

INTERCOMPARISON OF FORECASTING METHODS FOR FLOOD WARNING IN THE RIVER CAM CATCHMENT

D.F. LEKKAS^{1*}
R.T. MAXEY²
H.S. WHEATER¹

¹ *Imperial College London
Dept. of Civil and Environmental Engineering
South Kensington Campus, SW7 2AZ London, UK*
² *Environment Agency, Anglian Central Area,
Bromholme Lane, Brampton
Huntingdon, PE28 4NE, Cambridgeshire, UK*

Selected from papers presented at the 8th Conference
on Environmental Science and Technology,
8 - 10 September 2003, Lemnos, Greece.

*to whom all correspondence should be addressed:
e-mail: d.lekkas@imperial.ac.uk

ABSTRACT

Currently flood warning in the catchment of the River Cam in Cambridgeshire relies on the issuing of alerts when the river level at the monitoring station at Byron's Pool, just upstream of Cambridge, reaches certain pre-determined levels. Warnings are shown to be fairly accurate, but there is very little lead time between the trigger being exceeded and the commencement of flooding. At present there is no method used that can forecast in advance when the trigger is likely to be reached.

Three conceptually different methods of forecasting if and when the trigger at Byron's Pool will be exceeded are presented. The first of these is a simple additive model, in which flows from the three tributaries that are gauged are summed to give a combined flow. The second method involves the derivation and application of two transfer function models capable of transforming river levels on the upstream tributaries to a level at the trigger site. These models are applied both with and without real-time updating techniques. The third method involves the calibration and application of a lumped rainfall-runoff model of the whole catchment to Byron's Pool. Two different calibration periods are used, and the results compared.

The results indicate that the simple additive model, while being better than no model at all, is very inaccurate, and fails to replicate the hydrograph shape and timing, most likely because of the influence of an ungauged tributary. The transfer function models perform well, especially when real-time updating is used. The rainfall-runoff model performs less well, struggling to reproduce the hydrograph shape.

The main conclusions are that for this site a hierarchy of models may be appropriate, with rainfall-runoff models providing an early indication of flooding, and transfer function routing models with updating providing a more accurate forecast, with the additive model as a back up. The importance of obtaining more data, including validation of ratings, and the future gauging of the ungauged tributary, is noted throughout this investigation.

KEYWORDS: Flood warning, forecasting, Additive model, Transfer Functions, NAM model

INTRODUCTION

The impact of flooding ranges from the loss of lives, through destruction of property and loss of agricultural production to disruption of transport and services. To minimise the effects of flooding two complementary approaches exist: a) flood protection works, including the design and construction of river banks, dams and flood storage areas, to protect flood prone areas; and b) flood warning. Effective flood warning can facilitate evacuation of people, property and livestock; amelioration through temporary flood proofing, opportune maintenance and early alerting of emergency services; and control by adjusting reservoir discharges or preparation of balancing ponds.

The effect of the flood warning depends on the accuracy of the forecast used, the conversion of this forecast into a warning for the area affected, and effective interpretation of the warning once it is received (Butts 1995).

All forms of flood forecasting use some form of trigger mechanism to anticipate when the river level at the flood risk area will exceed the flooding threshold. The ability of the trigger to forecast the correct outcome is termed its accuracy. The consistency with which it does this is termed its reliability. A successful trigger is one, which offers sufficient lead-time of warning, with acceptable accuracy and reliability (Cadman and Moore 1998). The simplest form of trigger is based on observations of water level at the site that is at risk from flooding. When the trigger reaches a predetermined level that is less than the threshold, a warning will be triggered.

Setting the trigger lower will increase the warning time, but lower the accuracy and reliability. Usually, a better solution is to have the water level trigger upstream of the risk area. This increases timeliness, and reduces false alarms, however there will be increased uncertainty between the trigger and the threshold. Other triggers, which increase timeliness but decrease accuracy, may be based on rainfall amounts or forecasts. The trade-off between accuracy, reliability and timeliness is the key issue when setting flood warning triggers.

