A Stochastic Space-Time Rainfall Model for Engineering Risk Assessment

by Michael Leonard

Submitted in fulfilment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

April 14, 2010

FACULTY OF ENGINEERING, COMPUTER AND MATHEMATICAL SCIENCES
School of Civil Environmental and Mining Engineering

THE UNIVERSITY OF ADELAIDE
AUSTRALIA

Chapter 1
Introduction

Humans have a long history dating back to ancient times attempting to understand rainfall and runoff processes. Typically, this has been in order to harvest water for economic purposes and mitigate the impacts of extreme events such as floods and droughts. At the turn of the century, the annual average direct cost of flooding damage in Australia was estimated at over $\$300$ million. It accounts for the highest proportion (29%) of costs attributable to natural disasters trumping storm-damage, cyclones, earthquakes and bushfires [*BTE*, 2001]. It is therefore of significant national interest that the variability of rainfall patterns in both time and space can be modelled reliably.

It is assumed here that the main impetus for considering rainfall patterns is for flood design estimation, where a level of risk is assigned to some statistic associated with the abundance of water. This topic is broad and covers a wide range of applications, with methods highly linked to the spatial and time scales of interest. Consider for example,

1. an estuarine environment subject to joint flooding from rainfall on a large catchment and elevated ocean levels;

2. a steep escarpment in an urbanised area having quick response times;

3. a small catchment highly sensitive to the occurrence and advection of frontal storms.

Furthermore, the statistic of interest will vary, ranging from the overall volume of water, to the peak flow rate at an outlet or the water level at any point along a channel. Depending on the scenario, the magnitude of the statistic will vary from frequent smaller events to infrequent larger events and in each of these cases spatial and temporal scales are critical. Various terms are used interchangeably to describe the measure of risk, for example, the annual exceedence probability (AEP), non-exceedence probability, frequency or average recurrence interval (ARI).

This chapter provides an introduction to the history of flood design methodology and the Australian context of climate, rainfall measurement and case studies used in this thesis.

1.1 THE RATIONAL METHOD

Streams are often the focal point for a design, but it is uncommon that a reliable streamflow gauge is available. This is the crux of the problem. Since rainfall is a more abundant source of data, it is convenient to base flood estimation techniques on transforming rainfall inputs to runoff. This is perhaps an obvious statement, but it is not an obvious practice.

The rational method of *Mulvaney* [1850] is the backdrop for many of the concepts still used in hydrologic design today. This method derives the flow as the product of a specially selected intensity, the area of a catchment, and an assumed coefficient,

$$Q_{AEP} = 0.278 C_{AEP} A I_{tc,AEP} \tag{1.1}$$

where the subscript AEP corresponds to the level of risk, Q_{AEP} is the peak flow-rate at the outlet of the catchment ($m^3 s^{-1}$), 0.278 is a unit conversion factor, C_{AEP} is a dimensionless runoff coefficient, A is the catchment area (km^2) and $I_{tc,AEP}$ is the rainfall intensity (mm/hr) at the critical duration, t_c. This equation embodies the concept of AEP neutrality that a rainfall event of a certain AEP can be converted to a runoff event of the same AEP. It also embodies the concept that only the most intense part of a storm, the peak rainfall burst, is needed to determine the peak flow. Also fundamental to this equation is the notion that one duration, t_c, will cause a catchment to generate the maximum flow. At this duration, the simple product of area with intensity bears the assumption that rainfall is uniformly distributed over the catchment. The requisite that the AEP subscript on the LHS of Eq. 1.1 is the same as the AEP subscript on the RHS is achieved by the supposed constant C_{AEP} which must act as a surrogate for multiple unparameterised factors (intercept losses, infiltration losses, areal reduction, inhomogeneous rainfall, temporal variability, storm advection, catchment geometry, depression storage, etc.).

1.2 DESIGN EVENTS

It is possible to view the design event approach as an extension of the rational method, where the coefficient C_{AEP} is now replaced by temporal patterns, areal reduction factors, a conceptual runoff model and a channel routing model. The design event method determines an entire hydrograph at all river sections, which is in contrast to having just the peak flow at the catchment outlet. This method is more versatile than the rational method

because it allows for the influence of storage characteristics within a catchment, such as detention basins. Typically, several design storm durations are considered at the AEP of interest (such as 1 hour, 12 hour, 24 hour duration rainfall bursts at a 1% AEP) and the intensities for each scenario are determined from intensity-frequency-duration (IFD) rainfall maps. The rainfall for the various storm durations is distributed according to a standardised temporal pattern for that duration and according to the zoning of temporal patterns across Australia.

Because rainfall is an intermittent process, a storm burst will not simultaneously cover the entire catchment. For this reason, an areal reduction factor corresponding to the catchment area, storm duration and ARI is applied to the design rainfall hyetograph to produce a realistic rainfall total that lands on the catchment. A loss model is required to mimic the antecedent soil moisture conditions of the catchment and produce a rainfall excess hyetograph from the incident hyetograph. The rainfall excess hyetograph is converted into a runoff hydrograph by employing a parameterised runoff-routing model that represents the reaches of the catchment as lumped systems. A baseflow is then added to produce the design hydrographs, from which the burst duration yielding the largest peak flow is selected as the critical design duration. While the structure of the design event approach is better for appreciating model errors, the method has many limitations. Three of these are briefly covered here, the idea of AEP neutrality, the use of rainfall bursts rather than entire storms and the use of a single temporal pattern.

1.2.1 AEP Neutrality

Let R denote a set of rainfall events, let θ represent all other variables (for simplicity consider them fixed values such as parameters, initial conditions, ground water, evaporation, etc.), let $g()$ be some nonlinear transformation (flow, water level, volume) and $h_{1\%}[\,]$ some function that ranks values to obtain the magnitude having a 1% chance of exceedence. The concept of AEP neutrality is then summarised as follows

$$h_{1\%}[g(R,\theta)] = g(h_{1\%}[R],\theta) \tag{1.2}$$

for some suitably chosen values of θ . That is, the transform of a 1% AEP rainfall event (RHS) can be made equivalent to the 1% AEP of a set of transformed events. This concept is akin to arguing $1 = 2 * \theta$ for a suitably chosen value of θ. This statement is true for a unique value of θ, is approximately true for a small range of θ values and for the remaining majority of values it is untrue. In the instance of conceptual runoff models it is usually the role of loss coefficients to be manipulated to achieve the stated aim of AEP neutrality. Since catchment models contain many nonlinear components and initial conditions, it is

difficult to estimate representative parameters that satisfy AEP neutrality [*Kuczera et al.*, 2006].

Although the parameters of Eq. 1.2 often have a physical interpretation, AEP neutrality is a mathematical convenience and it is not strictly necessary that θ values have a realistic interpretation. This observation is acknowledged in the design guidelines Australian Rainfall and Runoff (ARR) [*Pilgrim*, 1987, Book II, §3], *Initial loss values obtained from fitting actual storms will be too high, compared with the values that should be used in assessing the median value for use in design*. Obtaining suitable values is essentially a calibration problem. It is not a given that traditional calibration is any more reliable than experienced judgement as several aspects can obscure the calibration process (i) significant spatial and temporal averaging, (ii) measurement error of inputs and outputs and (iii) structural model error [*Thyer et al.*, 2009]. The emphasis here is on the *concept* of conceptual runoff models and is best testified to by the many unsuccessful attempts to regionalise model parameters for application to ungauged catchments.

1.2.2 Intensity Frequency Duration curves

The most fundamental aspect of IFD methodology is that it refers only to a burst of rainfall for some fixed time period. Thus the design event is not an entire storm but merely a design burst. The strongest limitation of using peak rainfall bursts for design events is that they ignore preceding rainfall which is necessary for estimating initial conditions for volume based designs and for estimating catchment losses when converting rainfall to runoff. For example, 2-D hydraulic models are being increasingly used and have accurate ability to model catchment storages, but they are biased toward low volumes when used with design bursts due to the absence of pre-burst rainfall. *Blaikie and Ball* [2005] conducted a study using 14 pluviometers from the greater Sydney region which were analysed for antecedent rainfall prior to a 24 hour period. This study reached several conclusions: that a significant portion of 5 day antecedent rainfall occurred on the day prior to the rainfall burst, that the antecedent rainfall depth increased with the storm burst rainfall depth and that the antecedent rainfall depth was not constant over the Sydney region.

1.2.3 Temporal patterns

There are generally no typical patterns [*Pilgrim*, 1987, Book II, §2.2b]. The average variability method (AVM) is a method to derive typical patterns. Figure 1.1 shows the AVM temporal patterns of the 24 hour duration storm burst for the eight different classification zones about Australia. It can be seen that there is significant variability in the temporal

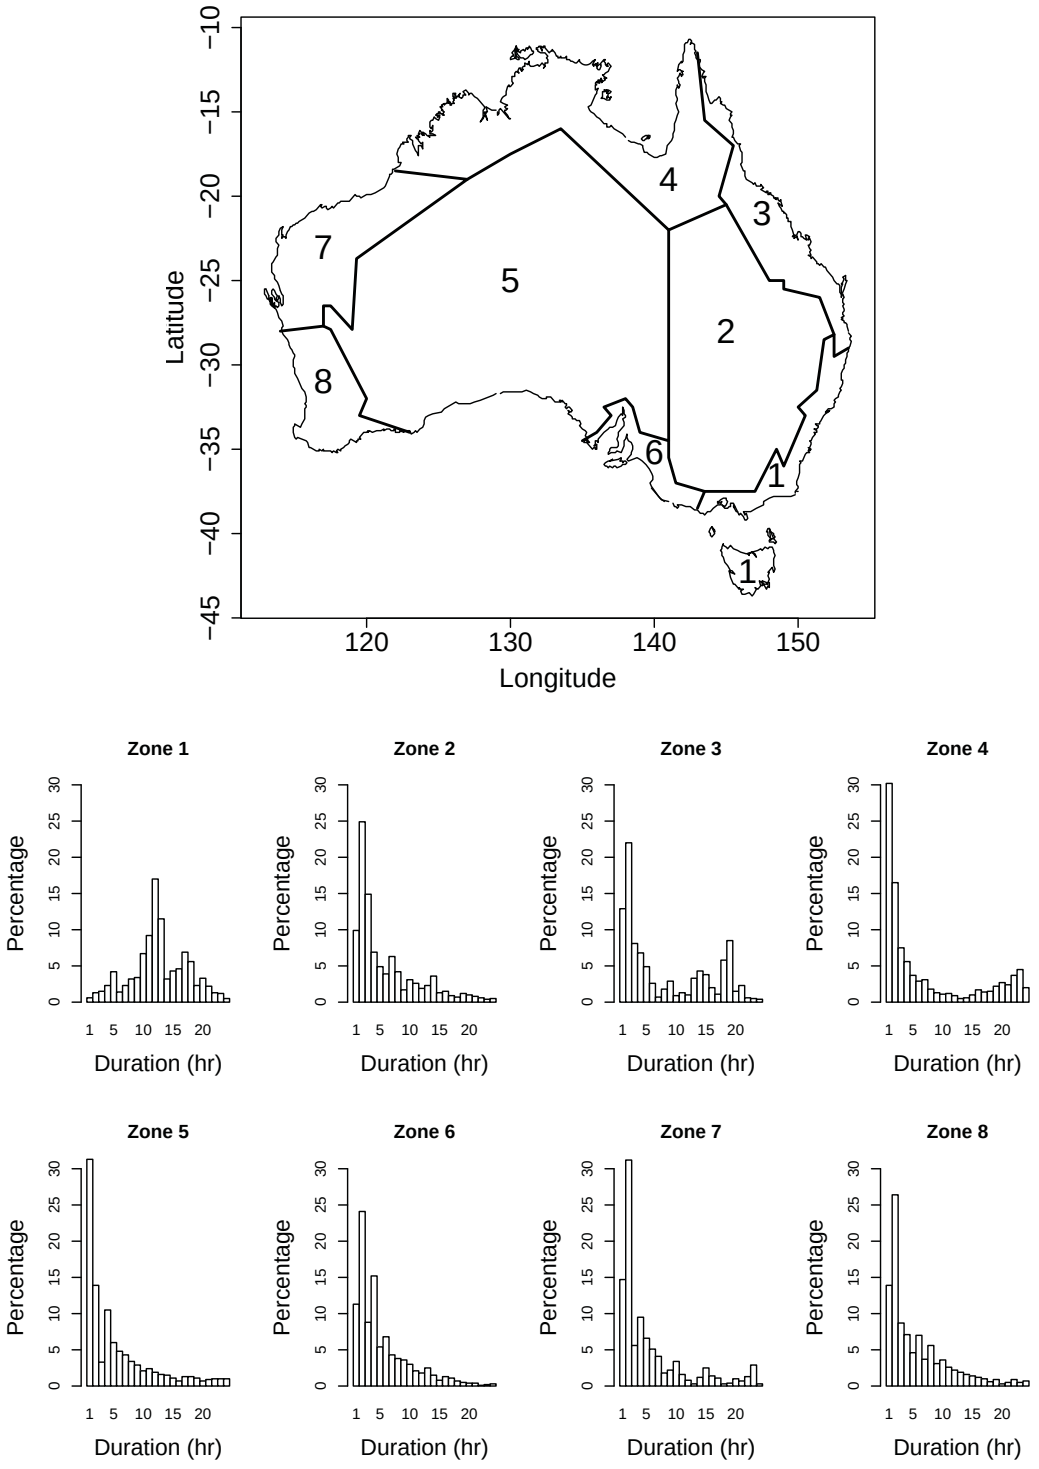

Figure 1.1 24 hour design storm temporal patterns for 8 different zones of Australia.

patterns between the zones, though the pattern is assumed to be the same at all points within the zone. The sole purpose of the AVM is to transform a peak rainfall into a peak flow of the same AEP. Any summary statistic such as an average or median is an abstract idea that simultaneously represents everything and nothing, so it is of no surprise then that the AVM yields atypical patterns that are unlike real rainfall patterns. This is further emphasized by the fact that the AVM requires corrections such as smoothing and filtering to provide consistency across durations and AEPs.

Since the 1987 version of ARR, numerous authors have compared the AVM temporal patterns to other rainfall patterns. *Ball* [1994] highlighted that different temporal patterns lead to different catchment responses and that the variability of different temporal patterns is important. *Green et al.* [2003] presented temporal patterns for use with probable maximum precipitation (PMP) estimates within the tropical region. They surveyed a wide range of techniques that were used to produce temporal patterns and acknowledged that there is *no definitive answer regarding the most appropriate temporal pattern to be adopted*. For the PMP design flood, they recommended the use of an unsmoothed AVM temporal pattern based on the top ten events of each duration. They also discussed options where, with the use of meta-data, the hydrologist may wish to choose just one of the top ten patterns or an alternative historical extreme event in the near vicinity of the catchment. *Sih et al.* [2008] compared AVM temporal patterns with flow estimates obtained using a collection of temporal patterns across regions of Australia. They noted biases in the flow estimates produced by AVM patterns when compared to replicates of temporal patterns and that the mean flow estimate from a set of temporal patterns produced less bias than the AVM pattern.

1.3 MONTE CARLO SIMULATION

An alternative to the design event approach is referred to as Monte Carlo Simulation (MCS) and it is capable of more reliable AEP estimates [*Rahman et al.*, 2002]. The idea of MCS is not to obtain a set of deterministic θ values for catchment processes that happen to reproduce a 1% AEP peak runoff for a 1% AEP peak rainfall burst (RHS of Eq. 1.2), but to allow θ to be random, covering all joint possibilities of significant catchment processes (using LHS of Eq. 1.2). Many replicates of parameter values θ are sampled from representative distributions and are combined with the rainfall to determine the AEP of the transformed statistic. The rainfall input should also be randomly varied along with the other inputs as it is likely that rainfall patterns from a range of AEPs will generate transformed values having a 1% AEP (e.g. a 0.5% AEP storm on a drier

catchment or a 2% AEP storm on a wetter catchment). The advantage of MCS is that it allows for variability of the inputs and therefore avoids bias in the output. Since the 1987 ARR version, there has been considerable development of the MCS approach. For example *Rahman et al.* [2002] and *Weinmann et al.* [2002] highlight that loss coefficients, temporal patterns, pre-burst rainfall and seasonality are important variables to randomise. *Kuczera et al.* [2006] highlight the initial level of storage basins is an important variable to consider when estimating flows. *Need et al.* [2008] show that for estuarine environments the ocean depth is an important variable to randomly vary. *Tan et al.* [2008] highlighted the importance of correlation when large systems are being considered. Their specific study used mean sea level pressure to associate water levels with rainfall amounts in the Gippsland region.

One of the attractions of MCS is that the concepts and processes involved are similar to those used in the design event approach. MCS is however more computationally demanding than the design event approach, where the computational burden depends on the complexity of the models being used. More complex models, larger sample sizes or multiple replicates for uncertainty analysis will increase the computational demand. This is especially the case when spatially gridded rainfall, runoff and hydraulic models are considered.

Another technique that is similar to MCS is referred to as continuous simulation. Like MCS, continuous simulation uses many events that are sampled randomly. The difference is that the events are not standalone but are embedded into an event sequence. The chief benefit over MCS is that other random variables within the system can be updated and do not need the user to specify a sampling distribution (e.g. soil moisture, initial water level). Unlike MCS, continuous simulation cannot be stratified to focus only those events which result in an extreme outcome. In other words, continuous simulation aims to reproduce the full distribution of rainfall and runoff events which can make it more computationally demanding than MCS.

1.4 CONTINUOUS SIMULATION

Continuous simulation refers to a process of converting catchment rainfall inputs into flood outputs by successively simulating sequences of storms in conjunction with a water budget model to continuously update soil moisture conditions. The rainfall inputs are generated stochastically, using a conceptual framework to jointly specify the incidence of storms and the properties of each storm. Differences between models arise from the way in which random variables are assigned to the conceptual storm properties such as

duration, intensity, temporal pattern and storm occurrence. Storm sequences and corresponding runoff are simulated for long time periods so that a complete distribution of outflow hydrographs is generated and used to evaluate levels of risk. *Heneker* [2002] provides a thorough evaluation and discussion of this approach.

The spatio-temporal variability of rainfall is widely cited as one of the main sources of uncertainty in understanding and modelling the hydrological dynamics of a catchment [*Chaubey et al.*, 1999; *Faures et al.*, 1995; *Wood et al.*, 2000]. To date, however, continuous simulation models have primarily focused on reproducing temporal statistics at a point (rain-gauge) or set of points that are assumed to sufficiently represent the orographic effects of the catchment. To convert point rainfall to a volume over a region an areal reduction factor is required correct for the patchiness of rainfall over that region. The spatial averaging that is implicit in the use of areal reduction factors can often lead to biases.

There is growing demand for the development of continuous simulation space-time rainfall models and the potential application of these models is diverse. Examples include environmental and land use decision making, erosion/sediment modelling, water quantity/quality modelling, spatially distributed groundwater recharge and ecological and riparian modelling. The need for spatially variable rainfall estimates is even greater in urban environments, due to the more rapid response times generated by impervious surfaces. Examples include stormwater and roofwater re-use, sewerage overflows and urban flooding. The development of space-time rainfall models for continuous simulation is the focus of this thesis and the background is discussed more thoroughly in Chapter 2.

1.5 THE AUSTRALIAN CLIMATE

Being a large land mass, surrounded in entirety by the Pacific, Indian and Southern oceans, Australia experiences a broad range of weather patterns. Australia is highly influenced by a number of air-masses originating over these basins [*Sturman and Tapper*, 2006], including:

- The Southern Maritime air mass - cool moist air in the Latitudes $35°$ to $55°$, bringing cloudy drizzle to Southern Australia throughout the year.
- The Tropical Maritime Tasman - warm moist air that originates in the Tasman Sea and bringing warm, cloudy and drizzly weather to the east coast of Australia.
- The Tropical Maritime Pacific - warmer than the Tasman maritime air mass and

originating further north in the western Pacific. It affects North Queensland most of the year, bringing heavy rainfall when associated with tropical cyclones.

- The Equatorial Maritime - very warm moist and unstable air affecting the north of Australia (as far South as 30° Latitude) with heavy rainfall and humidity during the summer monsoon season.
- The Tropical Maritime Indian - sourced in the Indian ocean, affecting Western Australia and similar to the Tropical maritime Pacific.

As a consequence of these processes, Australia experiences moderate rainfalls around the majority of the western, southern and eastern coastlines, high rainfall totals along the northern coast due to the summer monsoon and a very dry interior due to the distance from the coastline and the the lack of relief in this region. This is shown in Figure 1.2 (a) which depicts the annual average rainfall over the Australian continent with values ranging from below $200\ mm$ in the centre of Australia to over $2000\ mm$ in northern Australia. The eastern seaboard yields higher rainfall totals than the western seaboard due to the effect of orographic lifting over The Great Dividing Range mountains. Figure 1.2 (b) and Figure 1.2 (c) show the monthly totals for two months, January and July, representing typical seasonal patterns that occur in summer and winter respectively. Figure 1.2 (b) shows that, during summer, the northern regions are inundated due to the summer monsoon, while the southern and south-western regions are very dry. By way of contrast, Figure 1.2 (c) shows that the southern and south-western regions are dominated by winter-rainfall while the northern regions are very dry. The eastern seaboard shows less variation between the seasons than the other parts of Australia.

The climate of Australia is summarised using a classification of climatic zones based on rainfall totals, seasonal incidence and temperatures. Figure 1.3 shows the Bureau of Meteorology classification including the locations of major urban centres [*BOM*, 1983]. With respect to the capital cities, Perth, Adelaide and Melbourne have winter-dominated rainfall, Sydney and Hobart have uniform rainfall across the year, Brisbane experiences sub-tropical rainfall and Darwin is tropical.

It is evident that rainfall is a highly variable process and, attributable to its geography as a large island nation, Australia is subject to a wide variety of climates.

1.6 RAINFALL MEASUREMENT

In addition to appreciating atmospheric processes in the Australian climate, the ability to model rainfall occurrence is dependent on the quantity, quality and nature of rainfall

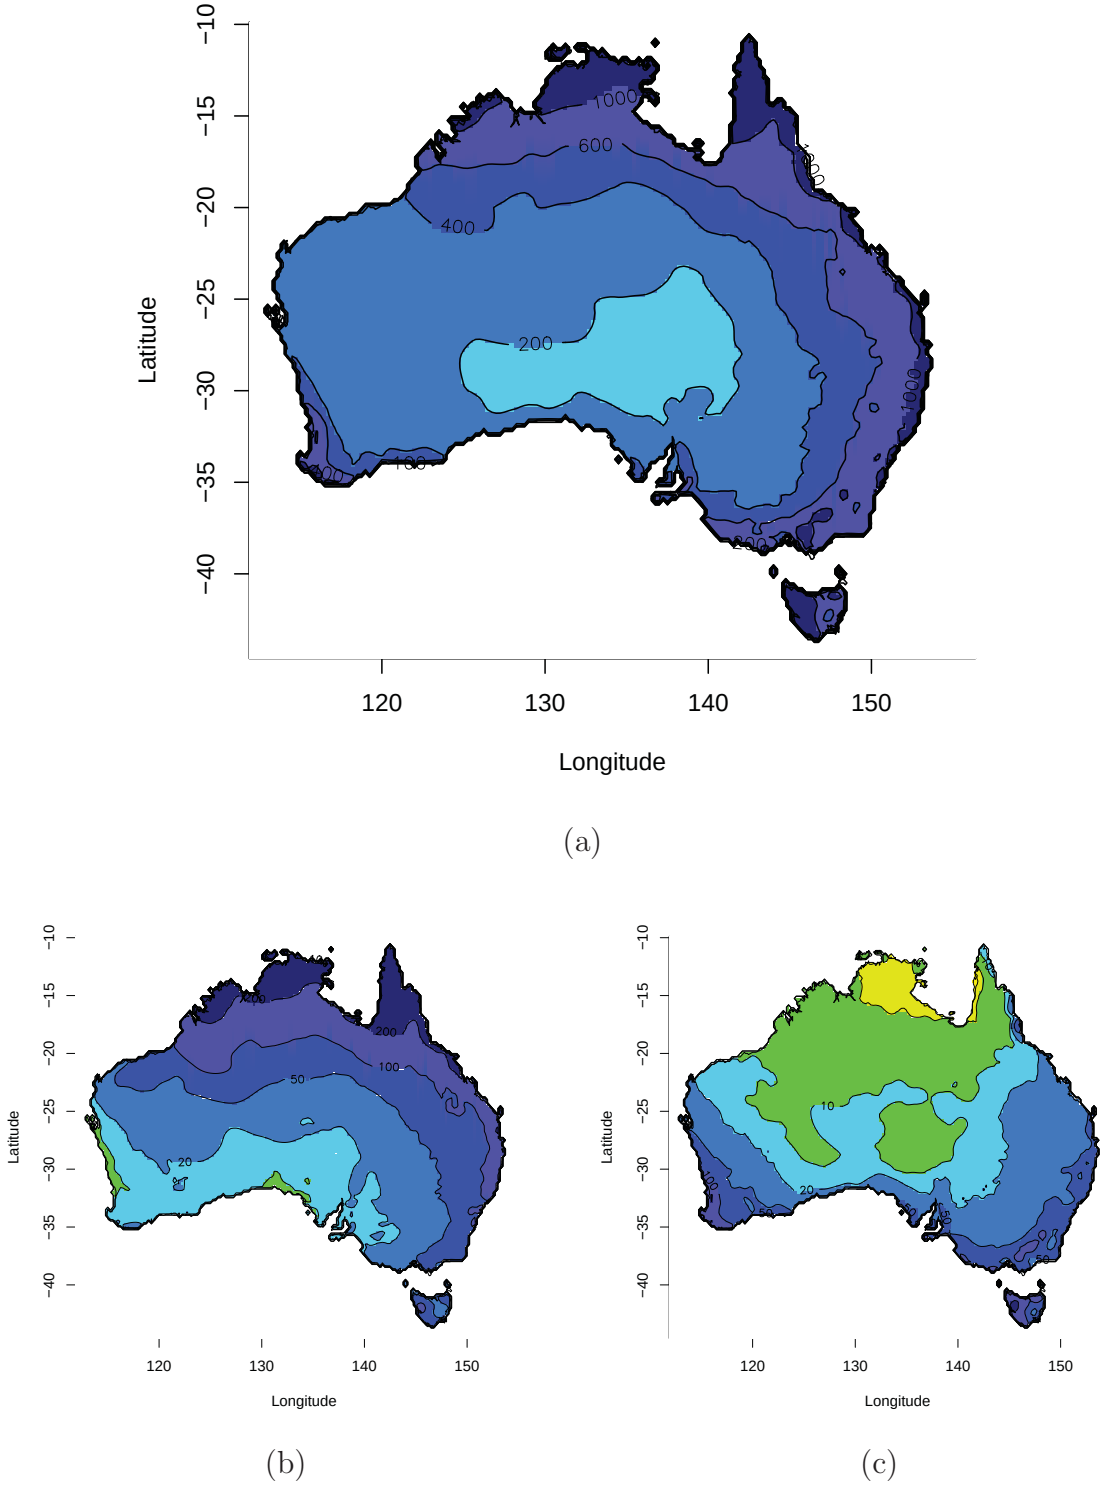

Figure 1.2 Mean rainfall in millimetres depth across Australia, (a) Annual, (b) January (Summer), (c) July (Winter).

Figure 1.3 Climatic Zones of Australia, from *Heneker* [2002], based on *BOM* [1983].

measuring devices. Rainfall, in the form of individual droplets, is fundamentally discrete in space and time, but the process of measurement typically requires averaging in space and time. For this reason, the most common unit of interest is rainfall intensity (mm/hr for point observations and $mm/hr/km^2$ for areal observations), which has a continuous, non-negative scale. Like an understanding of the rainfall process itself, measurements of rainfall require an appreciation of temporal and spatial scales. Traditionally, measurements have been made at a single point in space using rain-gauges at either daily or sub-daily timescales. More recently, modern radar and computing technology, via reflection of electromagnetic waves, has made inferential estimates of rainfall possible over large regions at sub-daily time scales. The temporal and spatial resolution of the radar is governed by the wavelength, examples include X-band, C-band and S-band radars. Figure 1.4 summarises the distribution of these measuring devices throughout the Australian continent.

1.6.1 Daily rainfall gauges

Daily rainfall gauges are the most conceptually simple and inexpensive mechanism for measuring rainfall quantities. A gauge, little more elaborate than a container, is read on a daily basis to give a depth of water over that day. Due to the simplicity of the apparatus,

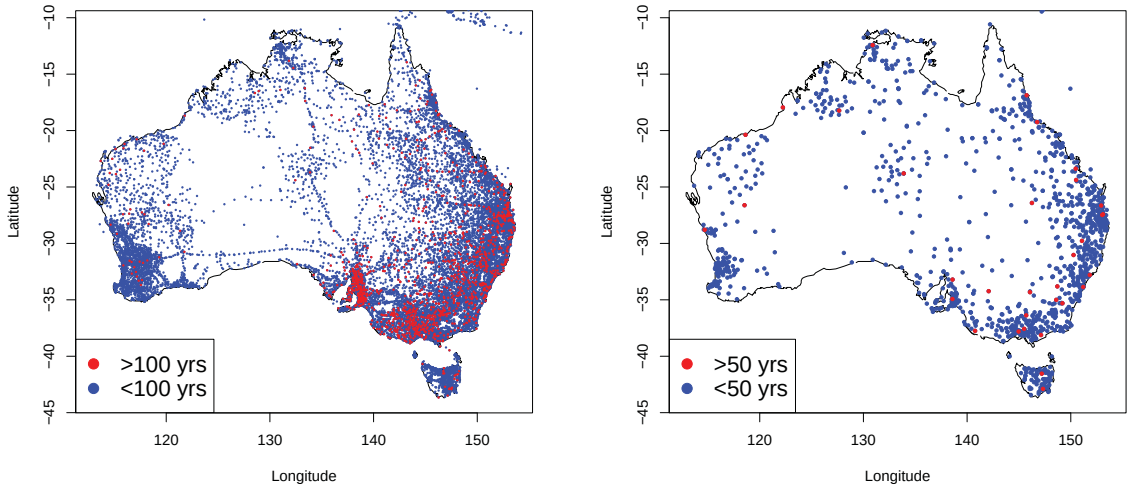

(a) (b)

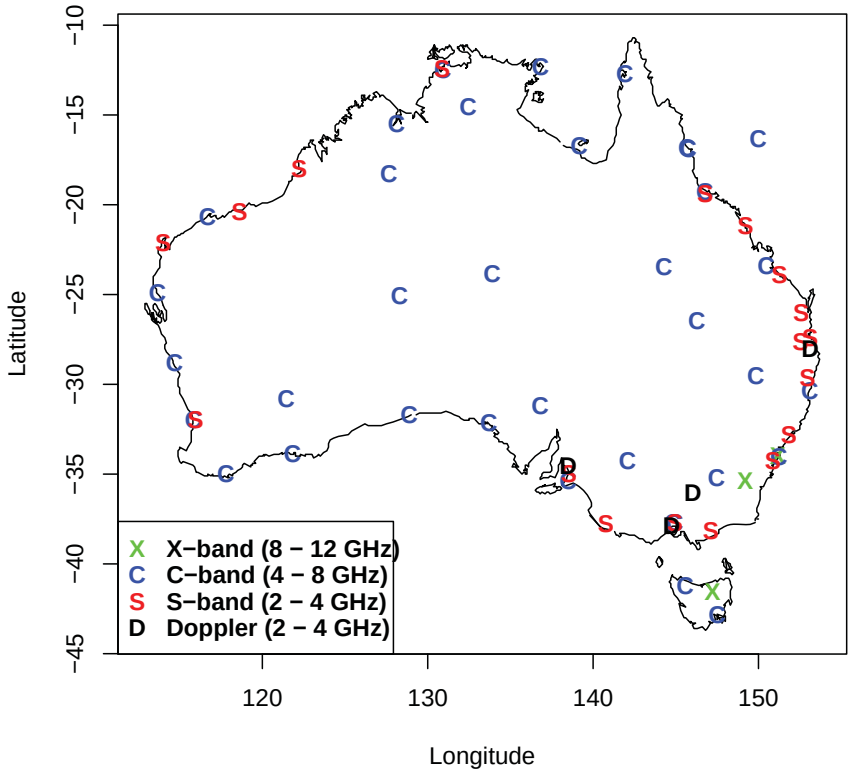

(c)

Figure 1.4 Distribution of rainfall measuring devices about Australia (a) Daily gauges, (b) Pluviograph gauges (c) Radar.

daily rainfall gauges are the most common source of rainfall observations. Over 7000 gauges are currently active across Australia and over $16,000$ in the Bureau of Meteorology archive have been active at one point or another during Australia's history. The oldest gauges have records that extend as far back as the mid 1800s, but it is not uncommon to have substantial periods of missing observations during this era. Nonetheless, there are approximately 1000 records in the Australian continent that have at least 100 years' of observations. Figure 1.4 (a) shows the distribution of daily rainfall gauges across Australia and verifies that there is a good distribution of gauges with over 100 years worth of record. State and private agencies own numerous records in addition to those from the Bureau of Meteorology, though they typically have a shorter period of observation.

