 New Approaches to Hydrological Prediction in Data-sparse Regions (Proc. of Symposium HS.2 at the Joint IAHS & IAH Convention, Hyderabad, India, September 2009). IAHS Publ. 333, 2009.

A high-resolution hierarchical space–time framework for single storm events and its application for short-term rainfall forecasting

JUAN QIN[1], MICHAEL LEONARD[2], GEORGE KUCZERA[1], MARK THYER[1], ANDREW METCALFE[2] & MARTIN LAMBERT[2]

1 *School of Engineering, The University of Newcastle, Callaghan, New South Wales 2308, Australia*

2 *School of Civil, Environmental & Mining Engineering, The University of Adelaide, Adelaide 5005, Australia*

juan.qin@studentmail.newcastle.edu.au

Abstract A new phenomenological hierarchical stochastic model is developed to robustly simulate rainfall fields consistent with 10-minute 1-km^2 pixel radar images. The hierarchical framework has three levels. The first level simulates a latent Gaussian random field conditioned on the previous time step. In the second level, first-order autoregressive (AR(1)) models are used to describe the within-storm variations of the level-one parameters that control the evolution of rain fields. The third level is designed for simulation of storm sequences. Calibration and validation of the first two levels using an observed storm event (typical in Sydney, Australia) demonstrate that this two-level model produces realistic sequences of rain images which capture the physical hierarchical structure of clusters, patchiness of rain fields and the persistence exhibited during storm development. A variety of important statistics are adequately reproduced at both 10-minute and 1-hour time scales over various space scales. Application of this model to short-term rainfall forecasting is also presented.

Key words stochastic space-time rainfall; hierarchical framework; high-resolution; block Toeplitz, circulant decomposition; generalized moments; parametric bootstrap; short-term forecasting

INTRODUCTION

The spatial and temporal variation of rainfall over small scales substantially affects the estimation of runoff, particularly for urban catchments where the time of concentration is typically small. A rainfall model that can adequately reproduce the spatio-temporal variation of important rainfield characteristics is highly desirable. In recent years, high-resolution stochastic space-time rainfall modelling has received increased attention from hydrologists. Various stochastic modelling approaches have been developed to characterize the rainfall variation in both time and space, including cluster-based models (e.g. Waymire *et al*., 1984; Mellor *et al*., 2000; Cowpertwait *et al*., 2002), multi-fractal models (e.g. Seed *et al*., 1999), and other random field models (e.g. Pegram *et al*., 2001). Most of these models, however, make simplifying assumptions on the temporal characteristics of rainfall fields and thus cannot be expected to reproduce important statistics, especially at small time and space scales. In particular, cluster-based models usually attempt to describe the physical structure of rainfield and thus lead to a large number of model parameters. Some common deficiencies associated with this class of model include an inability to generate sub-hourly rainfall due to simplified cluster mechanisms, and an inability to reproduce statistics at time scales not used for calibration (e.g. Cowpertwait *et al*., 2002). The space-time rainfall forecasting model developed by Mellor *et al*. (2000), for example, despite being conditioned upon the latest observed radar images and allowing the spatial covariance structure to change over time, has difficulty in dealing with storms with a less well-defined structure such as convective events. Compared to cluster-based models, the main attraction of multifractal models is that their self similarity property enables parsimonious parameterization of rainfall fields over a wide range of scales. However, they are not without disadvantages, such as the issue of imperfect scaling of space–time rainfall, making some subjective assumptions on space–time cascade generators (e.g. Over & Gupta, 1996; Krajewski, 1996). So far these models have not been used for long-term simulation and applications have only been restricted to single storm events. These current state-of-the-art scale invariant models are still at the stage of theoretical development and there is still a long way to go in their practical application (Mellor *et al*., 2000). Apart from the above two

primary approaches, other theories of random field, e.g. Gaussian Markov Random Field (GMRF), Spatial-Temporal Auto-Regressive Moving-Average (STARMA) and other Gaussian random fields, have also been applied in spatio-temporal rainfall modelling. Among them, models of GMRF and STARMA share the Markovian property in spatial correlation which typically produces inaccurate marginal variances. Other Gaussian random field models have their own limitations as well. For instance, the "String of Beads" model (Pegram & Clothier, 2001) requires rain intensity fields to be lognormally distributed.

To overcome these limitations, an improved rainfall model should therefore be:

1. capable of generating images conditioned on previous time steps;
2. able to adequately reproduce the temporal variation of important rainfield characteristics by allowing model parameters to evolve over time;
3. capable of dealing with different types of storms, e.g. frontal and convective systems;
4. flexible with regard to the distribution of rainfall intensity; and
5. capable of modelling storm sequences for long-term simulations, not only for the simulation of single storm events.

In this paper, we develop a new model framework that satisfies (1)–(4), with (5) being reserved for future work.

A THREE-LEVEL HIERARCHICAL FRAMEWORK FOR RAINFALL FIELDS

Statistical characteristics of rainstorms derived from radar rainfall images (Qin *et al*., 2006) suggest that the modelling of space–time rainfall is challenging because the parameters controlling the stochastic evolution of the images are themselves evolving during the storm according to a stochastic process. Moreover, there are different storm types implying the parameters of the stochastic process vary between storms. This hierarchical structure is expected to be critical in the development of a high-resolution model. In this section, a new hierarchical framework is proposed for simulating high-resolution rainfall fields. The core idea is to simulate temporally varying Gaussian random fields that are conditioned on the rainfall field at the previous time step. To do so, the rainfall data must first be transformed to remove skewness and to accommodate the patchiness of rainfields using a latent Gaussian variable approach (for details, see Qin *et al*., 2006).

