High Resolution Space -Time Modelling of Rainfall: The String of Beads Model

by

Antony Neil Clothier

Submitted in partial fulfilment of the
requirements of the degree of
Doctor of Philosophy in Engineering,
Civil Engineering,
University of Natal, Durban

1st September 2003

Abstract

The purpose of this study was to develop a rainfall model, continuous in space-time, which captures both the spatial and temporal structure of rainfall over a range of scales varying from 1km to 128km pixels at temporal resolutions ranging from 5 minute up to 1 year. Such a model could find application in a variety of hydrological fields including the management of flash flood scenarios where it could be used in combination with runoff models as a training tool in the operation of flood control structures, the assessment of flood risk, the management of water resources in an area through the simulation of long rainfall sequences and as a short term rainfall forecasting tool, to name a few.

The *String of Beads* Model (SBM) is a high-resolution space-time model of radar rainfall images. It is a stochastic model that takes advantage of the detailed spatial and temporal information captured by weather radar and combines it with the long term seasonal variation captured by a network of daily raingauges.

The alternating wet-dry process, or event arrival and duration, is modelled as a one-dimensional process, while the detailed wet process is modelled as a three-dimensional (two space and one time) process at 1km, 5 minute spatial and temporal resolutions respectively, over an area of 16000km^2, consistent with the observed radar data. The three-dimensional rainfall events distributed on a one-dimensional time line, is analogous to a "String of Beads".

The SBM makes use of a combination of power law numerical filtering techniques and well-known time series models to achieve an efficient algorithm that can be run on an ordinary personal computer. Model output is in the form of image files which, when viewed as an animated sequence, are difficult to distinguish from observed radar rainfall images. Apart from the realistic appearance of these images, when calibrated to daily raingauge data for the region, analysis of the simulated sequences over periods of up to ten years, reveal convincing rainfall statistics for a wide range of spatial and temporal scales. It can be used both as a simulation tool and as a short term forecasting tool.

In simulation mode, it can quickly produce long sequences (tens of years) of 128 x 128 km rainfall images at five minute, one kilometre resolution. Such simulations can be used as input to distributed and semi-distributed hydrological models to produce "what if" scenarios for applications in water resources management and flood risk assessment amongst others.

In forecasting mode, the SBM has proved effective in producing real time forecasts of up to two hours making it a useful tool for flood warning and management, particularly in steep or urban catchments which react quickly and often give rise to flash floods. It can also be used in a combined simulation-forecasting mode to quickly produce many short term "what if" scenarios which can be used to assess the risk of possible growth or decay scenarios in real time.

Preface

This thesis is compiled entirely by myself, Antony Neil Clothier, under the supervision of Professor G G S Pegram at the University of Natal, in partial fulfilment of the degree of Doctor of Philosophy in Engineering (Civil).

The University of Natal offers a mechanism whereby it is possible to upgrade a study from Masters status to Doctorate status. In order to take advantage of this opportunity, a Masters thesis is submitted and graded and then it is banked by the faculty and the study upgraded to Doctorate level with proposed additional content. The Masters degree is then awarded, but not conferred. On completion of the study, a new thesis is submitted. This Doctorate thesis can contain some or all of the content of the Masters thesis. If the Doctorate is awarded, the Masters thesis becomes void as does the Masters Degree.

For this reason, part of the work was successfully submitted in 1998 as a dissertation in fulfilment of the degree of Master of Science in Engineering and the upgrade to Doctorate level was recommended. In addition, the contents of the masters thesis were published with very little alteration in reports to the Water Research Commission WRC K5/752 (Pegram and Clothier, 1999) and WRC K5/1010 (Clothier and Pegram, 2001) as well as in papers published in the Journal of Hydrology (Pegram and Clothier, 2001a) and Hydrology and Earth System Sciences (Pegram and Clothier, 2001b) *without reference* to the Masters thesis, since it will become void if the Doctorate is awarded. As a result, there are sections of this thesis which are almost identical to sections of these publications. Reprints of the journal publications are included in the final pages of this document.

Antony Neil Clothier Date:.....

As the candidate's supervisor I have/have not approved this thesis/dissertation for submission.

Professor Geoffrey GS Pegram Date:.....

Table of Contents

Acknowledgements

The author gratefully acknowledges the following organisations and people:

Geoff Pegram for his inspired supervision and guidance throughout the project, for the encouragement he offered in times of despair and for being so receptive to different ideas, but most of all, for providing the opportunity for me to take part in the development of his vision,

Sue Clothier and Pat Clothier, my mother and father and two of my dearest friends, for all the opportunities they have provided for me in my life. I would particularly like to thank my Mum for encouraging me to embark on this project and for her essential financial and emotional support throughout, especially after my Dad died in 1996,

Water Research Commission of South Africa for their generous financial support, facilitation over the past two years,

The WRC steering committee members:

Dr George Green (Chairman) of the Water Research Commission,
Mr Hugo Maaren of the Water Research Commission (deceased),
Dr Deon Terblanche of the South African Weather Services,
Prof. Bruce Kelby of the University of Zululand,
Prof. Dennis Hughes of Rhodes University,
Dr Mark Dent of Computing Centre for Water Research, University of Natal,
Mr Stefan van Biljoen of the Dept. of Water Affairs and Forestry,
Mr Steve Lynch of the University of Natal,
Mr Peter Visser of the South African Weather Services,

for their active interest in the development of the project and for the encouragement and valuable ideas that they offered during the steering committee meetings,

Acknowledgements...