The main advantage of daily rainfall records is their coverage across all areas of the Australian continent and their comparatively long periods of observation with respect to other recording devices. A limitation of daily observations is that many hydrologic processes of interest, especially in urban areas, have time-scales less than one day. An example is flash-flooding events associated with individual storms. Another limitation of many daily gauges is the presence of biases, accumulations and outages in the observed record, which are prevalent, despite the simplicity of the measuring system. Examples include (i) when a gauge has moved location, (ii) when a human incorrectly reads or records the total (iii) when the gauge has been overtopped (iv) when a nearby building or tree blocks rainfall and (v) when readings are accumulated over multiple days, typically weekends, where only the accumulated total is known but not which day(s) the rain fell on. While most accumulations are flagged, *Viney and Bates* [2004] highlighted that some are unflagged and therefore not obvious and potentially misleading. This can occur when a recorder mistakenly indicates that a multi-day rainfall accumulation occurred solely on one day and that the other days had zero rain.

1.6.2 Pluviograph rainfall gauges

In Australia, pluviograph rainfall gauges typically provide information on rainfall intensity at six-minute increments. Initially, siphon gauges and continuous recording charts were used to measure intensity, though after a while tipping-bucket gauges and digital storage became more common-place. Tipping-bucket gauges operate by counting the number of tips of two, $0.2\ mm$, buckets mounted on a see-saw below the gauge funnel. The time-series of bucket releases is then converted to intensity over six minute intervals. Currently there are approximately 1300 pluviograph gauges in the Bureau of Meteorology archive, as shown in Figure 1.4 (a). The limited spatial coverage of pluviograph gauges is

a serious impediment to hydrological science in Australia, though sometimes additional gauges owned by state and private agencies can provide a useful supplement. Such additional networks of gauges mostly use continuous updating of measurements via radio telemetry for applications such as flash-flooding detection and real-time radar calibration.

While pluviograph gauges provide invaluable information at sub-daily timescales, their poor coverage in time and across Australia is a significant limitation. Figure 1.4 (b) shows that very few gauges have more than 50 years of observations on record. In addition to this, pluviograph gauges often experience outages or data corruption, for example due to blockages from leaf-litter. In contrast to daily rainfall gauges, the total rainfall depth over the outage period is often not recorded, which implies that aggregate statistics from pluviograph measurements will not conserve the total rain depth and can become biased. With networks that report observations via radio telemetry, accumulations of rainfall can be common due to instances of poor reception during the remote data transmission. Other minor issues associated with pluviograph gauges include incorrect digitisation of older chart records [*Deidda et al.*, 2007] and bias arising from the tipping motion during extreme heavy rain [*Humphrey et al.*, 1997].

1.6.3 Radar rainfall reflectivity

Following the Second World War, radar technology was put to civilian use in remotely measuring weather activity. Radars are mounted on a rotating pedestal set above tree height and scan volumes of the sky. The radar sends out electromagnetic pulses and, based on the timing and strength of the reflected pulse, the location and intensity of rainfall activity can be estimated. *Battan* [1962] is a good introductory text on this topic. The type of radar is identified by the wavelength of the incident pulses. Weather radar is typically one of: S-band, $2-4\ GHz$ frequency ($8-15\ cm$ wavelength), C-band, $4-8\ GHz$ frequency ($4-8\ cm$ wavelength) or X-band, $8-12\ GHz$ frequency ($2.5-4\ cm$ wavelength). Smaller wavelength radars are able to detect more detailed structure in a weather system, but they also have increased attenuation, which means that they make less reliable measurements over increasing distance. For this reason, S-band radars are regarded as best for measuring storm systems over larger areas (e.g., of $100\ km$ radius). Figure 1.4 (c) shows the distribution of operational weather radars about Australia. Note that some are dedicated to inferring wind measurements, some are dedicated to inferring rainfall measurements and some are dual-purpose. This figure shows four radars, recently built, that concurrently estimate wind and rainfall by measuring the Doppler-shift from pairs of reflected waves. These radars were commissioned as part of a recent $\$62$ mil-

lion upgrade, which will produce a total of six new Doppler, S-band radars (Sydney and Tamworth Doppler radars not yet installed).

Radar has significant advantages in its ability to provide information on the detailed spatial structure of storms not otherwise available from gauges. While spatial detail is a strength, radar instrumentation is subject to numerous outages and has not been operational for as long as gauges. Therefore it is not as good for estimating long-term temporal statistics. Importantly, radar is an inferential measuring device, where *reflectivity* is the actual measurement. Numerous forms of error can occur in its conversion to *rainfall intensity*. These errors include (i) systematic errors such as ground clutter, anomalous propagation, bright band contamination and inappropriate reflectivity conversion and (ii) random errors such as sampling errors, quantisation and variability in the reflectivity relationship [*Jordan and Seed*, 2003]. Rain-gauge measurements are relied upon to provide a 'ground truth' for radar observations, however there exists a discrepancy in scale (in comparing a $1\ km \times 1\ km$ radar grid to a point measurement). For this reason, the calibration of radar images is an area of active research requiring advanced correction algorithms. For example, see *Lanza et al.* [2001].

1.6.4 Summary of rainfall data

Figure 1.5, showing the number of recording instruments over time, summarises the history of rainfall measurement in Australia. Daily rainfall gauges, having the simplest technology, are the earliest and most common source of rainfall measurements. Pluviograph gauges, initially recording on drum-charts, were installed from the 1920s onwards and offer improved resolution in time compared to daily gauges. They are, however, far less numerous. Radar is a relatively modern technology, having been established for weather reporting since the 1960s.

In contrast to gauge-based measurements, radar remotely infers rainfall amounts and requires considerable amounts of calibration to provide best estimates. Whereas radar has poor temporal content and detailed spatial content, (networks of) rain-gauges have poor spatial content but have comparatively detailed temporal content. In addition to these primary sources of rainfall-related data, there is a wide catalogue of observations from ground-based instruments for other atmospheric processes linked to the rainfall-runoff process. These include, but are not limited to, pressure, wind, temperature, evaporation, humidity, soil moisture and groundwater observations. Another source of data is remote satellite observations, dating from the first weather monitoring satellite launched in 1960 by the National Oceanic and Atmospheric Administration. Environmental monitoring

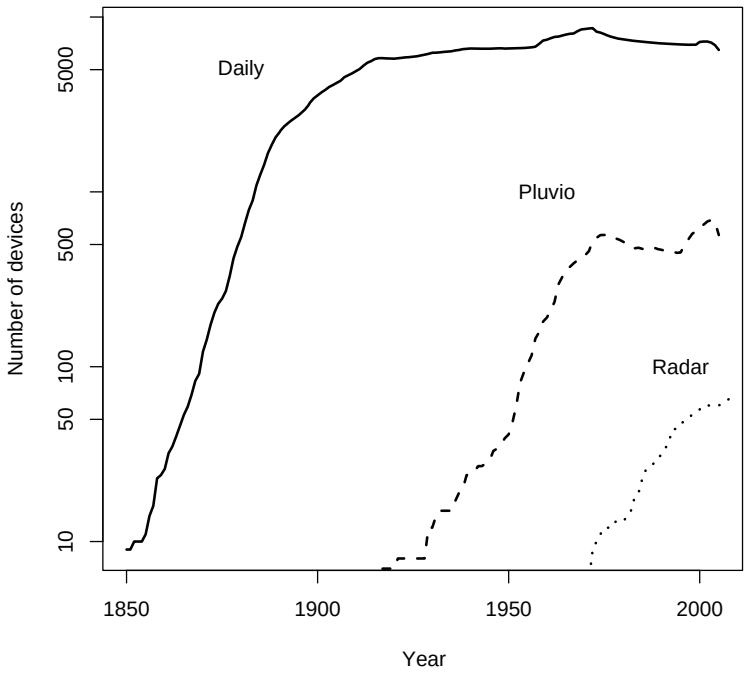

Figure 1.5 Number of commissioned gauges/radars in time.

satellites are similar to radar in that they measure reflected intensities that are mapped to properties such as moisture content in various layers of the atmosphere. However, these measurements are at a coarser spatial resolution than ground-based measurements.

1.7 THESIS OBJECTIVES

There has been considerable growing interest in space-time models of rainfall occurrence. There are many models based on radar imagery [*Northrop*, 1998; *Seed et al.*, 1999; *Mellor et al.*, 2000; *Mackay et al.*, 2001; *Allcroft and Glasbey*, 2003; *Pegram and Clothier*, 2001a], though they are mostly focused on the ability to reproduce single storm events. There are also numerous space-time models based on rain-gauge data that have been developed for continuous simulation [*Cowpertwait*, 1995; *Shah et al.*, 1996; *Willems*, 2001; *Zhang and Singh*, 2006]. This thesis is focused on space-time models based on rain-gauge networks in preference to radar imagery for several reasons:

1. Rain gauges are widely available, which enables gauge based models to be applied more widely.

2. There are significant issues of bias and error associated with radar-rainfall imagery such that there are few archived datasets that have had thorough and reliable quality control applied to them. These issues arise from the difficulty in merging gauge based rainfall measurements with radar sequences.

3. Radar archives are not as long as rain gauge records and are less likely to capture extreme events.

4. Rain gauge based algorithms should, in principle, be less computationally demanding than radar based models.

5. While rain gauge based models will have a simpler spatial structure than radar based models, there is an argument that this is compensated for by the runoff process. In short, rainfall is spatially aggregated during the runoff process when it lands on a catchment.

For reasons further outlined in Chapter 2, the Spatial Neyman-Scott Rectangular Pulse (SNSRP) model of *Cowpertwait* [1995] is used as the basis for the developments outlined in this thesis. In short, (i) it is based on temporal Neyman-Scott models of rainfall which have demonstrated a good reproduction of temporal statistics based on numerous case studies (ii) the SNSRP model has itself been applied to several case studies and (iii) other models require large numbers of subdaily gauges and cannot easily incorporate daily rainfall totals. This last point is especially significant given Australia's limited availability of sub-daily data. The primary objective of this thesis is therefore to develop a rainfall model that:

- is able to be simulated for long periods of time at a sub-daily time increment,
- is able to be simulated over large regions at a high spatial resolution ($1\ km \times 1\ km$),
- is able to reproduce key statistics and in particular intensity frequency curves used for flood designs,
- is able to reproduce important storm features given the information content of the data,
- is able to account for significant sources of variability such as seasonal and climatic variation,
- is accompanied by feasible computational methods, and
- is appropriate to Australian climatic conditions.

1.8 THESIS STRUCTURE

There are ten chapters in this thesis.

Chapter 2 provides a literature review so that the work of *Cowpertwait* [1995] is placed into context. Chapter 3 discusses the approach of *Cowpertwait* [1995] separately. There are no new results in Chapter 3. Rather, it serves to provide background information needed to understand the formulation and calibration of the model.

Chapter 4, 5 and 6 present pragmatic and general developments to the calibration and simulation methodology (as distinct from the conceptual formulation of the model). Chapter 4 addresses the simulation technique of the SNSRP model. It is shown that a spatial buffer is a computationally demanding means for simulating the model and that an analytic alternative is developed which is much more efficient. Chapter 5 highlights a calibration issue arising from independent monthly calibrations. Constraining the range of one parameter (β) is demonstrated as a practical solution. Chapter 6 investigates alternative calibration techniques that do not rely on analytic derivations of model properties and uses the Launceston region as a case study. A simulation-based calibration is demonstrated to be more flexible than analytic derivations, however it is not pursued in subsequent chapters owing to its computational burden.

Chapter 7 presents an extension to the conceptual formulation by introducing a storm radius parameter. This development arises from an inability of the original SNSRP model to correctly reproduce correlations at long distances. A case study involving 76 gauges across metropolitan Sydney is used to demonstrate this extension.

Chapter 8 addresses temporal heterogeneity. A methodology for analysing extreme values taking into account both seasonal and climatic partitions is presented. The utility of this method is demonstrated in a flood design for a hypothetical detention basin. The concept of partitioning extremes is then applied to the SNSRP model as a means for model verification. A case study involving the Burke region in New South Wales is used because it has both a strong climatic signal and a strong seasonal signal.

Chapter 9 formulates an inhomogeneous SNSRP model. This model couples an inhomogeneous spatial process of rain cells with the Neyman-Scott temporal process. The model is applied to the Avon Basin in Western Australia and is shown to explain variability in spatial correlations better than the homogeneous SNSRP model.

Chapter 10 provides concluding remarks and recommendations for ongoing research.

1.9 OVERVIEW OF AUSTRALIAN CASE STUDIES

A summary of case studies that the SNSRP model has been applied to is shown in Table 1.1. The top portion of the table shows existing studies in the United Kingdom, New Zealand and Italy while the lower portion summarises Australian case studies presented in this thesis. Existing applications of the SNSRP model have predominantly been to temperate, winter-dominated climates such as those in Europe and New Zealand. As the model has a European heritage, it is well suited to these conditions. A characteristic feature of these climates is that storm events are not able to be easily distinguished from one another. *Frost et al.* [2004] fitted the temporal counterpart of the SNSRP to individual gauges about Australia and showed that, amongst other limitations, it has some difficulty reproducing the distribution of dry-spells between storm events. This observation is due to the more event-based occurrence of storms in some parts of the Australian climate which gives rise to long dry-spells. This statistic is not specifically addressed in this thesis, but is strongly related to statistics that are, including the dry portion, skewness and coefficient of variation (CV).

For ease of comparison, the statistics shown in Table 1.1 are non-dimensionalised and are averaged over all months of the year. The case studies are summarised as follows:

- Launceston, located in north-eastern Tasmania, has similar statistics to existing case studies in the literature.

- While Sydney's rainfall is uniform throughout the year and is considered temperate, it has a mixture of generating mechanisms: winter is dominated by temperate southerly/south-westerly systems and summer is dominated by sub-tropical systems. While it has a similar dry portion compared to previous studies, it has a high degree of variation and skewness.

- Bourke is located inland near the border of New South Wales and Queensland. It experiences most of its rainfall from isolated but heavy storm events and has a high variability, skewness and proportion of dry days all year round.

- The Avon Basin case study, located in Western Australia, experiences a mixture of climates owing to its large area and geography. These range from temperate conditions along the coastline to a more arid climate inland. The coastline therefore experiences regular wetting during winter months some of which reaches inland, while during summer months the inland region is wetted infrequently. The portion

Table 1.1 Summary statistics for varied case studies which the Neyman-Scott model has been applied to.

Catchment	Location	No. Gauges [Daily,Sub-daily]	Ave. 1 hr Skewness[c]	Ave. 24 hr CV	Ave. 24 hr Dry Portion	Climate	Reference
Thames Basin[a]	U.K.	[0,23]	13.2	2.28	0.56	Temperate / Cool	*Munroe* [2005]
Auckland[b]	N.Z.	[0,1]	14.6	2.73	0.66	Temperate, Winter Dominant	*Cowpertwait* [2004]
Arno Basin	Italy	[0,9]	22.43	2.94	0.70	Temperate, Winter Dominant	*Cowpertwait et al.* [2002]
Launceston	Aust.	[35,6]	11.48	2.71	0.68	Temperate, Winter Dominant	Section 6.3
Sydney	Aust.	[52,24]	19.91	3.92	0.67	Temperate, Uniform Rain	Section 7.6
Bourke	Aust.	[16,1]	23.25	6.05	0.93	Arid	Section 8.5
Avon Basin	Aust.	[422,10]	18.58	3.92	0.78	Temperate (coast) / Arid (inland)	Section 9.5

[a] A smaller case tudy (6 gauges) for the same basin is reported in *Cowpertwait* [1995]

[b] This study is related to a larger case study of Auckland City (12 gauges) that is not reported in literature

[c] Non-dimensionalised skewness

of dry days is quite high in this region and ranges from 0.60 during winter months up to 0.92 in summer. The skewness and variation are also high in this region.

Section 1.6 outlined the availability of daily and sub-daily rainfall data about Australia, and in particular that long sub-daily records are uncommon. For this reason, daily rainfall gauges are jointly used with sub-daily rainfall gauges in this thesis. This is shown in Table 1.1 where previous studies have relied solely upon sub-daily gauges. A notable exception is the study of *Cowpertwait et al.* [1996b] who developed a regionalised model for any point within Great Britain. A secondary aspect of the case studies in this thesis is that larger numbers of gauges and larger regions have been considered for model calibration, and in this context, daily rainfall gauges are invaluable. It is a feature of the SNSRP model in comparison to other models that it can easily accommodate data at different aggregates in the calibration.

The use of large numbers of gauges brings into question the quality of the observed records and also the homogeneity (similarity) of the data over the defined region. Regarding the quality of the observed gauges, several safeguards have been employed. Firstly, all of the case studies make use of a minimum number of data-points to ensure that individual statistics are not too variable. This specification is arbitrary and varied for each study depending on the availability of gauges. Typically a condition of at least 10 to 20 years worth of available record is specified *for every month* in the record. The requirement for every month is a result of the independent calibration of monthly parameters. The estimation of statistics relies upon the presence of appropriate flags for missing data and accumulations of data. For the latter, the total amount of rainfall over the accumulation is considered reliable but the sequence in which it occurred is not. Where a data-point is flagged as accumulated it was completely removed from the calculation of statistics that require information on the pattern in which the rainfall occurred (such as the auto-correlation, cross-correlation, dry-portion, etc.). Statistics that require only a total accumulated depth can make use of the accumulated data, for example, statistics at higher aggregates (such as monthly totals) and calculation of mean rainfall amounts (e.g. daily). Cross-correlations require the period of two gauges to be overlapping in time. Where gauges do not overlap considerably (say less than 250 data points) the cross-correlation statistic was not estimated. This is acceptable since there are many permutations of gauge-pairs that can be used to estimate cross-correlations. Where rainfall observations are completely missing, these periods were excluded from the estimation procedure and in the event that a significant portion of a month is missing (say $>$3 days) the entire month of data was removed. The procedure of *Viney and Bates* [2004] was not applied to the daily records to check for

unflagged accumulations, which could be a source of bias in the distribution of extreme values if a gauge had a significant number of unflagged accumulations. The question of homogeneity of a study region is important, as this is a requirement of the existing SNSRP model formulation. To date this requirement has been acceptable since the model has been primarily applied to smaller case studies that by nature of their size have a greater chance of being similar. Over larger regions it becomes less feasible to maintain this requirement. Chapter 9 provides a formal treatment of this topic for the SNSRP model.

Chapter 2
Literature Review

The variability in space and time of precipitation processes affects both the infiltration dynamics of the soil surface of a catchment and the production and successive propagation of surface runoff. This is also fundamental to the rate of accumulation of water above and below ground levels and the rate of evapo-transpiration in the area. The entire circulation of water in a basin is thus governed by the spatial and temporal distribution of rainfall.

[*Bacchi and Kottegoda*, 1995]

Given rainfall's important role as the driving force of the water dynamics within a catchment, hydrology has always been concerned with accounting for its inherent variability. Throughout the history of hydrology this has not always been possible, and limiting assumptions have been applied to both spatial and temporal rainfall analyses in step with the understanding and technology of that era. Techniques restricted to the spatial domain have not been unique to hydrology, for example spatial interpolation and disaggregation have been employed widely in other areas. A similar observation is true for techniques restricted to the temporal domain, including time-series analysis, cluster models, Markov models and event models. Intermediate attempts to include both spatial and temporal aspects have typically employed the same concepts as temporal models, but in a multi-variate setting. These include regionalised temporal models and multi-site models. By way of contrast, space-time models aim to simultaneously accommodate both spatial and temporal sources of variability. They can be considered in two sub-categories depending on the underlying source of data, either radar imagery or gauge networks. Radar-based models tend to be well described in the spatial domain but poorly account for temporal development, while gauge-based models have good temporal performance but are limited in their spatial characterisation.

Whether temporal or spatial, most approaches to rainfall modelling are typically statistical and conceptual by nature, as distinct from physically deterministic models. Physical models are generally based on the Newtonian dynamics of multiple state variables (e.g. pressure, temperature, water content), as represented by a system of non-linear partial differential equations [*Kessler*, 1969]. Layers of the atmosphere (and soil) are represented as grids of discretised cells each containing the state variables to be updated. Physical atmospheric models suffer from several limitations: (i) sensitivity of the partial differential equations to initial conditions, (ii) inability to specify boundary conditions due to a lack of observations, and (iii) requirements of immense computing power [*Georgakakos and Krajewski*, 1996]. In contrast to physical models, conceptualisations of the dominant processes, underpinned by stochastic sampling, offer a more viable approach for simulating long-term rainfall records. However, the idealisation of the dominant processes requires calibration to observed data and extensive comparison between simulated and observed records in order to verify the appropriateness of the model. Thus the emphasis of conceptual models lies in the preservation of correlations and key statistics across all scales of interest, as distinct from modelling the *true* underlying representation of the system. In this light, conceptual modelling requires the balance of model limitations (due to reductionist assumptions and idealisation) with respect to the perceived benefits of increased model complexity. The history of conceptual modelling is a reflection of this balance as governed by parameters such as data availability, computational power and scientific insight.

2.1 RAINFALL HOMOGENEITY

Prior to considering models of rainfall, it is essential to consider the similarity of rainfall processes either over a space or in time. Even early mathematical approaches in hydrology required an appreciation of this concept, for example [*Thiessen*, 1911] who calculated average rainfall totals over a region. A similar appreciation is also required in the time domain as rainfall records are now known to have many long-term variations, for example the El Niño Souther Oscillation phenomenon [*Bjerknes*, 1969].

2.1.1 Spatial homogeneity

In order to conduct a spatial analysis of rainfall, it is necessary to define the boundaries of a region within which there are similar hydro-climatic properties. If this does not hold, then so-called representative statistics calculated from within that region will have little

meaning. The term *homogeneous* is used in a general sense to refer to whether the underlying rainfall process in a region can be pragmatically considered to be the same for any statistic of interest. Though not an exhaustive assessment, it is convenient at certain points to refer solely to key statistics when assessing homogeneity. Related to homogeneity is the concept of *stationarity* over a region, but this term is reserved for implying that the expected mean and expected variance are constant over the region (weak stationarity). The term *heterogeneous* is used in a general sense as the opposite to homogeneous. In the context of Poisson models, the term *inhomogeneous* is used in the same sense as heterogeneous but is often reserved for the rate parameter of a Poisson distribution (e.g. rate of storms, rate of cells).

Assessment of the homogeneity over a region is a subjective problem because of the abstract basis on which a region is defined. This holds true whether using physical or statistical arguments. From a physical perspective, the climate is a gradually varying phenomenon (with the exception of strong geographic influences such as mountains) which makes it difficult to determine distinct boundaries. In addition to this, the scale of interest is important, as local geographic variability can cause discrepancies to the homogeneity of a region defined at a higher scale (e.g. a large valley in the middle of a hilly region). From a statistical perspective, the variability in observed records and the array of potential statistics used for comparisons inhibit conclusions over regional similitude. As an example of the ambiguity of a region, the Bureau of Meteorology have several maps that can be used to define a climatic region [*Stern et al.*, 2006]. These maps differ due to the scale of interest and also the combinations of observed data used in the analysis (rainfall, temperature, humidity). In a hydrologic context, *Hosking and Wallis* [1997] provide useful discussion on the homogeneity of a region with respect to analysis of extreme values.

2.1.2 Temporal homogeneity

Changes in the rainfall process over time are an equally important when considering homogeneity. Seasonality, while fitting this description, is such an obvious source of variability that its inclusion in hydrologic modelling is an unquestioned necessity. Modelling techniques typically account for this by allowing parameters to vary across seasons, for example by using new parameters each season, or by using harmonic functions to describe the parameter variations in time. Seasonality is typically straightforward to model because the 12-month seasonal period is well defined. Of more subtle significance is inter-annual variability, which has escaped critical investigation until recent decades because of the complexity of the governing processes, the necessity of long observation records and the

irregular period of its fluctuations.

Long-term climatic variability is linked to fluctuations in the temperature of the world's ocean basins and circulation patterns in the atmosphere. The El Niño Southern Oscillation phenomenon is one such system that operates on the scale of multiples of years and the Inter-decadal Pacific Oscillation is an example that operates across multiple decades. These systems are conceptualised by referring to the climate as being in a particular *state* corresponding to the general condition of the system and its expected behaviour. For example, the El Niño state of the Southern Oscillation is associated with warmer waters and convective activity in the East Pacific giving lower rainfalls over Eastern Australia. In contrast the La Niña state is associated with warmer waters in the West Pacific and higher rainfalls in Eastern Australia. An intermediate neutral state is said to act when neither state is clearly identifiable. In the context of streamflow extremes, *Franks and Kuczera* [2002] show that there is a significant difference between climate states and that this challenges the traditional modelling assumption that extreme values are independent and identically distributed.

In terms of modelling the temporal heterogeneity of rainfall, a similar approach to modelling seasonal variation can be adopted: use of varying parameters with respect to each state. [*Thyer and Kuczera*, 2000] demonstrate an example of this using a Hidden Markov Model, which determines the rainfall distribution and transition probabilities between hidden states that can only be inferred from the observation record. Another approach is to regress observed rainfall totals on climate indices [*Whiting*, 2006], which for simulation would require a time-series model of the climate indices.

2.2 Single-site Rainfall Models

Due to the limited spatial extent of rain-gauge observations, single-site rainfall is solely concerned with the temporal statistics of rain at a point in space. Having reached a mature stage of research, single-site rainfall provides the historical context for many of the developments in stochastic rainfall modelling and much of the foundation from which space-time rainfall modelling has emerged. The various conceptualisations of single-site rainfall are highly associated with the scale of interest. For example annual and monthly models can be transformed to Gaussian distributions with relative ease and for many locations they are free of no-rain observations. If these conditions are satisfied then time series models can be employed. At the daily scale, the wet-dry intermittency and skewed distribution of totals become significant characteristics of rainfall and requires a different approach. Markov models and non-parametric models are popular at this scale. Sub-daily

rainfall is characterised by even greater runs of no-rain and rainfall totals that are highly variable and correlated in time. Event-based models and Poisson-cluster models that have separate variables for each component of a storm are common at this scale. Temporal disaggregation, for example from monthly to daily or from daily to sub-daily is a common counterpart to rainfall models to increase their resolution in time while preserving the statistical properties of the model at higher aggregates.

2.2.1 Event-based models

Event-based models seek to characterise a rainfall record as a succession of storm *events* dispersed by periods of no rainfall. This formulation requires the definition of what constitutes an event, which is a subjective process, but is necessary so that storms can be distinguished within observed rainfall records. This is in contrast to other models (such as Poisson-cluster models) that allow storms to overlap in time. Many event-based models are alternating renewal models that alternate between successive wet and dry events where random variables are assigned to the event lifetime, L, and inter-event duration D, and the event intensity, X, as shown in Figure 2.1. A slight variation is to model the total storm depth in preference to the event intensity.

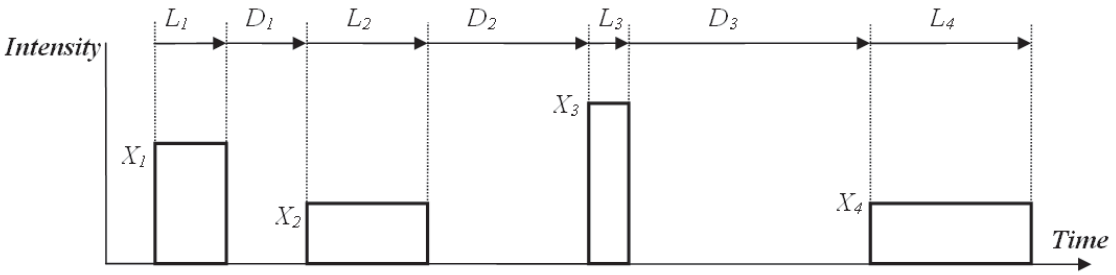

Figure 2.1 Schematic diagram of an alternating renewal model of storm events.

Green [1965] developed early alternating renewal models in preference to Markov models because they offer a more flexible representation of event and inter-event lifetimes. Specifically, the distributions of duration can take on forms other than a geometric distribution. Whereas *Green* [1965] used daily rainfall to calibrate wet and dry sequences, *Grace and Eagleson* [1967] used ten-minute data for calibration and included the event depth in their model. The authors defined the independence of storm events using a minimum of two-hours of no-rain, thus observed storms were able to have intra-event dry periods less than this amount. *Grace and Eagleson* [1967] used linear regression to model the conditional relationship between the event depth and event duration, a Weibull distri-

butions for the event and inter-event lifetimes and a beta distribution for the residuals of the event depth.

More recently, [*Heneker et al.*, 2001] developed an alternating renewal model fitted to six-minute rainfall data and using an independence criterion of two hours between storm events. Modified-exponential distributions were used for the event duration and inter-event duration and the event intensities were modelled using a Pareto distribution with parameters conditioned on the event duration. The authors used Bayesian calibration techniques that were able to account for missing observations, which is a significant advantage because of their prevalence in pluviograph records [*Heneker*, 2002]. Bayesian calibration provides a natural framework for obtaining the predictive distribution which allows for parameter uncertainty to be quantified. The model allows for seasonality by partitioning the observed record with respect to each month and using separate parameters in each month. To allow for long-term persistence in observed rainfall records, the model was extended to include climatic partitions, each having their own set of parameters. This extension however requires an additional model to simulate the overarching climatic process in which the rainfall model is embedded, for example, Hidden Markov Models [*Thyer and Kuczera*, 2000].

The model of *Heneker et al.* [2001] employs a disaggregation scheme to distribute the storm rainfall throughout the duration of the event. The total duration and rainfall depth of each storm is represented using a dimensionless mass-curve and the disaggregation proceeds via nested jumps sampled from a truncated log-normal distribution. The model shows a good reproduction of numerous statistics across a broad range of scales from 6-minute aggregates up to annual aggregates. The model also shows a reasonable reproduction of extreme rainfall bursts.

A common limitation of the alternating renewal framework is that it only allows for the definition of one event type. This can produce considerable difficulty in fitting distributions to event properties when the observed data show structured relationships between the event variables (i.e. more than one storm type). This problem is common and *Sansom and Thompson* [2003] addressed an alternative to alternating renewal models which introduced a Markov framework that models multiple types of events and the transition between them.

2.2.2 Poisson-cluster models

The Poisson distribution is derived as the discrete number of independent events on a continuous domain according to some rate parameter. Corresponding to this, a point-

process has a Poisson distribution for the number of events that are placed with uniform probability over the domain of interest. For single-site models, the domain is the time-line, though point-process models are equally common in the spatial domain. A marked point-process is one where properties are assigned to each event according to specified distributions. *Cox and Isham* [1980] present the general theory for point-process models.

Rodriguez-Iturbe et al. [1984] developed a Poisson White Noise model, which is a type of marked Poisson model. The formulation is shown in Figure 2.2 (a) having a Poisson process to specify the rainfall bursts in time and instantaneous rainfall intensity as the sole marked process. The second-order moments of the model were determined and fitted to aggregated rainfall totals. A later improvement, termed the Poisson Rectangular Pulses model, uses an additional marked process to represent the duration of the rainfall burst [*Rodriguez-Iturbe et al.*, 1984, 1987a,b]. As shown in Figure 2.2 (b) the rectangles are allowed to overlap, which is in contrast to event-based models. Where storms overlap the resultant process is the aggregate of the pulse intensities. Variations of the models can be achieved by specifying different distributions for the marked processes. Despite this however, these models demonstrate inadequate comparisons to observed statistics, in large because the arrival of rainfall is not uniform in time but *clustered* by association with storm events.