Description of the Three-Level Hierarchical Model

Figure 1 presents the three-level hierarchical structure of the model. In Level 1, the transformed rain fields y_{t+1} conditioned on the previous time step y_t are simulated, and are then transformed back into real rainfall fields r_{t+1} through the transformation coefficient γ using equation (1). The parameters of the multivariate normal model, $\mu_{t+1|t}$ and $\Sigma_{t+1|t}$, can be derived using equation (2) through some intermediate variables that are directly linked to Level 2. The mean and standard deviation of y_t and y_{t+1} ($\mu_{y_t}, \sigma_{y_t}, \mu_{y_{t+1}}, \sigma_{y_{t+1}}$) can be obtained using equation (3) by sampling γ, μ_R and WAR from the level-two AR(1) models. The spatial correlations of y_t and y_{t+1} ($\rho_{y_t}, \rho_{y_{t+1}}$) can be obtained using equation (4) by sampling α and β from its level-two AR(1) model. The temporal correlation between y_t and y_{t+1} (ρ_t) and the advection velocity (υ, θ), respectively, are sampled from their AR(1) models. The parameters of these AR(1) models in Level 2 are classified into different sets and fitted using Level-3a probability models conditioned on the storm gross characteristics (d, t_d, t_s) generated by a Level-3b probability model.

$$r = \begin{cases} y^{1/\gamma}, & y > 0 \\ 0, & y \le 0 \end{cases} \tag{1}$$

where r represents the observed rainfall data, y denotes the transformed latent Gaussian variable, and γ is the transformation coefficient.

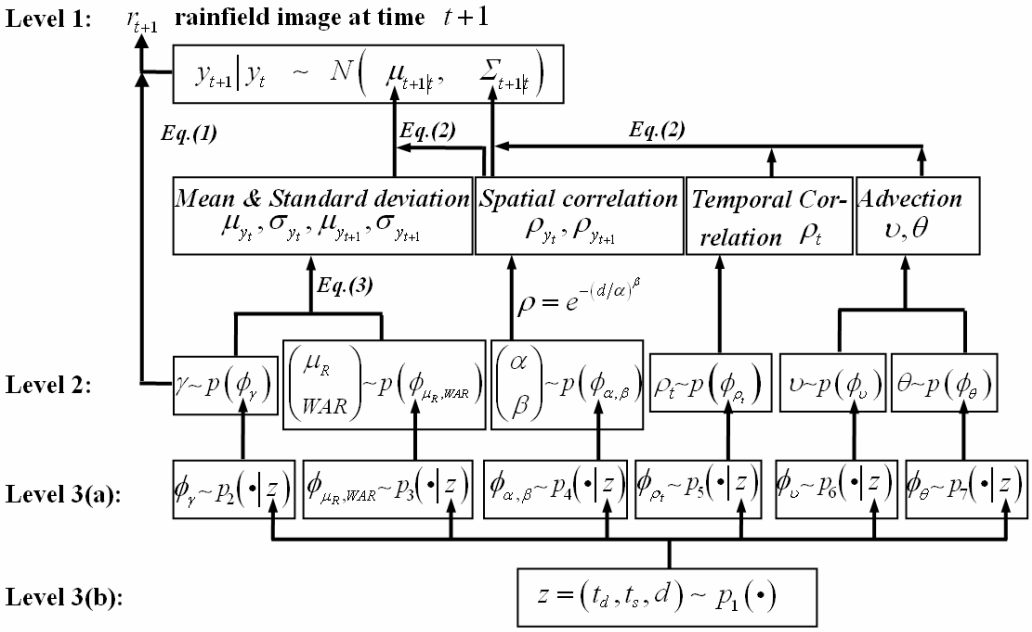

Fig. 1 Structure of the three-level hierarchical model.

$$y_{t+1} \mid y_t \sim N(\mu_{t+1} + \Sigma_{t+1,t}\Sigma_{t,t}^{-1}(y_t - \mu_t), \Sigma_{t+1,t+1} - \Sigma_{t+1,t}\Sigma_{t,t}^{-1}\Sigma_{t,t+1}) \tag{2}$$

In equation (2), μ_t is the mean areal intensity for image t, $\Sigma_{t,t}$ denotes marginal spatial covariance of image t.

$$\begin{cases} \mu_R = \int_0^\infty y^{1/\gamma} p(y \mid \mu_y, \sigma_y^2) dy \\ \mu_y + z_{y=0}\sigma_y = 0 \end{cases} \tag{3}$$

where $z_{y=0}$ represents the normal score corresponding to $y = 0$ and has an exceedence probability equal to the wet area ratio (WAR), and μ_y and σ_y, respectively, are the mean and standard deviation of y:

$$\rho = \exp\left(-\left(\frac{d}{\alpha}\right)^\beta\right) \tag{4}$$

where the power exponential class of spatial correlation functions are described by two parameters, a correlation length scale (α) and a shape parameter (β).

The first two of the three levels are discussed in detail in the following two subsections. The level 3 models are left for future work.

Level 1: Conditional rainfield model: $y_{t+1} \mid y_t$

Due to the computational limitations of the Cholesky decomposition method, the conditional space–time rainfall model is built in a Lagrangian system with the advective velocity vector used to convert Lagrangian coordinates to Eulerian ones. Due to the exclusion of velocity, the space–time covariance in the model now becomes separable, and can be written as the product of a purely spatial covariance matrix and a purely temporal covariance matrix. Noting that $Cov(y_{t+1})$ is block Toeplitz and can be embedded in a circulant matrix, the circulant embedding and the Fourier transform technique can then be employed to efficiently generate random fields (Dietrich & Newsman,

1997). Details of the formulation of the conditional rainfield model, $y_{t+1} \mid y_t$, can be found in Qin *et al*. (2008). Due to space limitations, only the final result of the derivation is given below.
The conditional random variable y_{t+1} given y_t can be expressed as:

$$y_{t+1|t} = F_{t+1}\sqrt{\Delta_{t+1}}\,\varepsilon_{t+1|t} \tag{5}$$

where y_t is some $(n \times 1)$ vector of correlated Gaussian random variables at time t, ε_t are uncorrelated $(n \times 1)$ Gaussian variables $\sim N(0, I)$, ε_t and ε_{t+1} are correlated in time according to some correlation ρ_t (the same for all pixels), F is a Fourier matrix, Δ is a diagonal matrix of eigenvalues of a circulant matrix of $Cov(y_{t+1})$, and

$$\varepsilon_{t+1|t} = \rho_t \varepsilon_t + \sqrt{1-\rho_t^{\,2}}\;\varepsilon_{t+1} \tag{6}$$

This is a new result. Equation (5) can be solved using fast Fourier transforms which results in remarkably efficient generation of conditional normal variables. It takes approximately one second to generate a rainfall image with a 256 × 256 lattice grid on a computer with a 2.0GHz CPU.

