Hydrological Extremes: Understanding, Predicting, Mitigating (Proceedings of IUGG 99 Symposium HS1, Birmingham, July 1999). IAHS Publ. no. 255, 1999.

Spatial–temporal stochastic rainfall modelling for hydrological design

PAUL J. NORTHROP, RICHARD E. CHANDLER, VALERIE S. ISHAM
Department of Statistical Science, University College London, Gower Street, London WC1E 6BT, UK

CHRISTIAN ONOF, HOWARD S. WHEATER
Department of Civil and Environmental Engineering, Imperial College, Imperial College Road, London SW7 2BU, UK

Abstract The concept of the design storm has its limitations because of its inability to take into account antecedent conditions and the consequent difference between the return periods of the storm and the flood it generates. Stochastic rainfall point models have been developed to model rainfall raingauge data at a single site and, through continuous simulation, provide a new design tool. In the case of many hydrological catchments however, the spatial variability of rainfall becomes crucial to the runoff generation. We present a model in which rain cell origins arrive according to a clustered point process in space and time. The model is fitted using a generalized method of moments. Results are presented for the simulation of individual storms over the Brue catchment in southwest England. These show that the spatial and temporal correlation structures are well reproduced. The model also provides a good fit to the observed marginal intensity and cumulative depth distributions.

INTRODUCTION

Knowledge of the flood with a given return period has long been a main element of the design of flood protection works. In some cases an *a priori* decision is made concerning the desired frequency of flood protection, and the associated flood magnitude is required. More generally, a cost-benefit analysis is undertaken, in which the desired frequency is selected to optimize the ratio of benefits of damage mitigation to the costs of a protection scheme.

Alternative approaches to flood frequency estimation are possible. Frequency analysis of flood flows depends on the availability of suitable site-specific flood data, or appropriate regional methods. The result is usually a simple estimate of flood peak. This is of limited value where a complete hydrograph is required, or where catchment change has taken place, or is proposed.

The main alternative approach is to adopt a rainfall–runoff modelling procedure. This can be relatively simple as in the case of a regionalized unit hydrograph, or more complex, where a site-specific conceptual model is applied. These rainfall–runoff methods produce a complete hydrograph, suitable for example for the design of flood storage schemes, and can in principle incorporate the effect of catchment change.

There are, however a number of problems with conventional rainfall–runoff methods, particularly centred on the concept of the design storm. Firstly, the relationship between the frequency of rainfall and of flows is complex. Alternate temporal and spatial distributions of rainfall can yield the same flood peak, and the state of antecedent wetness on a catchment is another important control on flood generation. In practice, empirical relationships between rainfall and flood frequency, adopting some measure of average catchment wetness, have been derived by simulation (NERC, 1975). However these are specific to the particular design procedure adopted.

Secondly, there is an issue of how temporal variability is included in the design storm. In UK practice the rainfall depth of given frequency is derived, usually from regional analysis, for a specified storm duration. Storm duration is chosen to be a critical duration for a given catchment; for more complex problems, some iteration of assumed values may be required to define this critical duration for a given location. A specified, symmetrical distribution of intensities is then applied to disaggregate the storm depth (NERC, 1975). It is obvious that this pre-specified distribution may not be the most appropriate for the particular application. Elsewhere, it is not uncommon to construct a symmetrical storm profile which preserves the full range of intensity–duration–frequency relationships, i.e. the central 10-min intensity satisfies the required frequency, as does the central 30 min, 1 h, etc. It is evident that this combination represents a much more extreme event than that required, and that the intensity distribution bears no relationship to observed storm properties.

Thirdly, there is a problem in specifying a spatial distribution of storm rainfall. This is commonly achieved by applying an Areal Reduction Factor (ARF) to reduce the storm depth and intensity distribution, to account for the translation of point rainfall statistics to the catchment scale of interest. This should be based on a comparative analysis between point and spatial rainfalls of the same frequency, which, as pointed out by Wheater *et al.* (1989), has not always been the case. In any event, the ARF includes no explicit representation of partial storm coverage, spatial structure of rainfall intensity, storm movement, etc.