Poisson-cluster models such as the Neyman-Scott and Bartlett Lewis models were originally adapted from models for the spatial distribution of galaxies, investigation of traffic accidents and prediction of computer failures [*Neyman and Scott*, 1958; *Le Cam*, 1961; *Bartlett*, 1963; *Lewis*, 1964]. Poisson-cluster models use a mother-daughter process, where the mother-event distributes events across the domain according to a Poisson process. Daughter-events are distributed within the domain relative to the mother-event and by this association tend to be clustered together. In terms of rainfall, the mother-process describes the Poisson arrival of storms and the daughter process describes the arrival and properties of rain-cells associated with that storm. *Rodriguez-Iturbe et al.* [1987a] show that clustered models perform better at reproducing the continuous rainfall process than non-clustered Poisson models.

The Bartlett-Lewis Rectangular Pulse (BLRP) model is one type of Poisson cluster model [*Rodriguez-Iturbe et al.*, 1987a], and is depicted in Figure 2.2 (c). The storm origins arrive according to a Poisson process, and a number of cells are associated with the storm according to some distribution. Each cell is defined as a marked-process having random variables to describe the cell lifetime, L, intensity, X and starting time S. The first cell coincides with the onset of the storm and subsequent cell starting times are distributed

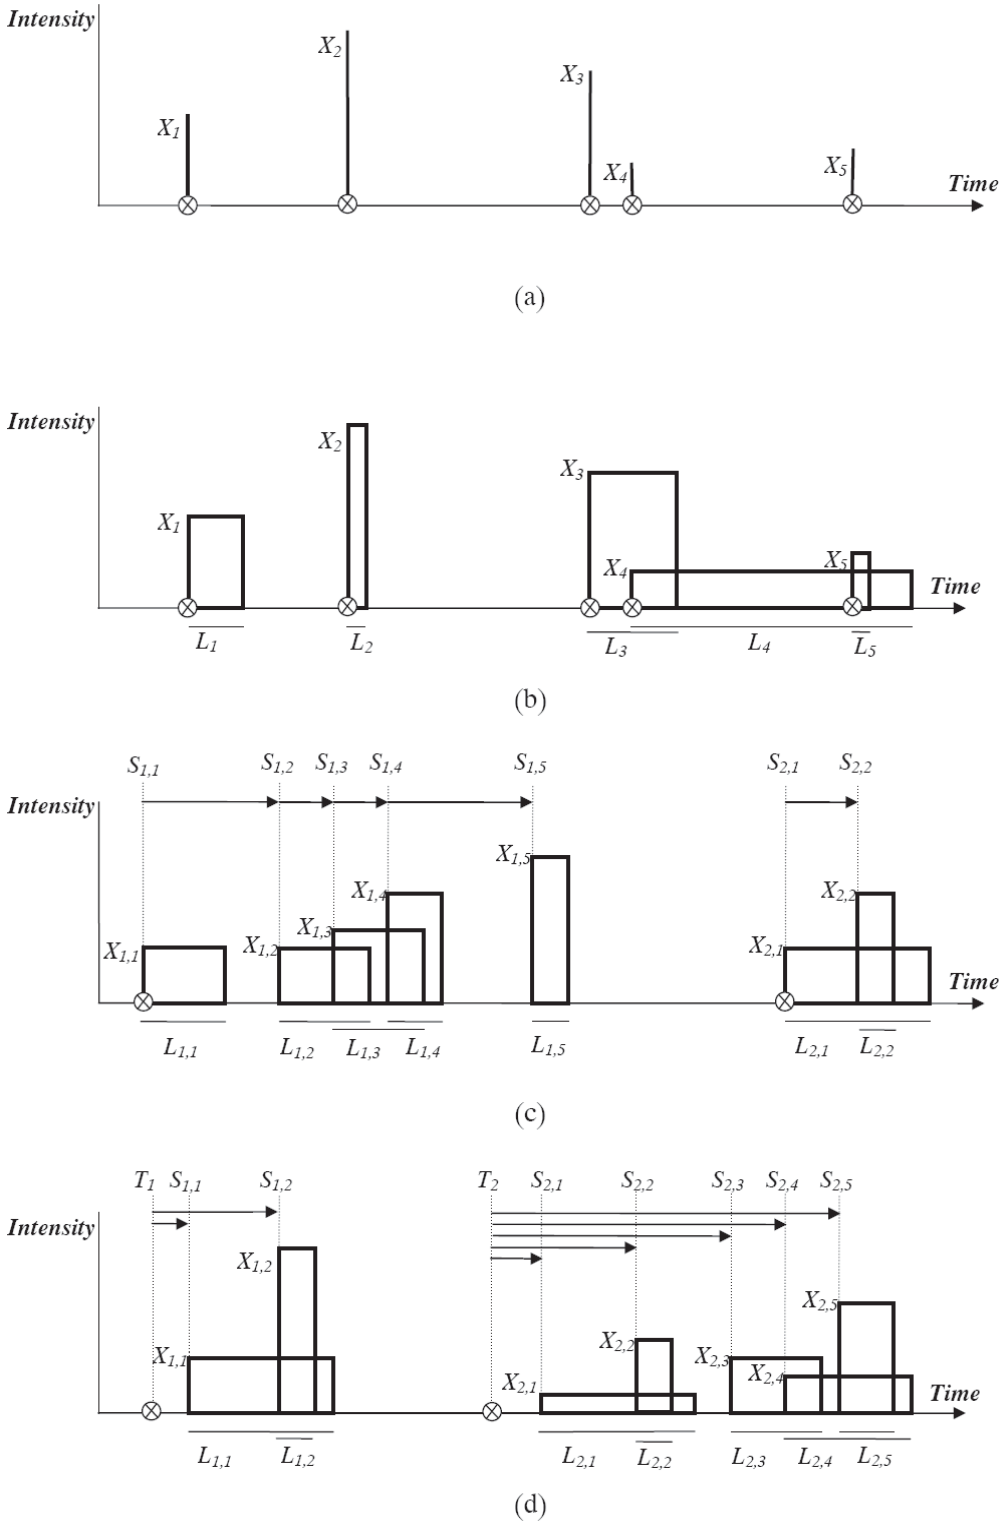

Figure 2.2 Schematic diagrams for various point-processes models (a) Poisson White Noise, (b) Poisson Rectangular Pulse, (c) Bartlett-Lewis Rectangular Pulse, (d) Neyman-Scott Rectangular Pulse. X represents cell intensities, L represents cell lifetimes, S represents cell start times and T represents storm start times.

according to an exponential distribution relative to the start of the previous cell, $S_{i+1} - S_i$. By way of contrast, the Neyman-Scott Rectangular Pulse (NSRP) model has the starting time of cells following an exponential distribution defined relative to the start of a storm $S_i - T$. Also, the first rain-cell does not coincide with the start of the storm, as shown in Figure 2.2 (d).

Apart from minor changes to the distributions of rain-cell properties, numerous extensions have been suggested for the BLRP model and *Onof et al.* [2000] provide a good summary of these. *Onof and Wheater* [1993] used a gamma distribution to incorporate parameter uncertainty into the rain-cell duration. While the proportion of dry intervals improved, discrepancies were observed in other statistics such as the auto-correlation, average event duration and extreme values. *Onof and Wheater* [1994b] improved the model by incorporating monthly seasonality into parameterisation, which improved the aggregate statistics across all time scales, but it overestimated the proportion of dry intervals. *Onof and Wheater* [1994a] applied the superposition of a high-frequency jitter process onto the rectangular pulse process of the BLRP to improve the reproduction of auto-correlations. When estimating the parameters for this model, several difficulties were noted and methods to reduce the overall number of parameters were suggested.

Regarding the NSRP, *Cowpertwait* [1994] developed a generalisation of the model to have multiple cell types which was later refined to the superposition of different storm types [*Cowpertwait*, 2004]. These extensions require additional parameters but provide a more flexible representation, for example, by allowing both intense short duration storms and lighter longer duration storms. *Cowpertwait et al.* [1996a] incorporated wet-transition probabilities into the model to improve the reproduction of historical dry spell sequences and dry-day proportions. *Cowpertwait* [1998] derived the skewness property of the NSRP formulation and used it in conjunction with hourly data to improve the reproduction of extreme values. Another statistic derived for the NSRP is the probability that an arbitrary interval is dry which has also been used in calibration to improve the reproduction of dry-spells [*Cowpertwait*, 1991].

Poisson-cluster models require the derivation of the model properties for a given formulation so that they can be calibrated to the equivalent statistics from aggregated rainfall data. Although the parameters of the cluster models often have physical interpretation, it can be difficult to calibrate these parameters because the model aggregates overlapping cells which makes it difficult to distinguish where cells start and stop. With the exception of *Chandler* [1997] who used spectral theory to develop a likelihood function, Poisson-cluster models are mostly calibrated using weighted least-squares comparisons

of the model properties across a range of time-scales [*Salas*, 1993].

2.2.3 Disaggregation models

There are many models not discussed in this review as they typically simulate at a daily or higher timescale, whereas the interest here is in sub-daily extremes of rainfall. A common technique to obtain sub-daily rainfall estimates is to use a daily rainfall model with a disaggregation technique. Disaggregation models often consider an entire rainfall record as a whole which is then split into rainfall sequences at a higher resolution according to some functional relationship. The attraction of disaggregation models is that they are conceptually simple, they provide a parsimonious approach to describing rainfall over a broad range of scales and, in many cases, they have provided good statistical fits. While there are numerous approaches to disaggregation, the discussion here is restricted to fractal-based (cascade) models.

The phenomenon of fractal patterns in rainfall is based upon empirical evidence about scale invariant properties of rainfall and upon an analogy with random multiplicative cascade models for fully-developed turbulence [*Zawadzki*, 1987; *Lovejoy and Schertzer*, 1990; *Menabde et al.*, 1997]. An atmospheric turbulent regime can be thought of as partitioning and concentrating energy, water and other fluxes into smaller and smaller regions. Scale invariance refers to statistical properties that are related by a scale changing operation involving only scale ratios, such as those observed in the energy spectra of rainfall. Multiplicative cascades are used to disaggregate rainfall from a larger scale to a finer scale by dividing the rain into branches according to random weights. Conserved cascades originally required that rainfall amount was conserved between scales by assuming that the weights at each scale summed to unity. Stochastic cascades were later developed by assuming that the weights sum to unity only on average, allowing better characterisation of the high variability of rainfall [*Menabde et al.*, 1997].

The scale invariant assumptions of rainfall posed by a fractal framework are seen to be in direct contrast with the observed hierarchical organisation of precipitation fields as modelled by conceptual models [*Kumar*, 1996]. While numerous authors have reported scale invariance for a significant range of scales [*Tessier et al.*, 1994; *Burlando and Rosso*, 1996] deviations from scale invariance have also been reported [*Olsson*, 1995; *De Lima and Grasman*, 1999; *Olsson and Burlando*, 2002]. [*Kumar*, 1996] used wavelet analysis of high-resolution temporal rainfall to demonstrate the presence of coherent structures coming to strongly oppose the scale invariant assumption. *Veneziano et al.* [1996] used empirical and analytical evidence to invalidate the scale invariant assumption, and they

advocated the use of multiplicative cascades by developing a multiplicative model that is scale dependent. Despite the controversy, cascade models are popular from a pragmatic point of view: that they provide a highly parsimonious framework for modelling rainfall at a range of scales. Disaggregation models also have several limitations, these include the difficulty of combining them with physical processes, for example, storm advection.

2.2.4 Regionalised single-site models

An often cited limitation of single-site models is that they can only be applied at the location of the rain-gauge used in the calibration. For example, high-quality rain-gauges may not be situated in or near the catchment of interest, or as another example, there may be several nearby gauges, but due to strong orographic effects, it would be preferred to have parameters corresponding to the exact location of interest. One method to overcome this is to calibrate the model to multiple sites and then spatially interpolate the parameters to produce a map of parameter values for any location of interest. This technique is not uncommon, as for example, intensity-frequency-duration parameters are mapped across Australia [*Pilgrim*, 1987].

Sansom and Thompson [2003] applied the thin-plate spline interpolation algorithm of *Hutchinson* [1995] to 49 separate parameters of a point rainfall model across 20 sites in central New Zealand. It was noted that this process is not straightforward as the parameters cannot be interpolated separately since certain relationships between them must be preserved. To overcome this an iterative process was used and additional spline surfaces were used to account for the relationships between the variables. The parameters were then determined for 563 grid points of size $6 \times 6\ km^2$. *Cowpertwait et al.* [1996b] conducted a regionalisation of the parameters for the Neyman-Scott model across England using a multiple regression technique. The study uses 112 gauges, where 85 gauges had daily data, 19 had hourly data and 8 had $1-$minute resolution. The regression terms included the elevation, the distance from the coast, a north Ordnance Survey grid reference and a binary indicator of the east-west position. Harmonic parameters were used to account for seasonality and reduce the overall number of parameters to be regionalised.

Gyasi-Agyei and Bin Mahbub [2005] provide a regionalisation approach that builds on the hybrid Bartlett-Lewis / jitter model. This model has been demonstrated for a wide number of sites about Australia. They demonstrate a robust calibration approach which makes the model stand out from other Poisson cluster models.

While regionalisation techniques based on interpolating parameters allow for a spatial dimension in the modelling of rainfall, the resulting model remains a single-site model

as it does not account for any space-time correlation. Nonetheless, for parameters having a physical interpretation, a comparison of values with respect to spatial coordinate can allow some inference about changes in the storm process across the region [*Sansom and Thompson*, 2003].

An alternative regionalisation technique is demonstrated by *Jennings et al.* [2009] that is used along with the model of *Heneker* [2002]. They use master-target approach where power-scaling parameters are used to modify the distributions of storm properties from the master (sub-daily gauge) to the target (daily gauge or very short sub-daily gauge). They demonstrate their approach for a number of sites about Australia ranging up to a separation distance of 190 km between the master and the target.

2.3 MULTI-SITE MODELS

Multi-site models provide the most common extension of single-site models in the spatial domain, often by using a multivariate equivalent of the single-site process but also incorporating spatial correlations.

Multi-site ARMA models are common at monthly or annual scales [*Srikanthan and McMahon*, 2001], although they have also been applied at the daily resolution [*Dalezios and Adamowski*, 1995]. With respect to daily rainfall, the Markovian framework is a more common approach for developing multi-site models [*Zucchini and Guttorp*, 1991; *Hughes and Guttorp*, 1999; *Wilks*, 1998; *Harrold et al.*, 2003]. *Thyer and Kuczera* [2003] developed a daily multi-site Hidden Markov Model for simulating regional climatic persistence and used Bayesian techniques to improve parameter estimation. *Sanso and Guenni* [2000] have applied a truncated Gaussian model to 80 rainfall stations using 10 day rainfall aggregations. This model utilizes a novel method for modelling the intermittency of rain, assuming that the data correspond to normal random variates that have been truncated and transformed. The truncated values correspond to the dry periods whilst wet periods are represented by the transformed normal variates. This model is also fitted using Bayesian techniques to handle dry periods, missing values and uncertainty in parameter estimates. [*Cowpertwait*, 1994] implemented a multi-site version of the NSRP model having some parameters common to all sites and calibrated to daily and sub-daily observed rainfall records.

Common limitations of multi-site models are that they are not continuous in space as the rainfall is modelled only at the gauge-locations. Also, unless techniques are employed to make parameters common across multiple sites, the number of estimated parameters can increase significantly with respect to the number of observed gauges.

2.4 SPACE-TIME MODELS

Taylor's hypothesis for fluid turbulence [*Taylor*, 1938], states that spatial correlations are related to temporal correlations through a velocity parameter. The implementation of radar imagery in the 1960s has verified this observation for precipitation fields and has provided numerous other insights into the space-time structure of rainfall [*Crane*, 1990]. Additionally, the technological advances in computing ability have extended the horizon of rainfall modelling from single-site rainfall to high resolution space-time modelling. There is a wide variety of approaches that seek to model space-time rainfall including empirical studies, radar calibration algorithms, short-term forecasting (now-casting) models, design-storm models and continuous simulation models. While continuous simulation models are the primary focus of this thesis, it is important to acknowledge the variety of approaches as they offer scope for further innovations in space-time rainfall modelling.

2.4.1 Radar-based models

Radar rainfall estimates are inferred from remote measurements of reflectivity and are subject to numerous errors and biases [*Jordan and Seed*, 2003]. For this reason, radar calibration and empirical studies of rainfall structure are two areas of on-going research. Although radar calibration techniques are distinct from simulation models, they provide useful techniques for handling the spatial error structure of rain-gauge and radar measurements. The most common technique for combining block (or pixel) radar estimates and point rain-gauge measurements is block co-Kriging [*Krajewski*, 1987]. More recent advances have provided improved estimates of spatial rainfall using artificial neural networks [*Matsoukas et al.*, 1999], block co-Kriged Bayesian Kalman filtering [*Todini*, 2001] and conditional merging [*Sinclair and Pegram*, 2004].

Using empirical observations of convective storms as a basis, *Waymire and Gupta* [1981] provided one of the first conceptual descriptions of space-time rainfall in a mathematical framework. The features of main concern are the size, shape and relative motions of rain-cells, cluster potential regions, rain-bands and storms. Rain-bands occur within storms and move in the same direction as the storm, sometimes faster than the storm itself. Rain-cells within the rain-bands are born, grow, decay and die, moving in a direction that does not generally coincide with the direction of storm movement. Also within the rain-bands, there are regions of high rainfall activity, called cluster potential regions, where rain-cells have a tendency to cluster. The storms are up to $250\ km$ wide and generally move with speeds of $20\ km/hr$ to $90\ km/hr$.

Mellor [1996] implemented the specification of *Waymire and Gupta* [1981] in the

Modified Turning Bands (MTB) model and were able to recreate the physical and statistical aspects of a storm using physically meaningful storm parameters. This model uses lagged correlations of radar images to determine the rain-cell velocity. Once the velocity is determined, the spatial extent and the temporal extent of fluctuations were independently determined and related to the lifetime and intensity of rain-cells. The average displacement of the respective edges can be used to estimate the storm speed, however there has been difficulty in estimating several storm parameters, such as the size and speed of rain-bands and cluster potential regions. The MTB model has been successfully applied to generate stochastic ensembles of storm evolutions, as is required for now-casting applications [*Mellor et al.*, 2000]. There are numerous other potential applications for the model, including rain-gauge design, expansion of historic temporal data sets into the spatial domain and analysis of extreme events conditioned upon a continuous temporal model [*Mellor*, 1996]. The MTB has not been applied in the context of long-term continuous simulation, in part due to the highly computational nature of the model.

Northrop [1998] developed a space-time Poisson-cluster model for calibration to radar rainfall observations. Storms arrive according to a Poisson process and give rise to a random number of elliptical rain cells aligned with the direction of storm advection. Each cell moves at the overall storm velocity for a random amount of time prior to terminating. Rain is deposited by the cell at a random intensity that is constant over the area of the cell and over its lifetime. Two variations of the model were investigated: (i) each cell within the storm is displaced from the storm centre according to an envelope of Gaussian probability and (ii) each cell is uniformly located within the circumference of an ellipse about the storm origin. The main properties of the model are derived using analytic and numerical integration techniques for comparison to radar rainfall estimates. The model was demonstrated for a single storm recorded at the Wardon Hill radar in southwest England.

Seed et al. [1999] developed a space-time model for application as a design-storm using fractal-based techniques that are more commonly used for disaggregation. The model uses a multiplicative bounded (multi-fractal) cascade to simulate seven different spatial scales of rainfall. The spatial model is extended into space and time by replacing the cascade fields of independent, identically distributed weights with fields of stochastic processes that are correlated in time. Specifically, the cascade weights at each level were linked to the previous time-step via an ARMA(1,1) process [*Over and Gupta*, 1996]. The model also included rainfall advection by means of a Lagrangian framework. The model was verified using rain fields produced by a monsoonal depression that passed over a weather radar at Darwin. The model was considered easy to calibrate and the parameters

were able to satisfactorily reproduce the statistical behaviour of rainfall over a wide range of spatial scales (down to $2 \times 2\ km^2$) and temporal scales (down to 10 min). There was some difficulty, however, in reproducing temporal correlations, since this depends upon spatial correlation, temporal development, and advection speed which are all likely to be non-stationary with an event.

Pegram and Clothier [2001a] developed the String of Beads model for sequences of storm events. The *beads* constitute a sequence of stationary, two-dimensional spatial images of rainfall that are threaded onto a one-dimensional time line, the *string*. The string is modelled using a Markov chain having a dry state and numerous storm states and it is calibrated from rain-gauge data. The beads are modelled as a three-dimensional Gaussian Random Field (GRF) having a power-law correlation in the two spatial dimensions and in time across the multiple images in the storm. The technique used to achieve this uses a convolution of the spectral density function with Gaussian noise in Fourier space to generate correlated images of $1\ km^2$ pixels [*Bell*, 1987]. For the original model, the mean, standard deviation and fluctuation rate of the rainfall were all conditioned on the wetted-area ratio of the radar images [*Pegram and Clothier*, 2001a]. A later development of the model improved the downscaling and temporal aspects of the model using two autoregressive processes to control the image mean flux and the pixel scale intensity [*Pegram and Clothier*, 2001b]. The techniques employed in the model have been explicitly developed with simplicity and computational feasibility for continuous simulation in mind. It has been demonstrated for two separate storms from the Bethlehem radar in South Africa and was able to reproduce important statistics.

Several authors have developed space-time analogues of spatial disaggregation processes [*Mackay et al.*, 2001; *Allcroft and Glasbey*, 2003]. However, these extensions are mostly from a statistical point of view and do not place a significant emphasis on the physical and temporal storm structures such as storm advection, storm growth and decay. *Mackay et al.* [2001] used a Markov Random Field technique that disaggregates the output from global circulation models having grid size $40 \times 40\ km^2$ down to a scale of $8 \times 8\ km^2$. This technique first considers the pattern of wet and dry pixels within radar images and then conditions the rainfall intensity of wet pixels with respect to their distance from the edge of the storm. Temporal dependence probabilities were included in the model formulation to allow for successions of disaggregated images to be modelled. *Allcroft and Glasbey* [2003] used a latent variable GRF approach to disaggregate rainfall images. This approach uses a continuous distribution for the rainfall process, but a truncation is applied to obtain dry-pixels and a transformation is used to model the skewed distribu-

tion of observed rain. Similarly to *Mackay et al.* [2001], temporal correlation parameters were included in the calibration so that sequences of images could be disaggregated, but detailed modelling of the temporal evolution of the storm was not considered.

2.4.2 Gauge-based models

Two of the first attempts at space-time models involved a relatively simple Poisson process for rain-cell occurrences in space and in time [*Eagleson et al.*, 1987; *Cox and Isham*, 1988]. *Eagleson et al.* [1987] considered three different shapes for the spatial distribution of the rainfall, (i) a Gaussian envelope, (ii) an exponentially decreasing envelope and (iii) a linearly decreasing envelope. Each model had isotropic correlation, assumed a homogeneous area and required only three parameters. The models were fitted to 428 storm days from 93 gauges over a 4 month period corresponding to the season of air-mass thunderstorms in Arizona. In terms of gross storm statistics, the model showed good agreement. However, the Poisson assumption of rain-cell arrival is limiting as detailed rainfall structure is better represented as a cluster process.

Cox and Isham [1988] used a highly idealised model, consisting of only one cell per storm, to investigate the analytic properties of space-time rainfall. Storms were defined to arrive as a Poisson process and were characterised with four random variables: velocity, V, lifetime, D, intensity, X, and radial coverage, R. In this sense the storm can be thought of as a disc of constant intensity rainfall over a circular area that moves across the region of interest. *Cox and Isham* [1988] provided an extension to the model where numerous cells are associated with the a storm and follow a Bartlett-Lewis process of arrivals in time. The storm velocity, V, is common to all cells, but all other variables (R, D, X) are separate for each cell. While the first model allowed for simplified analytical treatment, the more realistic cluster-based process was relatively inaccessible analytically. First and second order moments were derived for both models, but neither model was calibrated to observed data.

Shah et al. [1996] proposed a space-time rainfall model that uses the Turning Bands algorithm [*Matheron*, 1973]. This model does not consider rain-cells or rain-bands, instead it uses a fractionally-differenced ARIMA line process to create a GRF that has correlations in space and time. The GRF can reproduce either isotropic or anisotropic correlation structures, and several transformations are applied to the underlying GRF to generate realistic rainfall patterns:

- a linear advection is applied to the field to mimic storm movement,
- a scaling term to decay rainfall with increasing distance from the storm centre, and,

- exponentiation of rainfall values to reproduce the skewed distribution of rainfall.

The authors used an empirical calibration approach by adjusting parameters until a good agreement was reached between observed statistics and simulated output from the model. The model was calibrated to hourly rainfalls for a set of ten storms across 3 gauges from the Upper Wye catchment, Wales, England. A grid of 250×250 m^2 was used to cover the 10.55 km^2 area of the catchment. The model successfully reproduced observed statistics, however the authors note that the model requires more extensive assessment using larger catchments and more variable rainfall regimes having stronger convective rainfall.

Sanso and Guenni [1999] developed a space time model for monthly rainfall calibrated to 80 rain-gauges irregularly located over a 250×300 km^2 area of Venezuela for a time span of 16 years. The model is based on the assumption that data follow a truncated and transformed multi-variate Gaussian distribution. A significant advantage of the model is that it has been designed for sparse networks with historical records limited by length and having periods of missing data. The model was implemented in a Bayesian framework and makes extensive use of Markov chain Monte Carlo techniques for evaluating missing data and parameter uncertainty. The main limitation of the model is that it was applied at a monthly resolution at which, the transformation of rainfall to Gaussianity is more reliable and there is less complication arising from periods of no rainfall.

Jothityangkoon et al. [2000] developed a cascade space-time rainfall model for calibration to daily rainfall gauges. The model uses a first-order, four-state Markov chain to generate a daily time-series of the regionally-averaged rainfall. Using a cascade model, the regional rainfall totals are then spatially disaggregated down to a scale of 12.5 km giving daily rainfall totals across all points in the region. While there is a temporal correlation in the rainfall total, the model does not incorporate any pixel-to-pixel temporal correlation of rainfall. The model was applied to a 400×400 km^2 region in the southwest of Western Australia, using 490 daily rainfall gauges over a period of 11 years. The authors note that the nonuniform density of rain-gauges introduced biases in the parameter estimates. Nonetheless, the model preserved the observed spatial patterns of long term mean rainfall at daily, monthly and annual aggregates across the region. However, some statistics such as the mean number of wet days and the mean wet spell length were underpredicted, which could have a significant impact on estimated flood magnitudes when combined with a hydrologic model. The authors suggest that this may be due to a lack of space-time correlations in the model. While the model works well at a daily time-scale, it is not applicable for generating rainfall fields at shorter time scales, since in that case,

storm movement and space-time correlations will need to be explicitly incorporated. The model can also be used to disaggregate output from global circulation models.

Willems [2001] developed a space-time rainfall model for application in small-scale urban catchments at a sub-daily timescale. The model has a hierarchical structure which has:

1. meso-scale storm onset including storm-lifetime and advection (the storm area is assumed to be larger than the observing network),
2. clusters of rain cells that arrive within the storm as a Poisson process, and,
3. detailed modelling of individual cells having bivariate Gaussian shape and linked to a diffusion equation to model cell growth and decay.

The authors developed heuristic algorithms to separate the rain-cells within the observed records at each gauge and then link identified cells across gauges in order to track their motion. Having separated the individual rain-cells, their lifetimes, areal extent and velocity are easily determined. The model was calibrated for the city of Antwerp, Belgium, having a dense network of 12, $1-$ minute gauges located within an area of $100\ km^2$. Over the 3 year period, 807 storms were identified, enclosing a total of 5940 rain cells. While the model differs from previous approaches in its detailed description of rain-cell evolution, its main drawback is that it requires the presence of a dense and high-resolution network of gauges. In addition to this, the heuristic algorithms have not been tested for a wide variety of rainfall regimes and may not be appropriate in other locations.

Cowpertwait [1995] developed a spatial generalisation of the NSRP model that has been applied to multiple case studies for catchments in England, Italy, and New Zealand [*Cowpertwait*, 1995; *Cowpertwait et al.*, 2002; *Cowpertwait*, 2004]. The model includes all of the developments of the temporal NSRP and therefore a wide range of properties and features are available for calibration, including:

- first and second order moments in time and space [*Cowpertwait et al.*, 2002],
- third order moments at a point [*Cowpertwait*, 1998],
- dry probabilities [*Cowpertwait*, 1995], and
- results for a mixture of multiple storm types [*Cowpertwait*, 2004].

The model is stationary in space over the simulation region, therefore scaling parameters are used to standardise the observed rainfall under the assumption of a constant

coefficient of variation (i.e. divide by the mean). This has the added benefit that the model will preserve the mean rainfall at each of the gauges used in calibration. However, some statistics do not scale with respect to the mean, for example the probability of a dry period cannot be altered by this scaling procedure.

2.5 SUMMARY OF LITERATURE

There is a long history of developing continuous simulation rainfall models for sub-daily simulation, but most of these methods have concentrated on rainfall at a single point in space. Regarding sub-daily space-time models there are two different categories depending on whether the underlying data is based on radar imagery or on networks of rain gauges. Of the radar models, only the String of Beads model has currently been used for continuous simulation [*Smithers et al.*, 2007] and all other models focus on a single storm. There are many reasons for the focus of radar-based models on individual storms including (i) the inherent complexity of the data source (ii) the computational burden and (iii) use in other applications such as now-casting and (iv) difficulty in comparing radar reflectivity estimates to rain gauge observations for bias correction. Also, since radar imagery only has a short observation record it is not easy to use it for simulating extreme rainfall events (e.g. 100 year ARI).

Gauge-based rainfall models provide an alternative means for generating replicates of sub-daily space-time rainfall. Their main benefit over radar-based models are that gauges are more abundant and have a longer history so that extreme rainfall events can be estimated more reliably. Furthermore, as they have less data quality issues when compared to radar rainfall estimates they are easier to use. In contrast to radar-based models, most of the gauge-based space-time rainfall models have been explicitly developed for the purpose of continuous simulation. The main drawback of gauge-based space-time models is that they do not provide means for combining daily rainfall and sub-daily rainfall estimates into the calibration procedure, thus their applicability is severely limited for practical reasons. Of these models, the SNSRP model stands out because it can incorporate data into the calibration across a range of time periods. Despite the advantages of the SNSRP model over other gauge-based models, space-time rainfall modelling is a relatively young research area and the SNSRP model has numerous limitations. Two of the main limitations addressed in this thesis result from the fact that the model has previously been applied to smaller regions where homogeneity is a safer assumption and where storms can be assumed to cover the entire region. To address these limitations requires the SNSRP model to be reformulated. Numerous other limitations of the model are ad-

dressed in this thesis, but require changes to the calibration and simulation methodology rather than the model itself. These changes explore issues of bias in the calibration, dependency on the derivation of analytic equations, efficiency of simulation and ability to handle climate variability.

Chapter 3
Model Formulation and Calibration

This chapter provides background on how the SNSRP model is implemented: model assumptions, model limitations and the constraints posed by calibration.

3.1 MODEL FORMULATION

Figure 3.1 shows a schematic of the NSRP model for an individual site. For the i^{th} storm, the arrival time of storm origins, T_i, occurs as a Poisson process having rate parameter, λ, per hour. The duration between storms is an exponential random variable having mean λ^{-1}. Each storm has a number of cells, C, associated with it following a Poisson distribution having mean μ_C. For the i^{th} storm, the j^{th} cell has properties: lifetime, L_{ij}, intensity, X_{ij} , and start time relative to the storm origin, $S_{ij} - T_i$. The cell lifetimes and start times relative to the origin both follow exponential distributions with parameters η and β respectively. Accordingly the mean cell lifetime is η^{-1} and the mean time until it starts is β^{-1}. The intensity, X_{ij}, remains constant throughout the lifetime of a cell and is taken as an independent Weibull random variable with survivor function, $Pr(X > x) = e^{-(x/\theta)^{\alpha}}$. Setting $\alpha = 1$ gives the exponential distribution as a special case. Moments of this distribution are given as $E[X^r] = \theta\Gamma(1 + r/\alpha)$ so that the mean cell intensity, μ_X, is given by $\mu_X = \theta\Gamma(1 + 1/\alpha)$. The temporal process of the Neyman-Scott model therefore has 6 parameters: $\lambda, \beta, \eta, \alpha, \theta, \mu_C$.