Level 2: AR(1) Models for the Level 1 model parameters: $(\mu_R, WAR, \gamma, \nu, \theta, \alpha, \beta, \rho_t)$

Investigation of the rainfall power-law transformation shows that μ_y, σ_y and γ are highly correlated with each other, while μ_R and WAR have low correlation with γ. Therefore, μ_R and WAR are the better choice for model parameters in terms of sampling. Given μ_R, WAR and γ, μ_y and σ_y can be obtained using equation (3). The time series analysis in Qin *et al.* (2006) suggests that μ_R and WAR are positively correlated, and that α and β are negatively correlated, and that two bivariate AR(1) models are appropriate for simulating these variables. Univariate AR(1) models can be used to simulate the remaining level-one parameters v, θ, γ and ρ_t.

DATA

The Kurnell radar images, provided by Australian Bureau of Meteorology (Seed, pers. comm., 2005), have 1-km spatial and 10-minute temporal resolution covering a 256 km × 256 km area centred about 20 km south of Sydney. These reflectivity data were converted to rainfall intensity by using the following climatological relationship for Sydney, reflectivity = 280rain$^{1.6}$, which has been calibrated on raingauge data. The following criterion separates the images into different storms: prior to the start of a new storm, the wet area ratio (*WAR*) must be less than 5% for more than two hours. Though intuitive, this criterion is arbitrary but nonetheless convenient.

The event (storm 37) observed on 18 December 2000, commencing at 02:55h and ending at 12:35h UTC, is used to demonstrate the calibration and validation of the high resolution space-time rainfall model. In this paper, meanR (μ_r) and stdR (σ_r) denote the mean and standard deviation over all pixels (both wet and dry) within an image; the overall meanR and stdR present the mean and standard deviation over all pixels through both space and time within a storm; while wet or conditional meanR and stdR represent the mean and standard deviation over all wet pixels through space and time.

CALIBRATION AND VALIDATION OF THE FIRST TWO LEVELS OF THE HIERARCHICAL MODEL

The method of moments and the maximum likelihood are two common approaches for model parameter calibration. Due to bias in maximum likelihood estimates of the latent variables, the method of moments is used for calibration. Qin *et al.* (2008) compared two different model calibration strategies, namely the method of moments and the method of generalized moments.

The results suggest that the method of generalized moments produces more realistic images and is able to reproduce aggregated statistics over a wide range of space and time scales. Details on calibration and validation of the first level of the model can be found in Qin *et al.* (2008).

Calibration of the first two levels

For level one, eight model parameters are calibrated. Among them, μ_y, σ_y and γ are highly correlated and are jointly estimated from the data; α is calibrated using generalized moments by fixing $\beta = 1.5$ (Qin *et al.*, 2008); υ and θ are estimated by maximizing the cross correlation between two consecutive images (Qin *et al.*, 2006) and the lag 1 temporal correlation of the observed rainfall data r is then the maximal cross correlation between two images. Once the lag 1 temporal correlation of the observed rainfall data is obtained, the approach of Bell (1987) can be applied to work out ρ_t in "*y*" space. For details please refer to Qin *et al.*, 2008. For level two, it was found that a bivariate AR(1) model is capable of simulating the time series of $\log \mu_R$ and log*WAR* for most storms, and that a univariate AR(1) model is appropriate for simulating the time series of γ, α, υ, θ and ρ_t, respectively.

Validation of the first two levels as a whole

The parametric bootstrap validation technique is used to verify the space–time model for the selected storm by comparing the observed and simulated statistics of interest. The test will focus on the following characteristics: (1) visual comparison of historical and simulated radar images; (2) basic summary statistics, including the unconditional (all pixels involved) and conditional (only wet pixels involved) overall meanR, stdR and *WAR* over a wide range of spatial and temporal scales (Wheater *et al*., 2000). The unconditional statistics provide an overall picture of the entire field, while the conditional statistics allows focusing on the rain events themselves. The space scales considered for model verification are 1 km, 2 km, 4 km, 8 km, 16 km and 32 km. The temporal scales are 10-minutes and hourly.

For a given storm, using the initial values of the time series of the estimated level-one parameters to initialize the corresponding level-two AR(1) models, 100 replicates are generated as follows: Firstly, the eight level-one parameters (μ_y, σ_y, γ, υ, θ, α, β and ρ_t) are sampled from the fitted level-two AR(1) models; secondly, transformed rain images in "y" space are generated one by one using the eight sampled parameters and then transformed back into real rainfall fields in "r" space using the power-law transformation. Then the sampling distributions of important statistics over a wide range of space–time scales are calculated and compared with the observed storm statistics. For each simulation, the initial image was marginally sampled and the subsequent images were generated sequentially by conditioning on the previous image.

Figure 2 presents a visual comparison of the historical and simulated radar images for storm 37 at the 10-minute scale. Although each replicate of the sequence of three consecutive images may look quite different from the observed one, the persistence of the clusters seems to be well preserved. More importantly, as demonstrated in Fig. 3, 100 simulations of storm 37 show that the seven important statistics averaged throughout the storm have been well reproduced for both 10-minute and 1-hour aggregation over spatial scales ranging from 1 km to 32 km. The reproduction of the unconditional class of statistics, such as the overall unconditional mean, overall unconditional standard deviation, the unconditional within-image standard deviation averaged over all images of the storm, and the wet area ratio, ensures that the statistics over the entire field domain (including both wet and dry pixels) at various space scales are maintained. The preservation of the conditional category of statistics, on the other hand, ensures the main features of the rain event itself (wet pixels only) are reproduced. This is a very encouraging result as it is well known how difficult it is for a stochastic rainfall model to reproduce statistics over various aggregated space and time scales, and particularly over those not used for model calibration (e.g. Seed *et al.* 1999; Pegram & Clothier, 2001; Cowpertwait *et al.* 2002).