In an attempt to overcome the problems of rainfall–runoff modelling outlined above, continuous simulation rainfall–runoff modelling has been proposed (Linsley *et al.*, 1975; Reed, 1994). This approach has the disadvantage of computational expense, and raises questions concerning the regionalization of appropriate rainfall–runoff models, but has several important advantages over the design storm approach. For example, it removes the problem of specifying a design storm duration and frequency, and any hydrograph characteristic of interest can be investigated from the synthetic flow time series. It does, however, require a capability to generate continuous rainfall input time series. Whilst considerable recent progress has been made in stochastic modelling of point rainfall (Rodriguez-Iturbe *et al.*, 1988; Onof & Wheater, 1993; Kakou, 1998), it is desirable to develop the capability for continuous spatial–temporal modelling to overcome the problems of generating spatially consistent rainfall input series. In this paper we present a spatial–temporal model developed with a view to capturing the fine-scale spatial structure of rainfall on a sub-daily level. The model is fitted to radar data, but can, in principle, be fitted to raingauge data. An alternative Generalized Linear Modelling (GLM) approach (Chandler & Wheater, 1998a,b) has been developed to model daily data from a network of raingauges. Future research may seek to combine the merits of these two approaches by using the spatial–temporal model to disaggregate the predictions of daily rainfall totals given by the GLM.

A SPATIAL–TEMPORAL STOCHASTIC RAINFALL MODEL

Our approach to rainfall modelling can be described as *stochastic–mechanistic*. We attempt to incorporate some physical knowledge about the rainfall process into a parsimonious stochastic model that is parameterized in terms of physically meaningful quantities. We consider a model based on a point process of localized areas of relatively intense rainfall or *rain cells*. It is assumed that rain cells arrive according to some random mechanism in space and time and that individual rain cells have random durations, spatial extents and intensities. Observational studies (Petterssen, 1956; Austin & Houze, 1972) have shown that there is a tendency for new rain cells to form in the vicinity of existing cells, so that rain cells tend to cluster within larger-scale structures that we will call *storms*. Storms themselves tend to cluster in similar manner. This hierarchical structure can be reflected by specifying a point process in which cells arrive in clusters to form storms. In turn, we could model clusters of storms or *rain events*, although the spatial scale of the data available to us means that this is unnecessary.

The model is constructed in continuous space and time but can be fitted to both rainfall radar data and raingauge data by evaluating the theoretical properties of the discretized model over the appropriate spatial or temporal scales respectively. The model is fitted to data from the Wardon Hill radar station in southwest England. We use images of resolution 2 km every 5 min over an area of 104 km × 104 km.

Once we have estimated the parameters of the model we are able to simulate realizations of rainfall fields in continuous space and time. A single parameterization can produce many different realizations for use in Monte Carlo simulation studies of catchment response.

Model description

We assume that storm centres arrive in a homogeneous Poisson process of rate λ in two-dimensional space (i.e. in some large region of interest) and time. Following each storm centre, $(\mathbf{u}_0, t_0)$, cell origins arrive in a temporal Bartlett-Lewis-type cluster, i.e. in a Poisson process of rate β starting with a cell located in time at the storm centre, t_0. The process of cell origins terminates after a time which is exponentially distributed with parameter γ. Thus, the number of cells per storm, C, has a geometric distribution with mean $\mu_C = 1 + \beta/\gamma$.

Each cell within the storm is displaced from the storm centre by a vector, which is drawn independently from a bivariate Gaussian distribution with mean $\mathbf{0}$ and covariance matrix

$$\Sigma = \begin{pmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{pmatrix} \tag{1}$$

The components of Σ vary from storm to storm so that, for example, distinct storms have different sizes and shapes. For example, when $\sigma_x = \sigma_y$, storms for which ρ is close to +1 or −1 will tend to have a banded structure, whereas storms for which ρ is near 0 will tend to be circular.

Each cell is elliptical, with a major axis of random length A_c, depositing rain at a constant intensity X on all points in space covered by its defining ellipse during its duration D. The elliptical cell assumption is desirable since spatial autocorrelation plots of radar data often have elliptical contours, even at small spatial lags. For simplicity we assume that each rain cell is a scaled version of the cell displacement density contours of the storm within which it is born (i.e. has the same eccentricity E and orientation Θ) so that, in particular, E and Θ are fixed given Σ. The assumption that the rain cell's intensity X is constant over its area and duration is made for simplicity. Although the results do not appear to be sensitive to this assumption, it is possible to consider other shapes for the cell intensity function or to insert irregularity in the form of a high-frequency "jitter" (Rodriguez-Iturbe *et al.*, 1987).

