HAL
open science

Spatial-temporal rainfall fields: modelling and statistical aspects

H. S. Wheater, V. S. Isham, D. R. Cox, R. E. Chandler, A. Kakou, P. J. Northrop, L. Oh, C. Onof, I. Rodriguez-Iturbe

To cite this version:

H. S. Wheater, V. S. Isham, D. R. Cox, R. E. Chandler, A. Kakou, et al.. Spatial-temporal rainfall fields: modelling and statistical aspects. Hydrology and Earth System Sciences Discussions, 2000, 4 (4), pp.581-601. hal-00304688

HAL Id: hal-00304688

https://hal.science/hal-00304688v1

Submitted on 18 Jun 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Hydrology and Earth System Sciences, 4(4), *581–601* (2000) © EGS

Hydrology & Earth System Sciences

Spatial-temporal rainfall fields: modelling and statistical aspects

H.S. Wheater[1], V.S. Isham[2], D.R. Cox[3], R.E. Chandler[2], A. Kakou[2], P.J. Northrop[2], L. Oh[1], C. Onof[1] and I. Rodriguez-Iturbe[4]

[1]Department of Civil Engineering, Imperial College, Imperial College Road, London SW7 2BU, UK
[2]Department of Statistical Science, University College London, Gower Street, London WC1E 6BT, UK
[3]Nuffield College, Oxford OX1 1NF, UK
[4]Department of Civil and Environmental Engineering, Princeton University, Princeton, N.J. 08544, USA
e-mail for corresponding author: h.wheater@ic.ac.uk

Abstract

The HYREX experiment has provided a data set unique in the UK, with a dense network of raingauges available for studying the rainfall at a fine local scale and a network of radar stations allowing detailed examination of the spatial and temporal structure of rainfall at larger scales. In this paper, the properties and characteristics of the rainfall process, as measured by the HYREX recording network of rainguages and radars, are studied from a statistical perspective. The results of these analyses are used to develop various models of the rainfall process, for use in hydrological applications. Some typical results of these various modelling exercises are presented.

Keywords: Rainfall statistics, rainfall models, hydrological design.

Introduction

For most hydrological systems, rainfall is the primary input, and hydrologists have long recognised the problems of interpolation of "point" raingauge data to estimate spatial rainfall fields. Indeed, the dramatic spatial heterogeneity of convective rainfall in some arid areas has led to the situation where, even with a relatively dense raingauge network (~10 km inter-gauge spacing), flood flows can arise with no rainfall observed (Wheater *et al.*, 1991a, b). Ambiguity in definition of spatial rainfall in such areas has been shown to be a major factor in failure of conventional rainfall-runoff analyses (Wheater and Brown, 1989). Localised convective rainfall is relatively rare in the UK, but quantitative insight into the internal structure of rainfall fields generated by the more typical frontal and convective systems is nevertheless lacking.

Since the implementation of the Dee Weather Radar Project in the 1970s (CWPU, 1977), recognition of the power of radar as a tool for measurement of spatial rainfall fields has led to the progressive implementation of a national network of weather radars in the UK. However, while the importance of these data for weather forecasting and real-time flood estimation has long been apparent, the wider implications for spatial rainfall analysis have received relatively little attention. The HYREX experiment has provided the opportunity, via both a dense network of raingauges and associated radar data, to explore a range of issues related to the characterisation of spatial rainfall and the implications for various aspects of hydrological practice.

Firstly the basic spatial properties of the raingauge data are studied. The HYREX raingauge network, of 49 0.2 mm tipping bucket raingauges, located in a grid of 28 2×2 km^2 squares in the catchment of the river Brue in South West England, is more fully described elsewhere in this volume (Moore *et al.*, 2000); maps showing the location of the Brue, and the raingauge network, are shown in Fig. 1. Although designed primarily to provide the basis for a detailed evaluation of 2 km resolution radar data, it provides a data set unique in the UK. Two earlier dense experimental networks, established at Cardington and at Winchcombe between 1957 and 1967, received limited analysis (Holland, 1967; Marshall, 1980) and have long since been discontinued; the data are believed to be no longer available. The HYREX network encompasses a modest elevation range, from 35 m to 190 m above sea level, and can provide insight into small-scale topographic controls on rainfall fields, as well as the internal structure of rainfall at the network scale.

Secondly, the development of models of spatial rainfall, with respect to both full spatial-temporal modelling and the multi-site problem of modelling correlated raingauge data are considered. Spatial-temporal models can provide a succinct summary of the data, and are a primary means of generalisation for hydrological applications. Here they are

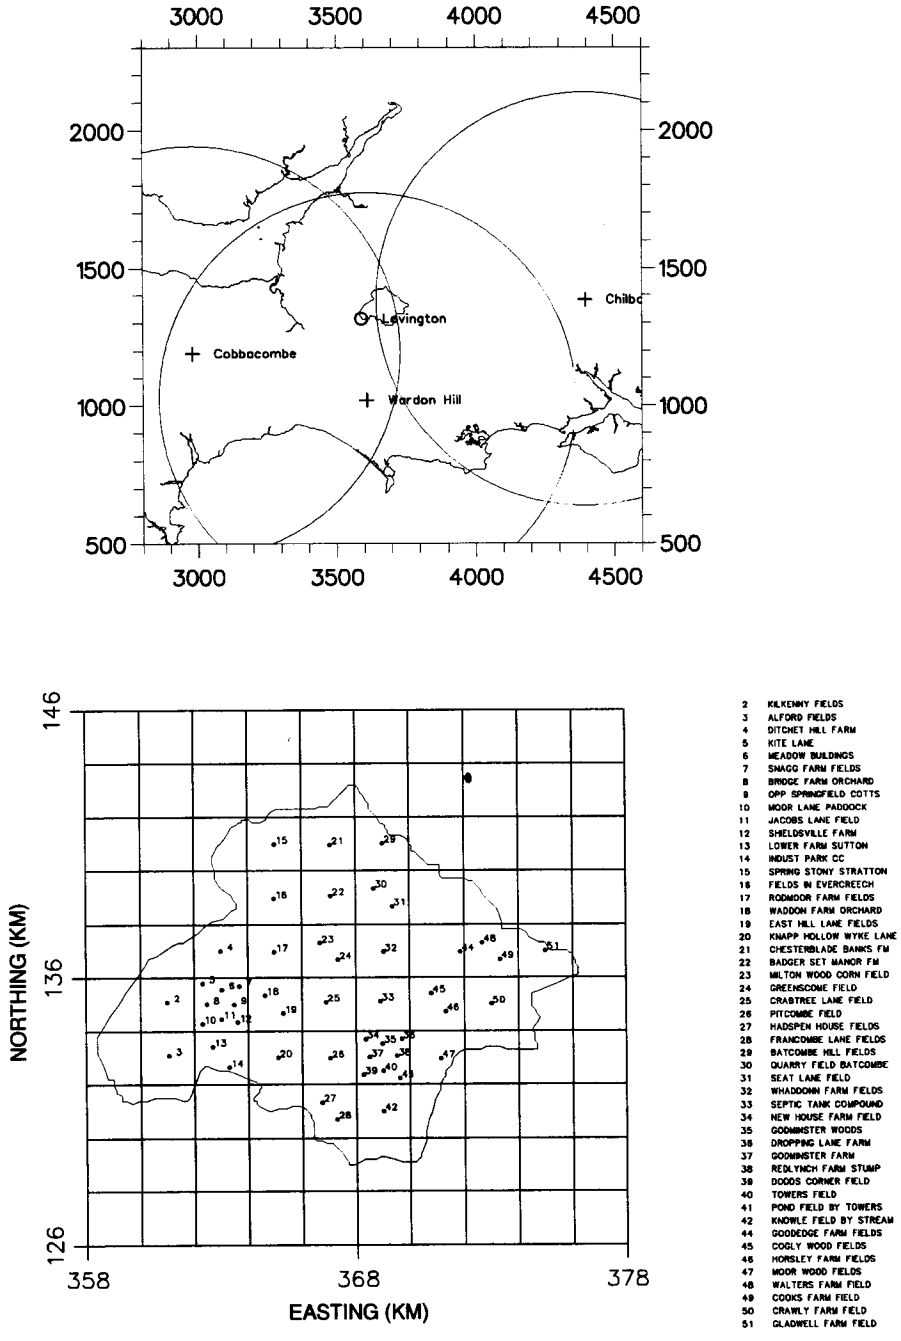

Fig. 1. *Map of the HYREX experiment. Top: location of Brue valley (outlined in centre of map, near Lovington) within South-West England, with locations and ranges of HYREX radar stations. Bottom: detail of Brue catchment, showing positions of HYREX raingauges.*

developed using individual radars which provide $2 \times 2\ km^2$ pixel resolution within a circle of radius 76 km and $5 \times 5\ km^2$ resolution to 210 km. Issues of radar calibration are discussed elsewhere (Wood *et al.*, 2000b). It is assumed for the present stage of model development that possible calibration errors will not bias seriously the spatial structure of the observed fields, once obvious errors have been removed. Multi-site modelling is based mainly on the raingauge network.

The implications of spatial rainfall distribution for large-scale hydrological modelling, focuses on the problems of disaggregation which arise when relating the large-scale average rainfall of a grid-square simulation in an Atmospheric General Circulation Model to parameterisations of the surface hydrology. Observed large-scale rainfall fields provide a means of testing algorithms and quantifying parameters for current GCM schemes, and analysis can point to potential methodological improvements.

Finally, the potential for a new family of design methods for hydrological application, is considered, in a preliminary

way, recognising the spatial and temporal variability in rainfall through the application of stochastic simulation methods.

Data analysis

To understand better the mechanics of the rainfall process, and to guide the development of statistical models, basic analyses of the data have been carried out routinely, on a monthly basis, throughout the HYREX experiment; these provide an archive of summary statistical information about the process. In addition, more detailed analyses of particular time periods or rain events have been used to address specific issues. The problem of calibrating radar measurements against data from the raingauge network has not been addressed, as this is being investigated elsewhere in the HYREX project.

ROUTINE ANALYSES

Most of the routine statistical analysis has been carried out using raingauge data from the Brue catchment. Primary quality control of the data has been carried out by the Institute of Hydrology and is reported elsewhere in this volume (Wood *et al.*, 2000a); this is supplemented by the data analysis reported here. For example, bad data may be identified by inspection of monthly cumulative hyetographs, although assessment of quality is more difficult in the summer months when the rainfall pattern is naturally more variable than in winter. If a gauge were suspect during a particular month, then it was excluded from any subsequent analysis for that month.

For each month, summary statistics have been computed at different levels of temporal aggregation for each of the reliable gauges, and the results tabulated (examples may be found in Wheater *et al.*, 1997). The mechanism of the gauges limits the temporal resolution of the raingauge data—the median time between bucket tips varies from around 5 minutes in the winter to 11–12 minutes in the summer, which means that for aggregation periods of less than about 15 minutes the data discretisation due to the tipping mechanism has a serious effect. This clearly has important implications for any temporal analysis which is carried out on these data. Seasonal variation in these summary statistics is much as expected – for example, the proportion of wet days is much lower in summer than in winter. Winter statistics are much more homogeneous across the network than those for summer months, which indicates less spatial variability in the winter.

Analysis of radar data has focused on the Wardon Hill radar station (rather than that at Cobbacombe), as this is closer to the Brue catchment and preliminary estimates of radar accuracy indicate that its performance is better (Wood *et al.*, 2000a). However, problems do occur and have been screened by visual inspection. Attempts to classify the radar images into different climate types have relied largely on subjective visual analysis (although some work on the development of an objective and statistically-based classification scheme has been carried out (Imrie, 1996)). An examination of the data over the period of the HYREX experiment reveals three distinct rainfall patterns: these correspond broadly to showers, rainbands and regions of widespread rain (more detailed examination suggests some kind of hierarchical organisation is taking place, with banded structures frequently visible within regions of widespread rain and 'rain cells' within the bands). The 'showers' category can be further subdivided into 'scattered' and 'widespread' showers. Sample images of each type can be found in Fig. 2. Notice the data error around the edge of the image for the event of 30/12/93, where nothing is recorded due to the rain being beneath the radar beam at this distance from the radar station – a common source of data error in winter (Collier, 1989).