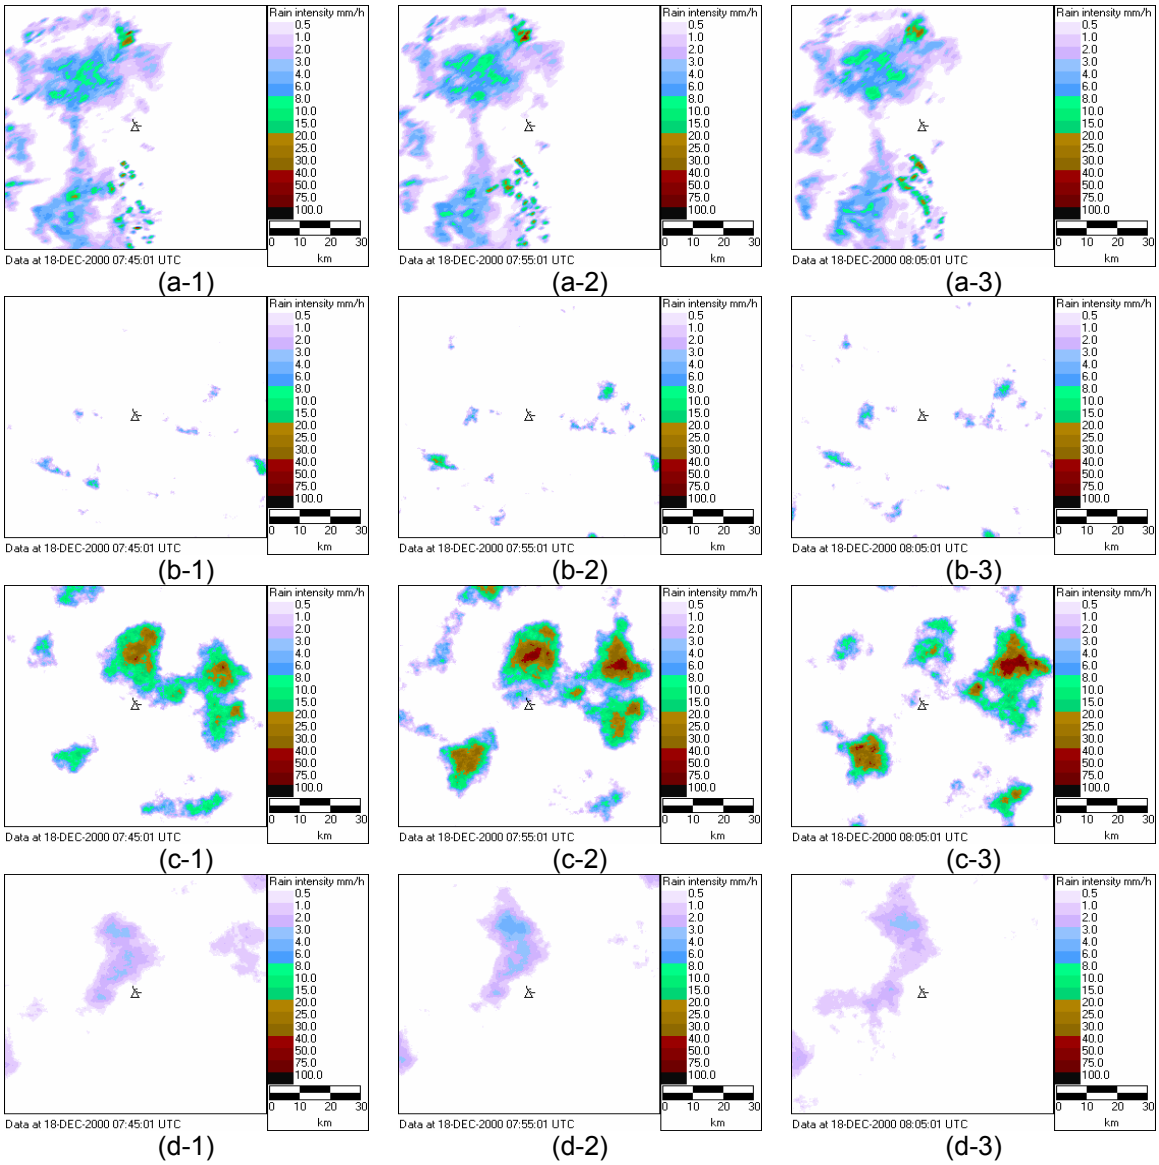

Fig. 2 Comparison of observed and simulated rainfall fields (with 1-km spatial resolution) for three consecutive images in storm 37 with images aggregated at 10-minute time scale: (a) observed sequence of three consecutive images; (b) (c) (d) three replicates of three consecutive images.

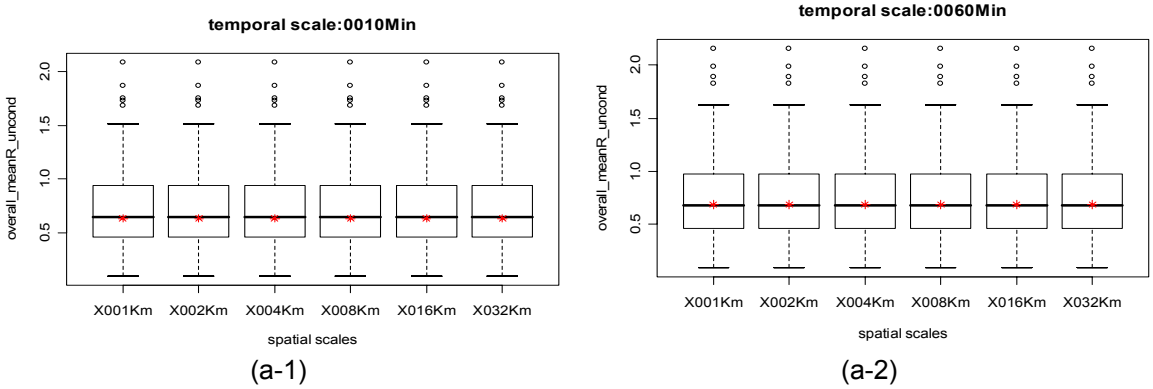

Fig. 3 Boxplots of important statistics over various space-time scales for 100 replicates of storm 37 generated by using level-two AR(1) models to generate the eight Level 1 model parameters that control the evolution of rain images. Red stars denote the observed values. See next page for parts (b)–(e).