The total rainfall intensity at a spatial location $\mathbf{u}$ at time t, $Y(\mathbf{u}, t)$, is the sum of the contributions from all cells active at $(\mathbf{u}, t)$. The variables A_c, D and X are assumed to be mutually independent between cells and independent of Σ. All cells within a storm and the storm centre itself move with the same random velocity V. Thus, all cells are born within the existing structure of the storm and the storm shape is not distorted by the motion of the storm. Cell clusters belonging to distinct storms are independent. We shall refer to this model as the Gaussian displacements spatial–temporal model (GDSTM).

A variant of this model which has also been studied is the random ellipse spatial–temporal model (RESTM), in which the spatial displacement of rain cells about the storm centre has a uniform distribution over an ellipse with a major axis of random length A. In the special case (of the RESTM and of the GDSTM) when $\beta = 0$, so that each storm consists of a single rain cell, the model is that of Cox & Isham (1988) with the generalization that the rain cells are elliptical rather than circular.

Model properties

We wish to fit the model to temporal sequences of radar images. Each radar image consists of a regular array of values, each value giving the average rainfall intensity over a pixel of dimensions h km × h km. Typically $h = 2$ or 5 and the temporal separation of images is 5 min. We need to derive properties of the spatially-averaged rainfall process. If $Y_{i,j}^{(h)}(t)$ is the average rainfall intensity over pixel (i, j), of dimensions h km × h km at time t, then

$$Y_{i,j}^{(h)}(t) = \frac{1}{h^2} \int_{(i-1)h}^{ih} \int_{(j-1)h}^{jh} Y(\mathbf{u}, t)\,\mathrm{d}u_1 \mathrm{d}u_2 \qquad (2)$$

The mean of the continuous rainfall process is given by the product of the cell arrival rate ($\lambda\mu_C$), and the expected area (μ_A), duration (μ_D) and intensity (μ_X) of a rain cell. Therefore, the mean of the spatially-averaged process is given by

$$E\left[Y_{i,j}^{(h)}(t)\right] = \lambda\mu_C\mu_A\mu_D\mu_X \qquad (3)$$

We have derived an approximate expression for the covariance function of the continuous space–time rainfall process. This is based upon approximations which are required to evaluate the expected area of intersection of two rain cells and to evaluate integrals that arise (Northrop, 1998).

The covariance function $c(\mathbf{u},t)=\mathrm{cov}[Y(\mathbf{0},0),Y(\mathbf{u},t)]$ of the continuous space–time rainfall process is given by

$$c(\mathbf{u},t)=\frac{\lambda\mu_C}{\eta}E[X^2]e^{-\eta t}E_{A_c}[A_c^2C(\mathbf{u}/A_c)]+\frac{\lambda\beta\mu_C\mu_X^2\psi\pi}{4\eta(\gamma+\eta)\varsigma}e^{-\gamma t}g(\mathbf{u}) \tag{4}$$

where $C(x)$ is the area of intersection of two discs of common unit radius whose centres are a distance x apart and $g(\mathbf{u})$ involves definite integrals of bivariate normal densities. In the interests of computational efficiency we make use of approximations in evaluating $C(x)$ and $g(\mathbf{u})$.

Aggregation of the covariance function over a radar pixel requires the numerical evaluation of a double integral. A full account of the derivation of these properties can be found in Northrop (1998).

Further properties

Other properties of the model are of potential interest for use in parameter estimation and assessing goodness of fit, particularly those that relate to the wet/dry pattern of rainfall, but most seem inaccessible to analytic study. For example, the aggregation of wet/dry properties over space creates difficulties. It is possible to derive an expression for the probability that an arbitrarily chosen point in space-time is dry, although it is only possible to extend this to consider an arbitrarily chosen pixel in the case where $\beta = 0$ (i.e. the EPPM). Properties of rainfall fields to which thresholds have been applied (i.e. rainfall intensities below a given level are set to zero) may also be of interest. Simulation from a fitted model can be used to assess the model's performance with respect to its ability to reproduce these properties.

PARAMETER ESTIMATION

We give a brief outline of a parameter estimation method and an example to illustrate the use of the model. The model is stationary in time and homogeneous in space. We therefore restrict attention to the interior of rain events in both space and time, i.e. without including the development and dissipation of rain events in time or their periphery in space. We assume that $V(=\mathbf{v})$, $E(=e)$ and $\Theta(=\theta)$ are fixed and common to all cells, an assumption which is reasonable over the relatively small temporal and spatial scales involved.