Routine analysis of the radar data has focused on the spatial variability of the rainfall. Again, data quality is an important consideration and these analyses are restricted to data at $2 \times 2\ \mathrm{km}^2$ spatial resolution lying within the small circle (radius 76 km) shown in Fig. 2. For each radar image, spatial summary statistics have been computed, and displayed as time series plots – see Wheater *et al.* (1997) for examples. These plots enable periods of potential interest to be isolated and have been used, together with a visual inspection of images, to identify a subset of rainfall events which is suitable for analysis from the HYREX radar data catalogue. These events were selected on the basis that they covered the inner radar circle for a reasonable period of time, and that there were no serious data quality problems within this circle.

ANALYSES USING RAINGAUGE DATA

In addition to the routine analyses described above, some more general properties of the data, relating to the spatial structure of the rainfall field and the influence of topography, have been investigated. It is important that the class of stochastic rainfall models to be considered has the ability to reproduce such properties of the data, at least over spatial and temporal scales of interest.

An extensive analysis of the topographic effects of the Brue area on the behaviour of the rainfall process has been carried out using raingauge data. Various monthly summary statistics were plotted against gauge elevation and regression relationships were sought. The only properties to show consistently a statistically significant elevation effect are the mean rainfall intensity and the proportion of wet intervals. A straight line fit to the data explains about half of the variation in each statistic at the hourly timescale. Plots of the two statistics against elevation are given in Fig. 3 for the months of June and December. The results for June

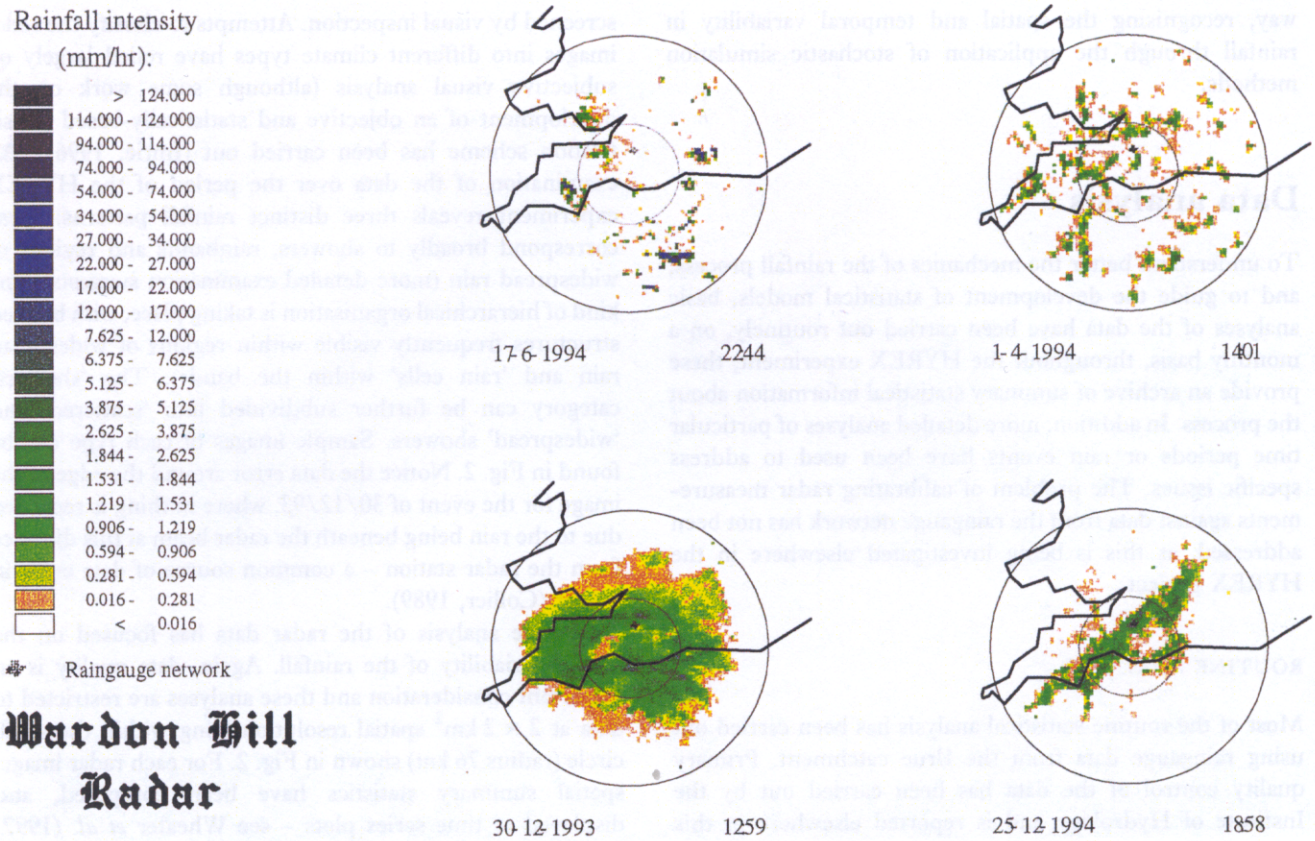

Fig. 2. *Different rainfall types observed with the Wardon Hill radar station. Clockwise from top left: scattered showers, widespread showers, rain band, stratiform rain.*

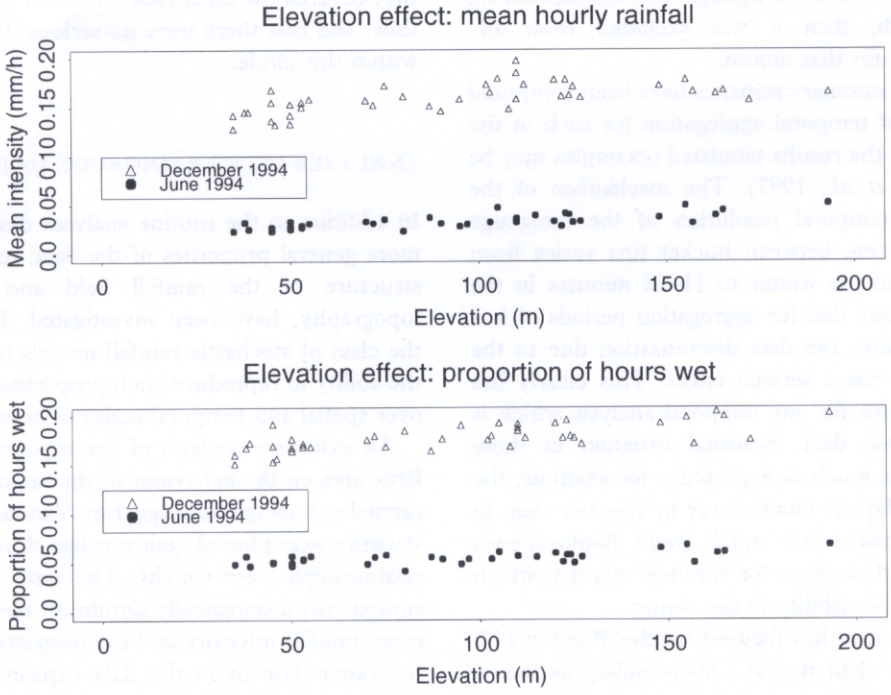

Fig. 3. *Effect of elevation on hourly rainfall statistics for months of June and December 1994. Top: mean rainfall intensity. Bottom: proportion of hours wet.*

indicate a 70% increase in the mean rainfall at an elevation of 190 m compared with that at 35 m, with an increase of 20% for the proportion of wet intervals. The December results indicated a 20% change in both statistics over the range of elevations in the catchment. These results show that relatively small changes in elevation can make a substantial difference to the rainfall. The main difference, occurring throughout the year, is in the proportion of intervals which are wet; it is only in summer that elevation has an effect on rainfall intensities during wet periods, for the 20% increase in mean rainfall in winter is explained by the increase in the proportion of wet intervals alone.

In addition to examining the effect of elevation on the rainfall behaviour, the possible presence of other geographical effects has been investigated by plotting monthly summary statistics against the grid co-ordinates of the gauges. No significant relationship was found – this is perhaps unsurprising given the small scale of the Brue raingauge network.

The spatial association of the rainfall process within the Brue catchment has been examined by looking at the effect of distance and orientation on the inter-site properties (again on a monthly basis). The Brue area was split into North-South and East-West transects, defined by the $2 \times 2\,km^2$ grid squares of the Wardon Hill radar station. For all pairs of raingauges within each transect, the probability of both sites being wet and the cross correlation of the rainfall intensity at zero time lag were calculated, using hourly data. Figure 4 shows the cross-correlations plotted against grid square separation, again for June and December 1994. The mean and standard deviation of the cross-correlation are calculated over all pairs of sites with the same grid square separation. A decreasing spatial association with distance is clear; it drops off slightly more quickly in a North-South direction than in an East-West direction (this is probably a topographic effect, due to the East-West orientation of the Brue valley). The seasonal effect is obvious: intense and localised summer storms result in a faster decay in the spatial correlation of the process with distance than do more uniform winter systems. Moreover, the spatial correlation structure is much less variable in the winter than in the summer, as indicated by the larger standard deviations across the network in the latter case.

Similar analyses for data at different levels of temporal

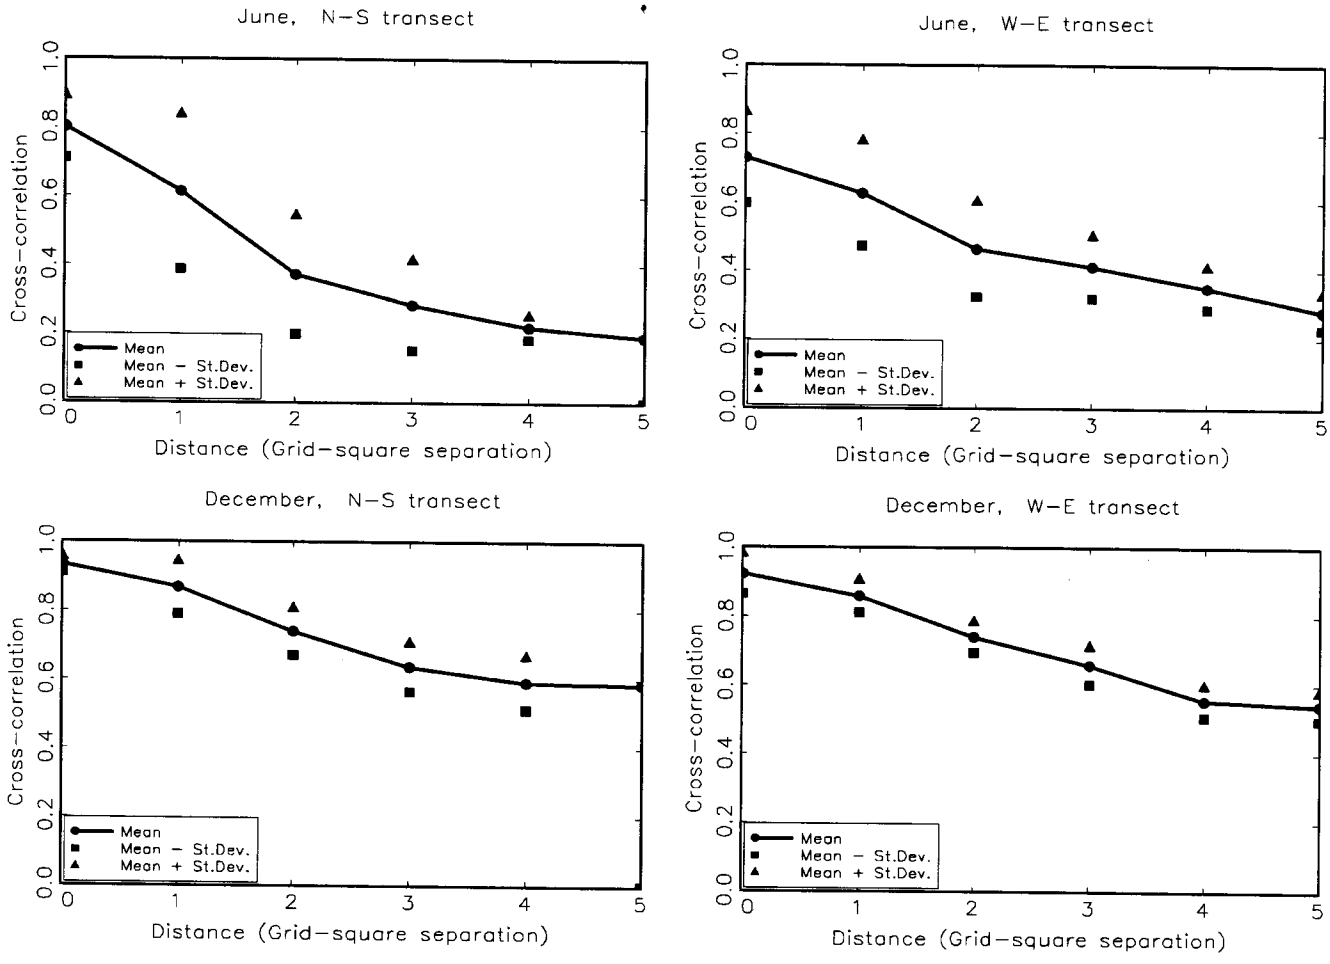

Fig. 4. *Spatial cross-correlations for hourly rainfall amounts from the HYREX raingauge network, showing effects of distance and direction. Top: June 1994, bottom: December 1994. Left-hand plots show results for North-South transects, right-hand plots show results for East-West transects.*

aggregation produce the same conclusions, although aggregation of the data over larger time periods makes the process more spatially homogeneous and reduces the effect of distance.