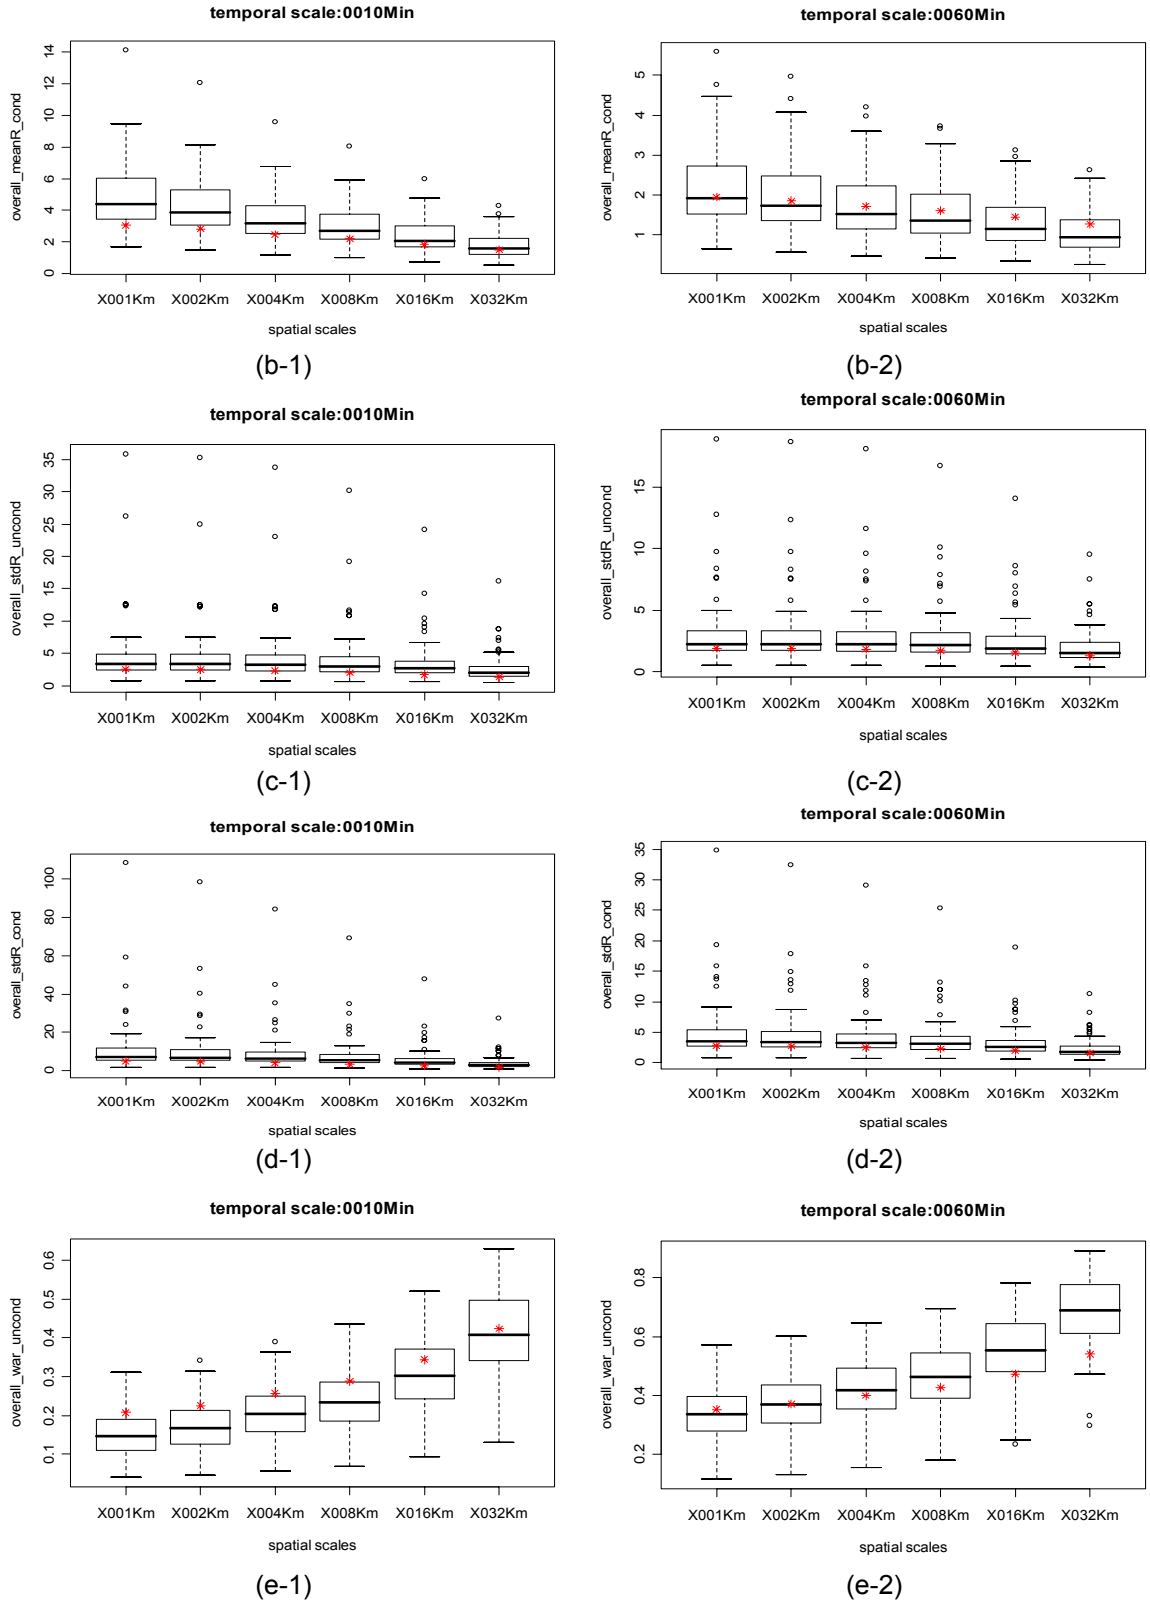

Fig. 3 Continued from previous page. Boxplots of important statistics over various space-time scales for 100 replicates of storm 37 generated by using level-two AR(1) models to generate the eight Level 1 model parameters that control the evolution of rain images. Stars denote the observed values.