The complex dependencies produced by the structure of the model mean that a maximum likelihood approach is not feasible (it is difficult to obtain a likelihood in useful form) and, indeed, not appropriate due to the fact that the cell intensity structure leads to short-term deterministic features. A more subjective approach is to select properties of the data which are regarded as being important and find values for the parameters that produce as close a fit as possible to these properties. This is in the spirit of the Generalized Method of Moments of Hansen (1982) in which a weighted sum of squared differences between observed values of selected properties and their model values is minimized numerically with respect to the parameters of the model. The properties used are chosen to enable the parameters of the model to be estimated reliably. The common cell/storm velocity $\mathbf{v}$ can be estimated directly from the data using a cross-correlation method, and e and θ can be estimated using the spatial autocorrelation function. The computing time required to estimate the remaining parameters is thus reduced. Properties not used for fitting are used to assess the adequacy of fit of the model.

The parameter estimates produced by this method are conditional on the sample properties selected for fitting and are therefore not unique. It is possible that quite different sets of parameter estimates may give very similar fits to the sample properties chosen. This parameter identification problem may be compounded further by difficulties in finding the global minimum of the sum of squared differences. A grid search of the parameter space, subject to certain physical constraints, (for example, that cells are of lesser spatial and temporal extent than storms) is conducted to find sets of initial estimates. It seems that the prior estimation of $\mathbf{v}$, e and θ alleviates the problem of multiple minima by distinguishing, for example, between large storms moving quickly and small storms moving slowly.

We present results from the fitting of the GDSTM to a one hour sequence (13 images) of radar images recorded by the Wardon Hill radar station on 6 February 1994. It is assumed that X and A_c follow exponential and gamma distributions respectively.

Table 1 gives the estimated values of the parameters of the model. The cell and storm dimensions are in broad agreement with those given by Austin & Houze, 1972.

Table 2 shows the observed and fitted values of a variety of properties over a range of levels of spatial aggregation. The fitted values for the proportion of h km × h km pixels in an image that are wet are estimated using a sequence of images simulated from the fitted model. The close agreement between many properties not used for parameter estimation indicates that the model is performing well.

Figure 1 illustrates the goodness of fit of the model in terms of the temporal and spatial autocorrelation functions (at the 2 km scale of spatial aggregation) and the space–time variance across scales of spatial aggregation h. There is good agreement between the observed and fitted values of these functions, the latter plot (in conjunction with Table 2) demonstrating that the model is able to reproduce properties of the empirical data over a range of pixel sizes.

Table 1 Parameter estimates: GDSTM fitted to the storm of 6 February 1994, 13:00–14:00.

Model property	Fitted value
Cell arrival rate (km^{-2} h^{-1})	0.00034
Expected cell duration (min)	39
Expected cell area (km^2)	23
Expected cell intensity (mm h^{-1})	0.54
Expected number of cells per storm	650
Expected storm duration (h)	2.7
Expected storm area (km^2)	310
x component of cell/storm velocity (km h^{-1})	36
y component of cell/storm velocity (km h^{-1})	34
Cell/storm eccentricity	0.83
Cell/storm orientation (degrees)	94

Table 2 Rainfall recorded by the Wardon Hill radar 6 February 1994, 1300-1400 Observed and (in parantheses) fitted values under the GDSTM.