While some of the statistical calculations reported above will assist in developing and comparing models for the rainfall process, perhaps the most striking feature to emerge from these analyses is the effect of topography on the rainfall. The Brue catchment encompasses a very small elevation range, yet the rainfall on the upper slopes can be as much as 70% more than that on the lower slopes in summer; moreover, the cross-correlations shown in Fig. 4 indicate that, at this local scale, the Brue valley influences the rainfall correlation structure. These findings clearly have implications for hydrological design at a small catchment scale.

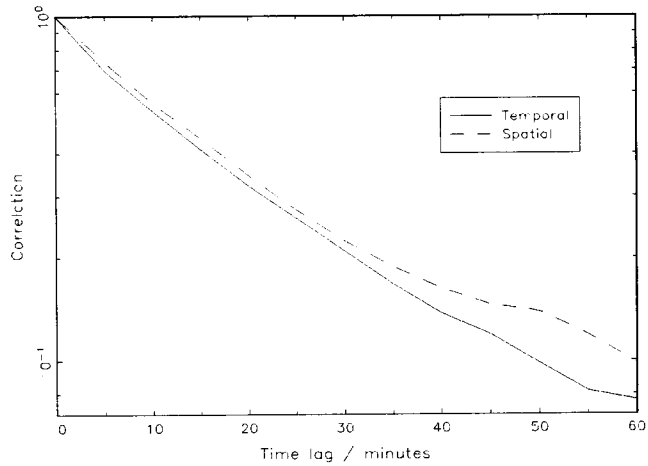

Fig. 5. *Estimated temporal and spatial autocorrelation functions for Wardon Hill radar event of 25/12/94, verifying the Taylor Hypothesis for time lags up to 30–40 minutes.*

ANALYSIS OF RADAR DATA

While the raingauge network allows the examination of the effects of topography on the rainfall process and to quantify its very local behaviour, other issues (particularly relating to the spatial variability of the rainfall on a larger scale) may be addressed using radar data.

One area of analysis has been the investigation of *Taylor's hypothesis* (Taylor, 1938), which relates spatial and temporal autocorrelation structure via the average velocity of the field. Explicitly, if $\rho(u,\tau) = \mathrm{Corr}[Y(0,0),Y(u,\tau)]1$ for a spatial-temporal field $Y(u,\tau)$ viewed as a stationary stochastic process, then the field is said to satisfy Taylor's hypothesis if

$$\rho(V\tau, 0) = \rho(0, \tau) \qquad (1)$$

for some velocity vector, V and all time lags, τ. Essentially, this expresses the notion that an observer moving at the same velocity as the field would not see any qualitative changes in its structure over time. While a fixed spatial field which moves with constant velocity (a 'frozen field') implies Eqn. 1, the converse is not true.

Figure 5 is a plot of $\rho(V\tau,0)$ and $\rho(0,\tau)$ against τ for the storm of 25/12/94 shown in Fig. 2. The velocity of the storm, V, was estimated as (22.8,−23.6) km h^{-1} by tracking the centroid of the storm. The figure shows the two lines beginning to separate substantially after around 30–40 minutes. Before this 'cutoff' Taylor's hypothesis appears to hold. This is in broad agreement with a previous study by Zawadski (1973). The extent of the validity of Taylor's hypothesis for stochastic spatial-temporal models will be investigated later.

Another informative analysis concentrates on how the spatial variability of the total accumulated rainfall intensity or *depth* deposited by a storm changes over time. The analysis uses a Lagrangian frame of reference: the storm is studied over a period of time through a circular 'data window', moving with the storm, and spatial statistics are computed over this window. For this analysis, attention is restricted to fairly slow-moving storms and to the 5 × 5km^2 resolution radar data, in order to be able to track the storm for a reasonable period of time as it evolves. An interesting feature of the data is that the spatial variance of the depth stops increasing after 3–4 hours in many of the storms analysed. Thus, in Fig. 6, where the spatial standard deviation of the rainfall depth is plotted against the mean on a log-log scale for a variety of different window sizes, the curves start to flatten out at the right-hand end of the graph, beginning with those for the smallest window sizes. Although the implications of this phenomenon are not yet understood fully, it may indicate the presence of some inhibition in the rainfall process – for example, an area of high rainfall intensity within a storm may inhibit further high intensity rainfall within that area at a later time. This conclusion is speculative, however, and the phenomenon bears further investigation.

Rainfall models

OVERVIEW

In the models described and used in this paper, the main observable features of precipitation processes, such as the clustering and movement of rain cells, are represented explicitly. However, the detailed deterministic behaviour of the physical processes is replaced by simple stochastic assumptions, whence a rather small number of physically interpretable parameters may be used to represent the rainfall process. Such models fit part way between the deterministic General Circulation Models (GCMs) of dynamical meteorology involving complex sets of differential equations (e.g. Mason, 1986), and the empirical statistical models based upon ideas such as multiscaling

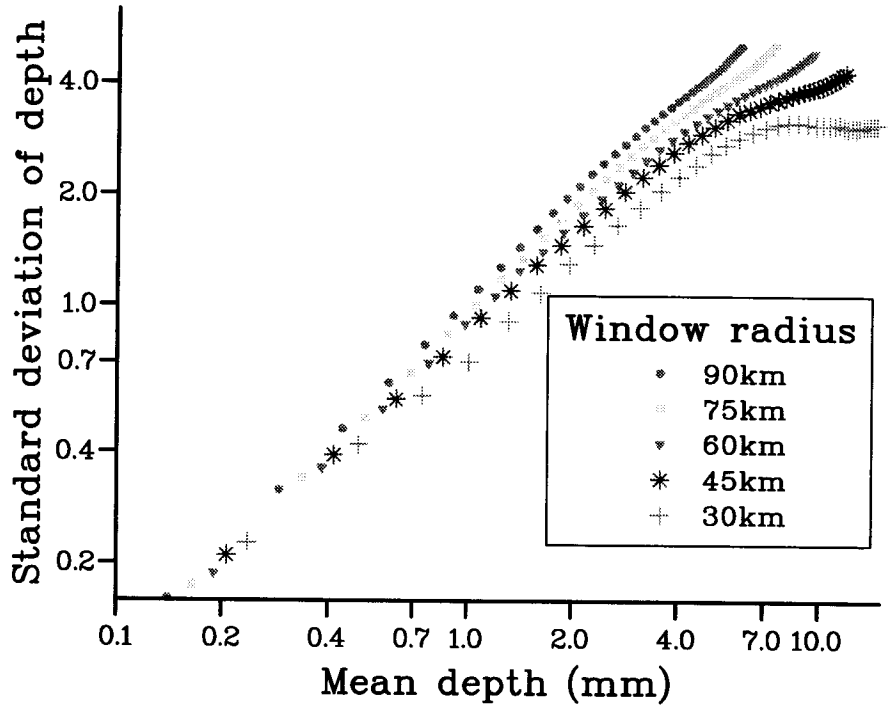

Fig. 6. *Evolution of spatial variability of a storm with time: log-log plot of spatial standard deviation of storm depth (= cumulative intensity) against mean storm depth for event of 7/12/93, at a range of different window sizes.*

processes and fractal cascades (e.g. Lovejoy and Schertzer, 1987; Gupta and Waymire, 1990).

Models for spatial-temporal precipitation based on stochastic point processes, go back at least as far as the fundamental work of Le Cam (1961). This approach developed rapidly in the 1980s through a series of papers (e.g. Waymire *et al.*, 1984). Such models are based on a hierarchical structure in which *rainfall fields* occur in a temporal Poisson process, *rain bands* (storms) occur within each field in a spatial Poisson process (the rate of which may reflect orography and seasonality), and *rain cells* occur in each storm, clustering in space and time. Typically the cells, storms and fields move: in the simplest models, all components have a common velocity.

Models for the temporal evolution of rainfall, both over continuous spatial regions and at discrete sets of locations are described. The former (referred to as *spatial-temporal models*) are appropriate where radar data are available, whereas the latter (*multi-site models*) are relevant to data from a network of raingauges. A crucial feature of the approach is that all models are constructed, and their properties determined, in *continuous* time (and, for the spatial-temporal models, in continuous space). To fit the models to empirical data, it is then a relatively straightforward matter to aggregate their properties, as appropriate, over disjoint spatial and/or temporal regions; for example, rain gauge data are usually aggregated temporally, while radar data are averaged spatially. Then, in assessing the adequacy of the fit of a model, it is important that it reproduces well the properties of the data at levels of aggregation other than those used in its fitting.

In the work reported here, for simplicity only one level of clustering of cells within storms shall be included, although there are indications that a better fit would be obtained by allowing a second level of clustering, stochastic stationarity in both time and space is assumed. Thus, in fitting the models, each month is treated separately, using data for a relatively homogeneous spatial region.

The spatial-temporal models that have been developed are spatial analogues of models that have been used successfully to represent the temporal process of rainfall at a single rain gauge (Rodriguez-Iturbe *et al.*, 1987, 1988) and generalise that investigated in Cox and Isham (1988). The multi-site models generalise similarly the models of Cox and Isham (1994). All of these models have the desirable feature that they preserve the structure of the single-site models in their marginal properties.

SPATIAL-TEMPORAL MODELS

The focus of the HYREX experiment has been upon the use of radar data, which is ideally suited to the development of fully spatial-temporal models: it has been in this area that the most significant contributions have been made during the project, and this aspect of the modelling is therefore discussed first. The basic spatial-temporal model of Cox and

Isham (1988) does not have any clustering of cells. Its specification is as follows:

- Cells arrive in a Poisson process in (two-dimensional) space and time.
- Each cell is circular in space with a random radius, duration and velocity.
- Throughout its duration the cell deposits rainfall at a constant rate (intensity) on all points in space covered by its defining disc.
- The total rainfall intensity at a point in space-time is the sum of the contributions from all cells active at that point.

It is assumed that all cells move with the same velocity, an assumption which is realistic to make only when studying a single weather system.

It is relatively simple to derive expressions for the mean and second order properties of the process. Additionally, expressions for properties such as the probability that an arbitrary point in space-time is not covered by a cell (i.e. is *dry*) can be derived.

Northrop (1996) introduced a generalisation of this model in which cells are elliptical rather than circular. Hereafter referred to as the elliptical cell Poisson process model (EPPM), this is likely to be more realistic, especially in the cases where banding is apparent in the radar images. These cells are also identifiable by the elliptical contours of their spatial autocorrelation plots. This model requires two extra parameters, the eccentricity and orientation of the cells, which are both assumed to be common to all cells.

It is clear that these models are highly idealised, for inspection of radar images indicates that a model in which rain cells are clustered within storms in both space and time is required. Cox and Isham (1988) also investigate a model in which circular rain cells are clustered within storms. The temporal clustering has a Bartlett-Lewis type structure: explicitly, for a storm centre at (u,t), cell origins will arrive at u in a temporal Poisson process starting from a cell origin at (u,t), the process terminating after an exponentially distributed time. All cells within a storm originate at the same spatial location. Since all cells move with the same velocity, the result is that a storm consists of a band of cells moving in the direction of alignment of the band. This is unrealistic, for empirical observation is that the direction of cell movement tends to be orthogonal to any banding which is present.