SHORT-TERM FORECASTING

Feasibility of short-term forecasting

The conditional simulation capability of the two-level rainfall model makes it potentially useful for short-term forecasting. It was shown that the transformed rainfall field $\boldsymbol{y}_{t+1}^{simu}$ can be generated conditioned on the latest simulated image $\boldsymbol{y}_t^{simu}$ according to equation (5). Generating conditional normal variables is extremely efficient because it involves independently sampling $\boldsymbol{\varepsilon}_t$ and $\boldsymbol{\varepsilon}_{t+1}$ from white noise and then twice applying discrete 2-D fast Fourier transforms to obtain $\boldsymbol{F}_{t+1}$ and $\sqrt{\boldsymbol{\Delta}_{t+1}}$. In forecasting applications, the only difference is that the conditional simulation is based on the latest observed image $\boldsymbol{y}_t^{obs}$ (note that $\boldsymbol{y}_t$ cannot be observed directly, here *obs* denotes $\boldsymbol{y}_t$ derived from an observed image). Once $\boldsymbol{y}_t^{obs}$ is obtained, the corresponding uncorrelated variables $\boldsymbol{\varepsilon}_t$ for the observed image ($\boldsymbol{\varepsilon}_t^{obs}$) can be readily obtained by equation (8), which is the inverse of equation (7) used for marginal simulations. Then the forecast of $\boldsymbol{y}_{t+1}$ conditioned on the latest observed image can be generated by conditional simulation using equations (9) and (10).

marginal simulation: $$\boldsymbol{y}_t = \boldsymbol{F}_t \boldsymbol{\Delta}_t^{1/2} \boldsymbol{\varepsilon}_t \tag{7}$$

$$\Rightarrow \quad \boldsymbol{\varepsilon}_t^{obs} = \boldsymbol{\Delta}_t^{-1/2} \left(\boldsymbol{F}_t^{-1} \boldsymbol{y}_t^{obs} \right) \tag{8}$$

$$\Rightarrow \quad \boldsymbol{\varepsilon}_{t+1|t}^{pred} = \rho_t \boldsymbol{\varepsilon}_t^{obs} + \sqrt{1-\rho_t^2}\, \boldsymbol{\varepsilon}_{t+1} \tag{9}$$

$$\Rightarrow \quad \boldsymbol{y}_{t+1|t}^{pred} = \boldsymbol{F}_{t+1} \sqrt{\boldsymbol{\Delta}_{t+1}}\, \boldsymbol{\varepsilon}_{t+1|t}^{pred} \tag{10}$$

The principle challenge for rainfall forecasting using equations (8) to (10) is estimation of $\boldsymbol{y}_t$ from the observed radar image r_t. According to the power-law transformation given by equation (1), it is straightforward to estimate $\boldsymbol{y}_t$ for wet pixels where $r > 0$. However, $\boldsymbol{y}_t$ is unknown for dry pixels. So the challenge is how to estimate $\boldsymbol{y}_t$ at dry pixels given known $\boldsymbol{y}_t$ at wet pixels. This question can be solved by applying the theory of Gaussian Markov Random Field (GMRF). According to GMRF theory, the density for one variable conditioned on the rest does not depend on all the other variables but only those in the Markov neighbourhood. This is often referred to the Markov property of a GMRF. Based on its small neighbourhood, $\boldsymbol{y}_t$ at each dry pixel can be estimated using single-site Gibbs sampling iteratively until the field converges.

GMRF theory and single-pixel Gibbs sampling method are employed to estimate $\boldsymbol{y}_t$ at each dry pixel in an observed rain image. Basically the procedure involves two steps: Firstly, fitting a GMRF to the transformed Gaussian rain field in terms of preserving the historic spatial correlations over a long range of scales. Secondly, based on the local dependence of the GMRF, applying the single-site Gibbs sampling approach to estimate $\boldsymbol{y}_t$ for each dry pixel conditioned on its small neighbourhood.

Validation of single-pixel Gibbs sampling y_t for dry pixels

In this section, a small example is given to verify the methodology for estimating $\boldsymbol{y}_t$ at dry pixels, which is the only challenge involved in short-term forecasting. The validation involves the following procedure: (1) generating a synthetic rain field (4 × 4) as true values of $\boldsymbol{y}_t$; (2) using the single-pixel Gibbs sampling approach to estimate $\boldsymbol{y}_t$ for dry pixels; (3) comparing the simulated $\boldsymbol{y}_t$ with those true values of $\boldsymbol{y}_t$.

As shown in Fig. 4, a 4 × 4 synthetic rain image is generated at time t by setting $\mu_y = 1.5$, $\sigma_y = 1.0$, spatial correlation parameters $\alpha = 1.0$ and $\beta = 1.0$. Figure 5 presents histogram plots of the estimate $\boldsymbol{y}_t$ compared with the true values. It can be seen for each dry pixel that the histogram of 100 simulations of $\boldsymbol{y}_t$ shows a smooth truncated bell shape. Importantly, all the true values are consistent with the histograms for the dry pixels.

0.65	*0.10*	*0.47*	*1.11*
-1.20	-0.25	-0.02	-0.47
-1.63	*0.80*	-0.00	-1.19
-0.14	*0.22*	-0.15	-2.06

Fig. 4 A 4 × 4 synthetic rain field with true values of $\boldsymbol{y}_t$**:** wet pixels are italic, while dry pixels are black.