Property		Level of spatial aggregation (km) $h=2$	$h=4$	$h=8$	$h=16$
mean (mm h^{-1})		1.78† (1.78)			
variance (mm^2 h^{-2})		3.65† (3.65)	3.18 (3.22)	2.59† (2.59)	1.92† (1.93)
$\hat{\rho}^{(h)}(0,0,t)$*	$t=1$	0.66 (0.62)	0.74 (0.71)	0.84 (0.84)	0.93 (0.91)
	$t=2$	0.45† (0.45)	0.51 (0.51)	0.61 (0.63)	0.77 (0.76)
	$t=3$	0.32 (0.33)	0.36 (0.37)	0.43 (0.45)	0.59 (0.58)
$\hat{\rho}^{(h)}(x,0,0)$	$x=1$	0.83† (0.82)	0.76 (0.72)	0.65 (0.63)	0.36 (0.41)
	$x=4$	0.45† (0.45)	0.22 (0.22)	−0.23 (0.01)	−0.24 (0.00)
$\hat{\rho}^{(h)}(0,y,0)$	$y=1$	0.86 (0.92)	0.81 (0.84)	0.76 (0.78)	0.66 (0.67)
	$y=4$	0.56 (0.57)	0.45 (0.47)	0.33 (0.20)	−0.03 (0.01)
$\hat{\rho}^{(h)}(x,0,1)$	$x=-1$	0.56 (0.52)	0.53 (0.51)	0.47 (0.46)	0.22 (0.27)
	$x=12$	0.72 (0.77)	0.80 (0.84)	0.78 (0.77)	0.51 (0.53)
$\hat{\rho}^{(h)}(0,y,1)$	$y=-1$	0.59 (0.57)	0.60 (0.60)	0.60 (0.63)	0.56 (0.54)
	$y=1$	0.69 (0.66)	0.79 (0.76)	0.82 (0.80)	0.70 (0.71)
$\hat{p}_h$		0.92 (0.93)	0.97 (0.97)	1.00 (0.99)	1.00 (1.00)

* $\rho^{(h)}(x,y,t)$ denotes the autocorrelation at spatial lag (hx, hy) km and time lag $5t$ minutes;
† denotes a property that was included in model fitting.

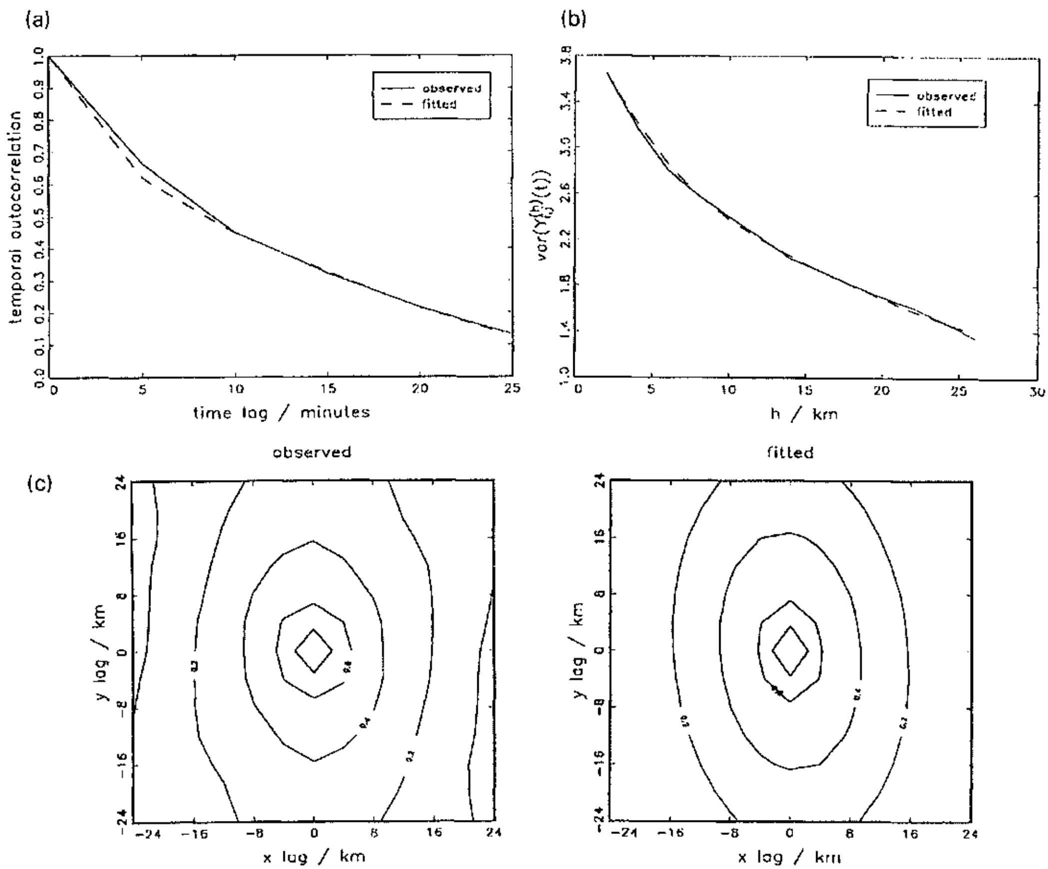

Fig. 1 Assessment of the fit of the GDSTM to the storm of 6 February 1994, 13:00–14:00. (a) temporal autocorrelation; (b) scaling of variance; (c) spatial autocorrelation.