This paper describes an investigation, more fully documented by Maxey (2002), of the performance of a number of types of trigger, of varying degrees of complexity, in order to see how well they perform in forecasting flooding in the Cambridge

area. A number of conceptually different models of the Cam were constructed and compared.

STUDY AREA

The catchment of the River Cam is located in the south east of England, and covers approximately 770 km² in area to Cambridge before entering the fenland river system. The major tributaries of the Cam are the Rhee, Granta and Bourn Brook (see Figure 1), which all join the main river just upstream of the level monitoring station at Byrons Pool, to the south of Cambridge.

River flows are comprised of two principal natural components, namely run-off resulting from rainfall or surface drainage, and baseflow, derived from chalk spring flows. Parts of the Rhee and Bourn Brook catchments are overlain by boulder clay, with the consequence that run off from these areas tends to be rapid. Extensive flooding in this area is not common. In recent years moderate flooding has tended to happen every ten years or so. The most severe flooding in recent years occurred in October 2001, when approximately 80 properties in the Cambridge area flooded. There is an extensive network of rainfall, river flow and river level sites in the catchment. The primary flow and level sites, including those that are (or potentially could be) used for flood warning, are shown on the map in Figure 1.

FLOOD WARNING PROCEDURES

Current flood warning practice is based on trigger levels at Byron's Pool. When the level is reached, an alarm is triggered. Warnings are disseminated via direct contact with households at risk, the media and the emergency services. This system has remained substantially unchanged for many years, apart from a few organisational and technical refinements. At present there is no method for anticipating the trigger, or procedure for early dissemination of warnings. The travel time of flows from the site to the main flood warning area in Cambridge is around 2 hours.

The trigger levels at Byron's Pool were redefined in September 2000 as part of the new Environment Agency flood warning system, which introduced new warning codes and terminology in response to the Bye Report on the floods of 1998 (which did not significantly affect the River Cam) (Bye and Horner 1998). The new trigger levels are given in Table 1.