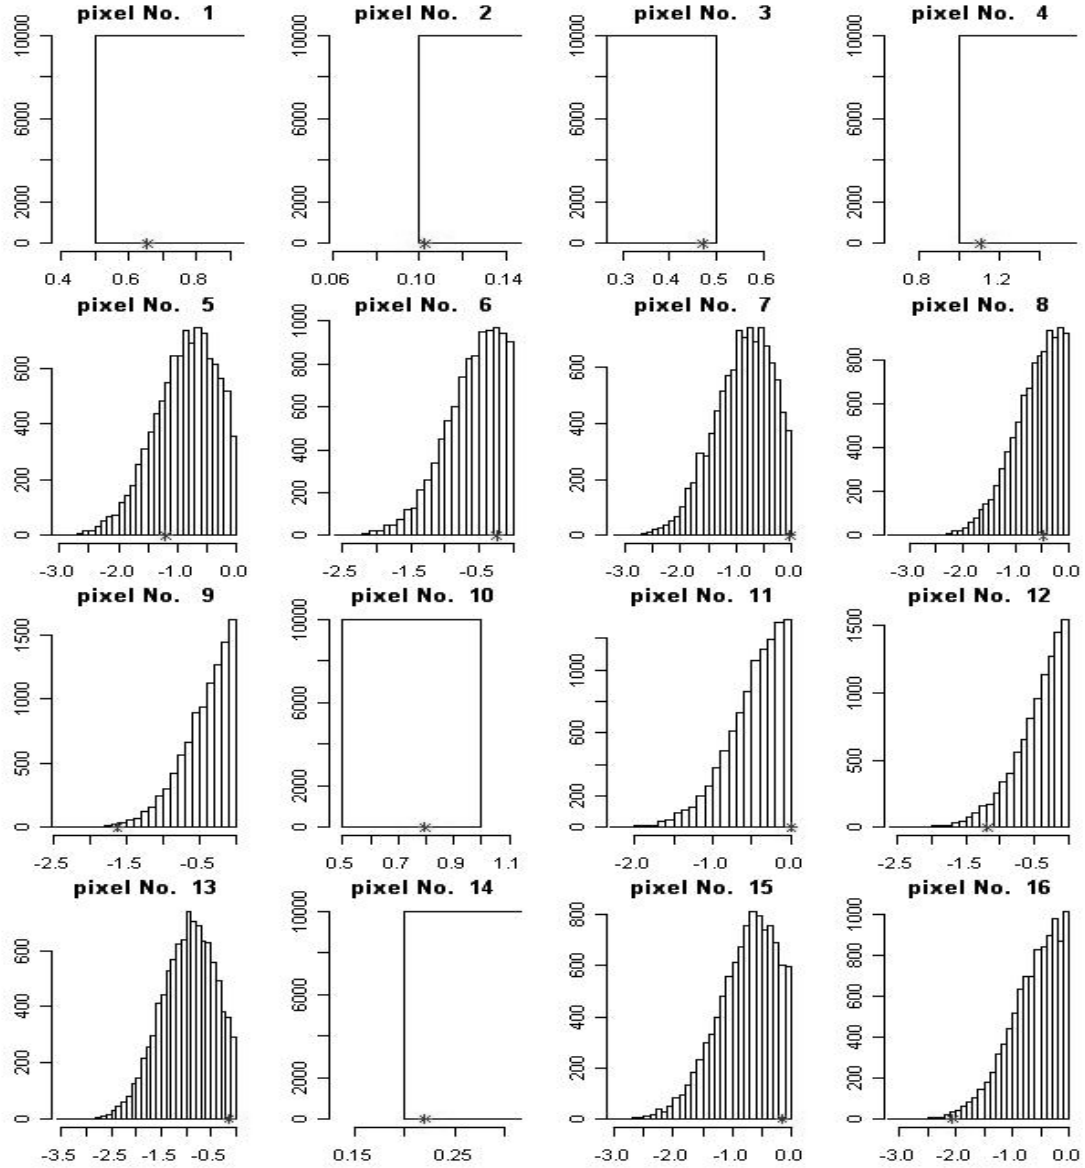

Fig. 5 Histogram plots of estimated y_t compared with the true values (denoted by stars).

An example of short-term forecasting

Once $\boldsymbol{y}_t$ has been estimated, it is straightforward to make short-term forecasts using the proposed conditional rainfall model. In this section an example of short-term forecasting for a 16 × 16 lattice is presented to demonstrate how uncertainties increase with lead-time using the stochastic model. Firstly, a 16 × 16 Gaussian rain field is generated as the initial image y_0. Given lead-time t, $y_1 \mid y_0$, $y_2 \mid y_0, \ldots, y_t \mid y_{t-1}$ are in turn simulated as a "true" observed sequence of rain fields. Secondly,

$\hat{y}_0$ for dry pixels of image y_0 is estimated by iterative single-site Gibbs sampling. Ten chains are used and after every 100 iterations along each chain (defined as one "loop", i.e. one loop = 100 iterations × 10 chains) convergence is checked. The last loop with 1000 realizations can be viewed as samples from the distribution of $\hat{y}_0$, labelled as $\hat{y}_0(i)$, $i = 1,2,\ldots,1000$. Thirdly, having obtained $\hat{y}_0$ for dry pixels in the initial image y_0, conditional simulation can then be used for short-term forecasting. Given lead-time t, 1000 ensembles of forecasts are generated as follows: $\hat{y}_1(i) \mid \hat{y}_0(i)$, $\hat{y}_2(i) \mid \hat{y}_1(i), \ldots, \hat{y}_t(i) \mid \hat{y}_{t-1}(i)$, $i = 1,2,\ldots,1000$.

The 90% forecasting uncertainty interval is defined as the range over which 90% of the simulated probability mass is found. The average of the 90% uncertainty interval over all pixels in an image is referred to as "average 90% uncertainty interval". Figure 6 depicts the increase of the average 90% uncertainty with lead time increasing. As can be seen, for a 10-minute lead-time, the uncertainty is relatively small, but increases as lead-time increases. The curve levels off after approximately 100 minutes, at which point the stochastic model is reporting the marginal variance independent of the observed image y_0 (in other words, the stochastic model no longer has any memory of the observed image). However, for lead-times of less than 1 hour, particularly 30–40 minutes, the stochastic model has useful skill. This is important for flash flood forecasting in urban catchments. For lead times beyond 1 hour, reliance would need to be placed on numerical weather prediction models which incorporate physical processes.