To overcome this problem, a modified version of this model (Northrop, 1997) has been investigated. The temporal clustering of cells is achieved using a Bartlett-Lewis structure as above. Additionally, spatial clustering is incorporated using a Neyman-Scott-type mechanism in which the displacements of the cell origins from the storm centre follow a bivariate distribution in space. A range of storm shapes (e.g. bands and large masses) can be produced by variation of the parameters of the spatial clustering distribution. An important modification to the model of Cox and Isham (1988) is to have the storm centre moving with the same velocity as the cells so that cells are born within the existing structure of the storm. Two spatial clustering distributions are considered:

1. A bivariate Gaussian (normal) distribution. The resulting model is referred to as the *Gaussian displacements spatial-temporal model (GDSTM)*;
2. A uniform distribution over a random ellipse. This gives rise to the *random ellipse spatial-temporal model (RESTM)*.

This model can be generalised by taking the cells to be elliptical in shape – in the simplest case, they are just scaled versions of the elliptical contours of the bivariate Gaussian distribution (GDSTM) or storm ellipse (RESTM) within which they were born.

Expressions for the mean and second order properties of the model can be derived (Northrop, 1996) although the second order properties require the numerical evaluation of integrals or the use of approximations. The same is true of the probability that an arbitrary point in space-time is dry. In addition, for the purposes of model fitting using the method of moments, the aggregation of properties over a radar pixel will require the numerical evaluation of a double integral for all the models considered here. Further details of the specification and fitting of these models are in Northrop (1997).

For the class of spatial-temporal model considered here, it can be shown theoretically that Taylor's hypothesis holds approximately for time lags which are small relative to the mean cell duration.

MULTI-SITE MODELS

In regions where radar data are not available, it is useful to develop spatial-temporal models that can be calibrated from a network of raingauges. In principle, it is possible to apply a full spatial-temporal model as described above, by deriving and using properties of the model at a discrete collection of spatial locations. This is complicated, however, and an alternative approach is to develop models in which between-site interactions are expressed directly in a manner which reflects the underlying spatial-temporal structure of the rainfall, but which avoids the complications of a full spatial-temporal model. These are referred to as *multi-site models*: they should have inter-site properties that depend on the distances between the sites and possibly on topography, and should also preserve the structure of single-site models for their marginal processes at each site.

The models considered are generalisations of those proposed by Cox and Isham (1994), whose motivation was as follows: when studying rainfall over a network of k sites, a storm or cell may be categorised according to the subset of sites which it affects. Excluding the empty subset, there are $N = 2^k - 1$ possible subsets of the sites, which may be

labelled $S_1,\ldots,S_N$. A storm (cell) affecting all of the sites in S_j and no others is said to be of *type j*. The basic idea underpinning the models of Cox and Isham (1994) is that cells arrive at the study area in a point process called the *master process*, and the type of each cell is determined by some random mechanism. This type of model structure offers great flexibility for building a wide variety of different multi-site models; however, it requires the probabilities of each of the 2^k -1 different types to be modelled, which in general will involve a large number of parameters unless simplifying assumptions are made.

An intuitively reasonable strategy for simplifying this problem is to assume that the probability of a cell hitting two sites depends in some way upon the distance between the sites. There are various ways of doing this, the easiest being to assume a simple algebraic form for this probability as a function of inter-site distance (e.g. an exponential decay). Alternatively, the choice of functional form for this dependence can be made more physically plausible by incorporating assumptions about the spatial structure of a storm. These assumptions reflect the structure of the full spatial-temporal models already described. Specifically, it is assumed that storms move with a constant velocity and that they hit the study area in 'fronts' which are perpendicular to the direction of movement. The centre of a storm thus traces out a straight path called the *storm axis*, which is fixed throughout the storm's lifetime. A storm gives rise to a cluster of rain cells, each of which moves with the same velocity as the storm, along its own axis which is displaced from the storm axis according to some random mechanism. Each cell has a random 'width' (measured at right angles to the direction of movement). A cell's type (i.e. the subset of sites which it affects) is thus determined by its axis position and width.

In the specific example developed and presented here, various choices have been made regarding the specification of various components of the model. These are detailed below. Unless otherwise specified, any random quantity in the model has been taken to follow an exponential distribution, realised independently and identically for each cell and independently of any other quantity in the model.

- The master process of storm arrivals (i.e. times at which storm fronts hit some arbitrarily chosen point within the study region) is Poisson in time.
- The axis position of each storm is random, distributed uniformly over some region enclosing the study area, and realised independently from storm to storm.
- The speed of storm and cell movement is random, fixed for each storm but realised independently and identically between storms according to a gamma distribution. The direction of storm and cell movement also varies randomly between storms, this time according to a discretised cardioid distribution (Fisher, 1993).
- Within storms, the temporal clustering of rain cells is achieved via a Neyman-Scott mechanism with an exponential scatter distribution. The number of cells in each storm follows a Poisson distribution.
- The spatial displacement of each cell's axis from the storm axis is drawn, independently and identically for each cell, from a double-sided exponential distribution whose probability density function is of the form

$$f(x) = \frac{\lambda}{2} e^{-\lambda|x|} \quad (-\mathit{infinity} < x < \mathit{infinity})$$

  for some positive parameter λ.
- Each cell has a random intensity and nominal duration. However, the time for which a cell is active at any site after hitting that site is obtained by scaling this nominal duration by a site-specific constant. This is intended to model topographic effects; its justification lies in the observed influence of topography upon mean rainfall through the proportion of wet intervals (as mentioned previously and in Fig. 3), for it is only if cells are active for different lengths of time at each site that this feature can be incorporated realistically into the model.

These multi-site models are reasonably parsimonious in their parametrisation, requiring a single extra parameter, the cell duration scalar, for each new site that is included in the study. The cross-correlation function of the rainfall intensity at a pair of sites has been derived (Kakou, 1997) as has the implied functional form of the probability of a cell hitting two sites. It turns out that, for individual storms, this probability decays approximately exponentially with inter-site distance for sites which are well-separated and which are not aligned along the direction of the storm's movement; for sites which are closer together, the dependence is no longer exponential.

SINGLE-SITE MODELS

The models described in the preceding sections were generalisations of models that have been used successfully to model the temporal evolution of rainfall at a single site. A first step towards improving the performance of these models involves studying ways in which the single-site models can be improved.

One of the most obvious ways in which the basic single-site models of, for example, Rodriguez-Iturbe *et al.* (1987), can be extended is by allowing for different types of storm. One way of achieving this, adopted by Rodriguez-Iturbe *et al.* (1988), is to randomise the cell duration parameter between storms – in this approach, storms have a common structure but occur at different timescales. The main advantages of such models, in practical terms, lies in their ability to reproduce well the observed probability of no rainfall at various levels of aggregation.

An alternative to the randomisation of the cell duration parameter for single-site models has been investigated. This allows for different types of storm using an inverse

relationship between the duration of an event and its intensity (the motivation being that intense convective events tend to be shorter-lived than shallower stratiform systems). By adopting an explicit functional form for the dependence between cell depth and cell duration, it is possible to overcome the problems of over-parameterisation typically associated with attempts to model different cell types explicitly (e.g. Cowpertwait, 1994).

The work is based on the Neyman-Scott and Bartlett-Lewis point process models (Rodriguez-Iturbe *et al.*, 1987) which are modified to allow raincells with stochastically dependent duration and intensity (Kakou, 1997). It is assumed that the cell duration, L, is exponentially distributed with parameter η, and that given $L = l$, the conditional intensity, X, is exponentially distributed with mean of the form

$$E(X|L = l) = f\ l^d \exp(-cl) \tag{2}$$

where f, c and d are positive constants. Depending on the value of d, this family of models covers a wide range of possible forms of dependence. The special cases with $d = 0$, corresponding to a negatively correlated X and L, and $d = 1$, resulting in a positive correlation when $\eta > c$, have been explored and the second order properties of the models have been derived. These properties involve only the distribution of X through $E\{XL\}$ and $E\{X^2L\}$, so the explicit distributional forms are unimportant once these moments are known (Kakou, 1997).

MODEL FITTING

One of the biggest problems with the type of model described above is the difficulty in estimating model parameters. Standard statistical techniques such as Maximum Likelihood estimation are not available here, for the complex dependencies in the models prohibit the formulation of a likelihood function. Further problems are created by the aggregation of the data, for the models themselves are built in continuous space and time. Two approaches to model fitting have been investigated. The first is a generalised *method of moments*, where parameters are chosen to minimise a weighted sum of squared differences between suitably-chosen model properties and their corresponding sample values; and the second is a *spectral approach* which treats the sample Fourier coefficients as data and formulates an approximate likelihood function in terms of these coefficients.

Generalised method of moments

In the absence of any more sophisticated techniques for parameter estimation, an intuitively reasonable procedure is to try to obtain as close as possible a match to observed features of the data. Denote the p-dimensional parameter vector of the model by $\Theta = (\theta_1,\ldots,\theta_p)$, and pick a set of k features $y_i (i = 1,\ldots,k)$ of the data for which model expressions can be obtained. If the model expression corresponding to y_i is $f_i(\Theta)$ then take $k = p$ and try to solve the set of simultaneous equations

$$y_i = f_i(\Theta) \quad i = 1,\ldots,k \tag{3}$$

for Θ.

The model expressions involved will be highly non-linear so that a unique solution to Eqn. 3 may not exist. A generalisation of the approach is to minimise the sum of squared differences between the $\{y_i\}$ and the $\{f_i(\Theta)\}$, weighting each term appropriately to avoid bias due to the differing orders of magnitudes of observed values of the features involved. So a value for $\hat{\Theta}$ is found which minimises the objective function

$$\sum_{i=1}^{k} w_i\ (f_i(\Theta) - y_i)^2 \tag{4}$$

where the $\{w_i\}$ are some suitably-chosen weights. There is now the possibility of having $k > p$, i.e. including more than p features in the minimisation.

The fitting features chosen should have relatively small sampling errors, not be highly mutually correlated and contain sufficient information that all the parameters of the model can be estimated reliably. Apart from these criteria, the choice of fitting features is essentially subjective and the estimates of certain parameters will, in general, depend greatly upon the set of features used for fitting the model – a parameter identification problem. Additionally, the parameter estimates may depend upon the initial estimates given to the numerical minimisation algorithm. Certain parameters may also be highly correlated in the sense that many different pairs of values of two parameters give very similar values for the objective function. This corresponds to the objective function being almost flat in some region of the parameter space, so that convergence to the minimum is difficult. Re-parameterisation of the fitting procedure using more stable parameters may alleviate this problem. A further feature of these models in general is that different sets of parameters may yield similar sets of fitted values, leading to multiple minima in the objective function.

Assessment of the fit of the model to the data is achieved by comparing the predicted values of features not used in the fitting procedure with those of the data. Simulation from the fitted model can be used to assess how well the model reproduces features for which it is not possible to obtain model expressions.

Spectral method

The method of moments suffers from a number of disadvantages. In particular, the choice of features to incorporate into the fitting procedure is subjective, and the parameter values obtained can be quite sensitive to the features used in the fitting – hence model comparison can be difficult. Moreover, the method makes inefficient use of available data, as only a few summary statistics are used in

the fitting. In an attempt to overcome some of these difficulties, a spectral method has been developed. This method uses the sample Fourier coefficients rather than the original data, and makes use of the fact that, for large samples, small collections of the Fourier coefficients have a joint distribution which is approximately multivariate normal (Brillinger and Rosenblatt, 1967). This enables approximate likelihood functions for the model parameters to be written in terms of small subsets of the sample Fourier coefficients. By combining all of these approximate likelihood functions, an objective function is defined which can be interpreted as a log quasi-likelihood (Chandler, 1997). This then provides a basis for objective model comparison procedures using standard statistical techniques such as likelihood ratio tests (McCullagh and Nelder, 1989).

The method has been developed for use in fitting single-site and spatial-temporal models. In both cases, the only theoretical properties required from the model are the mean and second-order spectral density of the process; the required expressions (corresponding both to the underlying model and to the aggregated domain of observation) have been derived for wide classes of models in both cases. The reliance on second-order properties is a potential disadvantage in distinguishing between models whose main difference is in their wet/dry interval properties. More details may be found in Chandler (1996b, 1997).

MODELLING RESULTS

This section presents a selection of results obtained when the models described in the previous section are fitted to HYREX data. Lack of space precludes a more complete enumeration of results, which may be found in Northrop (1996), Kakou (1997) and Chandler (1996b, 1997).