The instantaneous amount of rainfall at time t, is the sum over all cell intensities that are alive at that time,

$$Y(t) = \int_{s=0}^{\infty} X(t-s)dN(t-s)\, d\, s, \tag{3.1}$$

where $dN(t-s)$ gives the number of cells starting at $t-s$ that overlap time t. As rainfall data are sampled over discrete time intervals, the instantaneous process is integrated for some level of aggregation, h (in hours), such that the l^{th} interval has a rainfall depth

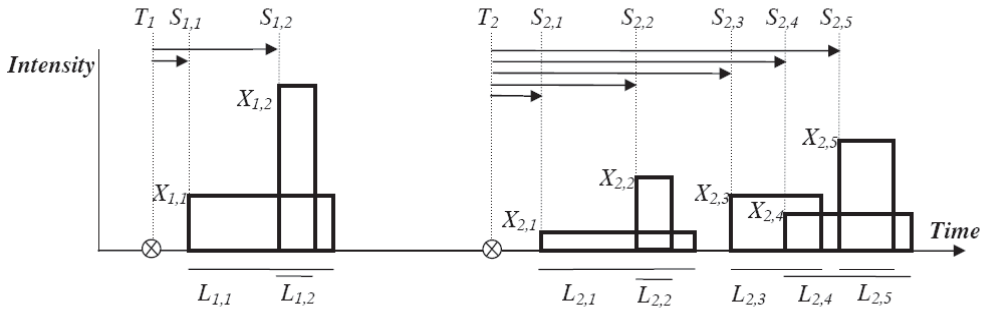

Figure 3.1 Schematic diagram of the temporal Neyman-Scott Rectangular Pulse model. X represents cell intensities, L represents cell lifetimes, S represents cell start times and T represents storm start times.

denoted as

$$Y_l^{(h)} = \int_{(l-1)h}^{lh} Y(t)\, dt. \tag{3.2}$$

Having defined the aggregated rainfall process corresponding to the model, various statistics are summarised, including first, second and third order moments and the dry-portion. The mean rainfall is the product of the storm rate, the mean number of cells, the mean intensity and lifetime of each cell and the level of aggregation,

$$\begin{aligned} \mu_h &= E[Y_l^{(h)}] \\ &= \lambda \mu_C \mu_X \tfrac{1}{\eta} h. \end{aligned} \tag{3.3}$$

From *Cowpertwait et al.* [2002], the auto-covariance of two intervals at at a given lag, τ, is given as,

$$\begin{aligned} \gamma_{0,h,\tau} &= Cov[Y_l^{(h)}, Y_{l+\tau}^{(h)}] \\ &= 2\lambda\mu_c E[X^2] A(h,\tau)/\eta^3 + \lambda(\mu_x \beta \mu_c)^2 A(h,\tau)/(\eta^3(\beta^2-\eta^2)) \\ &\quad -\lambda(\mu_x\mu_c)^2 B(h,\tau)/(\beta(\beta^2-\eta^2)), \end{aligned} \tag{3.4}$$

$where,$

$$\begin{aligned} A(h,0) &= (h\eta + e^{-\eta h} - 1) & A(h,\tau) &= \tfrac{1}{2}(1-e^{-\eta h})^2 e^{-\eta h(\tau-1)} \\ B(h,0) &= (h\beta + e^{-\beta h} - 1) & B(h,\tau) &= \tfrac{1}{2}(1-e^{-\beta h})^2 e^{-\beta h(\tau-1)}. \end{aligned}$$

When $\tau = 0$, Eq. 3.4 gives the expression of the variance. The third order moment was derived by *Cowpertwait* [1998] and is given by,

$$\zeta_h = E[Y_l^{(h)} - \mu_h]^3 = \frac{6\lambda\mu_C E[X^3]\left(\eta h - 2 + \eta h e^{-\eta h} + 2e^{-\eta h}\right)}{\eta^4} + \frac{3\lambda\mu_X E[X^2]\mu_C^2 f(\eta,\beta,h)}{2\eta^4\beta(\beta^2-\eta^2)^2} + \frac{\lambda\mu_X^3\mu_C^3 g(\eta,\beta,h)}{2\eta^4\beta(\eta^2-\beta^2)(\eta-\beta)(2\beta+\eta)(\beta+2\eta)}, \quad (3.5)$$

where the functions $f(\eta,\beta,h)$ and $g(\eta,\beta,h)$ are given by:

$$\begin{aligned} f(\eta,\beta,h) = \; & -2\eta^3\beta^2 e^{-\eta h} - 2\eta^3\beta^2 e^{-\beta h} + \eta^2\beta^3 e^{-2\eta h} + 2\eta^4\beta e^{-\eta h} + 2\eta^4\beta e^{-\beta h} \\ & +2\eta^3\beta^2 e^{-(\eta+\beta)h} - 2\eta^4\beta e^{-(\eta+\beta)h} - 8\eta^3\beta^3 h + 11\eta^2\beta^3 - 2\eta^4\beta \\ & +2\eta^3\beta^2 + 4\eta\beta^5 h + 4\eta^5\beta h - 7\beta^5 - 4\eta^5 + 8\beta^5 e^{-\eta h} - \beta^5 e^{-2\eta h} \\ & -2h\eta^3\beta^3 e^{-\eta h} - 12\eta^2\beta^3 e^{-\eta h} + 2h\eta\beta^5 e^{-\eta h} + 4\eta^5 e^{-\beta h} \end{aligned} \quad (3.6)$$

$$\begin{aligned} g(\eta,\beta,h) = \; & 12\eta^5\beta e^{-\beta h} + 9\eta^4\beta^2 + 12\eta\beta^5 e^{-\eta h} + 9\eta^2\beta^4 + 12\eta^3\beta^3 e^{-(\eta+\beta)h} \\ & -\eta^2\beta^4 e^{-2\eta h} - 12\eta^3\beta^3 e^{-\beta h} - 9\eta^5\beta - 9\eta\beta^5 - 3\eta\beta^5 e^{-2\eta h} - \eta^4\beta^2 e^{-2\beta h} \\ & -12\eta^3\beta^3 e^{-\eta h} + 6\eta^5\beta^2 h - 10\eta^3\beta^4 h + 6\eta^2\beta^5 h - 10\eta^4\beta^3 h + 4\eta\beta^6 h \\ & -8\eta^4\beta^2 e^{-\beta h} + 4\eta^6\beta h + 12\eta^3\beta^3 - 8\eta^2\beta^4 e^{-\eta h} - 6\eta^6 - 6\beta^6 \\ & -2\eta^6 e^{-2\beta h} - 2\beta^6 e^{-2\eta h} + 8\eta^6 e^{-\beta h} + 8\beta^6 e^{-\eta h} - 3\eta^5\beta e^{-2\beta h}. \end{aligned} \quad (3.7)$$

The probability that an interval is dry (which is the same statistic as the portion of dry intervals) was derived by *Cowpertwait* [1995] and is given below. Eq. 3.8 can be evaluated using routine numerical integration .

$$Pr\{Y_l^{(h)} = 0\} = exp\left[-\lambda\int_0^\infty (1 - p_h(t))dt - \lambda\int_0^h (1 - p_h(0))dt\right], \quad (3.8)$$

where $p_h(t)$ is the probability of no rain in the interval $(t, t+h)$ given a storm origin at time zero,

$$p_h(t) = exp\left[\mu_c e^{-\beta(t+h)} - \mu_C e^{-\beta t} - \frac{\mu_C\beta(e^{-\beta t} - e^{-\eta t})}{\eta - \beta}\right]. \quad (3.9)$$

To extend the model in the spatial domain, *Cowpertwait* [1995] assumed that each cell covers a circular area having a random variable, the cell radius, R_c. The cell radius follows an exponential distribution with parameter ϕ_C having mean value ϕ_C^{-1}. The cell intensity (cell height) is constant over the entire area, such that raincells can be thought of as cylinders. As for the temporal model, the total rainfall at a point in space is the sum of all raincells overlapping that point. The spatial position of cell-centres, Z_c, is distributed

according to a two-dimensional Poisson process with rate parameter φ_c (per km^2) which is related to the parameter μ_C as

$$\mu_C = \frac{2\pi\varphi_c}{\phi_c^2}. \tag{3.10}$$

This formulation is what has been referred to in earlier chapters as the SNSRP model and it offers the simplest extension of the temporal NSRP model into the spatial domain, requiring only one additional parameter, ϕ_c. Figure 3.2 depicts the spatial process of cell centres and radii at some instant in time and over some region of interest, the target region. By definition, the model is said to be homogeneous over this region as all statistics at any point within the region will be the same. In particular, the first and second order moments are constant over the region, since the model is stationary. Rainfall data are however non-stationary as there are often significant trends in the mean. To remedy this discrepancy, the observed data are standardised at each site, the procedure for which is discussed further in Section 3.2.

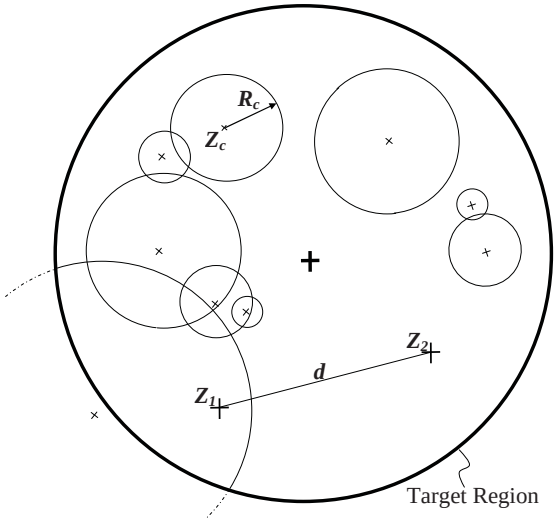

Figure 3.2 Schematic diagram of spatial cell locations and radii showing overlapping rain-cells in time and space.

As shown in Figure 3.2, consider two generic points within the region, Z_1 and Z_2 separated by a distance d. The cross covariance between two points at this distance is given as

$$\gamma_{d,h,\tau} = \gamma_{0,h,\tau} - 2\lambda\mu_c\{1 - P_{Z_1|Z_2}(\phi_c, d)\}E[X^2]A(h,\tau)/\eta^3, \tag{3.11}$$

where the aggregated auto-covariance $\gamma_{0,h,\tau}$ is given in Eq. 3.4 and $P_{Z_1|Z_2}$ is the

probability of a cell overlapping point Z_2 given it has overlapped Z_1. *Cowpertwait* [1995] derives this expression as,

$$P_{Z_1|Z_2}(\phi, d) = \frac{2}{\pi} \int_{y=0}^{\pi/2} \left(\frac{\phi d}{2 \cos y} + 1 \right) \exp\left(\frac{-\phi d}{2 \cos y} \right) dy, \tag{3.12}$$

which can be evaluated efficiently using numerical algorithms. The list of properties derived for the model are summarised below where the dependent parameters are given in parentheses. The parameter θ is not required as a dependent parameter as the standardisation procedure requires all sites to be simulated with the same mean, μ_h. Arbitrarily setting $\mu_h = 1$, an expression for θ can be obtained in terms of all other parameters. Therefore, the SNSRP model requires only 6 parameters to model the pattern of rainfall occurrence and 1 scale parameter (θ_i) at each site, $i = 1 \ldots M$, to preserve the mean at that site. Seasonal variation is accounted for by using independent sets of parameters for each month, thus there is a total of 12(6+M) parameters for the SNSRP model.

$$\begin{array}{lll}
\textit{Variance} & \sigma_h^2(\lambda, \beta, \eta, \mu_C, \alpha) & = \gamma_{0,h,0} \\
\textit{Coeff. of Variation} & \nu_h(\lambda, \beta, \eta, \mu_C, \alpha) & = \sigma_h / \mu_h \\
\textit{Lag-1 Auto-correlation} & \rho_h(\beta, \eta, \mu_C, \alpha) & = \gamma_{0,h,1} / \sigma_h^2 \\
\textit{Skewness} & \kappa_h(\lambda, \beta, \eta, \mu_C, \alpha) & = \zeta_h / \sigma_h^3 \\
\textit{Dry Portion} & \psi_h(\lambda, \beta, \eta, \mu_C) & = Pr\{Y_l^{(h)} = 0\} \\
\textit{Lag-0 Cross-correlation} & \rho_{i1,i2,h}(\beta, \eta, \mu_C, \alpha, \phi_c) & = \gamma_{d,h,0} / \sigma_h^2
\end{array} \tag{3.13}$$

Following *Cowpertwait* [2004], a generalisation of the model is obtained from the superposition of n independent $SNSRP^{(p)}$ processes, $p = 1 \ldots n$. In this way, mixtures of storm types, each having different parameters, can be modelled for a given month. When multiple storm types are used, the total rainfall at a given location for some interval l is the sum over each $SNSRP^{(p)}$ process that overlaps that point,

$$Y_l^{(h)} = \sum_{p=1}^{n} Y_l^{(p,h)}. \tag{3.14}$$

Following from this, the properties for the resultant moments are given as,

$$\mu_h = \sum_{p=1}^{n} \mu_h^{(p)} \tag{3.15}$$

$$\gamma_{d,h,\tau} = \sum_{p=1}^{n} \gamma_{d,h,\tau}^{(p)} \tag{3.16}$$

$$\zeta_h = \sum_{p=1}^{n} \zeta_h^{(p)}. \tag{3.17}$$

The auto-correlation and variance can be obtained from Eq. 3.16 as $\gamma_{0,h,\tau}$ and $\gamma_{0,h,0}$ respectively. The skewness can be obtained by standardising Eq. 3.17 using the overall standard deviation, as per Eq. 3.13. The resultant dry probability from a mixture of $SNSRP^{(p)}$ processes is given as,

$$\psi_h = \prod_{p=1}^{n} \psi_h^{(p)}. \tag{3.18}$$

In the event of multiple storm types, the expression for the parameter, $\theta_{h,i,k}$ (at site i and in month k), in terms of other parameters is given as,

$$\theta_{h,i,k} = \hat{\mu_{h,i,k}} \left\{ \sum_{p=1}^{n} \frac{\lambda^{(p)} \mu_C^{(p)} \Gamma(1 + 1/\alpha^{(p)})}{\eta^{(p)}} \right\}, \tag{3.19}$$

where $\hat{\mu_{h,i,k}}$ is the observed mean hourly rainfall at site i during month k. The generalised rainfall model requires an additional 6 parameters for each additional storm type, however it is possible to make some parameters common to both types of storms and reduce the total number of parameters. The need for multiple storm types and the selection of which parameters to make common is discussed with respect to the calibration procedure.

3.2 CALIBRATION STATISTICS

Consider $i = 1 \ldots M$ sites of observed rainfall at some aggregate, h, and having respective years of record, $j = 1 \ldots N_i$. For each month $k = 1 \ldots 12$, let the observed rainfall depth in each interval of the month $l = 1 \ldots 31(24/h)$, be denoted $x \equiv x_{h,i,j,k,l}$. Due to the presence of corrupt and missing observations in most rainfall records, an indicator function is introduced to mask unwanted values in the calculation of statistics. In addition to observations that are either missing or corrupt, for notational convenience all months are denoted 31 days but months of duration 28, 29 and 30 days are assumed to have a period of 'missing' observations at the end of the month. The indicator function is therefore defined as,

$$I \equiv I_{h,i,j,k,l} = \begin{cases} 1 & if\ valid\ data\ point \\ 0 & if\ corrupt/missing\ data. \end{cases} \tag{3.20}$$

For notational convenience it is useful to respectively define the total number of valid observations in a given month at an individual site, $n_{h,i,k}$, and the total number of valid observations across across all sites, $n_{h,k}$. These are given as,

$$n_{h,i,k} = \sum_{j=1}^{N_i} \sum_{l=1}^{31(24/h)} I \qquad \text{and} \qquad n_{h,k} = \sum_{i=1}^{M} n_{h,i,k}. \tag{3.21}$$

The model outlined in Section 3.1 is stationary in both time and space, however observed rainfall data is neither of these. To capture inherent seasonal variation, statistics are calculated on a monthly basis to allow separate parameter estimates for each month. Also, several statistics are calculated at multiple aggregates so that the variability of the rainfall process is represented across multiple time-scales (e.g. daily and sub-daily). The between-site variation is modelled by non-dimensionalising the data at each site using the mean rainfall. Thus the sample mean is first computed for each site i and month k as,

$$\hat{\mu}_{h,i,k} = \frac{1}{n_{h,i,k}} \sum_{j=1}^{N_i} \sum_{l=1}^{31(24/h)} I \cdot x. \tag{3.22}$$

All subsequent statistics are calculated by dividing the observed rainfall estimates by the estimate of the mean one-hour rainfall depth, $\hat{\mu}_{1,i,k}$. Thus, the model parameters, once estimated, will yield the property $\mu_1 = 1$ and the non-stationarity in the mean can be recovered from a simulation by a multiplication with the observed value of $\hat{\mu}_{1,i,k}$ at each site. While several statistics, such as correlations and the dry probability, are not dependent on the scale of the data, this process assumes that higher-order moments such as the variance and skewness are observed to scale with the mean. Because the data are non-dimensionalised, all subsequent temporal statistics are pooled across the region to yield one estimate of each statistic for the entire region rather than individual estimates at each site. Further discussion of the homogeneity of a given region and the appropriateness of using single estimates for an entire region is given in Chapter 9.

For the purpose of non-dimensionalising the data, the pooled mean of the region is given below as the sum of means at each site weighted by the number of observations,

$$\hat{\mu}_{h,k} = \frac{1}{n_{h,k}} \sum_{i=1}^{M} \hat{\mu}_{h,i,k} n_{h,i,k}. \tag{3.23}$$

The standardised moments pooled across all sites in the region and across all years of observation are given below in Eq. 3.24 and Eq. 3.25 for the variance and skewness

respectively.

$$\hat{\sigma}^2_{h,k} = \frac{\sum_{i=1}^{M}\sum_{j=1}^{N_i}\sum_{l=1}^{31(24/h)} I \cdot x^2}{\hat{\mu}^2_{1,k} n_{h,k}} - h^2 \tag{3.24}$$

$$\hat{\kappa}_{h,k} = \frac{\sum_{i=1}^{M}\sum_{j=1}^{N_i}\sum_{l=1}^{31(24/h)} (I \cdot x/\hat{\mu}^3_{1,k} - h)^3}{\hat{\sigma}^3_{h,k} n_{h,k}} \tag{3.25}$$

Given the standard deviation, the coefficient of variation is calculated as,

$$\hat{\nu}_{h,k} = \hat{\sigma}_{h,k}/h. \tag{3.26}$$

The cross-correlation between two sites, i_1 and i_2 requires concurrent years of overlapping data, denoted $N_{i1,i2}$. While two gauges may have lengthy records, these may have been taken during differing years and decades such that a cross-correlation between the gauges cannot be computed. The mutual length of the records is taken as $N_{i1,i2} = min(end_1, end_2) - max(start_1, start_2)$ and the indicator function for each record can be used to account for any periods where only one gauge has observed data. Whereas there is only one estimate over the entire region for the temporal statistics, the cross-correlation is a spatial statistic and it changes as a function of the distance. There are multiple estimates of the cross-correlation owing to the number of pairs of gauges. No attempt is made to average the correlation estimates with respect to distance, as each pair of gauges will give differing distances. While individual gauges may have long records, it is possible that some pairs of sites have short periods of overlap which leads to highly variable estimates. The Fisher transformation can be used to determine the uncertainty in a correlation estimate given the number of data used and a threshold can be applied if the estimate is deemed too variable (> 250 data points was considered acceptable using this criterion). The statistic to estimate the cross-correlation is given below in Eq. 3.27. The numerator estimates the covariance, while the denominator standardises with respect to the standard deviation at each site. The standard deviations are estimated from the overlapping portion only. Even though each record may have additional data that could improve the estimate of the standard deviations, as this data is from differing periods the

role of climatic variability could lead to biases in the estimate.

$$\hat{\rho}_{i1,i2,h,k} = \frac{\sum_{j=1}^{N_{i1,i2}} \sum_{l=1}^{31(24/h)} I \cdot (x_{h,i1,j,k,l} - \hat{\mu}_{h,i1,k}) I \cdot (x_{h,i2,j,k,l} - \hat{\mu}_{h,i2,k})}{\sqrt{\sum_{j=1}^{N_{i1,i2}} \sum_{l=1}^{31(24/h)} I \cdot (x_{h,i1,j,k,l} - \hat{\mu}_{h,i1,k})^2 \sum_{j=1}^{N_{i1,i2}} \sum_{l=1}^{31(24/h)} I \cdot (x_{h,i2,j,k,l} - \hat{\mu}_{h,i2,k})^2}} \tag{3.27}$$

The auto-correlation of an individual site can be computed as a special case of Eq. 3.27 by letting the second site be a replicate of the first but at some lagged distance in time. Since the auto-correlation is at distance zero it is pooled across all sites to give one estimate for the region as a whole.

To calculate the dry portion of a record (or dry probability), an indicator function is introduced based on whether the observation x in an given interval is non-zero,

$$D \equiv D_{h,i,j,k,l} = \begin{cases} 1 & if\ x = 0 \\ 0 & if\ x > 0. \end{cases} \tag{3.28}$$

The dry days are the portion of the dry intervals in the total number of valid observations, given by the product of the indicator and the number of valid cells:

$$\hat{\psi}_{h,k} = \frac{1}{n_{h,k}} \sum_{i=1}^{M} \sum_{j=1}^{N_i} \sum_{l=1}^{31(24/h)} I \cdot D. \tag{3.29}$$

Cowpertwait [1998] notes that the dry portion is a biased statistic because the observation record has some minimum threshold below which small rainfall volumes are indicated as zero. For daily rainfall records from the Australian Bureau of Meteorology this is $0.1\ mm$ and for pluviograph records, $0.01\ mm$. Additional networks, from other agencies (e.g. flash-flooding alert systems or radar calibration systems) are likely to have different thresholds. As it is not possible to derive the portion of rainfall above a threshold that is other than zero, this bias cannot be directly accounted for when fitting parameters. *Cowpertwait* [1998] recommended omitting this statistic from calibration and checking via simulation whether a parameter set adequately matched this statistic. The recommendation was not followed in this thesis and the statistic is used in calibration despite its known bias. The reason for this is based on preliminary observations that large discrepancies were observed for this statistic for some months ($\sim$ 10% to 20%) and especially in the event of higher dry portions. By way of contrast the bias for a threshold less than $0.1\ mm$ reported for two sets of example parameters in *Cowpertwait* [1998, Table 5] is

less than 5%. This suggests that the benefit (in terms of ease of calibration) of including the statistic for the dry portion can outweigh the negative effect of its bias.

The calibration proceeds by first selecting suitable daily and sub-daily gauges within a defined region covering the catchment of interest. Not all of the gauges need to be selected, as short gauges should be omitted in the event that they yield highly variable statistical estimates, in particular, the estimate of the mean that is used in the non-dimensionalising. The non-dimensional statistics are calculated at several aggregates, such as $1\ hr$, $6\ hr$ and $24\ hr$. The selection of the aggregate levels is arbitrary but should adequately cover the dominant time-scales at which the rain processes operate (e.g. cells$\sim$ sub-daily and storms $\sim$ daily). $1\ hr$, $6\ hr$ and $24\ hr$ were chosen in order to be consistent with prior studies. Pluviograph data is used to obtain sub-daily aggregates, however its use at a daily aggregate is limited because the daily gauges provide a more dense and longer record of observations. Also, the presence of significant numbers of missing observation in pluviograph records means that they are often unreliable at a daily aggregate.

The model properties are calibrated to the estimated statistics using a comparison termed the 'method of moments'. The method of moments is a least-squares technique between observed and estimated statistics. The calibration proceeds by minimising an objective function defined in terms of the squared differences between observed statistics (Eq. 3.24, Eq. 3.26, Eq. 3.25 and Eq. 3.29) and modelled properties (Eq. 3.13). The form of this function is given as,

$$F_1 = w_\psi(\psi_{24} - \hat{\psi}_{24,k})^2 + \sum_{h=1,6,24} w_\nu(\nu_h - \hat{\nu}_{h,k})^2 + w_\kappa(\kappa_h - \hat{\kappa}_{h,k})^2 + w_\rho(\rho_h - \hat{\rho}_{h,k})^2, \quad (3.30)$$

where w_ψ, w_ν, w_κ and w_ν are weights. Values of the weights were set to be inversely proportional to the typical ranges of each statistic: skewness - (0,50); coefficient of variation - (0,10); correlation - (0,1) and dry portion - (0,1). *Cowpertwait et al.* [2002] used non-dimensional terms instead of applying weights, however as it is possible for correlations to have a value close to zero this method was found to destabilise the objective function and was thus avoided. Another factor is that the skewness statistic is much more variable than the others, so it was less desirable to ensure that this statistic was fitted as accurately as the others. The statistic for dry proportion was included at the daily aggregate only, as this was sufficient to ensure that the calibrated dry portions are realistic. Therefore, 10 different statistics per month are used in calibrating the temporal properties of the model. The calibrated parameters corresponding to the temporal properties of the model are, $\lambda, \beta, \alpha, \mu_C, \eta$. They are given the broad constraints $\lambda, \beta, \alpha, \eta > 10^{-5}$ and $\mu_C > 1$.

To calibrate the spatial attribute of the model there is only one parameter ϕ_c which is calibrated to the cross-correlogram using a least squares objective function, which is

$$F_2 = \sum_{h=1,24} \sum_{i1=1}^{M} \sum_{i2=1}^{M} (\rho_{d,h,0} - \hat{\rho}_{i1,i2,h,k})^2. \tag{3.31}$$

A weight is not required for this function because the model is calibrated following a step-by-step procedure that does not compare the fit of the cross-correlations relative to the temporal goodness of fit. Instead, the parameters obtained in the temporal calibration are assumed to be fixed in the spatial calibration. Whereas previous calibrations have used only the cross-correlations at an hourly aggregate, Chapter 7 highlights the need to include the $24\ hr$ cross-correlations into the calibration. Cross-correlations between two sites and lagged in time are not considered because they quickly decay to zero and because they are adequately captured by the lag-one auto-correlation function and the lag-zero cross-correlation function.

By comparing the quality of the calibrated model to the observed statistics one can assess whether a good fit has been achieved. If a good fit is not possible, then Eq. 3.15 to Eq. 3.18 can be used to determine the model properties from two (or more) superposed SNSRP processes. This effectively doubles the parameters required for calibration, however it is beneficial to make several of these common to both storms. A range of subsets of parameters should be tested to determine the relative reduction in the sum-of-squares to indicate whether or not a parameter is made common to both storms. Preliminary investigations have determined that β and η were the most effective parameters to be made common across both storms, and to a lesser extent α and μ_C. Whether or not ϕ_c is made common depends on the outcome of the calibration of parameters to the temporal properties. An overall reduction in the sum of squares of $> 50\%$ was considered necessary before the benefit of using two processes was considered a viable alternative. This criterion is arbitrary, hence the improvement should also be assessed using visual comparisons of the plotted statistics. While it is possible to use three or more storms, at no stage has this been investigated due to the substantial increase in parameters.

Having determined the parameters for the stationary model, the scaling parameter θ can be calculated using Eq. 3.19 for each observation site. This gives the parameter set necessary for simulating the model as a multi-site model. While the cells are able to distribute rain to locations other than observed sites, it is not possible to determine a scaling parameter at these sites to give the corrected rainfall totals. However, it is possible to interpolate the scaling parameters over the domain of the region using standard spatial techniques such as kriging, thin-plate splines and regression. These techniques make al-

lowances for explanatory variables such as distance from the coast and elevation which can be used to reduce the variance of the interpolated scaling field. The number of additional parameters depends on the form of the equation, but a separate field should be used for each month to capture the seasonal non-stationarity of the scaling field. While the inclusion of additional explanatory variables is important, detailed investigation was considered beyond the scope of this thesis and the primary aim to address specific limitations of the model.

3.3 CALIBRATION SUMMARY

The Neyman-Scott model at a single point can be obtained as a special case from the *Cowpertwait et al.* [2002] model. This enables a step-wise calibration first calibrating the temporal parameters in isolation and then calibrating the cell radius parameter to hourly cross-correlations. Also, a step-wise iterative procedure allows for a comparison of the fit obtained when introducing additional parameters representing different types of storms. A further reason for a step-wise approach is that the model parameters are correlated. This can cause a joint optimization procedure to pursue a minimal increase in the goodness of fit at the expense of realistic interpretation or feasible implementation of the parameters (e.g. excessive numbers of short duration cells).

A step-wise procedure for calibration of the model is given below.

1. At each site, $i = 1, ..., M$, the rainfall data are non-dimensionalised by dividing by the estimated hourly mean, $\hat{\mu}_{1,i,k}$ (the choice of aggregate is arbitrary). Temporal statistics are calculated from pooled data across the region, including the pooled coefficient of variation, auto-correlation and skewness. The sole spatial statistic used in calibration is the cross-correlation between all gauges in the region. The statistics are calculated at multiple aggregates to ensure that the calibrated parameters capture the rainfall process at a broad range of timescales. The number of aggregation levels is arbitrary but should at least cover sub-daily and daily timescales.

2. A least squares objective function is defined between the observed statistics and the model properties for all temporal statistics and aggregates.

3. For each site, $i = 1, ..., M$, the scale parameter of the Weibull distribution is allowed to vary so that the mean of the process at each site is reproduced exactly. This scale parameter, $\theta_{1,i,k}$, can be expressed in terms of the other model parameters. Assuming the generalised model of *Cowpertwait* [2004] there are n storm types

so that the superposition of each storm type results in a mean equal to one (for the non-dimensionalised data). For the generalization of n storm types there are 5 parameters per storm type to be estimated. These parameters are constant across the region. Denoting the p^{th} storm type with a superscript, the parameters to be estimated are ($\lambda^{(p)}$, $\beta^{(p)}$, $\eta^{(p)}$, $\mu_c^{(p)}$, $\alpha^{(p)}$).

4. The model is first calibrated using one storm type. The improvement of the calibration (in a least-squares sense) by using multiple storm types is considered only in the event of a poor fit using one storm type. This criteria is arbitrary owing to the arbitrary number of statistics used to calibrate the model. A reduction of at least 50% in the sum of squares was used in this study before two storm types were considered preferable to one. If multiple storm types are used, it is likely that several parameters will be able to held fixed across the storm types without significantly affecting the quality of the calibration. The selection of how many and which parameters to fix is tested using several calibration attempts, by progressively increasing the number of parameters common to both storm types.

5. Having obtained a good fit to the temporal statistics, the calibration of spatial parameters is considered. The cell radius parameters $\phi_c^{(p)}$ are calibrated via least squares to the hourly cross-correlogram assuming storms cover the entire region.

6. To simulate the model continuously in space the scaling value θ is interpolated for all points within the region to model the non-stationarity in the mean.

Chapter 4

Simulation Involving a Regional Boundary

4.1 INTRODUCTION

A simulation of the SNSRP model proceeds by sampling, from calibrated distributions, the arrival of storms, number of rain cells and cell properties over some target region. However, with this formulation it is possible to have rain cells with centers lying outside of a target region, yet having radii large enough to cover points within the region. This introduces a boundary effect that significantly reduces the simulated rainfall depth at points within the region if these cells are ignored. It is essential to account for this boundary effect because the size of a rain cell is often of a similar magnitude to the size of the target region, meaning that all points within the region (and not just those near the outer perimeter) will be affected. One approach for avoiding this is to wrap the effects of rain cells across opposing points on the boundary of the region. This approach is efficient and effective with the exception that spurious cross correlations will be observed for points separated by large distances. An alternative approach is to mitigate the boundary effect by implementing a buffer around the target region. This approach is demonstrated to substantially inhibit the computational efficiency of the model, hence an algorithm is proposed that avoids the use of a buffer region. This algorithm directly simulates the number of cells that occur outside of the target region yet are known to intersect it. Whereas the buffer algorithm is approximate due to the finite size of the buffer, the direct algorithm is exact with respect to reproducing rainfall depths within the region. The method is not specific to the Neyman-Scott model as it could be applied to similar cluster-based models having the same spatial structure.