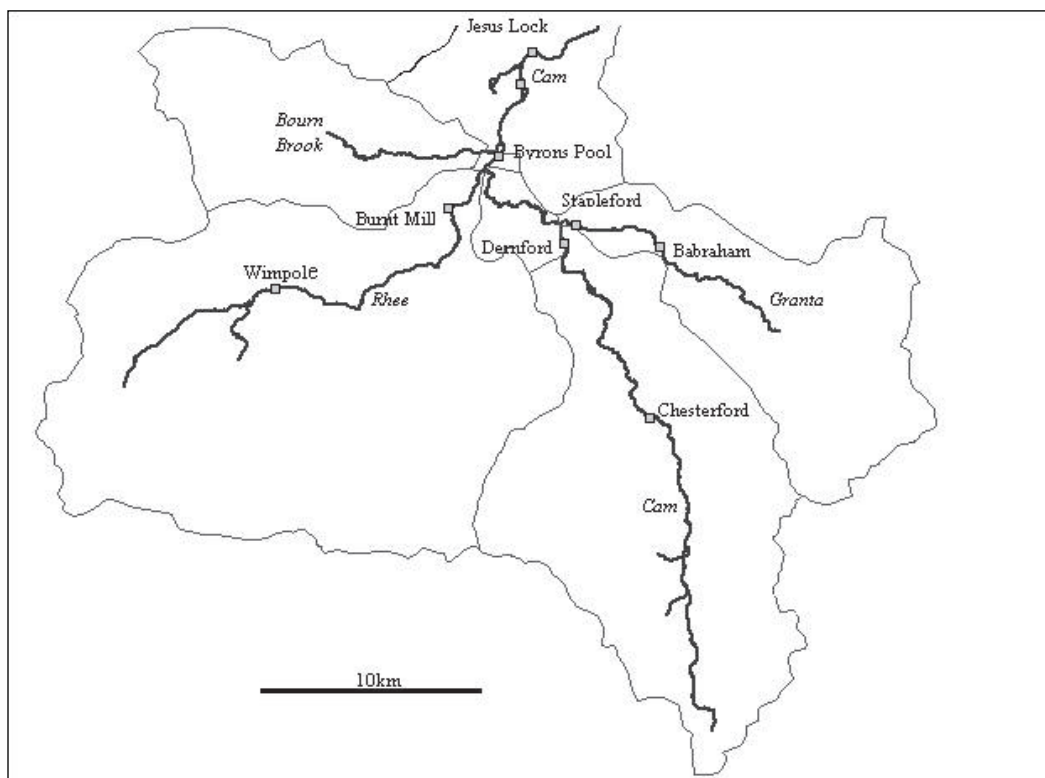


Figure 1. River Cam and its tributaries

These triggers have been derived empirically, using past events to see what level tends to correspond to a particular degree of flooding. In general they perform satisfactorily (Maxey 2002), and there are no immediate plans to change them. The aim of the project was to study how well the reaching of the trigger could be anticipated, using different methods. In the following sections, the methods that were used to predict flood warning triggers are presented.

Simple Additive model

In a flood event it is common for ad hoc real-time forecasts of flow to be made by means of the summing of contributing upstream tributaries in order to estimate the flow downstream of the con-

fluence. The method is intuitive, but it has not been properly evaluated before. These types of forecast are usually made in order to give the forecaster and other duty staff an indication as to the scale of the event; they are not generally used as a method of pre-empting the trigger level exceedence, in order to issue a warning early. However, there are occasions when, due to instrument failure at the downstream gauge, this is the only way of ascertaining what the flow is. In these circumstances, a warning must be issued by forecasting a flow based on the contributing upstream flows. As was described in Section 2 above, the River Cam flow at Byron's Pool is made up of the flow from four separate tributaries (see Figure 1), the Cam, Bourn Brook, Rhee and Granta. Of

Table 1. Flood Warning Triggers at Byron's Pool

Warning	Level (mAOD)	Flow (cumecs)	Meaning
Flood Watch	8.35	17	Flooding may be expected
Flood Warning	8.50	22	Major roads, isolated properties
Flood Warning Update	8.80	31	Property flooding
Severe Flood Warning	9.00	38	Major flooding, danger to life

mAOD: metres above ordnance datum (mean sea level); cumecs: $\text{m}^3 \text{sec}^{-1}$

these, the Cam, Rhee and Granta are gauged at the gauging stations at Dernford, Burnt Mill and Stapleford respectively (in addition to other gauges further upstream). The Bourn Brook, which makes up 10% by area of the total catchment to Byron's Pool, is presently not flow gauged.

When making a simple flow forecast for Byron's Pool, the three gauged tributaries are used, the Bourn Brook being ignored. Assuming the contributing flows have been calculated accurately (not necessarily the case at high flows), this should approximate the flow contributions of those tributaries fairly well, given that the times of travel from the three gauges to Byron's Pool are fairly similar, being of the order of 3 to 5 hours. The lead time of any forecast would thus be of that order.

Transfer Function Modelling

The application of linear Transfer Function (TF) models as the basis of real-time flood forecasting systems is becoming increasingly popular, with many operational systems in place world-wide. The main attractions of TF models are their rapid development times, minimum information requirements, ease of real-time implementation, data assimilation and increased forecasting performance (Lees, 2000; Young and Tomlin, 2000). The general TF model is a linear dynamic model that can be identified and estimated given input-output time-series data collected from the system to be modelled. The most commonly utilised form of TF is a multi-input single-output (MISO) dynamic autoregressive exogenous model (DARX) model,

$$y_t = \sum_{i=1}^r \frac{B_i(z^{-1})}{A(z^{-1})} u_{i,t-\delta_i} + \frac{1}{A(z^{-1})} e_t \quad (1)$$

where δ_i represents the pure time delay between upstream and downstream and

$$A(z^{-1}) = 1 + a_1 z^{-1} + a_2 z^{-2} + \dots + a_n z^{-n}, \quad (2)$$

$$B_i(z^{-1}) = b_{0,i} + b_{1,i} z^{-1} + b_{2,i} z^{-2} + \dots + b_{m,i} z^{-m}$$

