moore, R.J., 1993 Real-time flood forecasting systems: perspectives and prospects. *British—Hungarian Workshop on Flood Defence*, 6–10 September 1993, VITUKI, Budapest, 121-166.

Moore, R.J., Austin, R.M. and Carrington, D.S., 1993 *Evaluation of FRONTIERS and Local Radar Rainfall Forecasts for use in Flood Forecasting Models.* R&D Note 225, Research Contractor: Institute of Hydrology, National Rivers Authority, 156pp.

Moore, R.J., Bell, V., Roberts, G.A. and Morris, D.G., 1994 *Development of distributed flood forecasting models using weather radar and digital terrain data,.* R&D Note 252, Research Contractor: Institute of Hydrology, National Rivers Authority, 144pp.

Moore, R.J., Hotchkiss, D.S., Jones, D.A. and Black, K.B., 1992 *London Weather Radar Local Rainfall Forecasting Study: Final Report*, Contract Report to NRA Thames Region. Institute of Hydrology, Wallingford, Oxon, UK, 124pp.

Moore., R.J. and Weiss, G., 1980 Real-time parameter estimation of a nonlinear catchment model using extended Kalman filters. In: Wood, E.F. (ed.), *Real-time forecasting/control of water resource systems*, 83-92, Pergamon Press, Oxford, UK.

Nelder, J.A. and Mead, R., 1965 A simplex method for function minimisation, *Computer Journal*, **7**, 308-313.

Szollosi-Nagy, A., *1976 Introductory remarks on the state space modelling of water resource systems*, Int. Inst. for Applied Systems Analysis, RM-76-73, 81 pp.