4.2 Algorithm for Simulation of Cells Using a Buffer Region

Consider a target region, of fixed radius, r_K, as in Figure 4.1. For any given point within this region it is desired to model the rainfall process using the SNSRP model. According to this model, cells associated with a particular storm arrive within the region as a Poisson process having a spatial rate $\varphi_c \pi r_K^2$, where φ_c is the rate parameter.

It is however also possible to have cells arrive outside of the target region, referred to as an outer region, yet having a cell radius, R_c, greater than the distance from the circumference of the target region to the center of the cell, R_{xy}. Consider therefore a buffer, having radius, $r_K + r$, where r is defined relative to the target region. It is possible to also simulate cells over this region according to a Poisson process with the same spatial rate parameter, φ_c. Cells that intersect the target region are accepted, otherwise cells that lie solely outside are rejected.

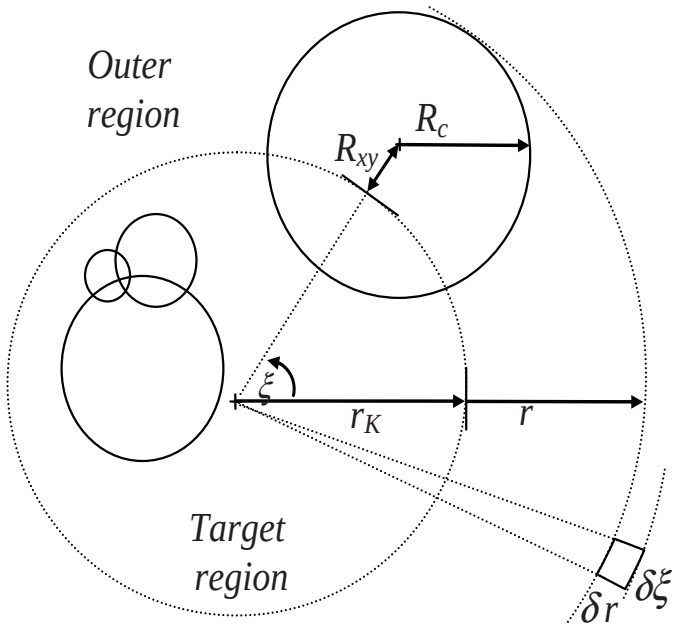

Figure 4.1 Schematic diagram of cells generated inside and outside the target region.

The arrival of cells according to a Poisson distribution requires their cells to be uniformly distributed over the region. Efficiently simulating uniformly over a circular region can be achieved using a transformation to polar coordinates. The equations for simulating

a point (x, y) uniformly within a circular region of radius $r_K + r$ are,

$$\begin{aligned} x &= (r_K + r)\sqrt{U}\cos\Xi \\ y &= (r_K + r)\sqrt{U}\sin\Xi, \end{aligned} \tag{4.1}$$

where $\Xi \sim Uniform[0, 2\pi]$ and $U \sim Uniform[0, 1]$.

To illustrate the impact of a buffer region, the Arno Basin case-study given by *Cowpertwait et al.* [2002] is used as an example. A circular target region of radius r_K = 65 km encompasses all of the rain gauges within this region. Cell radii are distributed exponentially with parameter ϕ_c. For the two example months of January and July reported values of this parameter for the Arno Basin are $\phi_c = 0.0446$ km^{-1} and $\phi_c = 0.0983$ km^{-1} respectively. An exponential distribution of cell radii having $\phi_c = 0.0446$ km^{-1} is capable of producing cells with large radii, for example, the 0.99 quantile gives a cell radius of 103 km. Even though such cells have a low probability of occurrence, they need to be included in a simulation otherwise the simulated rainfall will be substantially lower than the observed. For this reason, the buffer region needs to be sufficiently large to accommodate cells that occur with large radii up to an equally large distance.

Consider different sizes for a buffer region, r, set to quantiles 0.0, 0.5, 0.8, 0.95, 0.99 and 0.999 of the exponential distribution of cell radii. The 0.0 quantile gives the case of having no buffer. Figure 4.2 shows the simulated proportion of mean rainfall depth for each buffer size for the two example months, where increasing buffer sizes give increasing proportions. The non-dimensional quantity $\phi_c(r_K + r)$ is used to standardize the cell size with respect to the size of the region. In order to obtain accurate estimates of the simulated proportion, the statistic was estimated from an arbitrarily long simulation of $120,000$ years length (limited only because of computational requirements). The point located at the center of the 65 km target region was used for the comparison, while points closer to the perimeter will demonstrate an even greater reduction in the simulated proportion.

Figure 4.2 gives a theoretical comparison for a point at the centre of the region and shows that the proportion will reduce significantly in the event of either smaller target regions or larger expected cell radii (smaller ϕ_c). The bias, denoted Ψ, can be obtained theoretically as the ratio of the number of cells landing within the target and buffer region that overlap the centre point to the number of cells that overlap the centre point for a

region extending infinitely, given as

$$\Psi = \frac{\int_{\xi=0}^{2\pi} \int_{r=0}^{r_K+r} \varphi_c r e^{-\phi_c r}\, dr\, d\xi}{\int_{\xi=0}^{2\pi} \int_{r=0}^{\infty} \varphi_c r e^{-\phi_c r}\, dr\, d\xi} \qquad (4.2)$$
$$= 1 - e^{-\phi_c(r+r_K)}(1 + \phi_c r + \phi_c r_K).$$

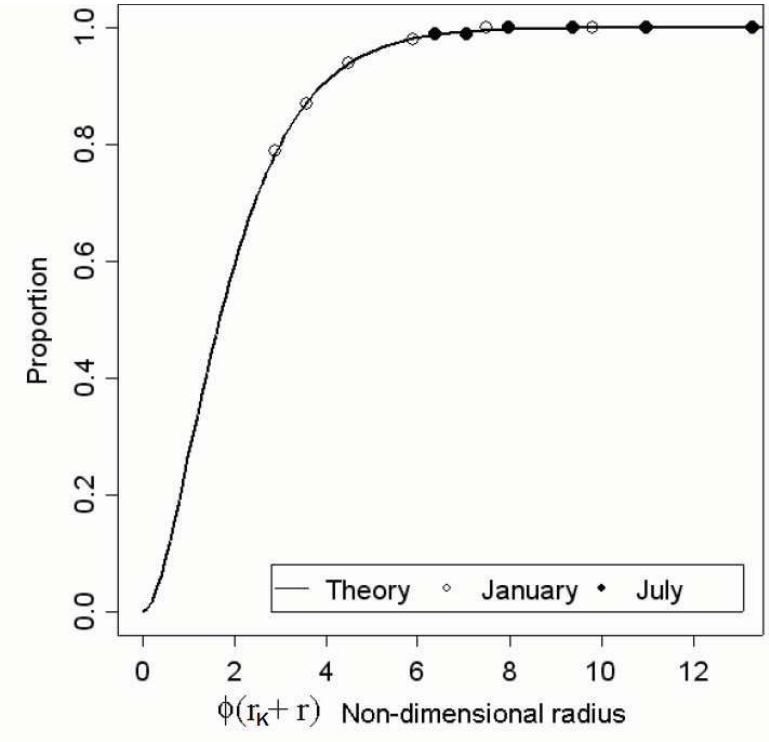

Figure 4.2 Proportion reduction in simulated rainfall depth, simulations for January and July compared with theoretical result.

Figure 4.2 shows that the size of the buffer region strongly affects the simulated statistics within the region, and that large buffer sizes are necessary to avoid reduction in the simulated statistics. For example, in January, the proportion of rainfall is 0.79 for the case having no buffer, and for a buffer set to the 0.99 quantile ($r = 103$ km), the proportion rainfall is 0.99. By comparing the two months, Figure 4.2 also shows that the proportion of simulated rainfall varies with each month due to the variation in the parameter ϕ_c.

The computational inefficiency introduced by the buffer is proportional to the ratio of the area of the target region to the area of the same region with the additional buffer. Assuming, that the 0.99 quantile is used to set the buffer size, the total simulation region is approximately 6.5 times larger for January and 3 times larger for July. For the Arno Basin case-study, an average month has a total simulation region 4 times larger than the

target region. This inefficiency is important given that the computational requirements for simulating the model are intensive. In order to address this problem, the following section develops an algorithm for directly simulating the number, location and radius of cells that land outside of the target region yet partially cover the region.

4.3 ALGORITHM FOR DIRECT SIMULATION OF CELLS

Consider a target region, of fixed radius, r_K, as in Figure 4.1. Cells arrive within this region according to a Poisson process with spatial rate parameter, φ_c. Consider also the same Poisson process occurring over an outer region that extends infinitely, with the exception that it is desired to retain only those cells that intersect the target region. This occurs when the cell radius, R_c, is greater than the distance from the circumference of the target region to the center of the cell, R_{xy}. The motivation then is to derive three distributions for the number, location and radius of cells, where each distribution is conditioned on the event that $R_c > R_{xy}$. Eq. 4.3 defines the discrete distribution of the number of cells in the outer region, N_o, that intersect with the target region. Having defined the distribution for the number of cells, it is necessary to define the properties for each given cell. The continuous distribution of cell centers, R_{xy}, is given in Eq. 4.4 for distances defined relative to the circumference of the target region. This distribution is conditioned on the event $R_c > r_{xy}$, as it is more likely for cells close to the target region to intersect it. The continuous distribution for cell radii is defined in Eq. 4.5. This distribution also depends upon the event $R_c > r_{xy}$, since the cell radius must be larger than the distance r_{xy}.

$$P_{n_o}(N_o = n_o | R_c > R_{xy}) \tag{4.3}$$

$$f_{R_{xy}}(r_{xy} | R_c > r_{xy}) \tag{4.4}$$

$$f_{R_c}(r_c | R_c > r_{xy}) \tag{4.5}$$

All other properties of the cell, the intensity, duration and starting time are independent of the location of the cell within the outer region.

4.3.1 Number of Outer Cells Intersecting Target

The following derivation relies on the Poisson distribution as the limit of a large number of Bernoulli trials. Consider an element, as in Figure 4.1, having radial increment, δr,

elemental angle, $\delta\xi$, and total radius, $r + r_K$, with cells arriving over that element according to a Poisson process having rate, φ_c. The probability of one cell landing inside the element is proportional to the elemental area, $(r + r_K)\delta r\delta\xi$, multiplied by the spatial rate. The probability of more than one cell landing in the element is of the order $(\delta r\delta\xi)^2$ and becomes zero in the limit. If a cell radius is distributed exponentially with parameter ϕ_c, then the probability that a cell, landing at distance r, intersects the target region is the survivor function, $e^{-\phi_c r}$. Thus the probability that a cell lands in the element and extends to the target region is given as,

$$p(r) \propto \varphi_c(r + r_K)\delta r\delta\xi \; e^{-\phi_c r}. \tag{4.6}$$

The proportionality factor required for Eq. 4.6 is obtained from the Poisson distribution as $e^{-\varphi_c(r+r_K)\delta r\delta\xi}$. This factor is omitted from the resulting derivation for notational convenience as the probability in Eq. 4.6 is linearized with respect to ξ when taking the limit $\delta\xi \to 0$, since this gives $e^{-\varphi_c(r+r_K)\delta r\delta\xi} \to 1$. It is clear from Eq. 4.6, that for a fixed elemental angle $\delta\xi$, the associated probability will vary with respect to the radius. Alternatively, for an arbitrary constant, $p(r) = p_o$, on rearranging Eq. 4.6, the elemental angle can be made to vary for a given radius to ensure that each element maintains this fixed probability, given as,

$$\delta\xi(r) = \frac{p_o}{\varphi_c(r + r_K)\delta r \; e^{-\phi_c r}}. \tag{4.7}$$

The set of elements, varying across all radii in the outer region and having elemental angles for a given radius specified by Eq. 4.7, define a set of Bernoulli trials. Each trial has a fixed probability, p_o, that a raincell lands in the elemental area at radius r and also intersects the target region. At a given radius the number of Bernoulli trials is given by $2\pi/\delta\xi(r)$, and the total number of trials, N_{tot}, can be obtained by integrating over all radii in the outer region,

$$\begin{aligned} N_{tot} &= \int_{r=0}^{\infty} \frac{\varphi_c 2\pi(r + r_K)}{p_o} e^{-\phi_c r} dr \\ &= \frac{\varphi_c 2\pi}{p_o}\left(\frac{r_K}{\phi_c} + \frac{1}{\phi_c^2}\right). \end{aligned} \tag{4.8}$$

A set of Bernoulli trials of size N_{tot}, from which n_o will be successful is distributed according to a binomial distribution with,

$$P_{N_o}(N_o = n_o) = \begin{pmatrix} N_{tot} \\ n_o \end{pmatrix} p_o^{n_o}(1 - p_o)^{N_{tot} - n_o}. \tag{4.9}$$

The Poisson distribution is derived as the limit to the binomial distribution as $N_{tot} \rightarrow \infty$ and is specified with a rate parameter, $\Omega = p_o N_{tot}$. Consequently, the number of cells in the outer region that have radius large enough to reach the target region is Poisson distributed with parameter Ω,

$$(N_o = n_o | R_c > R_{xy}) = \frac{\Omega^{n_o}}{n_o!} e^{-\Omega}, \tag{4.10}$$

where the parameter $\Omega = \varphi_c 2\pi \left(\frac{r_K}{\phi_c} + \frac{1}{\phi_c^2} \right)$ is obtained from Eq. 4.8.

4.3.2 Cell Centre Conditioned on Intersecting Target

The probability that a cell intersecting the target comes from an element with radial increment, δr, and elemental angle, $\delta\xi$, is the ratio of the expected number of cells with centre in the element that intersect the target, to the expected number of outer cells that intersect the target. This is defined as,

$$\begin{aligned} &P(R_{xy} = r_{xy} | R_c > R_{xy}) \\ &\quad = \frac{\varphi_c (r_{xy} + r_K) e^{-\phi_c r_{xy}} \, \delta r_{xy} \, \delta\xi}{\int_{r_{xy}=0}^{\infty} \int_{\xi=0}^{2\pi} \varphi_c (r_{xy} + r_K) e^{-\phi_c r_{xy}} \, dr_{xy} \, d\xi}. \end{aligned} \tag{4.11}$$

Evaluating the denominator for all angles $0 < \xi \leq 2\pi$, and all radii $(0 < r_{xy} \leq \infty)$, and taking the limit $\delta r \delta\xi \rightarrow 0$ this becomes the continuous distribution,

$$f_{R_{xy}}(r_{xy} | R_c > r_{xy}) = w\phi_c e^{-\phi_c r_{xy}} + (1 - w)\phi_c^2 r_{xy} e^{-\phi_c r_{xy}}, \tag{4.12}$$

where the resulting expression in Eq. 4.12 represents a mixture of an exponential distribution, $Exp[\phi_c]$, and a gamma distribution, $Gamma[2, \frac{1}{\phi_c}]$, and where the mixture ratio $w = \phi_c r_K / (\phi_c r_K + 1)$ depends upon the radius of the target region and the parameter for cell radii. For larger target regions and smaller expected cell radii, the distribution in Eq. 4.12 will tend toward an exponential distribution. The simulation of a point within the outer region requires a random angle, ξ, which can be independently sampled from a uniform distribution $\Xi \sim Uniform[0, 2\pi]$, giving the coordinates of the point as,

$$\begin{aligned} x &= (R_{xy} + r_K) \cos \Xi \\ y &= (R_{xy} + r_K) \sin \Xi. \end{aligned} \tag{4.13}$$

4.3.3 Cell Radius Conditioned on Location

Given that a cell has landed in an element at a distance r_{xy} from the edge of the target region, for the cell to intersect the region, its radius must be greater than the distance r_{xy}.

Given that the cell radius is exponentially distributed and that r_{xy} is the constant defined by Eq. 4.12, the remaining distance $R_c - r_{xy}$ is also exponential, based on the standard properties of this distribution. The distribution of the cell radius is therefore given as,

$$f_{R_c}(r_c|R_c > r_{xy}) = \phi_c e^{-\phi_c(r_c - r_{xy})}. \tag{4.14}$$

Simulation of the direct algorithm proceeds as previously outlined for the arrival of storms and occurrence of cells over the target region. For the cells that land outside the region yet intersect it, the number of cells is sampled using Eq. 4.10, their location is sampled using Eq. 4.12, and their radius is sampled using Eq. 4.14. All other cell properties such as the starting time, lifetime and intensity are independent of the location of the cell.

4.4 Results and Discussion

To demonstrate that the bias of the proposed algorithm is practically negligible, 100 replicates of 1,000,000 years length were simulated for the Arno Basin case-study. The resulting median proportion of rainfall was 0.99987, with a standard deviation of 0.0008, which arises because of the finite simulation length. This proportion is not statistically significantly different from a value of 1.0 at the 90% level.

To illustrate the efficiency of the direct simulation algorithm a comparison was conducted with the buffer algorithm for the Arno Basin case-study. Results for the direct algorithm are presented in Figure 4.3 as a proportion of the requirements for the buffer algorithm for a range of values of the non-dimensional quantity $\omega = \phi_c r_K$. The results were obtained using values of ϕ_c reported for the 12 months of the year for the Arno Basin case-study and for a range of target regions $r_K = 20, 40, 60, 80, 100$ km. For each month, the radius of the buffer region, r, was specified as the constant corresponding to the 0.99 quantile of the cell radii for that month. The results are reported for two variables, the memory usage and simulation time. Because the results of the direct algorithm are reported as ratios of the buffer algorithm, they can loosely be considered independent of the computational platform. The efficiency is largely independent of the number of simulation years, and a length of 300 years was used in order to measure run-times with sufficient accuracy. It is important to note that the measured efficiencies depend to some extent on the details of implementation in the computer code, since there is often a trade-off between memory usage and computational time. Therefore, Figure 4.3 should be regarded as an indication of the order magnitude of the efficiency and as a qualitative indication of the trend in efficiency with respect to various attributes.

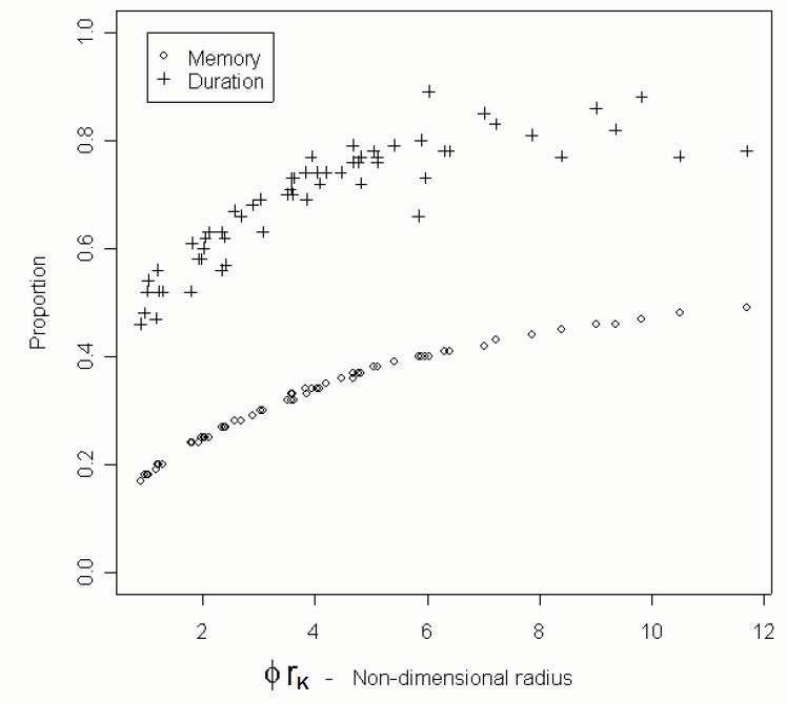

Figure 4.3 Comparison of direct-algorithm and buffer-algorithm for proportion of memory usage and simulation run time. Results reported for a range of sizes of the target region using different parameters across 12 months.

Figure 4.3 shows that the proportion of memory usage across the various months and target regions respectively follow two curves, with scatter attributed to random variation of the simulating process and the variability of parameters other than ϕ_c for each month. To interpret Figure 4.3, consider a parameter value of $\phi_c = 0.1$. For for a small radius of the target region, $r_K = 20\ km$, the buffer region is comparatively large to the target region, hence the direct simulation method is considerably more efficient. The direct method gives a proportion of memory usage at 0.25 and a proportion of duration at approximately 0.6. For a target region having a large radius of $r_K = 100\ km$, the buffer region is not as comparatively large, hence the efficiency of the direct method is reduced. At this radius, the proportion of memory usage is 0.4 and the proportion of duration is approximately 0.8. For this latter case, it is important to note that even a small reduction in the duration of the simulation is significant since a simulation will always maximize the available resources, for example by increasing the simulation length or by computing additional replicates. Similar comparisons could be made from Figure 4.3 by considering the variation with respect to the parameter ϕ_c with respect to a fixed target region.

4.5 CONCLUSION

The need to simulate cells outside of a region of interest was demonstrated in order to avoid significant boundary effects. An efficient method was derived for directly simulating these cells without the use of a buffer region. The performance of the algorithm varies with respect to the parameter values and the size of the target region, but in all cases it was demonstrated to provide considerable improvement. The significance of the direct method is due to the computationally intensive nature of the model and the desire to increase simulation lengths and the number of replicates of the model.

Chapter 5
Calibration Involving a Monthly Boundary

5.1 INTRODUCTION

The calibration technique of using separate parameters for each month implies that there is a boundary at each month and that the rainfall regime switches over at this boundary. This is of concern because it is possible for a storm origin occurring within a given month to generate raincells that do not land within that month but within the following month. The most obvious example is for a storm occurring on the last day of the month, however given certain parameter values this may not be the sole case. The main parameter in question is β, which controls the displacement of raincells relative to a storm origin. The mean value for cell displacement is given as β^{-1} hours, thus, if a typical value of beta is 0.01, then the average cell displacement is a period of 4 days and the 0.95 quantile is 12.5 days. This scenario is not uncommon when calibrating the Neyman-Scott rainfall model and it represents a conceptual limitation of the model. That is, while intuitively one may expect that rain may begin within a period of hours of storm inception, the model is capable of generating large displacements between the storm origin and the first occurrence of rain associated with that origin. This scenario comes about because the location of the storm origin is hidden and does not itself result in a burst of rainfall. Thus, in the observed record the first appearance of the storm coincides with the first rainfall cell and the location of the origin becomes immaterial. However, the exception to this is at the location of a monthly boundary.

While the rainfall model is conceptual, a calibration proceeds by obtaining parameters that give the best fit in statistical terms only. Therefore it is possible that parameters giving the best fit do not have the most realistic interpretation. The parameter β, controlling the dispersion of cells, is strongly related to statistical properties such as the skewness and

auto-correlation and in order to obtain accurate fits it is possible to generate physically meaningless values. For example, *Cowpertwait et al.* [2002] reports a value $\beta = 0.00446$ for one of the months, which has a mean displacement of approximately 9 days and at the 0.95 quantile this displacement is 28 days. In other words, a storm on the 4^{th} of January will have, on average, 5% of the rain associated with it occurring at the start of February. Storms beginning closer to the end of the month will have even higher portions of rainfall landing outside the month of origin. This will significantly affect the simulated statistics as they will not match the observed statistics that were used to calibrate derived properties of the model.

A simple remedy for this is to wrap rainfall across a monthly boundary, analogous to a clock going from 12 back to 1. Following the prior example, if rainfall associated with a January storm lands on the 1^{st} of February it is instead made to land on the 1^{st} of January. Statistically, this solution is feasible as the modelled rainfall process within a given month is stationary and the locations of storm origins are independently and uniformly distributed. However, the solution is physically undesirable as it could be detected with the eye that storms appear to abruptly stop at the end of a month and a suitable statistic could be easily developed to show this.

An alternative solution is to calibrate the parameters across all months simultaneously such that the amount of rainfall lost to the following month is balanced with the amount of rainfall gained from the preceding month. Technically, this will make the timeseries of rainfall within a month non-stationary, as the rain at the start of a month will be a mixture of the properties for the two months where the mixture proportion decreases further into the month. This artefact is however not undesirable since it is likely to mimic the true rainfall process and provide some remedy (though unquantified) to the arbitrary binning of observed data into 12 rigid periods of approximately 31 days over which the rainfall is assumed to be stationary.

5.2 Derivation of Monthly Bias

To derive the amount of rainfall falling outside of a monthly boundary it is necessary to consider two cases (i) cells landing entirely beyond the monthly boundary and (ii) cells partially landing beyond the monthly boundary. Figure 5.1 shows a schematic diagram for these two scenarios.

For the case shown in Figure 5.1 (a) it is necessary to consider all possible locations for a storm origin on the interval $s = [0, t]$ taken relative to the end of the month. The probability that a cell lies beyond the end of the month is given by the survivor function

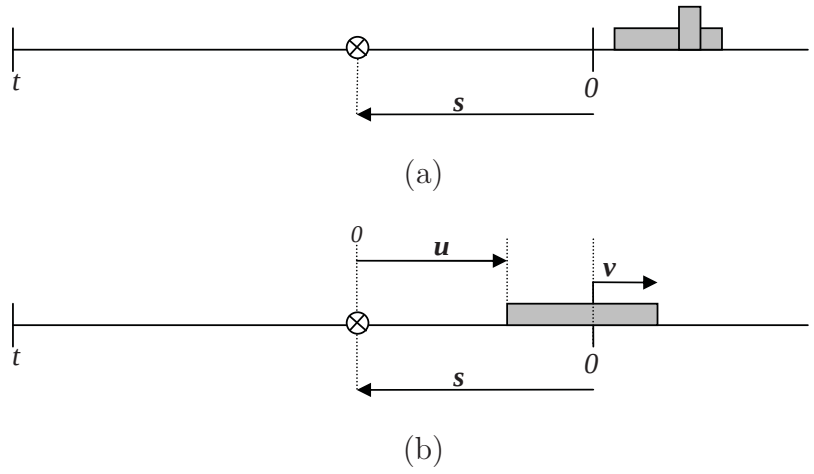

Figure 5.1 Schematic diagrams of rainfall landing outside the month boundary, (a) raincells entirely beyond the boundary and (b) raincells partially beyond the boundary.

$e^{-\beta s}$ and the total amount of rain for this cell is obtained as a product with expected values for: the rate of storms λ, the mean number of cells per storm μ_c, the mean cell intensities μ_X and mean cell duration η^{-1}. For the case shown in Figure 5.1 (b), all storms on the interval $s = [0, t]$ must again be considered, but the starting time of a given cell must occur prior to the month boundary. Therefore, all cell displacements on the interval $u = [0, s]$ relative to the storm origin are considered. The probability that rain lands beyond the boundary is then the probability that a cell lands prior to the boundary $\beta e^{-\beta u}$ multiplied by the probability that that cell does not terminate prior to the end of the month, given as the survivor probability, $e^{-\eta(s-u)}$. The total rainfall is then multiplied by the rate of storms λ, expected number of cells μ_C, expected intensity μ_X and expected duration of the cell beyond the monthly threshold. The expected lifetime of the cell beyond this threshold is obtained as

$$E[V] = \int_{v=0}^{\infty} v\eta e^{-\eta v}\, dv = \frac{1}{\eta}, \tag{5.1}$$

which is a standard result that the mean of a survivor function (relative to some threshold) is the same as the mean of the exponential distribution relative to zero. Using these two cases, the total amount of rain lost to a subsequent month, ϵ, is given as,

$$\epsilon = \frac{\lambda\mu_C\mu_X}{\eta} \int_{s=0}^{t} e^{-\beta s}\, d\, s + \frac{\lambda\mu_C\mu_X\beta}{\eta} \int_{s=0}^{t}\int_{u=0}^{s} e^{-\beta s} e^{-\eta(s-u)}\, d\, u\, d\, s. \tag{5.2}$$

Evaluating the integrals gives the result,

$$\epsilon = \frac{\lambda \mu_C \mu_X}{\eta} \left[\frac{(1 - e^{-\beta t})}{\beta} + \frac{(1 - e^{-\beta t})}{(\eta - \beta)} + \frac{\beta(e^{-\eta t} - 1)}{\eta(\eta - \beta)} \right]. \tag{5.3}$$

It should be noted that a simplification has been used in the derivation: that the duration of a subsequent month extends infinitely rather than having a finite range. This is unlikely to affect the equation significantly due to the low probability of rainfall landing two or more months after a storm origin. Ignoring then the effect of rainfall displaced by two or more months, for months $k = 1, \ldots, 12$, the bias is determined as the portion of rainfall simulated in a given month,Ψ_k, is then,

$$\Psi_k = \frac{\mu_{k,h} - \epsilon_k + \epsilon_{k-1}}{\mu_{k,h}}, \tag{5.4}$$

where December precedes January ($\epsilon_0 \equiv \epsilon_{12}$ and where the $\mu_{k,h}$ is calculated from Eq. 3.3 using a monthly aggregate level, $h(k) = 24\{31, 28, 31{,}30, 31, 30, 31{,}31, 30, 31, 30, 31\}$. Eq. 5.4 was used to inspect the parameters reported by *Cowpertwait et al.* [2002] for the Arno Basin case-study. Table 5.1 reports the the mean amount of rainfall for each month (non-dimensionalised such that $\mu_1 = 1$), the mean amount of rainfall that lands in the subsequent month, ϵ, and the bias in the rainfall mean, Ψ, for that month. The parameters β and η are reported for comparison with the bias because they control (i) the probability of a cell landing outside of a month and (ii) the probability of a cell within the month having a lifetime that extends beyond the month boundary. The percentage of the bias contributed by the latter case is also reported in Table 5.1.

Table 5.1 shows that the bias in the mean can be significant, ranging from 0.74 up to 1.29 for this case-study. It is also evident that large amounts of rainfall occurring after a monthly threshold are related to smaller values of the β parameter. The percentage of the bias controlled by the η parameter (when cells partially extend beyond the month threshold) is quite small with respect to the case where entire cells are located beyond the month boundary. This suggests that β is the dominant parameter of interest. To clarify this, the bias is shown in Figure 5.2 with respect to the ratio of parameters β_k/β_{k-1} between the current and preceding month. This figure shows that when this ratio is higher, it implies that the current month gains rainfall from the prior month relative to the amount it loses to a subsequent month. The figure shows some scatter due to variation in the other parameters between months (and for the same reason a ratio of one may not correspond exactly to the case of no bias).

5.3 CASE-STUDY

A mechanism to remedy this is to jointly calibrate pairs of months together to obtain parameters, for example, say January and February parameters, such that the parameters for February allow for rain contributed by January. As the parameters for January depend on December, December on November and so on, this procedure requires a joint calibration of parameters across all 12 months. This is a significant increase in the parameter space. Properly determining model properties under this scenario is a complicated task as the resulting model is non-stationary in time (with rain at the start of the month being more similar to the prior month). Instead, an approximation is used in this section whereby non-stationary effects are ignored and the properties of the monthly rainfall process are considered as the mixture of rainfall between the current and preceding month. The mixing proportion is the amount of rainfall contributed by the previous month with respect to the total amount of rainfall within that month. The mixing proportion, w_k is thus,

$$w_k = \frac{\epsilon_{k-1}}{\mu_{k,h} - \epsilon_k + \epsilon_{k-1}}. \tag{5.5}$$

A joint calibration using this approach was performed for the Arno Basin case-study. The results are shown in Table 5.2. From this table it can be seen that parameters were obtained such that the bias reduces to a value close to one. This does not correspond to the case of zero rainfall crossing the monthly threshold, as ϵ is non zero for all months.

Table 5.1 Bias in the mean rainfall due to rain landing outside of a month which that storm is associated with.