South African Weather Services at Bethlehem for providing excellent radar and raingauge data which were used in the analysis and for critically appraising the products of the research,

Merab Menabde then of the Physics Department of Auckland University for his guidance in the early stages of the research. The information he offered regarding the power-law filtering of white noise was crucial to the development of the model,

Alan Seed of the Australian Bureau of Meteorology in Melbourne for fuelling my interest in the topic and for sharing his contributions to the spatial and temporal modelling of rainfall in South Africa. His generous support regarding the radar database management is also acknowledged with thanks,

Duncan Forrest for the insight he offered regarding image data storage and manipulation.

List of Tables

List of Figures

List of Symbols

WAR_i instantaneous wet area ratio

WAR_i * normalised instantaneous wet area ratio

$WAR^*_{i(t)}$ normalised instantaneous wet area ratio at time t

$\overline{WAR_i}$ daily average of the WAR_i

WAR_d daily wet area ratio

WAR_m monthly wet area ratio

N_w is the number of *wet* pixels (rainfall rate in excess of 1mm/h) on the image

N_d is the number of *dry* pixels (rainfall rate less than 1mm/h) on the image

IMF_i instantaneous image mean flux (mm/h)

IMF_i * normalised instantaneous image mean flux

$IMF^*_{i(t)}$ normalised instantaneous image mean flux at time t.

IMF_d daily image mean flux (mm/day)

IMF_m monthly image mean flux (mm/month)

D_W duration of wet spell (hours)

M_{war*} mean of the WAR_i*

S_{war*} standard deviation of the WAR_i*

$\hat{M}_{war*}$ approximation of the mean of the WAR_i*

$\hat{S}_{war*}$ approximation of the standard deviation of the WAR_i*

M_{imf*} mean of the IMF_i*

S_{imf*} standard deviation of the IMF_i*

ε_{war*} difference^2 between $\hat{M}_{war*}$ and M_{war*}

δ_{war*} difference between $\hat{S}_{war*}$ and S_{war*}

β_{space} exponent of the radially averaged, two-dimensional power spectrum in space

β_{time} exponent of the one-dimensional power spectrum in time

D_w is the duration of the wet spell

D_d is the duration of the dry spell

$\Lambda(\mu, \sigma)$ represents the lognormal distribution with parameters μ and σ

m Maximum likelihood estimator of μ

s Maximum likelihood estimator of σ

μ_w represents the mean of the logs of D_w

σ_w represents the standard deviation of the logs of D_w

μ_d represents the mean of the logs of D_d

σ_d represents the standard deviation of the logs of D_d

Z_{WAR} percentage point of the standard normal distribution corresponding to WAR

Ne_m number of events in any month

Me_m Mean of the number of events in any month

Se_m Standard deviation of the number of events in any month

B_{adv} event advection direction

C_{adv}. event advection speed

$\overline{R}_e$ Average rainfall depth (mm) for an event

$\phi_1 \ldots \phi_5$ are autoregressive parameters of the pixel scale time series

Chapter 1

Introduction

1.1 OBJECTIVES OF THE RESEARCH

This research was funded by the Water Research Commission of South Africa (WRC) as a follow on research topic to work by Pegram and Seed (1998) regarding the feasibility of stochastically modelling the spatial and temporal distribution of rainfields. The final objective was to make use of the rainfall images measured by the MRL5 weather radar in Bethlehem, South Africa in conjunction with daily raingauge data for the region, to develop a rainfall model which was able to simulate radar rainfall images for application in water resources management and real time simulation of flooding scenarios.

In the initial stages of research, the model was to be a high resolution space-time model of rainfall for individual rainfall events. These simulations were to be used in training exercises for real time operation of major flood control mechanisms. In a particular case, the gates on the Vaal dam were used in conjunction with radar data in 1996 (Terblanche et al., 2001) to control a significant flooding event and thereby limit the resulting damage which would otherwise have been incurred downstream. It was thought that with an effective training tool such as a stochastic rainfall model which mimics radar images, the techniques learned in this event could be enhanced and used in the control of floods in other catchments. An immediate application is in real time flood simulation for urban catchments for the development and modelling of flood control mechanisms (Han et al., 2000) and in conjunction with hydrological runoff models for the prediction of potential flash-flood scenarios in steep mountainous terrain and other catchments which respond quickly. As research progressed, it became apparent that it would be possible to extend the ideas into a continuous space-time rainfall model for the region producing simulations many years long and thereby making it useful also as a tool in the management of water resources.