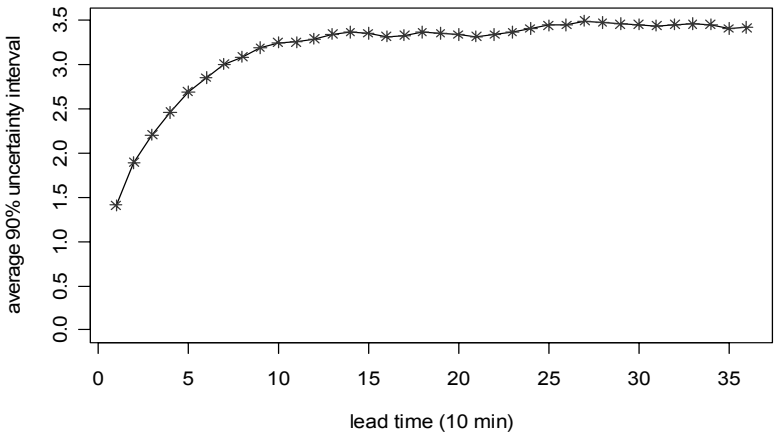

Fig. 6 Plot of forecast uncertainty against lead-time.

CONCLUSION

A new high-resolution space–time model for simulation of rainfall fields consistent with 10-minute 1-km^2 pixel radar images has been developed based on generating latent Gaussian random fields conditioned upon the latest image. A power transformation is applied to the rainfall field to ensure the transformed field is Gaussian. This model adopts a separable space–time correlation function so that the spatial correlation, modelled by a powered exponential covariance function, has the block Toeplitz structure. This allows the circulant embedding technique to be used for fast and accurate simulation. The simulation process is non-stationary over time, allowing for the parameters, which control the stochastic evolution of rainfields, to change dynamically from image to image within a storm. The simulated field is advected by velocity parameters that evolve through the storm's lifetime. Because this is a phenomenological model without explicit incorporation of storm physical structures, it is capable of dealing with different storm types.

The performance of the first two levels of the model was evaluated by comparing the characteristics of interest between a typical observed storm event and synthetic storms. Visual comparison suggested that the two-level rainfall model is capable of producing realistic sequences

of rain images that capture three major features of rainfield, i.e. the physical hierarchical structure of clusters, patchiness of rain fields, and the persistence exhibited during storm development. A number of important statistics, namely the unconditional overall mean, unconditional overall standard deviation, overall wet area ratio, conditional overall mean, and conditional overall standard deviation, were adequately reproduced at both the 10-minute and the 1-hour time scales over space scales ranging from 1 km up to 32 km. Many of these statistics were not subjected to calibration. This is a significant achievement and an advance in capability over current models.

The conditional structure of the rainfall model makes it feasible to perform short-term forecasting at high space-time resolutions. The main challenge is to estimate $\boldsymbol{y}_t$ for dry pixels in an observed image. GMRF theory in conjunction with the iterative single-site Gibbs sampling technique was used to estimate $\boldsymbol{y}_t$ at dry pixels and validated for a 4×4 rain field. A short-term forecasting case study was then undertaken on a 16×16 rain field. Not surprisingly, the uncertainties of forecasts increase with lead time. For lead-times of less than 1 hour, the stochastic model shows fair forecasting skill. That may be important for flash flooding forecasting in urban catchments.

To date the first two levels of the hierarchical model have been developed and tested with encouraging results. There is scope for improvement of the level-two model and extension of the short-term forecasting to large rainfields. However, the main challenge lies with the level-three component of the model.

Acknowledgements This work was funded by the Australia Research Council (DP0452180). The assistance of Alan Seed is gratefully acknowledged.

REFERENCES

Bell, T. L. (1987) A space–time stochastic model of rainfall for satellite remote sensing studies. *J. Geophys. Res*. **92**(D8), 963.

Cowpertwait, P. S. P., Kilsby, C. G. & O'Connell, P. E. (2002) A space–time Neyman-Scott model of rainfall: Empirical analysis of extremes. *Water Resour. Res*. **38**(8), 1131.

Dietrich, C. R. & Newsam, G. N. (1997) Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix. *SIAM J. Sci. Comput*. **18**(4), 1088–1107.

Krajewski, W. F. (1996) Effects of the radar observation process on inferred rainfall statistics. *J. Geophys. Res. – Atmospheres* **101**(D21), 26493–26502.

Mellor, D., Sheffield, J., O'Connell, P. E. & Metcalfe, A. V. (2000) A stochastic space–time rainfall forecasting system for real time flow forecasting, I. *Hydrol. Earth System Sci*. **4**(4), 603–615.

Over, T. M. & Gupta, V. K. (1996) A space–time theory of mesoscale rainfall using random cascades. *J. Geophys. Res. – Atmospheres* **101**(D21), 26319–26331.

Pegram, G. G. S. & Clothier, A. N. (2001) High resolution space–time modelling of rainfall: the "String of Beads" model. *J. Hydrol.* **241**, 26–41.

Qin, J., Leonard, M., Kuczera, G., Thyer, M., Metcalfe, A. & Lambert, M. (2006) Statistical characteristics of rainstorms derived from weather radar images. In: Proc. 30th Hydrology and Water Resources Symposium (4–7 December, 2006, Launceston, Australia).

Qin, J., Leonard, M., Kuczera, G., Thyer, M., Metcalfe, A. & Lambert, M. (2008) A High resolution spatio-temporal model for single storm events based on radar images. In: Proc. 31st Hydrology and Water Resources Symposium (15–18 April, 2008, Adelaide, Australia).

Seed, A. W., Srikanthan, R. & Menabde, M. (1999) A space and time model for design storm rainfall. *J. Geophys. Res.* **104**(D24), 31623–31630.

Waymire, E., Gupta, V. & Rodrigues-Iturbe, I. (1984) A spectral theory of rainfall intensity at the meso-beta-scale, *Water Resour. Res.* **20**, 1453–1465.

Wheater, H. S., Isham, V. S., Onof, C., Chandler, R. E., Northrop, P. J., Guiblin, P., Bate, S. M., Cox, D. R. & Koutsoyiannis, D. (2000) Generation of spatially consistent rainfall data. Report to the Ministry of Agriculture, Fisheries and Food (2 vols), also available as *Research Report no. 204*. Department of Statistical Science, University College London, Gower Street, London WC1E 6BT, UK, (http://www.ucl.ac.uk/Stats/abstracts.html).