Month	β	η	μ	ϵ	Ψ	% partial cells
1	0.00446	0.738	744	217.4	0.74	0.6
2	0.051	0.929	678	20.7	1.29	5.2
3	0.061	1.09	744	17.3	1.00	5.3
4	0.014	1.2	720	72.3	0.92	1.2
5	0.0513	1.55	744	20.1	1.07	3.2
6	0.00888	2.1	720	112.9	0.87	0.4
7	0.0185	2.63	744	54.4	1.08	0.7
8	0.0189	2.62	744	53.3	1.00	0.7
9	0.0195	2.05	720	51.8	1.00	0.9
10	0.00906	1.42	744	110.9	0.92	0.6
11	0.00973	1.05	720	103.6	1.01	0.9
12	0.0439	0.871	744	23.9	1.11	4.8

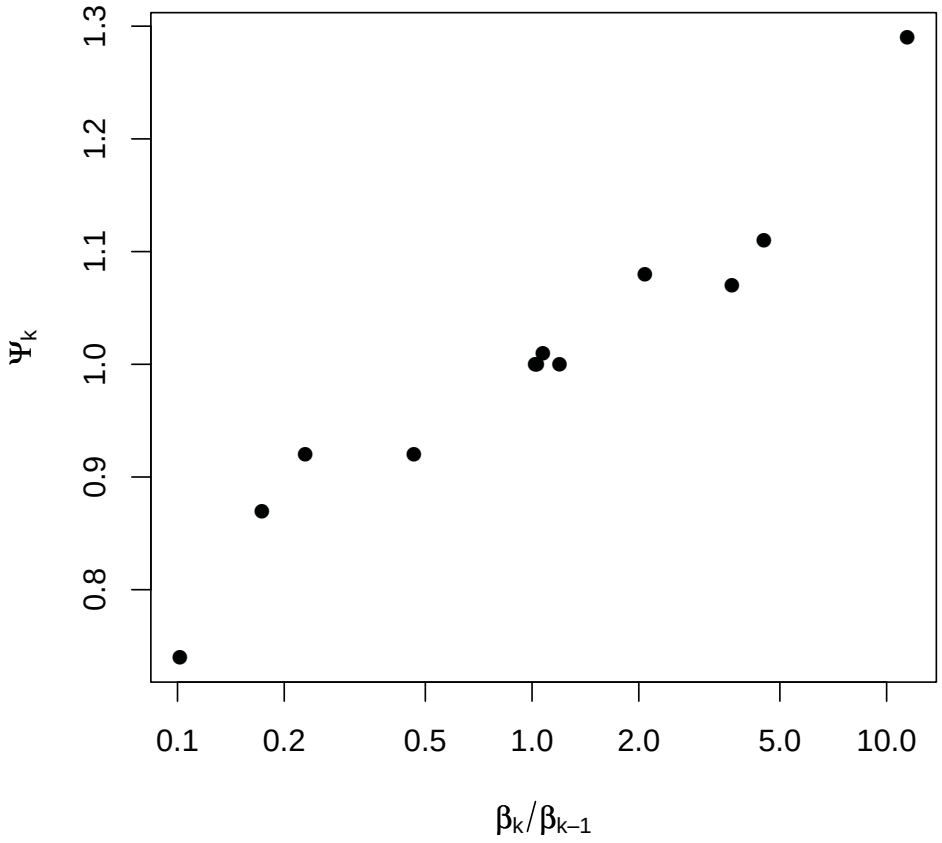

Figure 5.2 Bias in the mean rainfall, Ψ_k, for month k with respect to the ratio of parameters β_k/β_{k-1} controlling cell displacement in each months.

Instead it shows that the variability in the amount of rainfall crossing each month has reduced significantly and approximately the same amount of rainfall is allowed to land in a subsequent month as is received into that month. Table 5.2 shows the corrected β parameters with respect to the original values and it can be seen that they have all increased, implying that storms are less dispersed. Attempts to force other calibration solutions where an overall increase to the β parameters was not necessary (i.e. by maintaining a non-trivial value of the mixing proportion) were unsuccessful. It can therefore be concluded that the solution to the calibration was found by reducing the mixing proportion of rainfall to a value close to zero yet similar across all months. Due to the larger value of β it can be concluded the trivial amount of rainfall that lands over a month threshold comes from the latter days in that month, which is a realistic outcome. Table 5.2 also shows that for the improved calibration, the percentage of rainfall attributed to the parameter η is still

small compared to the influence of the β parameter.

Table 5.2 Bias in the mean rainfall after constrained parameter estimation.

Month	β prior	β corrected	μ	ϵ	Ψ	% partial cells
1	0.004	0.116	744	9.3	1.01	7.5
2	0.051	0.093	678	11.6	1.00	6.7
3	0.061	0.097	744	11	1.00	6.5
4	0.014	0.083	720	12.7	1.00	4.6
5	0.051	0.083	744	12.6	1.00	4.6
6	0.009	0.100	720	10.3	1.00	3.6
7	0.019	0.104	744	9.9	1.00	3.3
8	0.019	0.191	744	5.5	1.01	4.7
9	0.020	0.089	720	11.6	0.99	3.2
10	0.009	0.084	744	12.4	1.00	4
11	0.010	0.083	720	12.7	1.00	5.4
12	0.044	0.074	744	14.5	1.00	6.5

The improved calibration requires other parameters to change to compensate for the necessary change in the β parameter. To verify that the β parameter is the most sensitive when the constraint of no bias is applied to the calibration Figure 5.3 shows the changes in all other parameters from their original value to their corrected value.

Figure 5.3 shows that β is the most sensitive parameter as it has the largest deviations from the 45° line. Note that there is no significance in the fact that most parameters appear above the 45° line as some parameters are inversely related to the total amount of rainfall while others are not. The smaller the value of β the larger the change in value, which in some instances was an order of magnitude. The other parameters show some variation, but not as significant.

5.4 CONCLUSION

The case-study demonstrates that the effect of a constraint on the monthly bias in the mean is to increase the parameter that controls the cell dispersion. Typical values for this parameter are on the order of $\beta \approx 0.1$ to ensure that significant amounts of rainfall do not land in subsequent months. The calibration technique to implement this constraint is cumbersome as it requires all 12 months to be jointly calibrated, which is a significant increase in the parameter space. While it is possible to perform this type of calibration for subsequent case-studies, a more direct approach is to check that the parameter β is

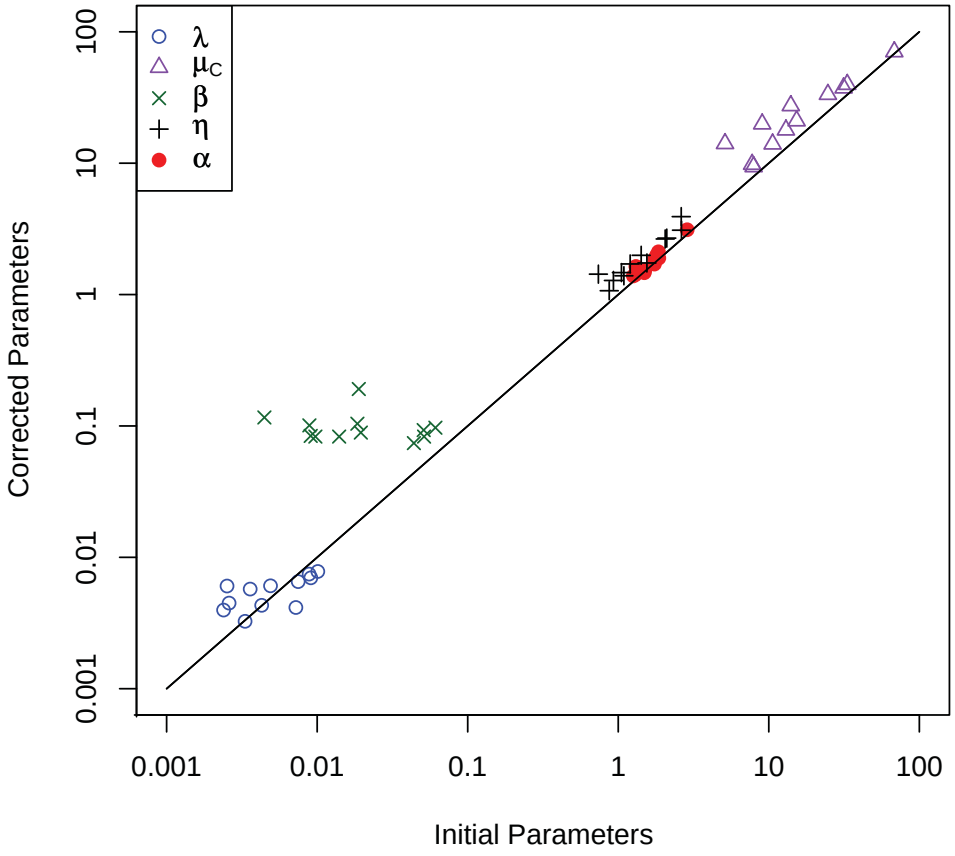

Figure 5.3 Change in parameters accounting for monthly bias. β is the parameter most sensitive to the bias correction.

not too small and if necessary provide a lower bound to the value taken by this parameter in the calibration. A secondary observation is to ensure that the variation in β is not too significant between the months.

Chapter 6
Calibration Using Simulated Moments

6.1 INTRODUCTION

Neyman-Scott and Bartlett-Lewis models are conceptual rainfall models based on the Poisson clustering of rainfall in time. They have been used extensively to model observed rainfall data at a single gauge, at a set of gauges and also radar rainfall estimates [*Northrop*, 1998; *Wheater et al.*, 2005]. The method for calibrating models that belong to the Poisson-cluster family has typically involved minimising the difference between observed statistics and the equivalent analytic properties of the model, known as the method of moments. This approach is adopted since likelihood functions are not tractable in terms of direct comparison to observed rainfall measurements at a given aggregate. *Chandler* [1997] has demonstrated a technique for calibrating a Poisson-cluster model using a likelihood function, by employing spectral theory, however this method is not common. No matter what calibration method is used, a central requirement is the need to specify the analytic properties of the model for the selected statistics used in the calibration. For a given model specification, deriving the corresponding properties can be a non-trivial and in some instances non-tractable task. This limits the flexibility for specifying alternative model formulations, as only those that yield simple expressions are considered.

Instead of expressing the analytic properties of the model, a Monte Carlo Simulation method (MCS) is proposed for use in calibration. For a given parameter set, the model is first simulated and statistics of interest are calculated. Longer records will yield statistical estimates that have less variability than those estimated from a shorter record. The simulated statistics may be then compared to their equivalents calculated from observed data, as per the method of moments. This procedure gives an indication of the goodness-of-fit for the given parameter values, and it can be employed with an optimiser to locate

parameters that model the data well.

This chapter considers the performance of calibration using MCS with respect to using known analytic expressions for the model. The following section outlines modifications to the SNSRP calibration procedure to ensure that it can be feasibly calibrated using MCS. A case study is subsequently presented for comparing the MCS calibration method to calibration using analytic expressions. The region surrounding Launceston, Tasmania, having 35 suitable daily rainfall gauges and 6 pluviograph gauges, is considered for this study.

6.2 CALIBRATION METHODOLOGY

Recall that the SNSRP model requires six parameters $(\lambda, \alpha, \mu_C, \beta, \eta, \phi)$ to simulate rainfall over a region, with a different parameter, θ, calculated at each site to scale the non-stationarity in mean rainfall across a region. While the model allows for the mean to vary, the coefficient of variation and the probability of a dry period are both assumed to be constant over the region, thus the region should be homogenous with respect to these statistics. Seasonality of the data is taken into account by fitting a different set of parameters for each month, thus for M sites, there is a total of 12×6 model parameters + $12 \times M$ scaling parameters.

As outlined by *Cowpertwait et al.* [2002], the above model assumptions enable analytical expressions to be derived for various properties of the model in terms of the model parameters. The properties used for calibration in this study are the coefficient of variation, lag one autocorrelation, skewness and lag zero cross correlations between each pair of sites. Since the model assumes the rainfall process to be stationary in space, the model properties are calculated for a non-dimensional case, where each observation site is scaled to have an hourly mean rainfall depth of $1\ mm$. By non-dimensionalising the model properties over a region, the observed data can be pooled across that region to yield a set of regional statistics used for calibration. The model properties are calibrated at multiple aggregates to ensure that the model performs adequately at a range of time scales. For this purpose, 3 aggregates at 1-hour, 6-hour and 24-hours are considered.

The MCS calibration exploits the ability to calibrate the model in a step-wise procedure. Specifically, the temporal parameters are calculated assuming that the set of regional statistics are equivalent to a set of statistics at a single observed point. That is, the model is first calibrated as a point rainfall model. Having obtained an optimal set of temporal parameters, the parameters that describe the spatial extent of the model are then calibrated. The benefit of this approach is that the majority of the parameters relate to the tempo-

ral process and it is significantly more efficient to simulate the temporal process without regard for the spatial distribution of rain.

Due to the finite length of a simulated record, there is sampling variability induced in the estimated statistics. Preliminary investigations revealed that a gradient-based optimiser was too sensitive to this variability, even for simulated records having length greater than 1,000,000 years. Therefore, a stochastic optimiser (shuffled complex evolution) was employed for the simulated calibration. This, significantly increased the computational burden of calibration, limiting simulations to 10,000 years. Whilst a length of 10,000 years may appear excessive given that the observed data is typically less than 100 years in length, it is equivalent to estimating the expected value of a statistic from say 100 replicates of 100 years each.

Preliminary results indicated that the skewness statistic estimated from a 10,000 year record was too variable, which caused difficulty for the optimiser to locate good parameter sets. To address this limitation, an alternative statistic, the L-skewness is considered [*Hosking and Wallis*, 1997]. The L-skewness is formulated as a linear combination of order statistics for the simulated data, and is therefore less variable than the product-skewness for a given sample size. For a ranked data set, X, where the j^{th} order statistic is denoted X_j, the L-skewness is defined as,

$$t_3 = \frac{6b_2 - 6b_1 + b_0}{2b_1 - b_0}, \tag{6.1}$$

where,

$$b_r = \frac{1}{n} \sum_{j=r+1}^{n} \frac{(j-1)(j-2)\dots(j-r)}{(n-1)(n-2)\dots(n-r)} X_j. \tag{6.2}$$

While the need to implement the L-skewness results from a limitation of the MCS calibration approach, it also demonstrates the main advantage of this approach. Specifically, the ability to select alternative statistics for use in calibration without recourse to deriving the same statistic in terms of the model parameters. Linear moments are a good example of this, since for the same reason that a likelihood function in terms of direct rainfall observations is not tractable, the properties of linear moments are not tractable (i.e. the rainfall distribution at a timescale of interest cannot be derived). For this investigation, the L-skewness was the only alternative statistic used in the calibration. It is important to note, however, that the simulated approach enables any statistic to be included provided it can be reliably estimated from the simulated data.

6.3 CASE-STUDY

The SNSRP model is calibrated to a network of gauges in the vicinity of Launceston, Tasmania, as shown in Figure 6.1. Tasmania experiences strong gradients in rainfall, particularly between the East Coast, the interior and the West Coast, but also between the northern and southern extent of the island. Launceston belongs to Rainfall District 91, as defined by the Bureau of Meteorology, which extends across to the north-west corner of Tasmania. Within this region, 35 daily rainfall gauges located with close proximity to Launceston were selected, in addition to 6 pluviograph intensity gauges (6 min). Several pluviograph records were included in the calibration despite their increased distance from Launceston because of their length of record and their similarity in terms of first and second order moments. However, not all of the daily gauges within the region were used due to the limited length of some records and because employing a larger region diminishes the assumption of homogeneity over that region. Excluding the remote pluviograph gauges, the simulation region is approximately $40 \times 40\ km^2$.

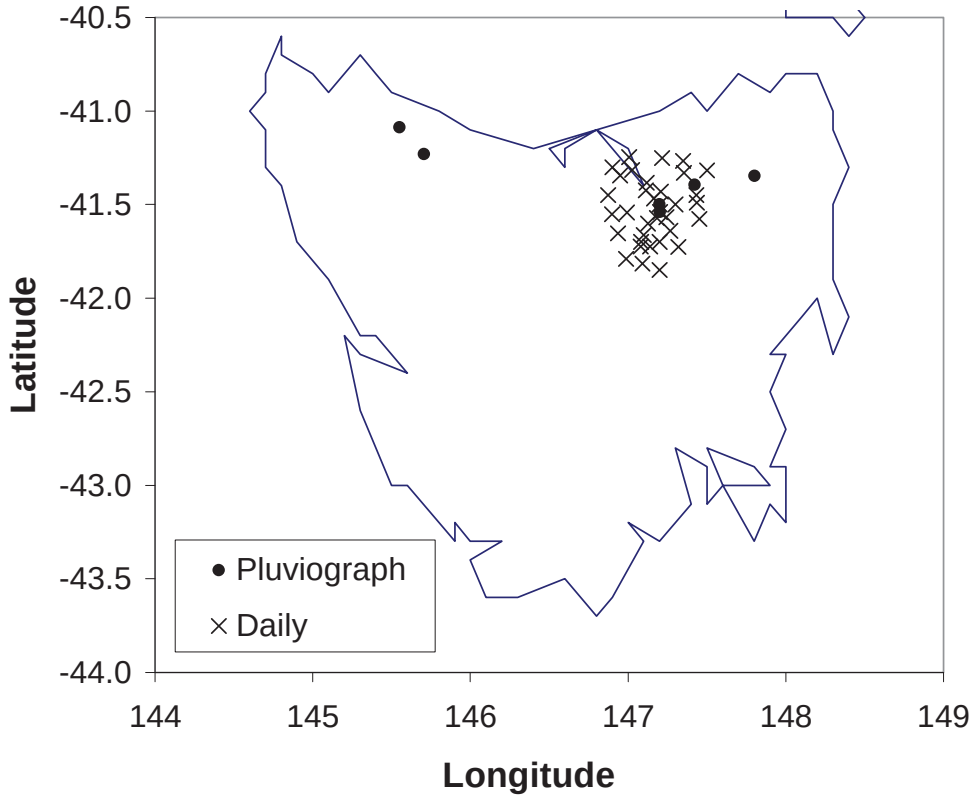

Figure 6.1 Spatial coordinates of observation gauges about Launceston, Tasmania.

The average observation length of the pluviograph gauges is 37 years, with dura-

tions ranging from 30 years to 68 years. The average observation length of the daily gauges is 59 years, with durations ranging from 31 years to 117 years. The annual rainfall is $720\ mm$ with winter rainfall approximately double the quantity of summer rainfall (monthly averages 41mm, January; $86\ mm$, July). The winter rainfall is also more variable than summer (coefficient of variation 0.32, January; 0.46, July).

The observed and fitted statistics used in the calibration of the SNSRP are presented in Figure 6.2, for each of the three aggregates, 1 hour, 6 hours and 24 hours. Figure 6.2 (a) shows the observed coefficients of variation ($\hat{\nu}_1$,$\hat{\nu}_6$,$\hat{\nu}_{24}$); Figure 6.2 (b) shows the observed autocorrelations ($\hat{\rho}_1$,$\hat{\rho}_6$,$\hat{\rho}_{24}$); and Figure 6.2 (c) shows the observed non-dimensionalised product-skewnesses ($\hat{\kappa}_1$,$\hat{\kappa}_6$,$\hat{\kappa}_{24}$). The model calibrated using analytic expressions (solid line) shows a very good agreement to all of the statistics, though the higher aggregates (24 hour and 6 hour) are fitted better than the 1 hour aggregate. The statistics calibrated using MCS show a similar agreement, however the fit is slightly poorer for all months (excepting December, in which it is marginally better). Again, the largest deviations are observed for the 1 hour aggregate, and for the skewness in particular. While the L-skewness was used for the simulated calibration, the equivalent product-skewness has been presented in Figure 6.2 (c) to facilitate comparison.

Figure 6.3 demonstrates the calibration of the model to hourly values of cross-correlation. The months January and July are respectively shown in Figure 6.3 (a) and Figure 6.3 (b) as example calibrations. It is evident from these figures that the model provides an adequate fit to the data, but this is difficult to conclusively verify given the low number of sub-daily gauges used in the calibration.

For each month, k, Table 6.1 shows the parameter estimates for the 6 parameters obtained using the gradient based optimisation of the analytic expressions. Table 6.2 gives the equivalent parameter values obtained using MCS with a stochastic optimiser. By comparing the two tables for any given month it is evident that the two methods have not converged to the same parameter values.

Whilst some of the parameters give highly similar values for any given month (for example the storm rate, λ, and the cell radius, ϕ_c), other parameters are significantly different (for example, μ_C). Despite this, it is evident that across all of the months, the calibrated parameter sets lie within similar ranges of each other. The ability of the model to yield parameter sets of similar quality (as observed in Figure 6.2 and Figure 6.3) yet having contrasting parameter values, suggests a limitation of the conceptual formulation: that it is over-parameterised. The observation of variability in parameter values for a near optimum fit to the data occurred irrespective of the calibration method. Therefore, the

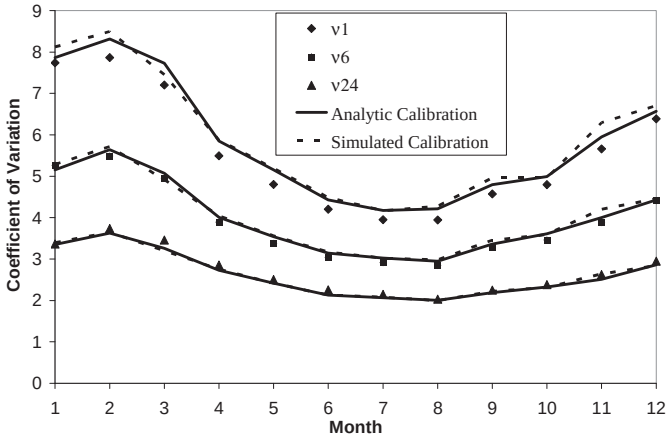

(a)

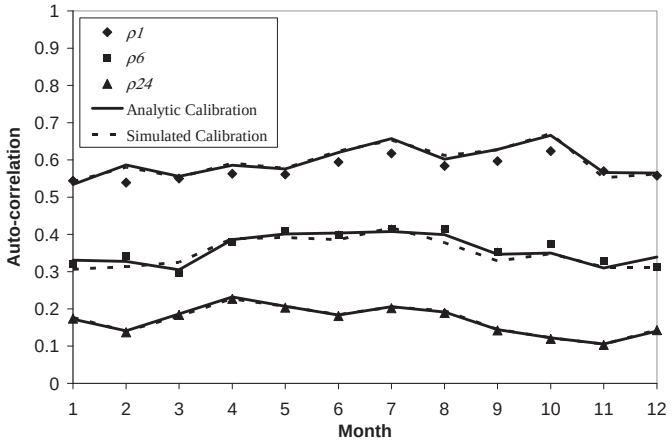

(b)

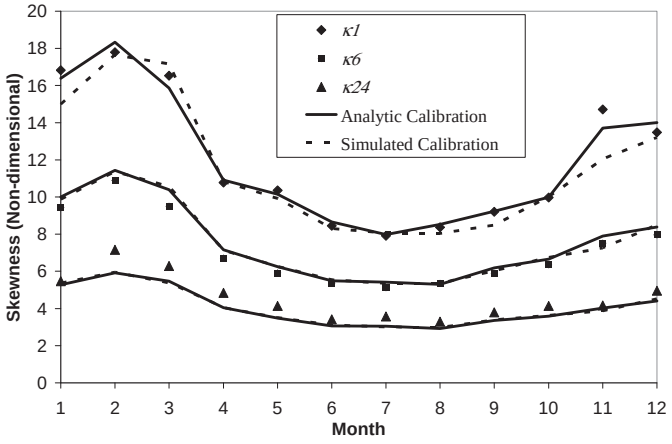

(c)

Figure 6.2 Observed statistics for aggregates 1 hr (diamond), 6 hr (square) and 24 hr (triangle) compared with model properties calibrated respectively using analytic expressions (solid) and MCS (dashed): (a) coefficient of variation, (b) auto-correlation, (c) skewness (non-dimensionalised).

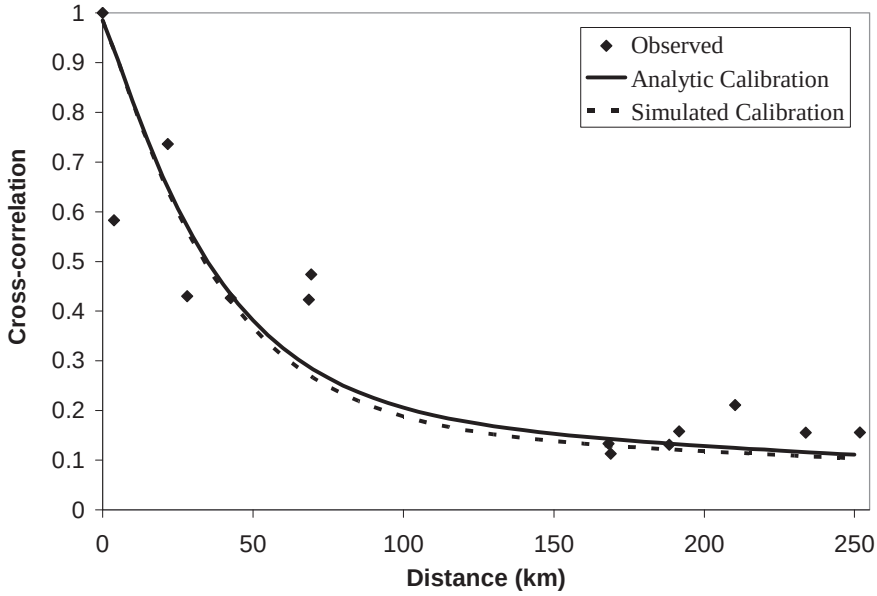

(a)

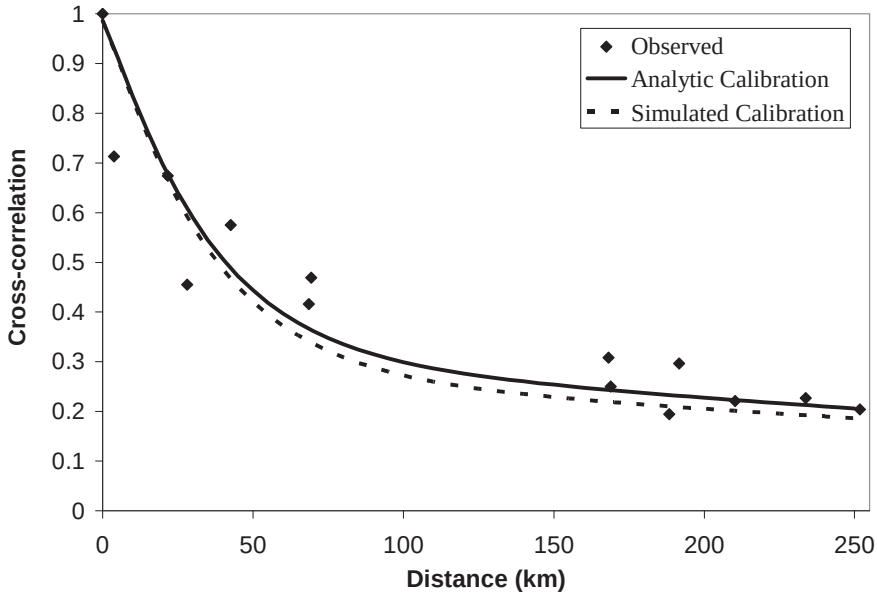

(b)

Figure 6.3 Comparisons of cross-correlations for 1 hour observations (diamond) model calibrated using analytic expressions (solid) and MCS (dashed) for two example months: (a) January, (b) July.

Table 6.1 Launceston parameter estimates for each month, via analytic expressions.

	$\hat{\lambda}$	$\hat{\beta}$	$\hat{\eta}$	$\hat{\mu_C}$	$\hat{\alpha}$	$\hat{\phi_c}$
k	h^{-1}	h^{-1}	h^{-1}	-	-	km^{-1}
1	0.0037	0.106	1.49	20.0	0.626	0.067
2	0.0035	0.131	1.28	24.7	0.536	0.096
3	0.0040	0.088	1.26	15.0	0.657	0.080
4	0.0046	0.083	1.30	22.0	0.722	0.052
5	0.0061	0.100	1.49	32.0	0.623	0.067
6	0.0082	0.115	1.32	35.0	0.587	0.074
7	0.0085	0.099	1.07	29.0	0.616	0.069
8	0.0092	0.108	1.37	36.0	0.572	0.111
9	0.0093	0.132	1.13	19.1	0.636	0.062
10	0.0084	0.161	1.05	27.0	0.543	0.079
11	0.0076	0.177	1.50	33.0	0.483	0.115
12	0.0053	0.137	1.47	27.0	0.556	0.098

Table 6.2 Launceston parameter estimates for each month, via MCS.

	$\hat{\lambda}$	$\hat{\beta}$	$\hat{\eta}$	$\hat{\mu_C}$	$\hat{\alpha}$	$\hat{\phi_c}$
k	h^{-1}	h^{-1}	h^{-1}	-	-	km^{-1}
1	0.0038	0.095	1.35	11.3	0.805	0.066
2	0.0036	0.126	1.25	17.3	0.604	0.096
3	0.0041	0.098	1.34	25.0	0.539	0.080
4	0.0046	0.086	1.29	21.6	0.733	0.052
5	0.0062	0.098	1.43	26.1	0.670	0.067
6	0.0085	0.109	1.21	22.8	0.679	0.074
7	0.0083	0.102	1.14	35.2	0.588	0.069
8	0.0095	0.100	1.20	21.5	0.681	0.110
9	0.0094	0.124	1.07	10.2	0.860	0.062
10	0.0086	0.159	1.01	26.0	0.539	0.079
11	0.0067	0.181	1.63	18.7	0.660	0.115
12	0.0058	0.123	1.33	14.6	0.675	0.098

difference between Table 6.1 and Table 6.2 is considered to be an artefact of the model rather than a deficiency in the MCS calibration technique. With further investigation into the behaviour of the model, this artefact could be mitigated by fixing some of the parameters as constant.

To verify the performance of the model, a simulation of 10 replicates, each having 54 years length, was conducted using the parameters in Table 6.2. A limited presentation of the results is provided, as there is a wide variety of statistics that could be employed to verify the model and 41 different sites at which a comparison could be made. For this reason, regionalised statistics were used as an indicator of the models overall performance, though they have the potential to conceal at-site variability. The selected statistics were pooled across each site by first dividing by the median of the simulated distribution for each replicate.

Two key statistics not included in the calibration were selected to demonstrate the models performance, the distribution of regional annual maxima and the distribution of regional annual totals. Figure 6.4 shows the standardised 1 hour and 24 hour maxima using a Gumbel axis. Comparing the 90% limits and median to the observed data shows that the 1 hour extreme values have good agreement but that the 24 hour extremes are underestimated in the upper tail.

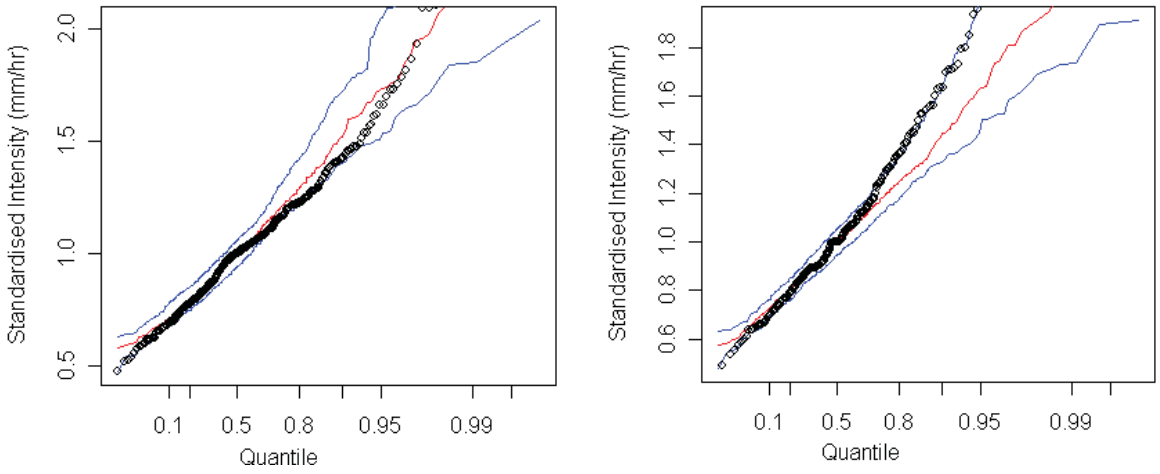

Figure 6.4 Comparison of standardized observed annual maxima with simulated median and 90% limits. Results presented for two separate aggregates: (a) 1 hour, (b) 24 hour.