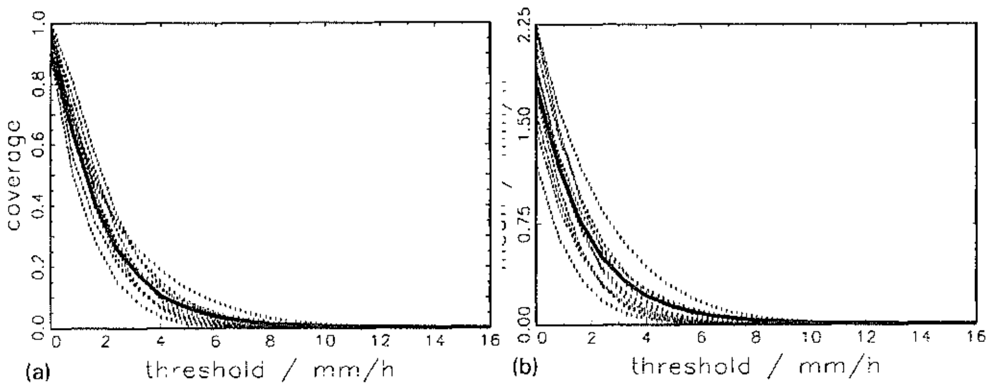

Fig. 2 Threshold analysis of the intensity field: (a) coverage *vs* threshold, (b) mean intensity of thresholded image *vs* threshold. Solid line is the empirical data curve, dotted lines are 10 independent simulations from the fitted model.

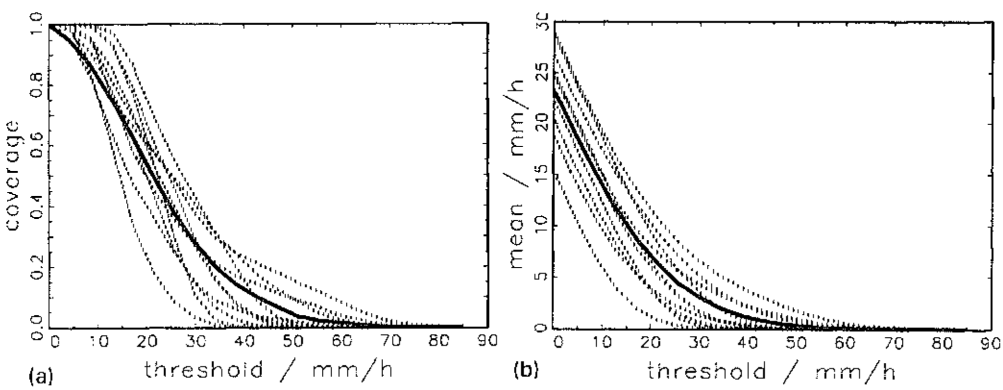

Fig. 3 Threshold analysis of the depth field: (a) coverage *vs* threshold, (b) mean depth of thresholded image *vs* threshold. Solid line is the empirical data curve, dotted lines are 10 independent simulations from the fitted model.

As a further assessment of model fit we investigate the effect of imposing thresholds on the empirical data and on data simulated from the fitted model. To impose a threshold of 1 mm h^{-1} on an image we simply set all rainfall intensities below 1 mm h^{-1} to zero. In particular we consider the proportion of pixels with a rainfall intensity over a given threshold, the *coverage*, and the mean rainfall intensity of images that have been thresholded. Figure 2 shows how the coverage and the mean intensity of empirical and simulated images is related to the threshold imposed. Each curve has been estimated by averaging the results from a 1-h sequence (i.e. 13 images) of data. The curves derived from the empirical data are generally consistent with the curves derived from the simulated data although the empirical data tends to have a greater coverage for large thresholds than the model. A further investigation examines the rainfall totals accumulated at each pixel over the 1-h sequence of data. We define the rainfall *depth* at a given pixel to be the sum of all the rainfall intensities observed at the pixel over a time period of interest. Figure 3 shows how the coverage and the mean depth of the empirical and simulated images is related to the threshold imposed. These plots exhibit broadly similar features to Fig. 2.