Here u_i , $i=1, 2, \dots, r$ are input variables; the triad $[n, m_i + 1, \delta_i]$ defines the model structure; z^{-1} is the backwards shift operator, *i.e.* $z^{-1}y_t = y_{t-1}$; $a_1 \dots a_n, \dots, b_{1,j} \dots b_{m,i}$ and are the model parameters; e_t is a stochastic error term. If $i=1$, then the

model is reduced to a single-input, single-output (SISO) TF model (Lees 2000). SISO models are typically used to model rainfall-runoff relationships, whereas MISO model in flow forecasting application where there are a number of gauging stations on upstream tributaries, as in the case examined here. The TF parameter estimates were obtained using a special optimisation procedure based on the SRIV (Simplified refined instrumental variable) method (Young, 1984).

The Error Prediction (EP) method (Ahsan & O'Connor 1994; Lekkas *et al.* 2001), which used past errors to inform an ARMA error prediction model to forecast future error, was used together with TF models as real-time updating procedure.

NAM model

The third method used is a rainfall-runoff model. NAM, which is an abbreviation of the Danish "Nedbor-Afrstromnings-Model" meaning precipitation-runoff model. It is a deterministic, conceptual, lumped type of model with moderate data input requirements (DHI, 1999). Based on rainfall and evaporation input data, the model produces catchment runoff, which is split conceptually into overland flow, interflow and baseflow components. Because it is a lumped model, each sub-catchment is treated as one unit. The parameters and variables thus represent effective values for the sub-catchment. Being a conceptual model, NAM is based upon physical structures and equations, together with semi-empirical ones. Thus, some of the parameters can be estimated from physical catchment data, but the final parameter estimation must be performed by calibration, employing concurrent input and output time series.

There are 12 key parameters in the NAM model. Three of these are fixed by the user, while the other nine can be subject to automatic calibration.

Hourly rainfall from telemetered tipping bucket gauges were used. Potential evaporation data were input to the model in the form of weekly Real Land Use data from the UK Meteorological Office Rainfall and Evaporation Calculation System (MORECS) (Hough and Jones, 1997). For calibration purposes the corresponding river flow record at Byron's Pool was used.

APPLICATION

The primary test of the effectiveness of each

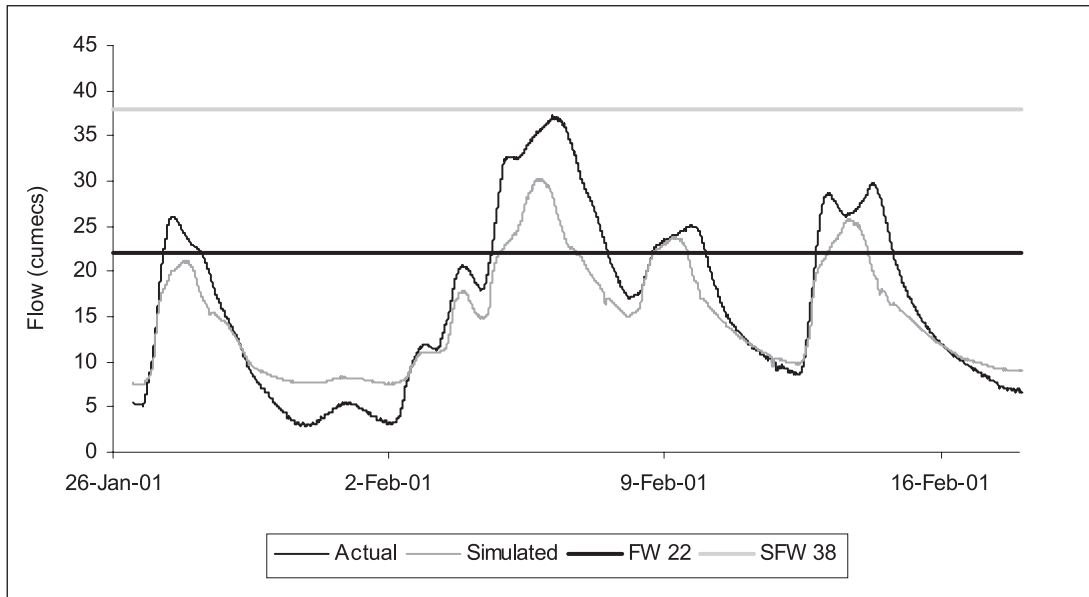


Figure 2. Combined flow at Byron’s Pool (3 hours forecast)