Figure 6.5 shows the standardised annual totals presented on a Normal axis. This figure shows good agreement in terms of the mean annual total, but the variability of the annual aggregate is under-simulated, and explanation can most likely be attributed to (i)

the method of independent monthly calibration, which removes variability due to rainfall cross-correlated between months and (ii) non-stationarity of the observed annual rainfall distribution due to climatic states not accounted for in the modelling.

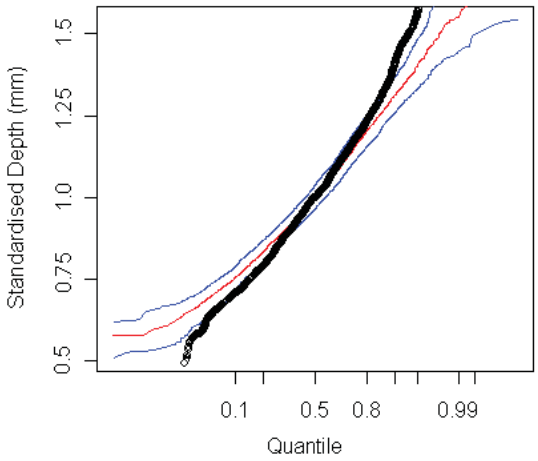

Figure 6.5 Comparison of standardised observed and simulated median and 90% limits of annual totals.

6.4 Conclusion

The method of moments is a common approach for calibrating conceptual models including those from the Poisson cluster family. In order to calibrate model parameters, this approach requires properties of the model (such as first and second order moments) to be compared to equivalent statistics from observed data for the specified region. A Monte Carlo approach was adopted to simulate the observed model properties, thus avoiding the need for deriving these same properties. The technique was demonstrated for the SNSRP model for 41 rainfall gauges located about Launceston, Tasmania.

The calibration was performed in two stages, first calibrating temporal parameters at a single point, then spatial parameters having optimised all other parameters. This approach exploited the ability to calibrate the model in a step-wise manner and enabled feasible computation. The temporal simulation required 10,000 year durations to produce reliable estimates of the model properties for a given parameter set. The spatial simulation was more computationally demanding as the cross-correlogram is comprised of estimated correlations between each pair of sites. A simulation of only 100 years was used, but this was sufficient as the goodness-of-fit was not as susceptible to the variability of an individual correlation estimate. A stochastic optimiser was employed to locate parameters

for the MCS calibration approach, as a gradient-based optimiser was overly sensitive to the variability in estimated statistics due to finite simulation length.

The product-skewness statistic was too variable to be reliably used for calibration, so the L-skewness was used instead. Whilst this was a necessity for the feasibility of the approach, the use of the L-skewness demonstrates the main advantage of using a simulated calibration: that additional statistics can be included in the calibration that would otherwise be beyond derivation for the specified model.

The results demonstrate that the simulation method gave parameter sets of comparable quality to parameters optimised using analytic equations, albeit marginally poorer in quality. Some of the calibrated parameters had similar values for both methods, while others showed significant contrast. The ability of the model to yield parameter sets of similar quality yet being from different neighbourhoods is not an artefact of the proposed simulation methodology but of the model.

The main disadvantage of the MCS calibration method is that significant computational resources are required. For this reason the MCS technique used here is not applied to further case-studies due to the necessity of more computationally demanding formulations and simulations over larger regions. Nonetheless, it provides a convenient platform for extending the model under different assumptions or using different distributions without the need to analytically rederive model properties. Also, should future users wish to introduce additional statistics (for example, dry portions above a certain threshold) this method is useful.

Chapter 7
Spatial Storm Extent

7.1 INTRODUCTION

The aim of this chapter is to address the cross-correlation statistic and spatial description of storms over relatively large areas (say $>$ 100 km $\times$ 100 km). This is achieved by characterizing the extent of a storm, which has not yet been considered as part of the SNSRP model due to smaller scale case studies in urban regions.

7.2 PROBABILITY OF STORM OCCURRENCE AT A POINT

In this section a superscript 'o' denotes the properties and parameters specific to the original formulation outlined in Chapter 3 to distinguish it from the extended model developed here. The SNSRP model performs well for small regions, but if the region is large, unrealistic cross-correlations at large distances occur as an artifact of the storm not having a defined spatial extent. To address this model limitation, an additional random variable, R_s, is introduced to describe the radial extent of a storm, as shown in Figure 7.1 (b). This radius is assumed to have an exponential distribution with parameter ϕ_s. The radius is defined with respect to a storm center, Z_s, which has coordinates that are uniformly distributed across the region. A parameter λ controls the Poisson arrival of storms. Whereas λ^o was previously the temporal rate of storms, λ is the spatio-temporal rate of storms (since storms no longer cover the entire region).

By introducing a storm extent the rainfall process is not independent of the storm center. That is, the storm boundary localises the rainfall process relative to the center of the storm. Nonetheless, the expected model properties averaged over many storms are independent of location due to the spatially uniform location of storm centers. In order to avoid boundary effects, simulation of storm centers should include those that fall outside target region yet have some portion of the storm that intersects it. The rainfall inside the

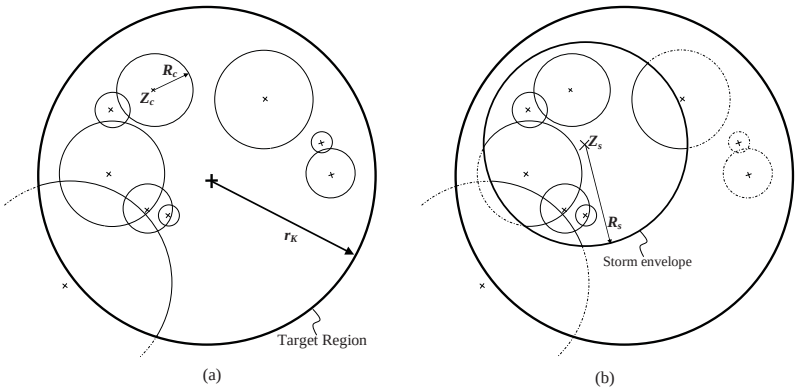

Figure 7.1 (a) Schematic diagram of Neyman-Scott process depicting overlapping rain-cells at a given point in time and with cell centers occurring as a two-dimensional Poisson process over some target region. (b) A storm extent is introduced so that cells and portions of cells landing outside of it are zero (indicated by dashed lines). Cells and storms with centers outside of the region also need to be included to avoid boundary effects.

storm boundary must maintain the uniform expected properties of the underlying rainfield. As a consequence, cells having centers outside of the storm boundary, yet having radial extent intersecting it, must be included. Also, any portion of a cell that lies outside of the storm boundary contributes zero rainfall. These scenarios are depicted in Figure 7.1 (b), where dashed lines indicate the regions contributing zero rainfall.

It is necessary to derive the probability that a storm intersects a point conditioned on it also intersecting the target region. This probability is denoted ψ. To obtain an expression for ψ, first consider a storm of fixed radius r. For the storm to cover an arbitrary point within the region the storm center should land within a distance r of that point (defining an area $A_1 = \pi r^2$), whereas for the storm to intersect the region its center should land within a distance $r + r_k$ of the regional center (defining an area $A_2 = \pi(r + r_k)^2$). The first criterion satisfies the second, since if a storm covers the point it also intersects the region. If the storm centers are uniformly distributed over the $\mathbb{R}^2$ plane, then the conditional probability ψ is the ratio of the A_1 to A_2, $\psi = \pi r^2/(\pi(r + r_K)^2)$. If the storm radius is a random variable that follows an exponential distribution, the probability is taken as the expected ratio of areas across all possible values of the storm radius, r. This is given as

$$\psi = \int_0^\infty \phi_s e^{-\phi_s r} \frac{r^2}{(r + r_K)^2}\, dr \tag{7.1}$$

where $\phi_s e^{-\phi_s r}$ is the probability density of a storm radius r, $r > 0$.

Using symbolic integration software, Eq. (7.1) is found to be equivalent to

$$\psi = 1 + \omega - \omega^2 e^{\omega} E_1(\omega) - 2\omega\, e^{\omega} E_1(\omega) \tag{7.2}$$

where $E_1(x)$ is the first order exponential integral function [*Abramowitz and Stegun*, 1964] and $\omega = \phi_s r_K$ which is the ratio of the radius of the target region to the mean storm radius.

From Eq. (7.1), substituting $r_K = 0$ returns the probability $\psi = 1$, which gives the case of storms covering a region that is a point. Increasing the region size, $r_K \to \infty$, shows that the conditional probability of a storm covering a point tends toward zero, $\psi \to 0$. Different shaped regions, such as a rectangle, can be derived by substituting the appropriate expression for the region area into the denominator of Eq. (7.1).

Because of the finite size of storms, when a storm occurs it does not cover every point within the region. The parameter ψ has the effect of reducing the rate of storms occurring over the entire region to give the rate of storms that overlap a single point only. If the arrival rate of storms over the region is λ (storms per hour per unit area), the rate of storms at a point is related by $\lambda^o \equiv \lambda\psi$. For a fixed target region, increasing the storm radius so that it always covers the entire region $\phi_s \to 0$ gives $\psi \to 1$, hence the two rates become equivalent (which is the case of the original model). As another example, assume the original model is calibrated with a storm rate λ^o (per hour, where storms implicitly cover the whole region) and the storm extent is later calibrated such that on average storms only cover half of the region $\psi = 0.5$. Then, twice as many storms need to be simulated over the region to obtain the equivalent temporal properties at a point $\lambda = 2\lambda^o$.

7.3 MODEL HAVING STORM EXTENT

This section provides a summary of the model properties incorporating storm extent, mainly emphasizing the differences with respect to the existing model. The instantaneous properties are discussed first and then the properties for the model aggregated at interval h are presented. The generalization in *Cowpertwait* [2004] holds for the extended model, though this is not explicitly demonstrated here.

The expected instantaneous rainfall intensity becomes

$$\mu_y = E[Y(t)] = \psi\mu_y^o \tag{7.3}$$

where ψ accounts for the fact that the storms do not cover the entire region. The

instantaneous covariance is defined as

$$C(\tau, d) = E\left[Y_1(t)Y_2(t+\tau)\right] - (\mu_y)^2 \tag{7.4}$$

where it is necessary to consider the expected value $E\left[Y_1(t)Y_2(t+\tau)\right]$. The expression for this is presented in Eq. (7.5) with three terms that are a superposition of the ways in which rainfall can cover the two points. Arbitrary overlapping of storms and cells is allowed since these processes are specified as rates. For example, if two cells c_1 and c_2 from the same storm cover points A and B respectively and a second storm contributes a cell c_3 that covers both points, this will contribute as cell c_1 and c_2 for as long as both cells are alive, c_1 and c_3 for as long as these two cells are alive, likewise for c_2 and c_3 and c_3 with itself. The result for all different combinations is given in Eq. (7.5). The first term is contributions from independent cells from independent storms, the second term is contributions from independent cells from a common storm and the third term is the contribution for the case of a common cell from a common storm.

$$\begin{aligned} E\left[Y_1(t)Y_2(t+\tau)\right] &= (\lambda\psi)^2\mu_x^2\mu_c^2/\eta^2 \\ &+\lambda\psi P(\phi_s, d)\mu_c^2\mu_x^2\beta\frac{\eta e^{-\beta\tau} - \beta e^{-\eta\tau}}{2\eta(\eta+\beta)(\eta-\beta)} \\ &+\lambda\psi P(\phi_c, d)P(\phi_s, d)\mu_c E[X^2]e^{-\eta\tau}/\eta \end{aligned} \tag{7.5}$$

All of the terms in Eq. (7.5) show that the storm rate λ is modulated by the probability of storm covering a point ψ. The first term corresponds to the quantity μ_y^2. The second term introduces the quantity $P(\phi_s, d)$ which is the conditional probability that given the storm covers one point it will also cover a second point a distance d away. The third term has the product of conditional probabilities $P(\phi_c, d)P(\phi_s, d)$ for storms and cells respectively. This occurs because the storm and cell processes are both exponentially distributed circular areas and because they are independent of each other. Substituting Eq. (7.5) into Eq. (7.4) gives

$$\begin{aligned} C(\tau, d) &= \lambda\psi P(\phi_s, d)\mu_c^2\mu_x^2\beta\frac{\eta e^{-\beta\tau} - \beta e^{-\eta\tau}}{2\eta(\eta+\beta)(\eta-\beta)} \\ &+\lambda\psi P(\phi_c, d)P(\phi_s, d)\mu_c E[X^2]e^{-\eta\tau}/\eta \end{aligned} \tag{7.6}$$

With the exception of $P(\phi_s, d)$ and ψ all other terms are the same as those in Eq. (3.12).

The properties that follow from the aggregated rainfall process are the aggregated mean intensity

$$\mu_h = E[Y_h(t)] = \psi\mu_h^o \tag{7.7}$$

and the covariance of aggregated intensities

$$\gamma_{d,h,l} = cov[Y_{1,h}(j)Y_{2,h}(j+l)] = P(\phi_s, d)\psi\gamma^o_{d,h,l} \tag{7.8}$$

The variance, autocovariance and coefficient of variation are temporal properties that can be obtained from these two equations. For the original model, all of the equivalent temporal properties can be obtained by substituting $\lambda^o \equiv \lambda\psi$. This follows naturally from the linear nature of expectations and is applicable to higher-order moments. This observation is applied in Section 7.4 to the skewness property which improves the fit to extreme values when calibrating the model [*Cowpertwait et al.*, 2002, Eq. (10)].

From Eq. (7.8) it is clear that the covariance of the storm-extent model is of an equivalent form but includes a multiplicative term $P(\phi_s, d)$. This term reduces the covariances according to the parameter ϕ_s. The covariance structure of the original model is obtained as a special case when the storm radius becomes infinitely large $\phi_s \to 0$ so that all points in the target region are covered by every storm. In the event of this and for all distances, d, $P(\phi_s, d) \to 1$. Figure 7.2 shows the influence of the storm extent parameter with respect to the cell radius parameter, where other parameters are held fixed at typical values. Figure 7.2 (a) shows that varying the cell radius parameter over typical scales 2 km $(\phi_c = 0.5)$ to 10 km $(\phi_c = 0.1)$ affects the cross-correlations over shorter distances, but the cross-correlations at large distances are not influenced significantly. Figure 7.2 (b) shows that the storm radius parameter over typical scales 20 km $(\phi_s = 0.05)$ to 100 km $(\phi_s = 0.01)$ affects the cross-correlations at larger distances.

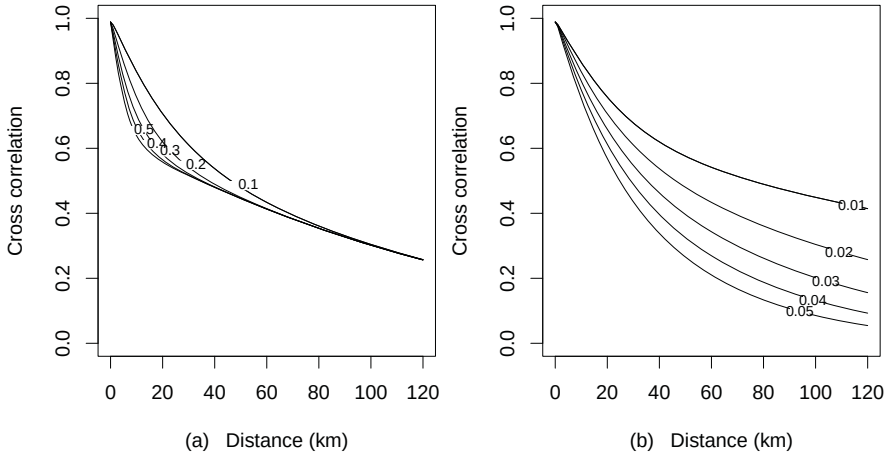

Figure 7.2 Example of variation in cross correlation at the 24 hour aggregate for two spatial parameters (other parameters fixed at typical values) (a) Cell radius parameter and (b) Storm radius parameter.

7.4 MODEL CALIBRATION

The original model allows cells to occur arbitrarily far apart and as a result independent cells arising from the same pattern of storms can cause a similar 'wet-dry' pattern at all points across the region. For this reason the original model causes points separated by a large distance to have spurious high correlations. This can be observed in Figure 7.3 for a model calibrated to the Sydney region for the month of January. The hollow circles are observed daily cross correlations, the dashed lines give the calibrated cross correlations of the original model and the solid line gives the cross correlations of the extended model. This figure shows that the correlation of the original model decays to a non-zero value for increasing distance, which is physically unrealistic. In comparison, the extended model which has an additional parameter to limit the storm extent allows for a more realistic decay. While it is possible to find parameters of the original model that improve the fit to the cross-correlation statistic, it is at the expense of other fitted statistics. This figure suggests that having calibrated the original model, the daily correlations are a suitable statistic that can be used to calibrate the parameter ϕ_s.

Additional steps that augment the existing calibration procedure are discussed below.

1. Having obtained a good fit to the temporal statistics, the calibration of spatial parameters is considered. The cell radius parameters $\phi_c^{(p)}$ are calibrated via least squares to the hourly cross-correlogram assuming storms cover the entire region. This is achieved by setting $\psi^{(p)} = 1$ and $P(\phi_s^{(p)}, d) = 1$ for all d. Use of several $\phi_c^{(p)}$ parameters for multiple storm types $p = 1, \ldots, n$ is considered relative to the sum of squares fit obtained using a common value of ϕ_c. Setting $\psi = 1$, the storm radius parameters $\phi_s^{(p)}$ are calibrated via least squares to the daily cross-correlogram. The length scale of $\phi_s^{(p)}$ will likely be an order of magnitude larger than $\phi_c^{(p)}$. For this reason, it will affect the cross-correlations at larger distances more significantly than shorter distances.

2. Having obtained $\phi_s^{(p)}$, the probability ψ is evaluated using Eq. (7.1) for a given region size. The equivalent arrival rate of storms in space and time λ is evaluated from the storm arrival rate of the original model using the relation $\lambda = \lambda^o/\psi$. This provides the full parameter set for simulating the model as a multi-site model. Interpolating the scale parameter allows for estimates at sites other than those used in calibration.

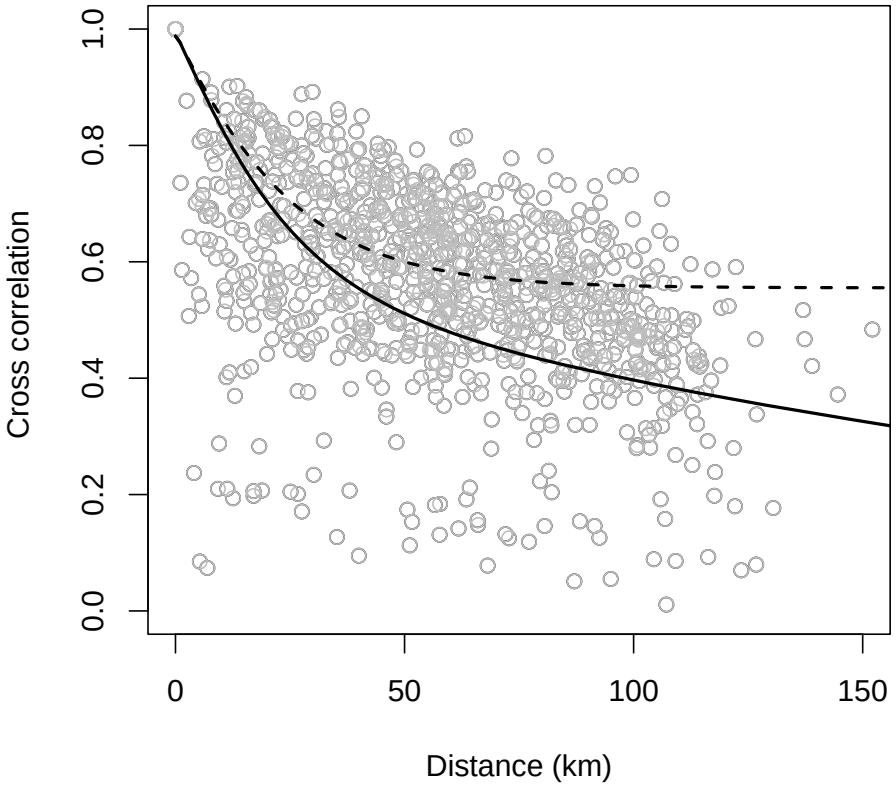

Figure 7.3 24-hour aggregate cross correlations for January where hollow circles are observed values, the dashed line is from the calibrated original model and the solid line is from the extended model having finite storm extent.

7.5 MODEL SIMULATION

As discussed in Chapter 4 a target region induces a bias when cells with centres landing outside of that region partially overlap the region. For several reasons this observation is also relevant to the storm envelope.

1. If cells that land with centres outside of a storm envelope (yet intersect the envelope) are not considered then there will be a reduction in the simulated rainfall total and also the properties of the storm will not be stationary. This undesired non-stationarity would be due to the fact that points closer to the outer boundary of the storm will have less chance of being overlapped by a raincell.

2. If storms that have origins landing outside of the region (yet intersect that region) are not included there will be a significant bias in the simulated totals.

3. If cells landing outside of the target region and belonging to storms that have origins

outside of the target region (where both intersect the region) are not included, a simulation will incur similar biases.

For a mean storm radius $\phi_c^{-1} = 100\ km$ the 0.95 quantile storm has a radius of $300\ km$. For this reason a buffer algorithm that includes all storm origins outside of a region is extremely inefficient. The direct simulation algorithm proposed for raincells in Section 4 is also valid for storms since they have the same theory underpinning them: a circular area having an exponentially distributed radius. Not only is the algorithm valid for cells and storms relative to the target boundary, it is also valid for cells relative to the storm boundary. Because of this an additional efficiency can be included in the simulation algorithm since there will be occasions in which the storm envelope is smaller than the target region and it is thus quicker to simulate cells relative to the storm. This statement is true even if the storm only partially overlaps the target region since the number of discarded cells simulated over the target region yet landing entirely outside of the storm would be higher than the number of cells over the storm area landing outside of the target region. Table 7.1 illustrates an efficient algorithm for simulating raincells.

Table 7.1 Algorithm for efficiently simulating storm and cell properties.

1.	Sample the number of storm origins occurring inside the region.
2.	Sample the number of storm origins occurring outside the region using Eq. 4.10.
3.	For all storms sample the storm-centre and storm radius, R_s.
4.	If $R_s < r_K$ then simulate all cells relative to the storm boundary. This includes all cells landing inside and outside of the storm. Discard cell portions that lie outside the storm envelope or the region boundary.
5.	If $R_s > r_K$ then simulate all cells relative to the region boundary. This includes all cells landing inside and outside of the region. Discard cell portions that lie outside the storm envelope or the region boundary.
6.	Simulate all other cell properties.

7.6 CASE-STUDY APPLICATION

7.6.1 Modelled region and calibrated parameters

To demonstrate the application of the model, the greater Sydney region was used as a case study. Two different data sources were used in the calibration: 24 pluviograph gauges of length greater than 20 years, and 52 daily gauges of length greater than 90 years. The spatial coordinates of these gauges cover an area approximately $100\ \mathrm{km} \times\ 140\ \mathrm{km}$

in size and are shown in Figure 7.4. A circular region of radius $85\ \mathrm{km}$ was used to simulate the model covering all of the gauges within this region. Calibrating the model proceeded on the basis that the simulation area was homogeneous, having a constant coefficient of variation and constant dry-probability. This requirement was not strictly met by the rainfall data in this region and the effect of this will be noted when comparing the model to observed data. For each site and for each month, the rainfall data were non-dimensionalised by dividing by the hourly mean and pooled statistics were computed for three aggregates, $1\ \mathrm{hour}$, $6\ \mathrm{hours}$ and $24\ \mathrm{hours}$. The temporal parameters were calibrated to temporal statistics including the coefficient of variation, lag-one auto-correlation and skewness. The cell radius parameter was calibrated to hourly cross-correlations, and daily cross-correlations were introduced for the calibration of the storm radius parameter.

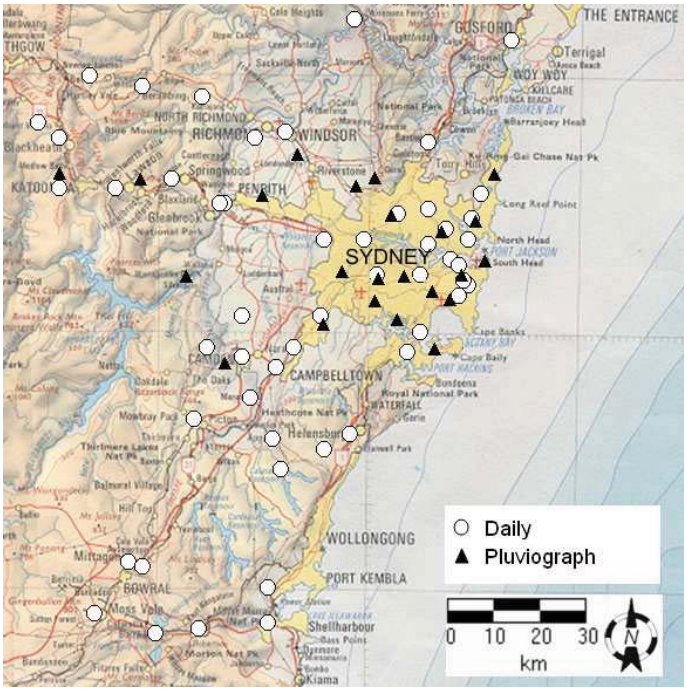

Figure 7.4 Coordinates of rainfall observation sites about greater Sydney region

Care needs to be taken when calculating the observed statistics to avoid biases arising from missing and accumulated data. The daily rainfall data contain numerous instances of accumulated totals where the total rainfall is known over some multiple of days but not the pattern in which it occurred. On average, 0.1% of the rainfall measurements were accumulations, with the worst-case site having 2.8% of its observations accumulated. The accumulated amounts are retained when calculating the mean rainfall and annual and

monthly totals, but are excluded from statistics such as the covariance which require the rainfall pattern within the accumulation to be known. The hourly rainfall data do not have accumulations, but do have missing observations. On average only 0.03% of the observations of the hourly data were missing, with the worst site having 0.07% missing data. These amounts are unlikely to significantly bias the estimated statistics.

For each month the model was fitted for two cases, only one storm type ($n = 1$) and a mixture of two storm types ($n = 2$). This decision was motivated by a difficulty obtaining adequate fits and realistic parameter values for several of the months. For the months having two storm types, reductions in the sum-of-squares varied between 60% and 97% when compared to using only one storm type.

Table 7.2 shows the calibrated parameters, with parameter sets having two types of storm processes selected for seven of the months. To avoid over-parameterization, the parameters $\beta^{(1)}$, $\mu_c^{(1)}$, $\phi_c^{(1)}$ and $\phi_s^{(1)}$ were made common to both storm types. An exception was made for the case where the second storm type reduces to a Poisson process, $\mu_c^{(2)} = 1.0$, in the months from May to August. Of these months, for May and June, the second storm type has a higher rate of storm occurrence but with shorter cell lifetimes than the first storm type. For July and August, the second storm type has less frequent storms with long cell lifetimes. The infrequent and long duration nature of these storms required the parameter $\phi_c^{(2)}$ to be re-estimated resulting in cells that have wide-spread coverage. This decision was motivated by an inability to characterise the cross-correlations using only $\phi_c^{(1)}$ for these months.

7.6.2 Interpolation of Scale Parameters

The spatial coordinates of these gauges cover an area ranging from Gosford to Wollongong (North-South extent), and from Bondi Beach to Katoomba (East-West extent). The elevation across the region is undulating from the coastline to approximately 150.6° longitude, then rising to $1000\,m$ above sea level due to the Blue Mountains. The mean annual rainfall over the region is $914\,mm$, and the monthly rainfall is slightly higher during summer months (up to $90\,mm$) and lower during winter months (down to $60\,mm$). The distribution of rainfall over the region varies with respect to season and elevation, as shown in Figure 7.5 for two example months, January and July. Figure 7.5 (a) and Figure 7.5 (b) show that January monthly totals (summer) are higher over the Blue Mountains but the July mothly totals (winter) get higher moving toward the coastline. This non-stationarity in the mean rainfall over the region is removed from the record by non-dimensionalising the data with respect to the mean hourly rainfall for each site and month.

Table 7.2 Parameter estimates for each month

$Storm^{(1)}$	$\hat{\lambda}^{o(1)}$	$\hat{\beta}^{(1)}$	$\hat{\eta}^{(1)}$	$\hat{\mu_c}^{(1)}$	$\hat{\alpha}^{(1)}$	$\hat{\phi_c}^{(1)}$	$\hat{\phi_s}^{(1)}$
Month	h^{-1}	h^{-1}	h^{-1}	-	-	km^{-1}	km^{-1}
1	0.00300	0.0578	1.79	24.4	0.571	0.112	0.010
2	0.00406	0.0630	1.17	23.5	0.786	0.120	0.015
3	0.00301	0.0803	2.04	36.4	0.562	0.141	0.012
4	0.00472	0.0616	2.95	77.8	0.391	0.166	0.017
5	0.00083	0.0490	0.23	39.2	0.500	0.147	0.010
6	0.00035	0.0539	1.41	60.0	0.878	0.154	0.014
7	0.00145	0.0083	2.13	49.1	0.408	0.100	0.012
8	0.00493	0.0263	1.60	19.1	0.400	0.140	0.011
9	0.00435	0.0875	2.30	29.1	1.00	0.168	0.017
10	0.00213	0.0942	2.07	97.2	0.423	0.139	0.016
11	0.00274	0.0587	1.99	44.4	0.481	0.131	0.013
12	0.00275	0.0587	2.51	20.7	0.725	0.134	0.010

$Storm^{(2)}$	$\hat{\lambda}^{o(2)}$	$^*\hat{\beta}^{(2)}$	$\hat{\eta}^{(2)}$	$^*\hat{\mu_c}^{(2)}$	$\hat{\alpha}^{(2)}$	$^*\hat{\phi_c}^{(2)}$	$^*\hat{\phi_s}^{(2)}$
Month	h^{-1}	h^{-1}	h^{-1}	-	-	km^{-1}	km^{-1}
2	0.00054	*0.0630*	2.15	*24.4*	1.927	*0.120*	*0.015*
4	0.00059	*0.0616*	1.71	*23.5*	0.592	*0.166*	*0.017*
5	0.05121	*0.0490*	2.81	1.0	0.694	*0.147*	*0.010*
6	0.03236	*0.0539*	1.68	1.0	2.684	*0.154*	*0.014*
7	0.00136	*0.0083*	0.06	1.0	1.057	0.01	*0.012*
8	0.00098	*0.0263*	0.06	1.0	1.185	0.01	*0.011*
9	0.00040	*0.0875*	2.74	*29.1*	2.092	*0.168*	*0.017*

* Values shown in italics have parameters fixed to be the same as storm 1

To obtain the non-stationary simulated totals for all points in the target region, the hourly rainfall totals used to non-dimensionalise the rainfall must be interpolated. Figure 7.5 (c) and Figure 7.5 (d) show the interpolated mean hourly rainfall totals across the entire Sydney region. The values were interpolated using ordinary kriging with an exponential covariance function fitted to the observed data. Every image for a given month produced by a simulation is multiplied by the respective monthly scaling field to give the corrected rainfall totals. A nugget effect was excluded from the kriging estimation so that

the interpolation field exactly matches the measured depths at the observation sites. The scaling field can appear 'peaked' because of this, especially when the exponential covariance decays quickly. Despite this, a benefit of this approach is that the mean rainfall total is preserved in the simulated statistics because there is no smoothing of the scaling field.