Spatial-temporal model (GDSTM)

The results reported here were obtained by fitting the GDSTM to a one hour sequence of 2 km resolution radar images from the event of 6th February 1994, using the method of moments to estimate parameters. The data used in the fitting were of 2×2 km^2 spatial resolution, and their quality appears reasonable—there are no indications from the radar images of any of the problems commonly encountered with radar data, and the mean radar/raingauge ratio for each radar grid square covered by the HYREX raingauge network is 0.68.

Table 1 shows the estimated parameters in the GDSTM model. These are in broad agreement with the spatial and temporal dimensions of rain cells and small mesoscale areas identified in observational studies such as Austin and Houze (1972). Table 2 shows the observed and model fitted values for features used in fitting the model. These properties have been chosen to enable the estimation of all the parameters of the model, and a very close match has been obtained. In addition to the properties used to fit the model, there is good agreement between the observed and fitted spatial and temporal autocorrelation functions (see Figs. 7a and 7b). Goodness of fit can also be assessed in terms of the spatial scaling properties of radar and model simulated data. One feature which the model does not reproduce is the decrease in spatial variability of accumulated storm depth reported previously (Fig. 6).

Figure 8 is presented to provide a visual assessment of the GDSTM. The figure shows three snapshots of radar data to which the model was fitted, together with a corresponding specimen realisation of the fitted model (the time labels for the bottom sequence are from the start of simulation). The

Table 1. Parameter estimates for Gaussian Displacements Spatial-Temporal Model fitted to event of 6/2/94.

Process property	Fitted value
Storm arrival rate (km^{-2}h^{-1})	3.44×10^{-4}
Expected cell duration (min)	38.80
Expected cell intensity (mm h^{-1})	0.537
Variance of cell intensity (mm h^{-1})	0.288
Expected cell area (km^2)	22.99
Cell/storm eccentricity	0.831
Cell/storm orientation (degrees)	94.35
Cell/storm velocity (x) (km h^{-1})	36.49
Cell/storm velocity (y) (km h^{-1})	33.76
Expected number of cells per storm	648.6
Expected storm duration (h)	2.707
Expected storm area (km^2)	307.3
Variance of storm area (km^2)	29.38

Table 2. Fitted and observed values of fitting features for Gaussian Displacements Spatial-Temporal Model fitted to event of 6/2/94. $Y^{(h)}$ denotes mean rainfall over a $h \times h$ km^2 pixel, and $\rho(x,y,t)$ denotes the space-time autocorrelation at spatial lag (2x,2y)km and time lag 5t minutes.

Property	Fitted	Observed
Mean	1.781	1.781
Var($Y^{(2)}$)	3.649	3.649
Var($Y^{(8)}$)	2.588	2.586
Var($Y^{(16)}$)	1.925	1.926
$\rho(0,0,2)$	0.447	0.450
$\rho(0,0,4)$	0.214	0.216
$\rho(2,2,1)$	0.860	0.861
$\rho(-2,-2,1)$	0.423	0.414
$\rho(-1,0,0)$	0.825	0.829
$\rho(-4,0,0)$	0.447	0.452
$\rho(-7,0,0)$	0.249	0.247
$\rho(-4,4,0)$	0.421	0.407
$\rho(0,8,0)$	0.415	0.396

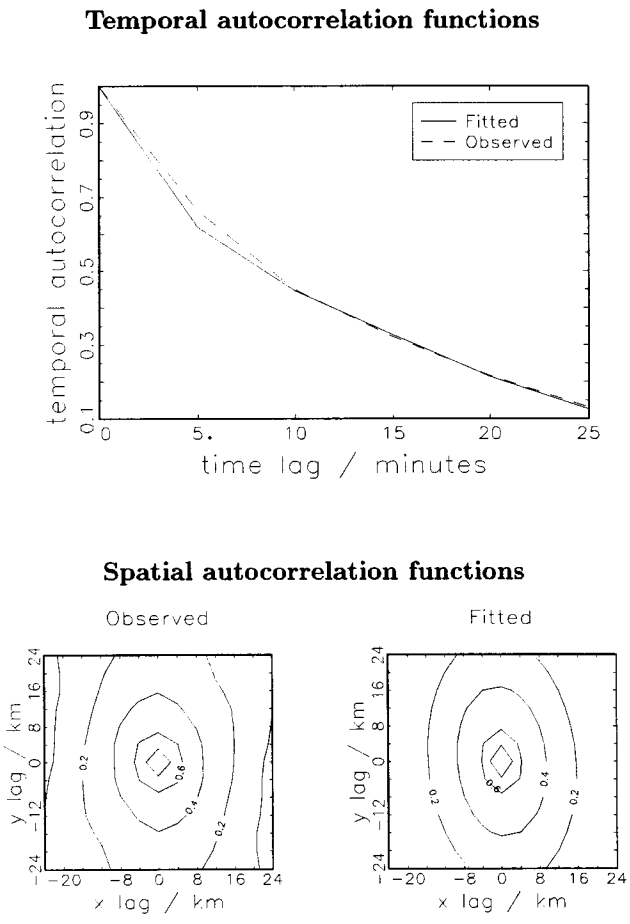

Fig. 7. *Event of 6/2/94: observed autocorrelations compared with those fitted using GDSTM. Top: temporal autocorrelations. Bottom: spatial autocorrelations.*

apparent similarity between the interior structures of each sequence is striking (the model is intended to provide a description of the interior of a rainfall event which can be regarded as statistically homogeneous; hence there is a lack of 'edges' in the simulated sequence). An exact correspondence is not expected owing to the stochastic nature of the model: a visual comparison of the two sequences should be made on the basis of whether or not they exhibit broadly similar features.

The simpler EPPM has also been fitted to this data sequence using the method of moments. The results are not reported here; the fit was significantly poorer, in terms of both the comparison of observed and fitted process properties, and a visual assessment of the data. The conclusion, from fits to this and other data sequences, is that the spatial-temporal clustering of the GDSTM is indeed necessary for the realistic modelling of rainfall fields, even within a single weather system.

Multi-site model

The results of fitting the multi-site model described above are now presented, again using the method of moments to estimate parameters. For illustrative purposes, results are presented for the month of December 1994, using hourly raingauge data from three sites which form a roughly equilateral triangle of side about 10 km see Fig. 1; the sites chosen are numbers 5, 29 and 42).

To model the rainfall at 3 sites, 13 parameters are used. To fit the model, the mean direction of storm and cell movement is first estimated from the observed cross-correlations between the sites. To estimate the remaining parameters, 18 features were selected for inclusion in the objective function Eqn. 4, 4 marginal features from each site and 2 cross-correlations between each pair. The estimated model parameters are given in Table 3.

The fitted model estimates nearly all the marginal statistics quite accurately, and fitted cross-correlations are close to observed ones. Table 4 shows some of the statistical properties of the data, and gives the corresponding values for the fitted model (the features used in the fitting are shown in **bold type**). From this table, it appears that the main areas of discrepancy between model and data are as follows:

1 There is a tendency for the model to underestimate the proportion of wet intervals at a site, particularly for the 6-hourly level of temporal aggregation. The discrepancy is especially noticeable at site 29, which is in the North-East of the Brue catchment – this may indicate a gauge location effect which has not been captured by the model.
2 The model overestimates the probability that two sites are simultaneously dry at large levels of temporal aggregation. This is perhaps to be expected from the previous comment. However, the problem is particularly noticeable for sites 5 and 29 taken together, once more indicating a geographical effect – these two sites are in opposite quadrants of the Brue valley (see Fig. 1).
3 The model underestimates the lag1 cross-correlation between all three pairs of sites for hourly data; conversely, there is a tendency to overestimate this feature at higher levels of temporal aggregation.

Despite these discrepancies between data and model, these results indicate that these models are capable of capturing most of the observed features of rainfall at a few sites; moreover, they are reasonably parsimonious in their parameterisation thanks to the simplifying assumptions regarding storm structure. There is clearly scope for further development of these models, and the features for which lack of fit is apparent from Table 4 may serve as pointers for such development. Indeed, the problems with reproducing proportions of wet intervals at a site were experienced by Rodriguez-Iturbe *et al.* (1987) in their original single-site model and resolved by randomising the cell duration parameter in Rodriguez-Iturbe *et al.* (1988): perhaps the multi-site models could be extended in a similar way.

Single-site models

Below are the results of fitting one of the dependent depth-

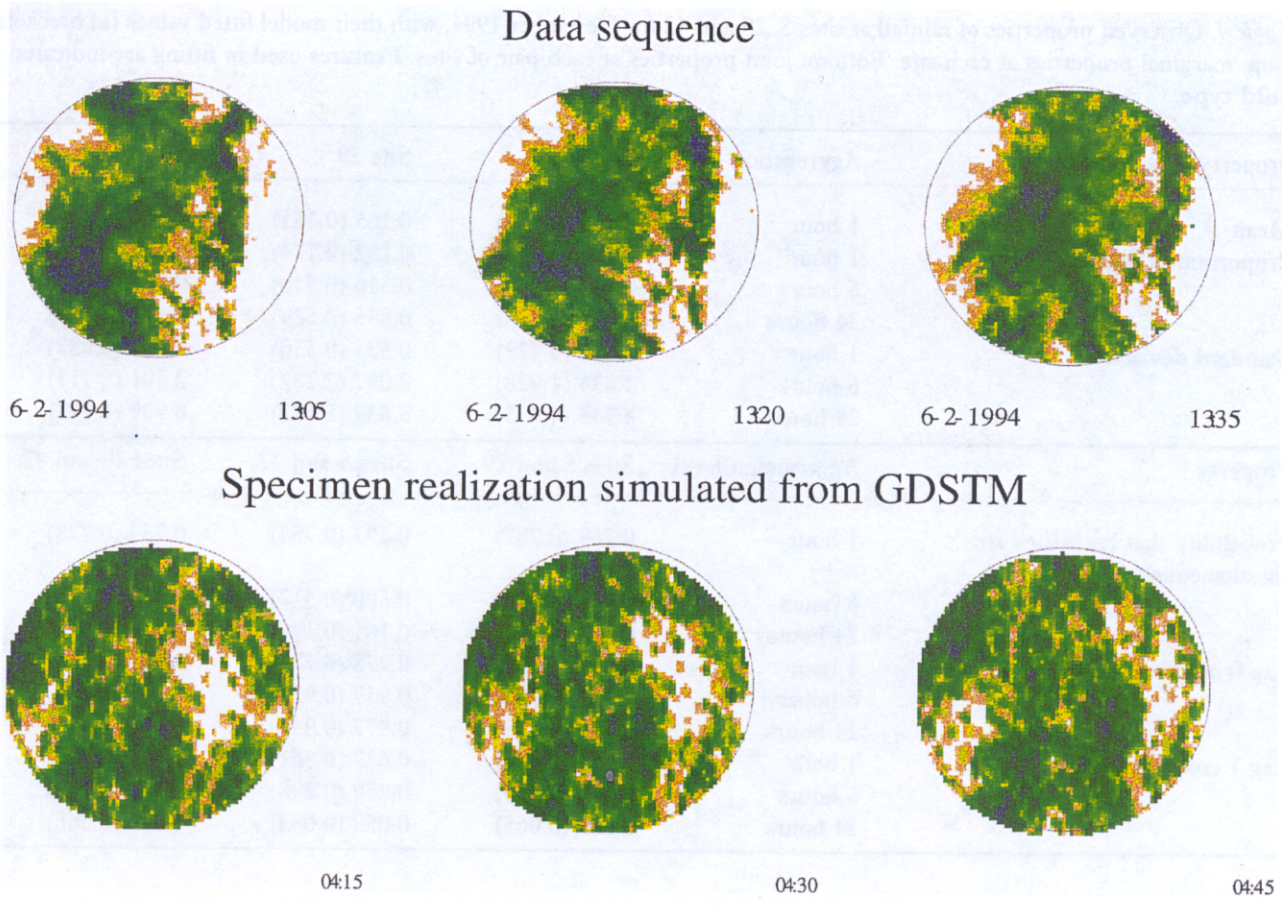

Fig. 8. *(a) Actual and (b) simulated sequences for GDSTM fitted to event of 6/2/94.*

duration models described previously to hourly data from gauge 19 for December 1994, again using the method of moments; this model is compared to the random parameter Bartlett-Lewis model of Rodriguez-Iturbe *et al.* (1988). In the dependent depth-duration model, a Neyman-Scott clustering structure was used and the value of d in the intensity-duration relationship Eqn. 2 was taken to be 0. This model has 5 parameters, whereas the random parameter Bartlett-Lewis model has 6.