method attempts to indicate is how well it predicts the trigger level exceedance at Byron’s Pool. One way of evaluating this is by a simplified version of the Critical Success Index (Bell, Carrington and Moore 2000). In Tables 2, 3 and 4 below, the Flood/ No Flood columns indicate whether a trigger was reached at any stage of the events. The Flood Warning/ No Flood Warning rows indicate whether the trigger was forecasted by the model. To calibrate the TF and NAM models the coefficient of determination (or efficiency) R_7^2 was used (for details see Young, 1984). All three methods are validated using a succession of events occurring during the period Jan- Feb 01 (see Figures 2, 3, and 4) and for an extreme event during October 2001. This latter event is not illustrated here, although results from it are included in the CSI tables.

Method 1

In order to apply the first, least complex method, combined hydrographs were constructed by the simple addition of the flows recorded at the upstream gauges at Burnt Mill, Dernford and Stapleford. The corresponding hydrographs for Byron’s Pool were then constructed using the gauged levels and the rating. It can be inferred from the results, some of which are plotted in Figure 2 that the combined flow estimate is generally lower than the actual flow, particularly at high flows. The values of combined flow for the

rising limbs appear to be slightly closer to the actual flows, which is encouraging. From the table 2 below, it can be seen that apart from one Severe Flood Warning, all Flood Warnings were forecasted.

In all cases the combined flow underestimated the actual flow, and the trigger was reached too late. It is acknowledged that there may well be considerable error in the value of the actual flow that is used, given that the rating is as yet untested. Nevertheless, the trigger levels have not been derived using flows, and therefore it is valid to use these Byrons Pool flow values as trigger flows, without worrying unduly about their absolute accuracy.

The most probable explanation for the lack of success of the model is that the flow from the Bourn Brook is being ignored. Despite having a small proportion of the total catchment area, it contributes a significant part of the run off in the early part of a flood event, due to its flashy nature, and is responsible for a large proportion of the first of the double peaks that are often seen at Byrons Pool. This portion of the hydrograph is

Table 2. Performance of the Additive method (includes Oct 01 result)

	Flood	No Flood
Flood Warning SFW	xxxx	-
No FW No SFW	X	-

totally lacking in the combined flow model. Thus the forecasting of trigger exceedence tends to occur several hours too late.

The above illustrates that, in this particular catchment, this method should not be used if alternatives are available. In particular it shows the importance of incorporating flows from the most responsive catchment, the Bourn Brook, in circumstances where the rising limb is of most interest.

Method 2

At the time that most of the modelling for this study was carried out (2001), it was not possible to accurately measure an instantaneous flow at any point on the River Cam downstream of the confluence of the main tributaries. It was therefore decided to model the river level at the Byrons Pool gauge. Because this is the parameter that is currently used as the flood warning trigger, and to which all outputs from the various forecasting methods have to be converted, this provided a direct method for anticipating the trigger already used. Furthermore the fact that levels are measured with a good degree of accuracy, whereas high flows are necessarily estimated, sometimes poorly, due to the inadequacy of the gauging station ratings at high flows, resulted in decreased uncertainty in the produced forecast. First a MISO TF model was selected and calibrated. The identified MISO TF was of the form

$$\hat{y}_t = \alpha \hat{y}_{(t-1)} + b_1 u_{b(t-1-30)} + b_2 u_{b(t-2-30)} + c_1 u_{c(t-1-30)} + c_2 u_{c(t-2-30)} + d_1 u_{d(t-1-30)}$$

where \hat{y} is the forecast level at Byrons Pool, u_b , u_c and u_d are the upstream levels at Wimpole, Chesterford and Babraham respectively and b_1 , b_2 , c_1 , c_2 and d_1 are the derived coefficients of the TF model.

The time step used was 15 minutes, which is the sampling rate used by the telemetry system and the interval generally archived. This means that the lag modelled from each upstream site to the downstream site was 7.5 hours (30 time steps).