Figure 4 gives a visual illustration of the ability of the model to reproduce the internal structure of rain events. We have included images from outside the time period used for parameter estimation, to illustrate the movement of the rain event across the area. The images have been simulated from the fitted model in a manner intended to mimic the movement of the data. This is achieved by defining a rain event area within which the model process is simulated and moving this area with the same velocity as that of the storms within it. The radar images and the realization simulated from the fitted model exhibit broadly similar features, although the assumption of constant

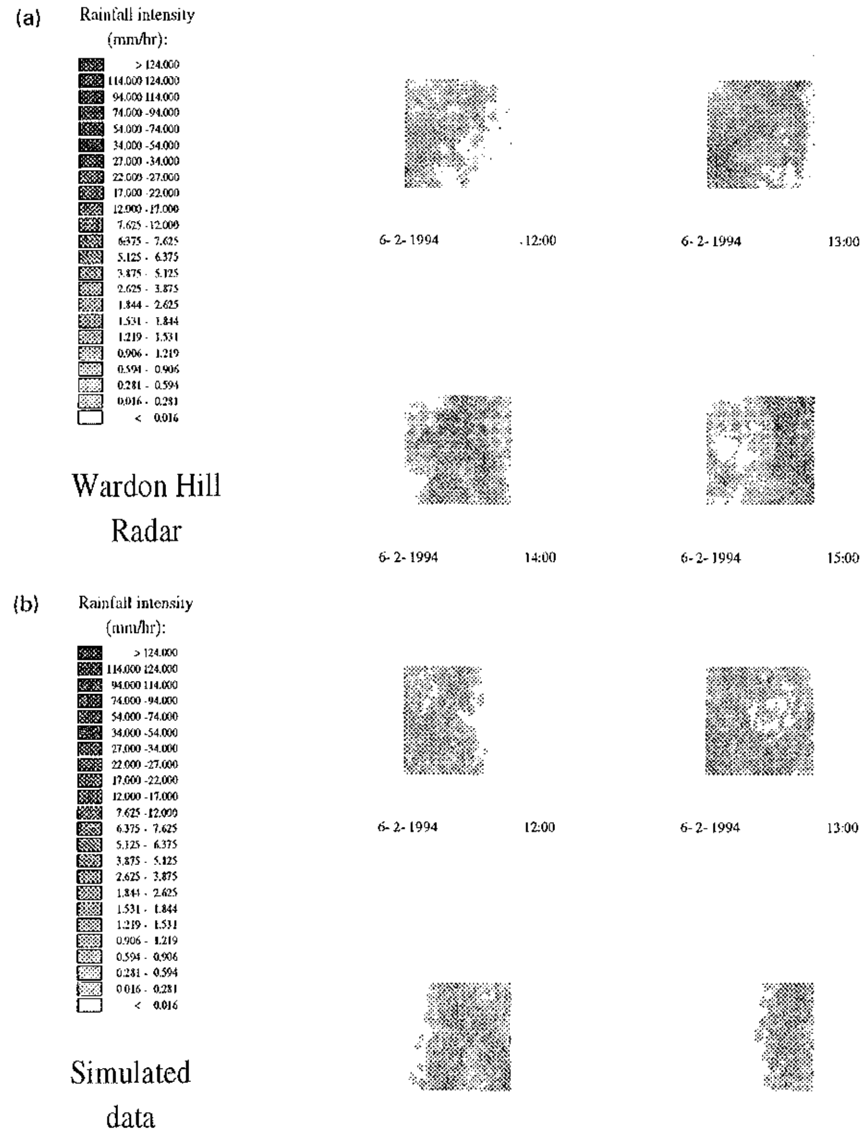

Fig. 4 Visual assessment of the GDSTM fitted to the storm of 6 February 1994, 13:00–14:00: (a) observed data, and (b) simulation. The temporal separation of the images is 1 h.

rainfall intensity over the area of a rain cell means that data simulated from the model is less smooth than the empirical data. The peripheral drizzle apparent in the empirical data is not present in the simulated data. We may have to address this feature if it becomes clear that the drizzle surrounding rain events is important hydrologically. It remains to be seen whether output from rainfall–runoff models is sensitive to this feature of the simulated data. Close agreement between the two sequences in terms of exactly *where* areas of relatively intense rainfall occur is not expected, just as close agreement between two realizations from the same stochastic model would not be expected.