Table 3. Parameter estimates for multi-site model fitted to hourly data from sites 5, 29 and 42, December 1994.

Parameter		Estimate
Storm arrival rate		0.020 h^{-1}
Mean rain cell intensity		0.97 mm h^{-1}
Mean temporal displacement of cells from storm origin		4.37 h
Mean no. of rain cells per storm		14.9
Mean speed of storm/cell movement		8.4 km h^{-1}
Std. dev. speed of storm/cell movement		2.8 km h^{-1}
Mean width of a rain cell		45.5 km
Mean distance of cell from storm axis		1.4km
Mean direction of storm/cell movement		202.5°
Concentration parameter of direction distribution		0.37
Mean cell duration	Site 5	34.6 min
	Site 29	41.2 min
	Site 42	49.9 min

The estimated parameters for both of these models are given in Table 5. In addition, an alternative parameter set for the Bartlett-Lewis model, derived using the spectral method, is presented. The fitting features used to obtain the moment-based parameter estimates were the mean and lag1 correlation of hourly rainfall amounts, and the variance of rainfall amounts and proportion of dry intervals at hourly and daily levels of temporal aggregation. The difference between the two sets of estimates for the random parameter model is worrying; however, quite different sets of parameters can result in similar performance, so these parameter sets must be assessed in the light of their ability to reproduce various features of the data.

Table 6 shows the observed values of various features of the data, together with their model fitted values for the various models and parameter values given in Table 5. First of all, on comparing the moment-based fits of the two different models, it is seen that both are generally in good agreement with the data, with some minor discrepancies (for example, the random parameter model overestimates the correlations in hourly data at lags 2 and 3). At levels of

Table 4. Observed properties of rainfall at sites 5, 29 and 42 for December 1994, with their model fitted values (in brackets). Top: marginal properties at each site. Bottom: joint properties at each pair of sites. Features used in fitting are indicated in **bold type**.

Property	Aggregation level	Site 5	Site 29	Site 42
Mean	1 hour	**0.133 (0.127)**	**0.165 (0.151)**	**0.176 (0.181)**
Proportion of wet intervals	1 hour	**0.164 (0.168)**	**0.202 (0.174)**	**0.165 (0.181)**
	6 hours	0.419 (0.308)	0.516 (0.310)	0.467 (0.400)
	24 hours	0.742 (0.528)	0.835 (0.529)	0.874 (0.530)
Standard deviation	1 hour	**0.465 (0.479)**	**0.533 (0.550)**	**0.653 (0.632)**
	6 hours	1.834 (1.928)	2.097 (2.282)	2.501 (2.711)
	24 hours	**4.745 (4.743)**	**5.537 (5.642)**	**6.999 (6.742)**
Property	**Aggregation level**	**Sites 5 and 29**	**Sites 5 and 42**	**Sites 29 and 42**
Probability that both sites are simultaneously dry	1 hour	0.769 (0.787)	0.793 (0.754)	0.753 (0.728)
	6 hours	0.419 (0.676)	0.500 (0.517)	0.411 (0.508)
	24 hours	0.132 (0.471)	0.161 (0.199)	0.132 (0.196)
Lag 0 cross-correlation	1 hour	**0.885 (0.811)**	**0.778 (0.772)**	**0.743 (0.721)**
	6 hours	0.952 (0.921)	0.957 (0.923)	0.902 (0.920)
	24 hours	0.971 (0.985)	0.977 (0.955)	0.947 (0.954)
Lag 1 cross-correlation	1 hour	**0.513 (0.336)**	**0.637 (0.561)**	**0.447 (0.358)**
	6 hours	0.193 (0.223)	0.259 (0.298)	0.261 (0.229)
	24 hours	0.025 (0.065)	0.059 (0.083)	0.036 (0.066)

Table 5. Parameter estimates for single-site models fitted to hourly data from gauge 19, December 1994. Top: dependent depth-duration model, fitted using method of moments. Bottom: random parameter Bartlett-Lewis model, fitted using both the method of moments and the spectral method.

Dependent depth/Duration model		
Parameter		Estimate
Storm arrival rate		0.027 h^{-1}
Mean number of cells per storm		8.3
Mean cell duration		5.4 h
Mean temporal displacement of cells from storm origin		1.36 h
c		0.79
f		3.47
Random parameter model		
Parameter	Moment estimate	Spectral estimate
Storm arrival rate	0.063 h^{-1}	0.021 h^{-1}
Mean number of cells per storm	20.65	9.04
Mean cell intensity	1.800 mm h^{-1}	1.307 mm h^{-1}
Mean of cell duration parameter	0.07 h^{-1}	2.860 h^{-1}
Std. Dev. of cell duration parameter	0.01 h^{-1}	1.728 h^{-1}
Cell duration/spacing ratio	15.7	0.0045

Table 6. Assessment of model fit for single-site models fitted to hourly raingauge data from gauge 19, for December 1994. H: Historical values. M_1: model with dependent duration and intensity as given in equation (2) with d = 0. $M_2^{(mom)}$: random parameter Bartlett-Lewis model, fitted using method of moments. $M_2^{(spec)}$: random parameter Bartlett-Lewis model, fitted using spectral method. Highlighted entries indicate features used in the moments fitting.

Level of aggregation		Mean	Prob. of no rain	Variance	Correlation		
					Lag 1	Lag 2	Lag 3
1 hour	H	**0.153**	**0.824**	**0.263**	**0.524**	0.279	0.188
	M_1	**0.153**	**0.824**	**0.264**	**0.524**	0.281	0.222
	$M_2^{(mom)}$	**0.153**	**0.824**	**0.264**	**0.524**	0.341	0.259
	$M_2^{(spec)}$	0.153	0.722	0.256	0.509	0.236	0.142
6 hours	H	0.916	0.556	4.271	0.274	0.186	0.021
	M_1	0.918	0.501	4.207	0.393	0.202	0.110
	$M_2^{(mom)}$	0.918	0.600	4.400	0.369	0.162	0.100
	$M_2^{(spec)}$	0.918	0.263	3.664	0.207	0.055	0.031
12 hours	H	1.832	0.387	10.805	0.255	0.127	0.061
	M_1	1.836	0.369	11.721	0.325	0.094	0.028
	$M_2^{(mom)}$	1.836	0.412	12.045	0.290	0.107	0.063
	$M_2^{(spec)}$	1.832	0.085	8.844	0.109	0.038	0.030
24 hours	H	3.665	**0.194**	**31.064**	0.112	0.081	−0.025
	M_1	3.672	**0.194**	**31.064**	0.204	0.018	0.002
	$M_2^{(mom)}$	3.672	**0.194**	**31.064**	0.220	0.071	0.041
	$M_2^{(spec)}$	3.665	0.011	20.242	0.109	0.038	0.030

aggregation which were not explicitly included in the fitting procedure, the dependent depth-duration model underestimates the probability of no rain whereas the random parameter model overestimates it; however, the magnitude of the discrepancy is approximately the same for each model. Hence, the dependent depth-duration model may offer an alternative to the random-parameter model, whose performance is roughly comparable.

Comparing the two different fits of the random parameter model in Table 6, it is clear that the spectral parameter set produces a much worse fit to these features of the data than does the moments set (the exact fit to the mean of the data in the spectral case was proved theoretically by Chandler (1996*a*) and is in fact used explicitly in the fitting procedure to ease computation). The poor fit to the wet/dry properties is to some extent expected, for the spectral method relies upon means, variances and covariances only. However, the evident lack of fit to the variances and covariances is somewhat surprising. Although this is a rather more extreme example than many encountered in fitting different

Table 7. Approximate log-likelihoods obtained using the spectral method to fit five different models to 15-minute data from gauge 19, for the months of June and December 1994. The 'best' models are those with the highest log-likelihoods.

Model	No. of parameters	Description	Log-likelihood	
			June	December
Poisson	3	No clustering of cells into storms	978.14	718.38
Bartlett-Lewis	5	Cells follow storm origin in a Poisson Process	986.05	739.25
Neyman-Scott 1	5	Symmetric scatter of cells about storm origin (Gaussian displacement)	985.75	738.34
Neyman-Scott 2	5	More cells at beginning of storm (Exponential displacement)	986.05	739.72
Random parameter Bartlett-Lewis	6	As Bartlett-Lewis, but storms can occur at different timescales	986.13	740.52

datasets, it would appear that in general this estimation method is not suitable for providing parameters for hydrological application (this comment regarding likelihood-based methods for this sort of model has been made by other authors such as Rodriguez-Iturbe *et al.* (1987)). In the next section the uses of the method are discussed in more detail.

Fitting using the spectral method

The spectral method has so far been used extensively in the fitting of single-site models (Chandler, 1997), and some preliminary work on the fitting of spatial-temporal models has also been done. The main area of interest has been that of model comparison, as it is here that the apparent objectivity of the method is particularly useful. In the single-site case, numerous different models have been fitted to data from the HYREX raingauge network. Table 7 shows the log-likelihoods obtained when the method was used to fit 5 different models to 15-minute raingauge data from June and December 1994. Rigorous procedures for model comparison, such as likelihood ratio tests, are available which allow for the different numbers of parameters in the models. The clustering models fit the data better than does the Poisson model and storms tend to be asymmetric with more intense activity towards the beginning of a storm than at the end. The same conclusions hold for all the data analysed. These results are in line with those reported by previous authors (e.g. Rodriguez-Iturbe *et al.*, 1988).

Spatial-temporal model fitting using the spectral method has so far been confined to two simple models: the EPPM already described and a simpler model with circular cells. Work has been restricted to an exploration of the parameter space for a few individual rain events, including that of 6/2/94 whose results were reported above for the GDSTM. The main conclusions are that elliptical cells provide a better fit to the data than circular cells, and that there is clustering in the data which is not captured by the simple Poisson models (Chandler, 1996b). This confirms the earlier results.

In general the spectral method suffers from being computationally expensive. This is particularly true of its application in a spatial-temporal context, where computation time is currently seriously restrictive. Recent theoretical developments (Chandler, 1996a) have, however, allowed its speed to be increased substantially for the fitting of single-site models and, in principle, may be applied to the spatial-temporal fitting problem in exactly the same way. The method provides a useful complement to the method of moments, particularly regarding model choice and comparison; and it makes more efficient use of available data, which allows model fitting to proceed based on shorter runs of data than does the moments method.

SUMMARY OF MODELLING WORK

The modelling work reported above makes a significant contribution to ability to generate realistic rainfall scenarios using a stochastic mechanism. In particular, the full spatial-temporal models have extended the preliminary work of Cox and Isham (1988) to the stage where the use of the models for hydrological design, on an event-by-event basis, is feasible. The availability of an extensive quantity of continuous radar data in the HYREX database has been invaluable in the development of the models, both in guiding their construction and in calibrating them against data. Continuing investigation of models for rainfall at a single site indicates directions in which the spatial-temporal models may be improved in the future.

Rainfall disaggregation

The models described in the previous section are ideally suited for use in problems where design rainfall is required; the stochastic nature of the models allows an ensemble of different realisations to be generated as input to models of the surface hydrology.: A rather different problem arises in linking the output of an atmospheric General Circulation Model (GCM) to the behaviour of surface processes. GCMs operate over grid squares, the size of which can be of the order of 200–300 km^2. The atmospheric component of such a model generates the mean rainfall intensity over each grid square at each timestep. To use this GCM output for hydrological applications, therefore, a disaggregation scheme is required to represent the distribution of rainfall at smaller spatial scales so as to provide a realistic input to hydrological models of surface processes.

Currently, methods of disaggregating GCM output assume that the *coverage* (that is, the proportion of a grid square which is experiencing precipitation in a fixed time interval) is a prespecified quantity depending only upon the rainfall type, and that where it rains the point rainfall depth is exponentially distributed (Warrilow *et al.*, 1986). Previous analyses of coverage in the UK have shown that this does not reflect the spatial structure of real radar fields (Onof and Wheater, 1996a) and that the time-series of consecutive coverages is highly correlated at short time-intervals and has a long memory (Onof and Wheater, 1996b), a feature which is not reproduced by the current scheme.