However, after thorough investigation it was determined that a SISO TF with a combined input $u_t = 0.5(u_b + u_c + u_d)$ was found to give the best performance. Summing the inputs can be seen intuitively to be reasonable for flow modelling; however here we are modelling levels, so there is not necessarily any physical basis for the

format of the combined function. The resulting TF has the following form $\hat{y}_t = \alpha \hat{y}_{(t-1)} + b_1 u_{t(t-1-30)}$. Different calibration and verification periods were used, incorporating the same events from 2000 - 2001 employed in method 1.

Figure 3 shows the results for one of the verification periods using a SISO TF with EP. The success of each model judged against whether or not it forecasts the breaching of a trigger (or falsely forecasts) is summarised for each actual (or forecast) event during the two verification periods in the CSI table 3 below. Initially, it would seem that using an updating function has very little effect on the model's performance, when this criterion is used. However, it can be seen from examination of the results that the runs without EP on the whole over-estimate the peaks, considerably in some cases. When this over-estimation occurs near to a threshold value, it could result in an incorrect warning being given. The overall fit of the non updated values is poor, although this tends to occur more during the falling limb of the peak, and between.

It can be seen that both options (with and without EP) correctly predicted all five of the Flood Warnings that would have been issued. The models wrongly predicted (over-predicted) one Flood Warning (but came fairly close) that did not occur until the following day.

The best performing set-up was the TF that incorporated EP. The basic simulation model provides a forecast lead time of 30 time steps (i.e. 7.5 hours). However, the better performing variant using Error Prediction gives a forecast with a lead time of 20 steps, ie 5 hours. This is the lead time that would be used for the model operators for the purposes of issuing warnings.

Using Error Prediction, the time of trigger level exceedence was forecast well in all cases, with no exceedences failing to be forecast. The modelled levels tended to be higher than those actually reached, and study of the hydrographs reveals that the model struggles to reproduce the detailed shape of the double peak that often occurs at this site. Given that the shape of the peak is dependant on the influence of the ungauged (and therefore unaccounted for) Bourn Brook, this should be expected at times. The timing of the peak is often wrong, but the most important criteria, namely the timing of trigger level exceedence and the value of the peak, were well forecast.




Figure 3. TF & TF+EP level forecast at Byron’s Pool using upstream levels.

The TFs used modelled level to level, ie upstream levels were used to forecast downstream levels. At no point were flows calculated. Given that the reliability of flow estimation at these particular sites diminishes at high flows, and also given that the trigger that is being forecast is a level (not a flow), this would seem to be a more accurate method than modelling flow to flow. But if flows can be calculated accurately it may be that a flow to flow model will prove more reliable, given that there would be a stronger physical basis to the model. However further work would have to be carried out in order to ascertain whether this is in fact the case.

The simulations produced in this project show a great improvement when updating routines are used and we suggest that updating mechanisms should be used wherever possible.

Method 3

The NAM model is evaluated by comparing the simulated runoff with the discharge measure-

ments. Graphical and numerical performance measures are used, the former including comparison of the simulated and observed hydrograph; and comparison of simulated and actual runoff. Another measure that was used is an evaluation of how well the model predicts trigger level exceedence, particularly with regard to timing. Two models were constructed, each using a different calibration period (Calibration 1: 1 Jan 1997 to 31 Dec 2000 & Calibration 2 : 1 Jan 1997 to 31 May 2001), one of which contained quite a large event Jan-Feb 01 event (Calibr. 2).