CONCLUSIONS

Simulation of rainfall fields over longer time periods, is required ultimately for input into distributed rainfall–runoff models. Therefore, the model has to be extended to deal with the "birth" and "death" of rain events in space and time and issues such as seasonality. Current research is promising and aims to use the parameter sets estimated from individual rain events to study the seasonal dependence of parameters. Continuous simulation will be achieved by constructing a temporal process of rain events arrivals and generating the parameters of a given rain event in a manner that reflects both the inter-dependence of individual parameters and their dependence on the season. Future developments will assess the effects of non-stationarity in space.

The radar data used so far covers a time period of five years and long-term daily raingauge records are required to investigate long-term performance. Model calibration using raingauge data will be investigated and the current model can be tested for consistency with long-term records. However, future research will also investigate the possibility of linking the spatial–temporal model with alternative approaches, in particular the GLM used by Chandler & Wheater (1998a,b) to model daily rainfall data from the west of Ireland. This model relates the daily rainfall total at a given gauge to various explanatory variables such as rainfall totals from previous days, the month of the year, gauge elevation and gauge location. Additionally, trend functions of different forms can be fitted and compared in an attempt to determine the presence/nature of long-term temporal variability in the area. A hybrid approach would ideally maximize the use of information from long-term data and sparse raingauge networks with a continuous spatial–temporal disaggregation to preserve the fine-scale detail.

Acknowledgement Support of NERC Grant no. GR3/10916 is duly acknowledged.

REFERENCES

Austin, P. M. & Houze, R.-A. (1972) Analysis of structure of precipitation patterns in New England. *J. Appl. Meteorol.* 11, 926–935.

Chandler, R. E. & Wheater, H. S. (1998a) Climate change detection using Generalized Linear models for rainfall – a case study from the West of Ireland. I Preliminary analysis and modelling of rainfall occurrence. *Research report no. 194, Department of Statistical Science, University College London.*

Chandler, R. E. & Wheater, H. S. (1998b) Climate change detection using Generalized Linear models for rainfall—a case study from the West of Ireland. II Modelling of rainfall amounts on wet days. *Research report no. 195, Department of Statistical Science, University College London.*

Cox, D. R. & Isham, V. (1988) A simple spatial–temporal model of rainfall. *Proc. Roy. Soc. Lond.* **A415**, 317–328.

Hansen, L. R. (1982) Large sample properties of generalized method of moments estimators. *Econometrica* **50**, 1029–1054.

Kakou, A. (1998) A point process based model for rainfall with dependent duration and intensity. *J. Hydrol.* (submitted).

Linsley, R. K., Kohler, M. A. & Paulhus, J. L. H. (1975) *Hydrology for Engineers*, 2nd edition. McGraw-Hill, New York.

NERC (Natural Environment Research Council) (1975) *Flood Studies Report.* Natural Environment Research Council, London.

Northrop, P. (1998) A clustered spatial-temporal model of rainfall. *Proc. Roy. Soc. Lond.* **A454**, 1875–1888.

Onof, C. & Wheater, H. S. (1993) Modelling of British rainfall using a random parameter Bartlett-Lewis rectangular pulse model. *J. Hydrol.* **149**, 67–95.

Petterssen, S. (1956) *Weather Analysis and Forecasting*, vol. 2. McGraw-Hill, New York.

Reed, D. W. (1994) Rainfall frequency analysis for flood design. In: *Coping with Floods* (Proc. NATO Advanced Study Institute, Erice, Italy) (ed. by G. Rossi, N. Harmancioglu & V. Yevjevich), Series E, Applied Sciences, **257**, 59–76. Kluwer Academic Publishers, Dordrecht, The Netherlands.

Rodriguez-Iturbe, I., Cox, D. R. & Isham, V. (1987) Some models for rainfall based on stochastic point processes. *Proc. Roy. Soc. Lond.* **A410**, 269–288.

Rodriguez-Iturbe, I., Cox, D. R. & Isham, V. (1988) A point process model for rainfall: further developments. *Proc. Roy. Soc. Lond.* **A417**, 283–298.

Wheater, H. S., Larentis, P. & Hamilton, G. S. (1989) Design rainfall characteristics for SW Saudi Arabia. *Proc. Instn Civ. Engnrs*, Part 2, **87**, 517–538.