The next section presents analyses of the dependence of coverage upon scale, using Wardon Hill radar data. Such analyses are important to enable a more realistic choice of the coverage in the existing disaggregation scheme. Subsequently, a different approach to the disaggregation problem is examined, which allows for the generation of disaggregated fields by assuming that rainfall fields have a Markovian spatial dependence structure. In what follows, a single temporal snapshot of the spatial rainfall field is referred to as an *image*.

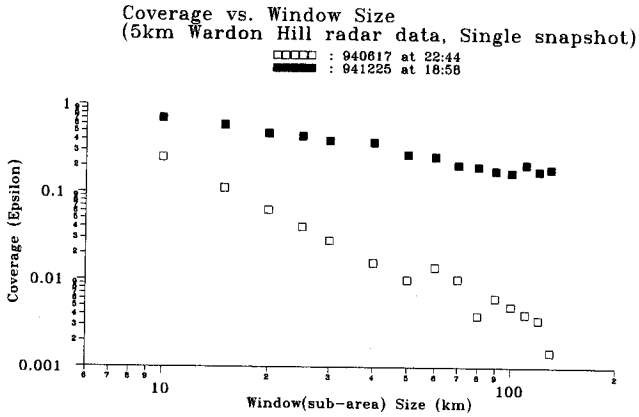

Fig. 9. *Log-log plot showing linear variation of coverage with window size for single snapshot from Wardon Hill radar station, events of 17/6/94 and 25/12/94.*

EMPIRICAL ANALYSES

The estimated rainfall coverage obtained from any image depends on two factors:

- the spatial resolution of the image, which is referred to as the *pixel size*;
- the size of the area studied, referred to as the *window size*.

Variations of the coverage with these two factors are investigated. In general, the coverage for a particular window of data is estimated by the proportion of pixels within that window which are wet; windows which are completely dry are excluded from the calculations.

Dependence of the rainfall coverage on the window size

Since GCMs are being used for a variety of grid-square sizes ranging from side lengths of a few hundred km to the meso-scale, it is essential to know how the rainfall coverage varies with the size of the grid-square over which it is estimated. This can be achieved, for any window size smaller than the range of a radar station, by defining a grid of windows inside the radar range, estimating the coverage over each window for which it is non-zero, and taking the mean value over all these windows. The estimated coverage is found to decrease with increasing window size, both for instantaneous rainfall fields and for fields which are temporally aggregated over a time-interval of a few days. The relationship invariably appears linear on a log-log scale – Fig. 9 is a typical example. This clear scaling pattern is maintained for most window sizes all the way up to the range of the radar; however, a deviation from this behaviour occurs for small window sizes approaching the pixel size. This is due to the constraint of unit coverage when the window size is the same as the pixel size.

Dependence of the rainfall coverage upon pixel size

The coverage ε which is used in the GCM is not dependent upon resolution since it is simply a proportion of the area which is wet; however, the coverage analysis above is based on data at the resolution of the radar pixels. It is therefore important to determine whether the coverages estimated from radar pictures exhibit a strong dependence on the pixel size.

Given radar data for either 2×2 or 5×5 km^2 pixel size, the resolution can be degraded by aggregation of the pixels so that a coarser-scale image can be obtained. In this manner, the coverage for a range of resolutions can be obtained (Fig. 10a). The variation of the rainfall coverage with the resolution is considerable and for the smaller scales, appears almost linear. The problem of how to provide an estimate for ε on this basis remains open.

However, this analysis is not wholly appropriate for determining the dependence of coverage upon pixel size. The problem is that radar estimates the true rainfall field for rainfall intensities above a certain threshold. When aggregating as above, rainfall intensities smaller than the given threshold will be obtained but have so far been counted as representing rain. The best way to mimic the mechanism of the radar and ensure a meaningful compari-

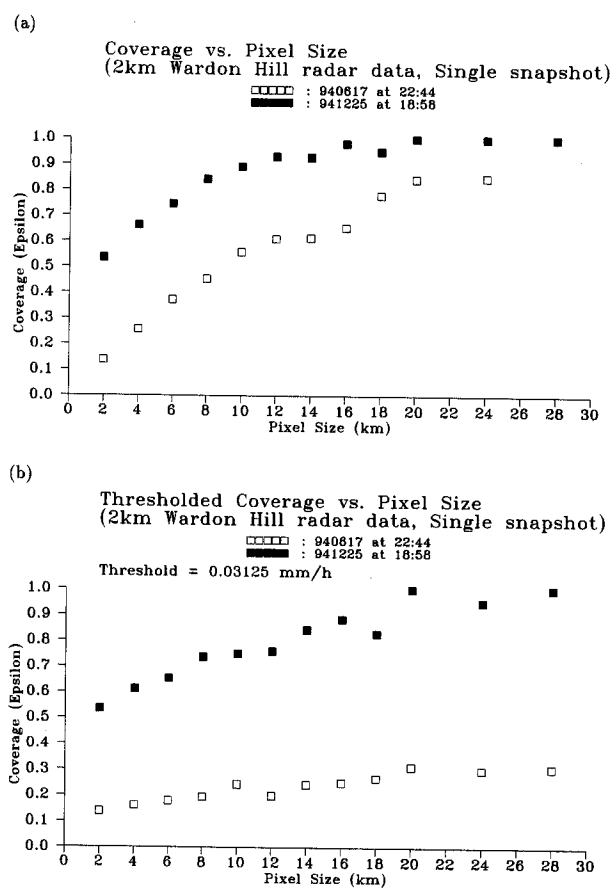

Fig. 10. *Variation of coverage with pixel size for Wardon Hill radar station, events of 17/6/94 and 25/12/94. (a) raw data, (b) data thresholded to mimic the behaviour of the radar station at coarser spatial scales.*

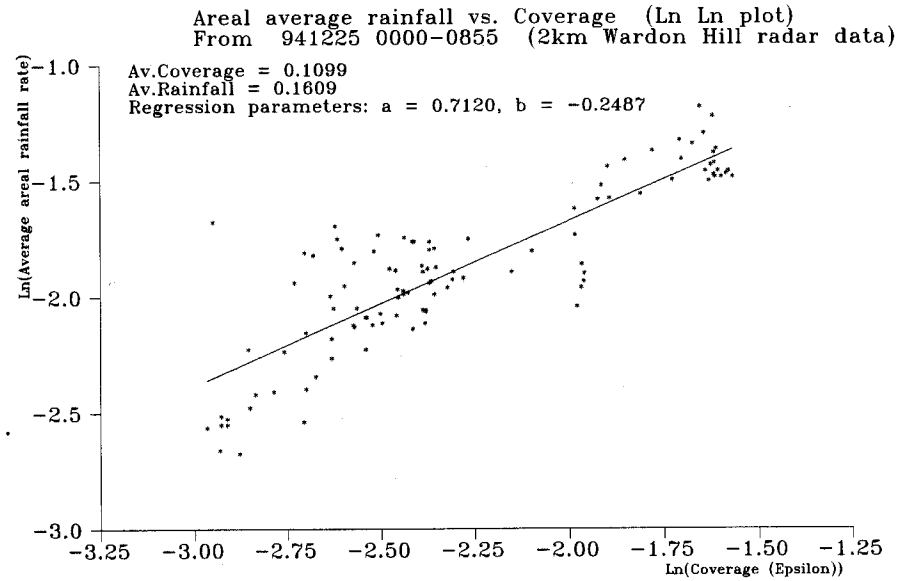

Fig. 11. *Log-log plot showing relationship between coverage and mean areal rainfall at a 2 × 2 km² pixel scale, Wardon Hill radar station, event of 25/12/94.*

son over the full range of pixel scales is for all such intensities to be set to zero.

With this new convention, the variation of the coverage with pixel size becomes much less, particularly for small scales (see Fig. 10b). This means that the 'true' coverage is well estimated by the coverage for small pixel sizes (2 × 2 or 5 × 5 km^2).

Inter-dependence of rainfall coverage and mean areal rainfall intensity

The atmospheric component of the GCM provides one value of mean areal intensity. An analysis of the radar data shows that this value is not independent of the coverage observed (Fig. 11). This dependence is probably best represented in the form of a linear regression over the logarithms of both quantities:

$$\ln(\varepsilon_t) = a \ln \overline{X_t} + b \qquad (5)$$

where ε_t and $\overline{X_t}$ denote respectively the coverage and the mean areal rainfall at time t. The squared multiple correlation (R^2) here is approximately 0.75. This relationship provides the basis for an improved disaggregation model which allows for temporal dependence by linking coverage to the mean areal rainfall intensity, which itself incorporates temporal dependence. Spatial memory, i.e. the fact that if it rains at time-step t at point M, there may be an increased probability of rain in the neighbourhood of M at time $t + 1$, is not addressed by this scheme. To deal with this problem, another approach, in which fields of disaggregated rainfall are generated, is considered.

MARKOV RANDOM FIELDS

An alternative approach to the disaggregation problem uses statistical image reconstruction techniques. The problems of image reconstruction and rainfall disaggregation share a common feature: in both cases, the objective is to reconstruct an unknown 'image' from imperfect data. In the case of image restoration, the imperfection is due to noise being added to the original image, whereas in the disaggregation problem, the imperfection arises because the data are available only as a very coarse-scale spatial average. Whereas in image processing there is a predetermined 'correct answer' to the reconstruction problem (the correct answer being the original image or 'true scene'), this is not necessarily the case in the rainfall disaggregation problem. In fact, it may be advantageous to have several different reconstructions of a rainfall field for design rainfall purposes.

The underlying concept behind statistical image reconstruction techniques is that of the *Markov Random Field*, in which the probability distribution for rainfall at a site ('site' being taken here to mean a pixel at a particular spatial scale) is specified conditionally on the pattern of rainfall in the *neighbourhood* of that site. As a simple example, consider the pattern of wet/dry pixels at a single time instant. An intuitively appealing neighbourhood structure may be obtained by defining two types of neighbours of site i; *type 1* neighbours are those immediately horizontally and vertically adjacent, and *type 2* neighbours are those immediately diagonally adjacent. Let X_i be a random variable taking the value 1 if site i is wet, 0 otherwise, and denote by $n_i^{(k)}$ (k = 1,2) the number of type k neighbours of

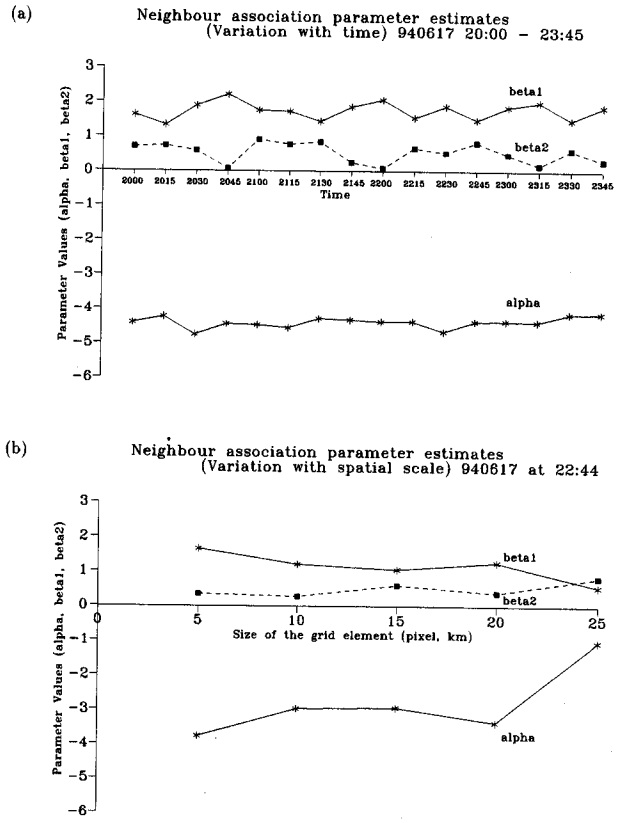

Fig. 12. *Estimated parameter values for simple Markov Random Field model of wet/dry rainfall pattern observed at Wardon Hill radar station, event of 17/6/94. (a) shows variation with time, (b) shows variation over different pixel sizes.*

site i which are wet (for example, $n_i^{(1)}$ is the total number of wet pixels lying immediately to the North, South, East or West of pixel i); then a Markov Random Field model may be specified as

$$\frac{Pr\{X_i = 1 |\ rest\ of\ scene\}}{Pr\{X_i = 0 |\ rest\ of\ scene\}} = \exp\left[\alpha + \sum_{k=1}^{2}\left(\beta_k n_i^{(k)}\right)\right]. \quad (6)$$

The specification of the model in terms of the conditional odds ratio is a consequence of the *Hammersley-Clifford Theorem*. For a fuller discussion of the theory of Markov Random Fields, see for example Besag (1974, 1986) and Isham (1981).