There was considerable variability in several of the parameters between the two calibrations. The parameters that were the most consistent are the lower zone storage capacity (L_{max}) and the overland and interflow routing time constant (CK1,2), suggesting that these are the best identified parameters. Both calibrations reproduced the overall shape of the hydrographs well, with similar R^2 values over 0.8. It proved difficult to adequately reproduce the overall water balance, with the

Table 3. TF level forecast (SISO) with and without updating (EP)

	TF		TF + EP	
	Flood	No Flood	Flood	No Flood
Flood Warning SFW	xxxxxX	xX	xxxxxX	xX
No FW No SFW	-	-	-	-

error being most significant on Calibration 2, which in other respects proved the best calibration. This was probably due to the poor (over-) simulation of low flows. It was obvious from inspection of the hydrographs that while Calibration 1 could reproduce the calibration period well it was less successful in the verification period, and totally inadequate in its attempt to forecast the October 2001 event. Calibration 2, however, reproduces the calibration period February 2001 well, as would be expected, and, more significantly, produces a much better simulation of the October 2001 event (for further details see Maxey 2002). However, neither calibration can reproduce accurately the double peaked nature of most of the flood events. This is unsurprising in a lumped model that is unable to properly represent the interaction of different tributaries reacting at different times. Ideally, more flood events are required in order to compare the two models more rigorously. Figure 4 shows results from the more successful model (Calibration 2).

Table 4.2 incorporates results from the calibra-

tion period to facilitate comparison with the results from Methods 1 and 2. Nevertheless, it does show that the calibration of the model was much improved by using this high flow period.

It can be seen that the model produced by Calibration 1 did not perform well enough to consistently forecast the flood warning triggers. It under-predicted the peak flows in all cases, sometimes by a considerable amount. Where the forecast peak was just under the trigger level, this meant that no trigger exceedence was forecast. When the extra high flow data are used in calibration, as expected the calibration period is simulated well, and the large event in October 2001 is better simulated. This suggests that as much high flow data as possible are required to properly calibrate the model.

The model presented is not sufficiently accurate to be used as the main prompt to issue a flood warning. However it could be used to give forecasters an idea of the approximate scale of a possible event, enabling them to be better prepared.

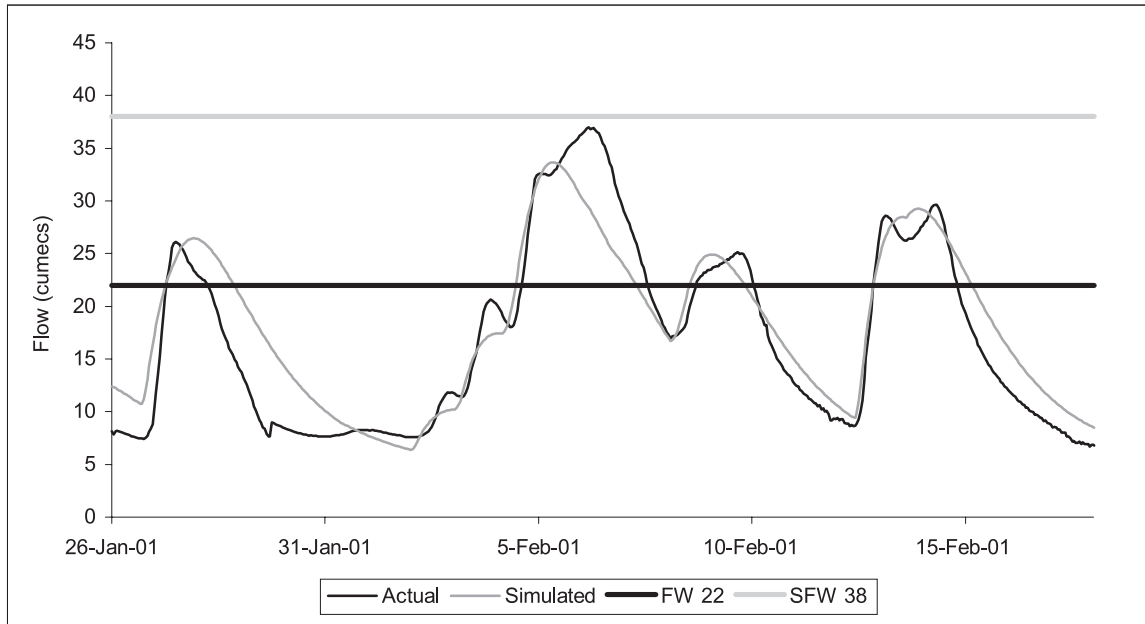


Figure 4. NAM flow forecast at Byron's Pool using rainfall input and Calibration 2.