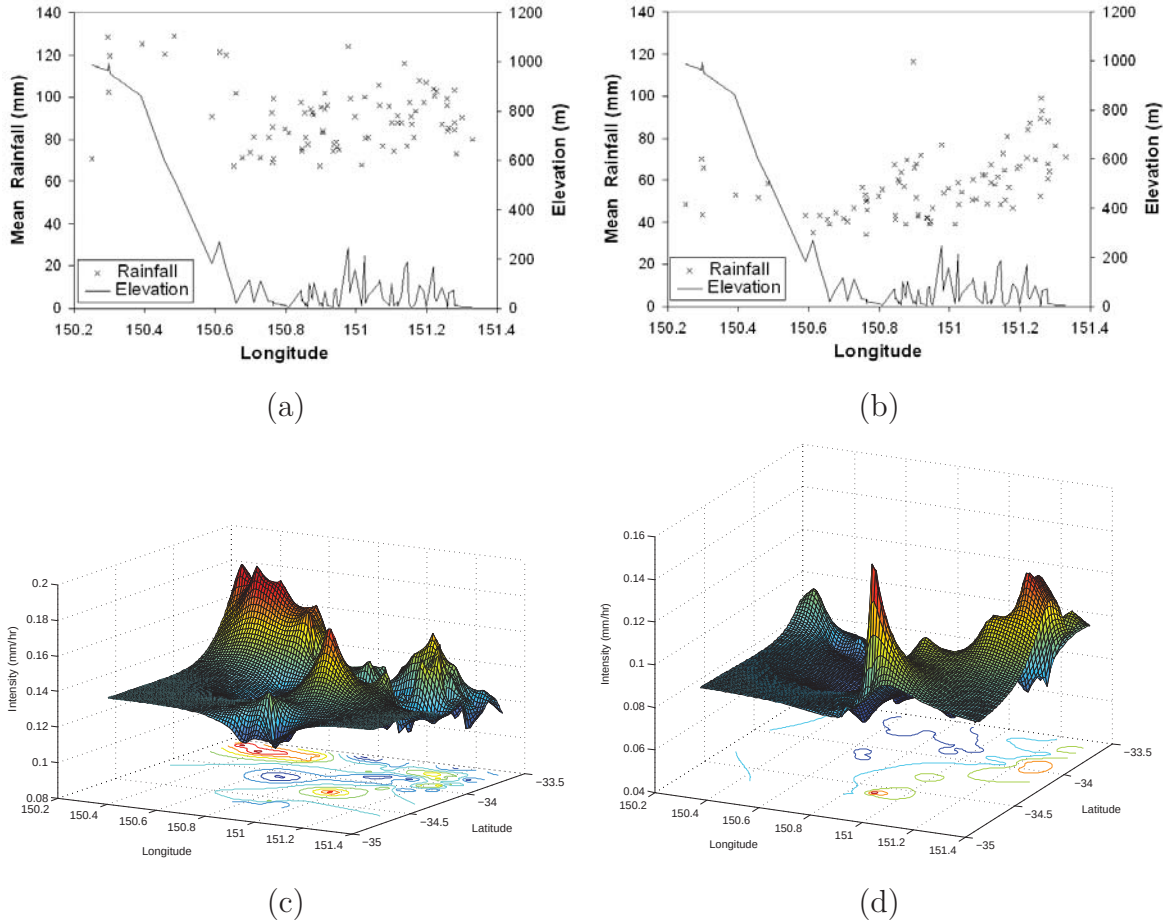

Figure 7.5 Elevation and variation in mean rainfall across Sydney (a) January monthly total (b) July monthly total (c) January interpolated scaling field (d) July interpolated scaling field.

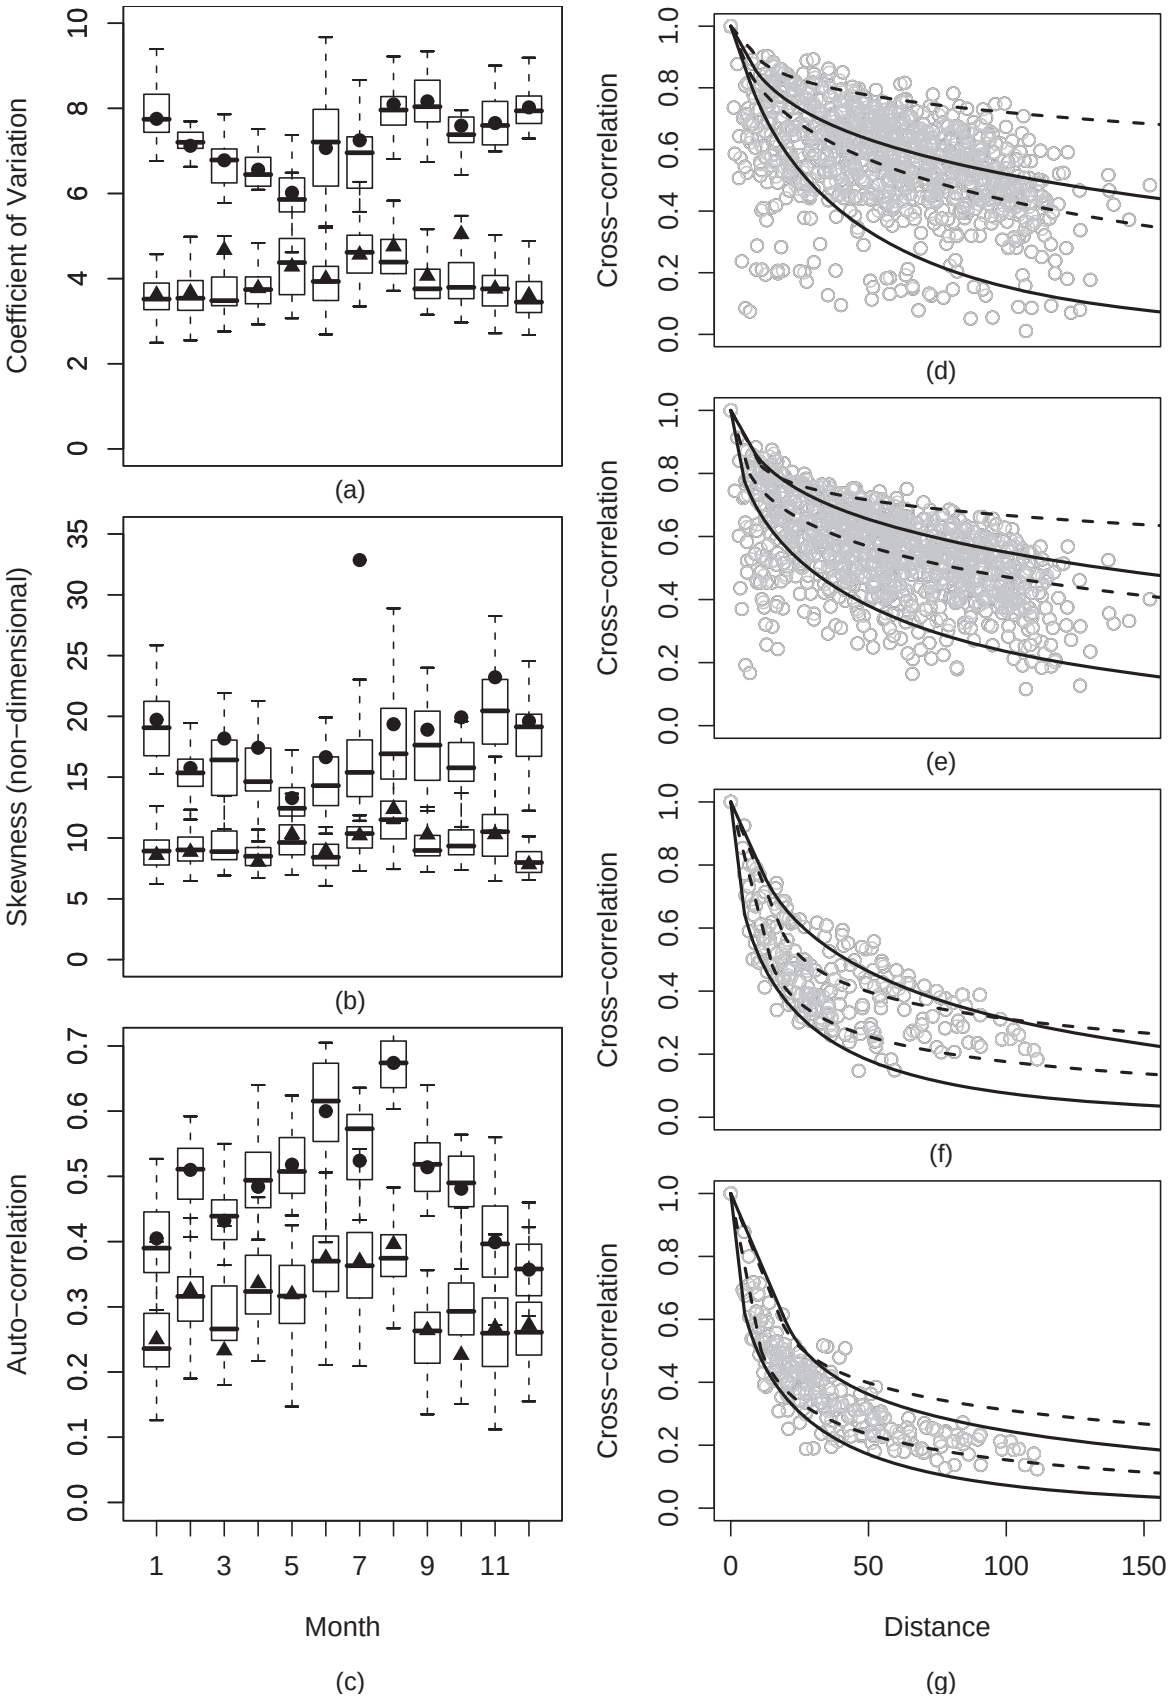

Figure 7.6 Caption footnoted on page 100.

7.6.3 Model Comparison to Calibration Statistics

Figure 7.6 shows a summary of the statistics that were used to calibrate the model[1]. Simulated results are also provided in Figure 7.6 for comparison. The model was simulated for a length of 95 years which is the median length of the observed daily rainfall records used in calibration. Where comparisons are made with sub-daily rainfall the simulation length was truncated to 24 years, which is the median length of the sub-daily records used in calibration. The variability in selected statistics was demonstrated using 100 replicates.

Figures 7.6 (a), 7.6 (b) and 7.6 (c) show comparisons for the temporal statistics used in calibration, respectively these are the coefficient of variation, skewness and auto-correlation. The statistics are shown for two aggregates (1 hour, 24 hours) where the symbols give the observed value and the box-plots show the variability in the simulated property for each month. The 6 hours statistics are not shown but have similar agreement.

Figures 7.6 (d) and 7.6 (e) show the fitted 24 hours cross-correlations with respect to two example months, January and July respectively. In these figures the dashed lines give 90% confidence limits for the original model and the solid lines give 90% limits for the extended model having the storm parameter. Variability in the simulated cross-correlation is due to the sample size of 95 years (as the record length increases the modelled cross-correlations collapse onto the line given by Eq. (7.8)). These figures show that the storm parameter reduces the cross-correlations for increasing distance. In these figures the hollow symbols show the observed cross-correlations which have much greater variability than what the model is able to reproduce. This feature is directly related to heterogeneity within the region. More specifically it is related to variability in the dry probability and that some sites have a higher frequency of rainfall than others. The geographic influences giving rise to this are in-part related to the distance from the coastline and the elevation, but the dry portion of the simulated record is the same for all sites since scaling the mean does not affect the number of zeros.

Figures 7.6 (f) and 7.6 (g) demonstrate the calibration of the model to 1 hour aggregate cross-correlations for two example months, January and July respectively. These figures show that the influence of the storm radius parameter is less significant at this time-scale. The observed hourly correlations are not as variable as the daily correlations but still show

[1]Figure 7.6 (page 99) Simulated values for statistics used in calibration. (a)-(c) temporal statistics, circles give observed 1-hour statistic, triangles give 24-hour statistic and boxplots give simulated values; (d)-(g) cross-correlations where hollow circles give observed values, dashed lines give 90% limits from 100 simulated replicates having storm extent covering the whole region, solid lines give 90% limits of replicates having finite storm extent. (d) 24-hour aggregate, January, (e) 24-hour aggregate, July, (f) 1-hour aggregate, January and (g) 1-hour aggregate, July.

more variability than the simulated replicates.

7.6.4 Model Comparison to Other Statistics

Figure 7.7 shows a selection of statistics not used to calibrate the model which are used to assess its performance[2]. The assessment here is by no means exhaustive as there are numerous statistics that can be used, a variety of aggregates and a large number of points within the region at which they should be assessed. Figures 7.7 (a) and 7.7 (b) show a comparison of cross-correlations of monthly totals for one simulated replicate and the observed data, for two example months January and July respectively. Unlike the daily cross-correlations these figures show that the model captures the variability of the observed data well.

Figure 7.7 (c) shows the at-site variability in the observed and simulated annual totals pooled across the region. Each site is standardised by dividing the distribution with respect to the median value, the distributions are then collated across all sites to generate a single regional distribution. The simulated values show the 0.05, 0.50 and 0.95 quantiles for each order statistic to indicate the variability between simulated replicates. Figure 7.7 (c) shows that the annual mean is matched, which is expected since the modelling methodology uses scaling values to force the simulated and observed mean to match at each site. However, Figure 7.7 (c) shows that the variability of the annual aggregate is significantly under-simulated. There are several reasons for this, (i) discrepancies in variability at the monthly aggregates, (ii) the method of independent monthly calibration, which removes any variability due to rainfall cross-correlated between months, (iii) heterogeneity of observed statistics such as coefficient of variation and (iv) non-stationarity of the observed annual rainfall distribution due to climatic states not accounted for in the modelling. The last point is well documented in literature, for example see *Srikanthan and McMahon* [2001]. Figure 7.7 (d) reinforces the observation made in Figure 7.7 (c) by comparing the standard deviation of the observed annual totals with respect to the standard deviation of one simulated replicate of annual totals at all sites. This figure shows that the annual variability is consistently underestimated at all sites within the region.

[2]Figure 7.7 (page 102) Cross-correlation of monthly totals, observed (black) and simulated (grey): (a) January (b) July. (c) distribution of annual totals standardised and pooled over the region, lines give 0.05, 0.5 and 0.95 simulated order statistics. (d) comparison of simulated and observed standard deviation in annual totals at 52 sites. Distribution of annual maximum values, lines give 0.05, 0.5, 0.95 simulated order statistics: (e) 1 hour aggregate, one maximum taken from entire region per year, (f) 24 hour aggregate one maximum taken from entire region per year, (g) 24 hour maximums at a selected site, (h) 24 hour maximums at another selected site.

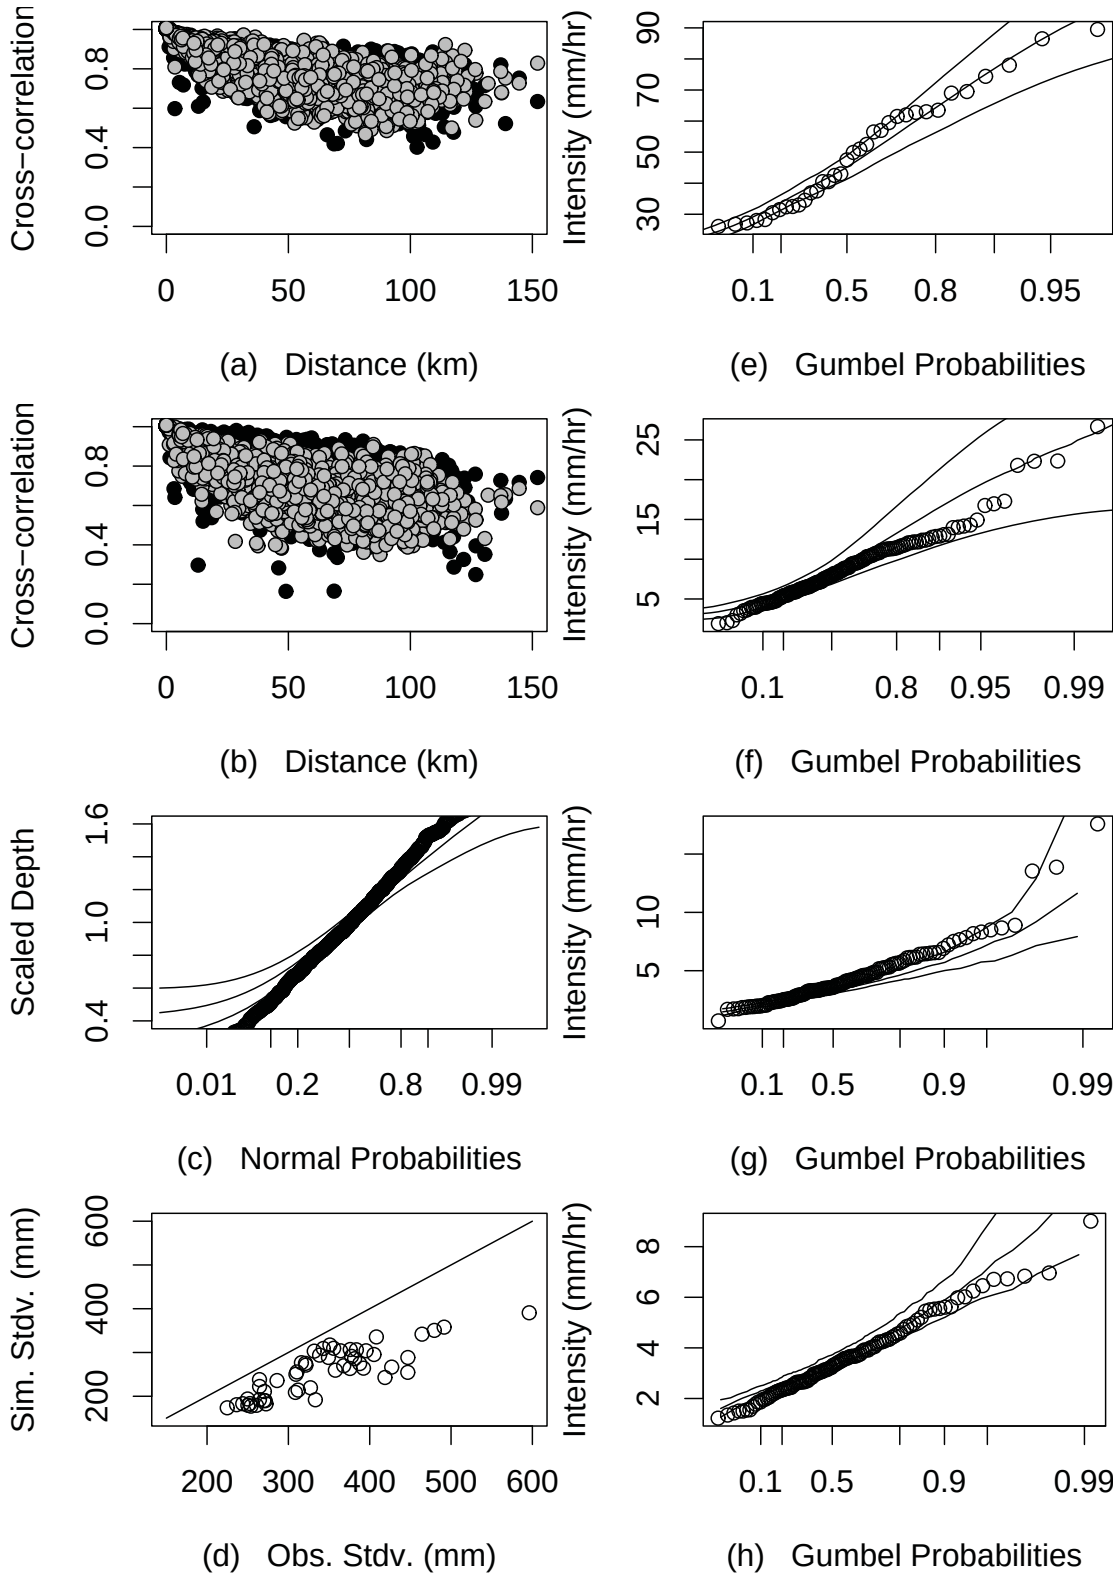

Figure 7.7 Caption footnoted on page 101.

Figures 7.7 (e) and 7.7 (f) compare the simulated and observed annual maximum frequency curves for the overall region at two aggregates, 1 hour and 24 hours respectively. For both aggregates the absolute maximum value within the region was selected for every year on record. There is the possibility that this statistic is not a good diagnostic of the model's regional performance if annual maximums consistently come from only a limited number of sites. This scenario was checked and is not the case: 35 of the 52 daily gauges contribute points to the distribution of annual maximums and 14 of the 24 sub-daily gauges contribute points to the distribution of annual maximums. Both of these figures show reasonable agreement to the observed data. Nonetheless, considering the region as a whole has the possibility of masking discrepancies at individual sites within the region. Figures 7.7 (e) and 7.7 (f) demonstrate this by comparing the observed and simulated annual maximums at two selected sites. Figure 7.7 (e) shows that the model under-simulates the maximums at one site and Figure 7.7 (f) gives an example at another site where the model over-simulates the observed maximums. This observation highlights the difficulty of assessing the performance of the model and the difficulty in obtaining a calibrated model that has a good overall fit to the region as well as for all sites in the region See Appendix A for comparisons of annual totals and extreme value statistics at each site.

7.7 VISUALISATION OF SPATIAL STORM PROCESS

The SNSRP model has a reasonable ability to reproduce temporal statistics due to its heritage as a point process model. By way of contrast, the spatial characterisation of the storm is more simplistic, which is in-part due to the limited spatial information-content of sparse gauge networks. To illustrate the spatial performance of the model, a simulated replicate of 12 hour duration is shown in Figure 7.8 over the region of Sydney. Alphabet letters are marked onto the figure to show examples of raincells that persist for more than one hour. The examples used are not exhaustive, as close inspection of the figures will reveal more cells that persist over multiple hours. The storm was simulated using the May set of parameters, which have the mean cell lifetime as $\eta^{-1} = 0.36\ hr$. As a result, only 6% of cells will persist for greater than one hour (thus appearing in at least two images) and only 0.4% of cells will persist for more than two hours (thus appearing in three images).

The spatial aesthetic of the storm is dominated by the circular shaped rain-cells. Nonetheless, numerous other shapes are generated by the storm that obscure the underlying circular structure. These are due to several processes (i) the non-stationary scaling

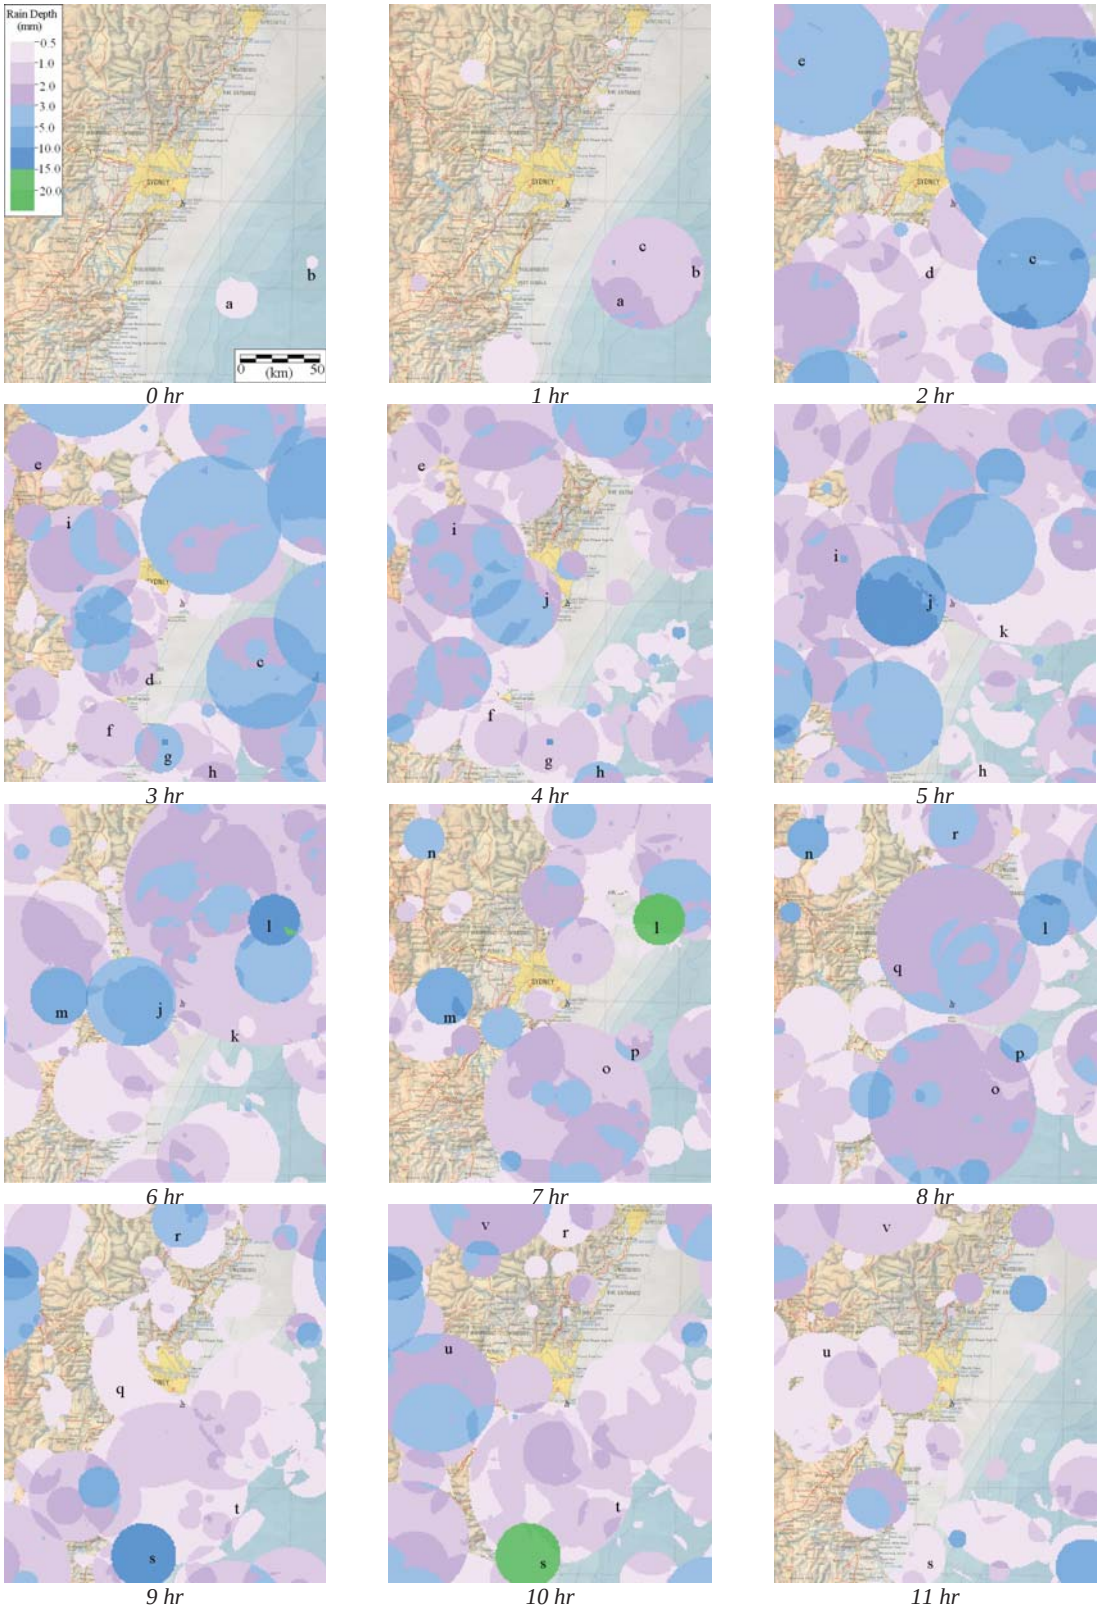

Figure 7.8 12-hour visualisation showing rainfall depths over Sydney of a storm in May. Marked letters provide examples of raincells that persist over multiple hours.

Figure 7.9 12-hour visualisation showing rainfall depths over Sydney of a storm in May. The envelope localises rainfall in the spatial domain.

field (as shown in Figure 7.5), (ii) overlapping raincells and (iii) the threshold levels used to produce the plots. The scaling field is a smooth and variable surface that follows the mean trend in the data. This shows in the images as curved lines that do not appear to follow any particular circular arc. One example is cell 'a' in the $0\ hr$ image which appears dented on one side since the scaling field has modulated the rainfall total at those points to a value less than 0.5 (the lowest threshold). The overlapping of raincells produces numerous shapes that appear as cell portions. For low numbers of cells the location of cell origins and radii that contribute to the overlap region are obvious (cf. cell 'a' cell 'b' and cell 'c'). Where a high number of cells overlap the resulting shape can be obscure and difficult to determine how many underlying cells are contributing or whether the shape is instead due to the scaling field (cf. points located about cell 'j'). The threshold levels also create an artefact in the plot as the overlap may not be evident unless the total rainfall is above the next specified threshold level. Multiple examples appear where a raincells seems to 'disappear', which is because it has fallen below the lowest threshold used for plotting. Some cells that persist over multiple images appear to grow and decay in their intensity (cf. cell 's'), where the model instead specified the intensity as a constant 'pulse'. This observation is however an artefact of the temporal aggregation into hourly 'bins'. For the example of cell 's', it overlaps three hourly bins (at $9\ hr$,$10\ hr$ and $11\ hr$). The cell entirely covers the interval at 10 hours but only partially covers the 9 hour and 11 hour intervals. As the rainfall depth is significantly lower in the latter interval it can be concluded that the cell did not persist for a long duration into this period.

Figure 7.9 shows the same storm as Figure 7.8, but including a storm envelope simulated south of metropolitan Sydney and having a radius of approximately $80\ km$. It can be seen that the storm envelope causes the rainfall to fall within a specific domain rather than over the entire catchment region. The rigid boundary of the storm is evident when there is a high number of raincells, but when there is a lower number of cells it is less obvious to detect. Comparing Figure 7.9 with Figure 7.8 shows that it is necessary to include rainfall contributions from cells occurring outside of the storm envelope yet have some extent that intersects the envelope. Also, when the size of the storm envelope is small relative to the region it can be seen that a high number of raincells are discarded from the underlying simulation of raincells, since the majority do not intersect the storm.

Visualising storm profiles in space is a useful analysis tool for understanding the behaviour of the model. While the spatial idealisation of the model may appear crude there are several counter-points that need to be reiterated. Firstly, rain-gauges are sparsely distributed in space (especially at sub-daily scales) and as such they are restrictive in the

amount of spatial detail that can be reproduced from this method of observation. Even if 100 sub-daily gauges were available this should be considered sparse when compared to a radar that can produce a 100×100 grid of observations at a $1\ km$ resolution. Secondly, a significant benefit of gauges is that they have significantly less bias and sources of error when compared to radar-rainfall estimates. For this reason, a simple gauge-based model may be more appealling than a sophisticated radar-based model that has bias induced by measurement errors. Thirdly, lumped catchment models integrate rainfall over the spatial domain to give runoff estimates into one-dimensional river reaches. The process of integration will, to a large degree, mask unwanted effects of simplifications in the spatial domain. This point is however largely untested due to the lack of studies using continuous simulation space-time models with distributed runoff models. An important future case-study would assess the performance of gauge-based and radar-based models using a runoff model and compared against observed streamflows.

7.8 CONCLUSION

This chapter demonstrated an improved Neyman-Scott rainfall model so that between site cross-correlations are modelled more realistically. Previously, separate rain cells from one storm could cover points in space that would not realistically be covered by one storm. An approach was developed to remedy this by letting the storm extent follow a circular region and have a random radius following the exponential distribution. It was shown that the properties of the extended model modulate the properties of the previous model with respect to two probabilities: ψ the probability that the storm covers a point given that it intersects the target region, and $P(\phi_s, d)$ the probability that the storm covers point A given that it covers point B a distance d away. The model was calibrated to 52 daily gauges and 24 sub-daily gauges from the metropolitan Sydney region according to various statistics at three aggregates, 1 hour, 6 hours and 24 hours, and allowing for multiple storm types. Results were presented for the statistics used in calibration and for several statistics not used in calibration. Statistics not used in the calibration were the cross-correlations at a monthly aggregate, the distribution of extreme values and the distribution of annual totals. Although the model is highly idealised it provides a parsimonious representation relative to other models and it can generate long, continuous realizations of rainfall in time for any number of points in space.