The parameters α, β_1 and β_2 in this model represent respectively some measure of the overall proportion of wet pixels, and measures of association in the vertical/horizontal and diagonal directions. Parameter estimation is straightforward, using methods such as the conditional likelihood approach suggested by Besag (1972). Figure 12 shows estimated parameter values from the event of 17/6/94 shown in Fig. 2, using the model specified in Eqn. 6. The variation of the parameter estimates with both time and spatial scale is shown. The stability of the estimates is encouraging, and is typical of the several events which have been analysed. In general, the estimates of β_1 and β_2 tend to be negatively correlated (presumably because of collinearity between $n_i^{(1)}$ and $n_i^{(2)}$ in the model at Eqn. 6, and β_1 tends to be lower in summer than in winter; moreover at coarser spatial scales there is less dependence in the rainfall field (indicated by larger values of α and smaller values of β_1 and β_2). All of these results are in line with what is expected – for example, the smaller values of β_1 in summer reflect the spottiness of summer rainfall relative to that in winter.

To apply this theory in a disaggregation context, an analogue of the *cascade algorithm* (Jennison and Jubb, 1988) can be used to reconstruct the rainfall field successively at finer and finer spatial resolutions. This has proved to be very successful at reconstructing the pattern of wet and dry pixels: the approach, and some results, are described by Onof *et al.* (1998) and in Chandler *et al.* (2000). This work will be extended in the future to include rainfall intensities rather than just the wet/dry pattern.

Applications

The results presented above indicate the complexity of spatial-temporal rainfall fields, but also point the way forward to new developments in the treatment of rainfall for hydrological applications.

The representation of rainfall in GCMs is clearly important in terms of the performance of the surface hydrology parameterisations and affects surface energy exchange and runoff production. The feedback between surface hydrology and climate is well established (Manabe, 1975); without rainfall disaggregation, runoff from major river systems such as the Nile cannot be represented (Abourgila, 1992) – a consequence of the highly nonlinear response of surface processes to rainfall inputs. The analyses of UK rainfall have pointed the way to simple improvements in current GCM procedures, for example, preserving temporal correlation of disaggregated fields (Onof and Wheater, 1996b) which are likely to be implemented shortly. However, problems of lack of spatial location of sub-grid-scale rainfall and associated lack of spatial memory in surface parameterisations remains a major difficulty. The work on Markov Random Fields provides potentially the basis of a method to overcome these difficulties.

More generally, the treatment of rainfall for hydrological design is primitive, and has thus far taken little account of the insights into spatial rainfall structure available from radar data. For example, current flood design practice in the UK (NERC, 1975) is based on uniform spatial distribution and a symmetric, unimodal temporal distribution of rainfall for small and medium-sized catchments (<500 km^2). For larger catchments, design guidance is unclear. The importance of these approximations is currently the subject of preliminary investigation, but one way forward could be based on the application of spatial-temporal models as

outlined above. This would require design to be cast in a framework of stochastic analysis, either based on alternative realisations of a design event, or from generated long sequences of continuous rainfall. A first step, considering stochastic generation of storm interior temporal distributions is reported by Onof *et al.* (1996). Since that time, spatial-temporal models have been developed to the point where they are suitable for continuous simulation, by incorporating a second level of clustering.

Acknowledgements

The authors are grateful to the NERC for their financial support (HSW, VSI, DRC, REC, CO, IRI) under the auspices of the HYREX project; also to the Greek State Scholarships Foundation (AK) and to the EPSRC (PJN) for research studentships. Additionally, they would like to thank Chris Collier for sharing his expertise on radar; Keith Browning (JCMM, Reading) for helping with the identification of events for analysis by checking the quality of radar data; and Sean Wood and Roger Austin at the Institute of Hydrology, for providing software support and database management facilities.

References

Abourgila, A.E., 1992. *Large-scale hydrological modelling of the Nile basin for the assessment of the impact of climate and land use changes*. MSc thesis, Department of Civil Engineering, Imperial College London.

Austin, P.H. and Houze, R.A., 1972. Analysis of the structure of precipitation patterns in New England. *J. Appl. Meterol*, **11**, 926–935.

Besag, J.E., 1972. On the statistical analysis of nearest-neighbour systems. In *Proc. Eur. meeting of Statisticians*. Budapest.

Besag, J.E., 1974. Spatial interaction and the statistical analysis of lattice systems. *J. Roy. Statist. Soc., Ser. B*, **36**, 192–236.

Besag, J.E., 1986. On the statistical analysis of dirty pictures. *J. Roy. Statist. Soc., Ser. B*, **48**, 259–302.

Brillinger, D.R. and Rosenblatt, M., 1967. Asymptotic theory of estimates of kth-order spectra. In *Spectral Analysis of Time Series*, (Ed. B. Harris), 189–232. Wiley, Chichester, UK.

Chandler, R.E., 1996a. A note on analytical solutions to the Whittle likelihood equation. *Technical report, no. 173, Department of Statistical Science, University College London*, London WC1E 6BT. http://www.ucl.ac.uk/Stats/research/abstracts.html.

Chandler, R.E., 1996b. The second-order spectral analysis of spatial-temporal rainfall models. *Technical report, no. 158, Department of Statistical Science, University College London*, London WC1E 6BT. http://www.ucl.ac.uk/Stats/research/abstracts.html.

Chandler, R.E., 1997. A spectral method for estimating parameters in rainfall models. *Bernoulli*, **3**, 301–322.

Chandler, R.E., Mackay, N., Onof, C. and Wheater, H.S., 2000. Bayesian image analysis and the disaggregation of rainfall. *J. Atmos. Ocean. Technol.*, to appear.

Collier, C.G., 1989. *Applications of weather radar systems: A guide to uses of radar data in meteorology and hydrology*. Ellis Horwood, Chichester, UK.

Cowpertwait, P.S.P., 1994. A generalised point process model for rainfall. *Proc. Roy. Soc. London*, **A447**, 23–37.

Cox, D.R. and Isham, V., 1988. A simple spatial-temporal model of rainfall. *Proc. Roy. Soc. London*, **A415**, 317–328.

Cox, D.R. and Isham, V., 1994. Stochastic models of precipitation. In *Statistics for the environment 2: Water related issues*, (Eds. V. Barnett and K. Turkman), 3–18. Wiley, Chichester, UK.

CWPU, 1977. *Dee weather radar and real time hydrological forecasting project*. Report by the Steering Committee, Central Water Planning Unit, Reading.

Fisher, N.I., 1993. *Statistical analysis of circular data*. Cambridge University Press, Cambridge, UK.

Gupta, V.K. and Waymire, E.C., 1990. Multiscaling properties of spatial rainfall and river flow distributions. *J. Geophys. Res.*, **95**, 1999–2010.

Holland, D.J., 1967. The Cardington rainfall experiment. *Meteorol. Mag.*, **96**, 193–202.

Imrie, C., 1996. *A classification of storms using radar measurements of precipitation*. Master's thesis, Department of Civil Engineering, Imperial College, University of London.

Isham, V., 1981. An introduction to spatial point processes and Markov random fields. *Int. Stat. Review*, **49**, 21–43.

Jennison, C.J. and Jubb, M.D., 1988. Statistical image resoration and refinement. In *Information processing in medical imaging*, (Eds. C. de Graaf and M. Viergever), 255–62. Plenum, New York.

Kakou, A., 1997. *Point process based models for rainfall*. PhD thesis, Department of Statistical Science, University College London.

Le Cam, L., 1961. A stochastic description of precipitation. In *Proc. of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, **3**, (Ed. J. Neyman), 165–186. Berkeley, California.

Lovejoy, S. and Schertzer, D., 1987. Physical modelling and analysis of rain and clouds by anisotropic scaling multiplicative processes. *J. Geophys. Res.*, **92**, 9693–9714.

Manabe, S., 1975. A study of the interaction between the hydrological cycle and climate using a mathematical model of the atmosphere. In *Proc. of the Conference on Weather and Food*, 10. MIT, Cambridge, Massachusetts.

Marshall, R.J., 1980. The estimation and distribution of storm movement and structure using a correlation analysis technique and raingauge data. *J. Hydrol.*, **48**, 19–39.

Mason, J., 1986. Numerical weather prediction. *Proc. Roy. Soc. London*, **A407**, 51–60.

McCullagh, P. and Nelder, J.A., 1989. *Generalized Linear Models (second edition)*. Chapman and Hall, London.

Moore, R.J., Jones, D.A., Cox, D.R. and Isham, V.S., 2000. Design of the HYREX raingauge network, *Hydrol. Earth System Sci.*, **4**, 523–530.

NERC, 1975. *Flood studies report*. Report of the Natural Environment Research Council (5 volumes).

Northrop, P., 1996. *Modelling and statistical analysis of spatial-temporal rainfall fields*. PhD thesis, Department of Statistical Science, University College, London.

Northrop, P., 1997. A clustered spatial-temporal model of rainfall. *Proc. Roy. Soc. London* **A454**, 1875–1888.

Onof, C., Faulkner, D. and Wheater, H.S., 1996. Design rainfall modelling in the Thames catchment. *Hydrol. Sci. J.*, **41**, 715–733.

Onof, C. and Wheater, H.S., 1996a. Analysis of the spatial coverage of British rainfall fields. *J. Hydrol.*, **176**, 97–113.

Onof, C. and Wheater, H.S., 1996b. Modelling of the time series of spatial coverages of British rainfall fields. *J. Hydrol.*, **176**, 115–131.

Onof, C., Mackay, N., Chandler, R.E. and Wheater, H.S., 1998. A rainfall disaggregation scheme for forecasting. In *Hydrology in a changing environment* (Eds. H. Wheater and C. Kirby), 107–116. Wiley, Chichester, UK.

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

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

Taylor, G.I., 1938. Statistical theory of turbulence. *Proc. Roy. Soc. London*, **A164**, 476–490.

Warrilow, D.A., Sangster, A.B. and Slingo, A., 1986. Modelling of land surface processes and their influence on the European

climate. *Technical report, no. DCTN38, UK Meteorological Office, Bracknell, Berks.*

Waymire, E., Gupta, V.K. and Rodriguez-Iturbe, I., 1984. A spectral theory of rainfall intensity at the meso-β scale. *Water Resour. Res.*, **20**, 1453–1465.

Wheater, H.S. and Brown, R.P.C., 1989. Limitations of design hydrographs in arid areas—an illustration from South-West Saudi Arabia. In *Proc. 2nd National BHS Symposium*, 349–356.

Wheater, H.S., Butler, A.P., Stewart, E.J. and Hamilton, G.S., 1991a. A multivariate spatial-temporal model of rainfall in South-West Saudi Arabia. I, Data characteristics and model formulation. *J. Hydrol.*, **125**, 175–199.

Wheater, H.S., Onof, C., Butler, A.P. and Hamilton, G.S., 1991b. A multivariate spatial-temporal model of rainfall in South-West Saudi Arabia. II. Regional analysis and model validation. *J. Hydrol.*, **125**, 201–220.

Wheater, H.S., Isham, V.S., Cox, D.R., Chandler, R.E., Kakou, A., Northrop, P.J., Oh, L., Onof, C. and Rodriguez-Iturbe, I., 1997. Spatial-temporal rainfall fields: modelling and statistical aspects. *Research Report No. 176, Department of Statistical Science, University College London* (http://www.ucl.ac.uk/Stats/research/rr176/176.html).

Wood, S.J., Jones, D.A. and Moore, R.J., 2000a. Accuracy of rainfall measurement for scales of hydrological interest. *Hydrol. Earth System Sci.*, **4**, 531–543.

Wood, S.J., Jones, D.A. and Moore, R.J., 2000b. Static and dynamic calibration of radar data for hydrological use. *Hydrol. Earth System Sci.*, **4**, 545–554.

Zawadski, I.I., 1973. Statistical properties of precipitation patterns. *J. Appl. Meteorol.*, **12**, 459–472.