Table 4.1. Performance of NAM model (Calibration 1)

	Flood	No Flood
Flood Warning SFW	xxx	-
No FW No SFW	xxX	-

Table 4.2. Performance of NAM model (Calibration 2)

	Flood	No Flood
Flood Warning SFW	xxxx	-
No FW No SFW	X	-

CONCLUSIONS

Widely varying methods were used in this study, each with the same ultimate aim: the timely forecasting of the exceedence of the flood warning trigger levels at the Byrons Pool river gauging station. The most obvious difference between the methods lies in the nature of the input data. The simple addition model required upstream flow data, to sum to give a downstream flow. The transfer function (TF) models took upstream level inputs and transformed them to a downstream level. The rainfall-runoff model used rainfall (and evaporation) inputs in order to calculate flow. By specifying the (level-defined) triggers as flows, effectively the methods that had a flow as their output were being used to forecast a level. Given that the outputs were essentially similar, it was possible to evaluate the results for each method in a similar way. A simplified version of the Critical Success Index was used in each case, although for the limited number of events evaluated, this method did not provide enough data for a full evaluation to be made, and reference had to be made to the hydrographs to provide a better picture. Where forecast hydrographs peaked just below a trigger level that had been exceeded, this was represented as a failure of the model; this may have suggested that the model was apparently less successful than it actually was. Analysis of the performance of the models in

terms of the trigger exceedence time and peak values proved useful. In many cases the trigger was forecast to be exceeded many hours after it did in fact occur, and there were often errors in the peak values (although it can be argued that, under the current local arrangements, once a warning has been correctly issued, a forecast peak level or flow, while valuable, is not essential). Given that the main criterion is the accurate simulation of levels, the best method trialled in this project would seem to be the TF method, incorporating real-time updating. A possible hierarchy of modelling might involve the use of rainfall-runoff techniques early in an event (using either conceptual or TF methods) to give an idea of the likely scale of the event, followed by use of the level-to-level TF method with updating, to produce accurate forecasts on which to issue warnings. The simple addition method would be kept as a back up, to be used when all else fails. However, there are several developments that could be made to this basic methodology, and other approaches to forecasting in this catchment, such as semi-distributed rainfall-runoff models, that could be considered.

REFERENCES

- Ahsan, M. and O'Connor, K. M. (1994), A reappraisal of the Kalman filtering technique, as applied in river flow forecasting. *Journal of Hydrology*, 161, 197-226.
- Butts, M. (1995), A Review of the Optimum Accuracy of Flow and Rainfall Forecasting, National Rivers Authority R&D Note 433
- Bye, P. and Horner, P. (1998), Easter 1998 Floods - Final Assessment by the Independent Review Team, Report to Board of the Environment Agency.
- Cadman, D. and Moore, R.J. (1998), A best practice guide to the use of trigger mechanisms in fluvial flood warning; report for the Environment Agency (Draft Final)
- Danish Hydraulic Institute (1999). MIKE 11 NAM Technical Reference Manual.
- Hough, M.N. and Jones R.J.A. (1997), The United Kingdom Meteorological Office rainfall and evaporation calculation system: MORECS version 2.0 - and overview. *Hydrology and Earth System Science*, 1, 227-239.
- Lees, M.J. (2000). Data-based mechanistic modelling and forecasting of hydrological systems, *Journal of Hydroinformatics*, 2, 15 - 34
- Lekkas, D.F., Imrie, C.E. and Lees, M.J., (2001), Improved non-linear transfer function and neural network methods of flow routing for real-time forecasting, *Journal of Hydroinformatics*, 3, 153-164.
- Maxey, R.T. (2002), Review and Intercomparison of Forecasting Methods for Flood Warning in the River Cam Catchment, MSc thesis, Imperial College, London
- Young, P.C. (1984), *Recursive Estimation and Time Series Analysis*, Springer-Verlag, Berlin
- Young, P.C. and Tomlin, C.M. (2000), Data-based mechanistic modelling and adaptive forecasting, from British Hydrological Society Occasional Paper No.12, (eds.) M. Lees and P. Walsh