The model is a stochastic space-time model of rainfall fields and as such it does not attempt to model the physical processes which influence the rainfield, but rather to mimic the statistics and appearance of weather radar rainfall images over a wide range of spatial and temporal scales.

1.2 LITERATURE

The *String of Beads* model has been developed using ideas from a variety of existing models which are cited in the text of the thesis. It is perhaps appropriate at this point to present a brief review of some of the work which was particularly pertinent to the development of this model.

1.2.1 Types of rainfall model

A general review of the more recent advances in rainfall modelling is given by Foufoula-Georgiou and Krajewski (1995) who consider the submissions to U.S. journals during the period 1991-1994. Foufoula-Georgiou and Krajewski (1995) divide rainfall models into two main categories, stochastic and dynamic models. They define dynamic models as being based on a set of partial differential equations describing conservation of mass, momentum and energy. This type of model belongs to the realm of cloud physicists and meteorologists and was not considered in this study since the *String of Beads* model is aimed more at the hydrological community who are really concerned with the statistics of rainfields rather that the four dimensional (three space and one time) physical behaviour of clouds. This leaves the stochastic models.

For the purposes of this discussion, stochastic rainfall models are separated into three main categories:

- Single point process models
- Multiple point process models
- Continuous space models

Each of these provides an insight into the behaviour of the overall rainfall field and will be considered in turn in the following paragraphs. The *String of Beads* Model falls into the last of these categories as a continuous space-time model of rainfall.

1.2.2 Single point process models

The most fundamental of rainfall models consider time series of rainfall data at a point. A variety of approaches have been proposed which mimic the behaviour of raingauge time series data and the earliest references to this type of model date back to the works of Le Cam (1961) and extend through to the present day. These models consider a range of temporal scales and they include the types listed below.

Single point binary wet-dry processes

Discrete binary wet-dry processes using Markov chains are proposed by Gabriel and Neumann (1992) and Katz and Parlange (1995) amongst others. These make use of transition probabilities between the wet and dry states in order to produce time series of rainfall state at various temporal resolutions. Explicit theoretical links between the discrete Markovian transition probabilities and the continuous alternating renewal models are presented by Small and Morgan (1986) and they also provide general relationships which can be solved numerically for the non-Markovian cases.

Daily rainfall distribution models at a point

Daily rainfall distribution models such as works by Woolheiser and Pegram (1979), Stern and Coe (1984), Rodrigues-Iturbe et al. (1988). These extend the ideas of the binary rainfall state, and suggest models which also incorporate the amount of rainfall which occurs in a particular discrete time interval. They are able to produce time series of rainfall at a point which aggregate to daily, monthly or yearly totals, consistent with the observed climate at that point.

High resolution time-series models at a point

High resolution time-series models such as the Neymann-Scott shot noise model developed by Cowpertwait (1994) from the work of Rodriguez-Iturbe et al. (1988) and Cox and Isham (1988). More detailed than the daily rainfall distribution models, these are able to reproduce continuous time series down to the smallest of useful temporal scales. Also included in these are disaggregation models (Güntner et al., 2001) which slice up daily rainfall totals into finer temporal scales.

Since the *String of Beads* Model is a continuous space-time model, it is not directly comparable to any of these models, however there are useful ideas which grew from

these works which have influenced the development of this work. Furthermore, point processes can be thought of as special cases of the more general continuous space-time process and although not attempted in this study, observations of these processes could serve as useful validation tests for any continuous space-time model.

Of those listed, it is the Neymann-Scott shot noise model developed by Cowpertwait (1994) which sparked the greatest interest in the context of the *String of Beads* Model. He adopts an interesting approach to modelling a high resolution time series of rainfall intensity data by considering the duration of an event and the event arrival as two separate random processes. Event arrival and Event duration are two independent Poisson processes. This is a useful concept and one which has been exploited in the *String of Beads* Model as will become apparent in the following chapters of this document.

1.2.3 Multiple point process models

Expanding into a second spatial dimension, the next breed of rainfall models considered networks of raingauges over large areas and they proposed various methods of interpolation to describe the spatial dependence structure between adjacent gauges. In much the same way as the point process models, these take the form of :

Binary wet-dry processes for an area

Using binary wet-dry processes in the discrete and the continuous form using discrete Markov Chains and Alternating Renewal processes, Harberlandt (1998) described the transition between wet and dry times over an area. These processes are dependent on a form of area-depth threshold to distinguish between wet and dry periods over an area which can make it difficult to compare results of different models. Pegram et al (2003) have also demonstrated the effectiveness of the Small and Morgan (1986) approach over an area by estimating an appropriate continuous alternating renewal process from a discrete Markov chain model for a network of pluviographs.

Discrete daily climate models

These are also dependent on a from of threshold in order to classify the daily state. Examples of these include works by Bardossy and Plate (1992) who use synoptic

atmospheric patterns to classify the daily state from which rainfall is predicted and Pegram and Seed (1998) who use a discrete Markov chain to switch between *Dry Day*, *General Rain Day* and *Scattered Rain Day* states. The state transition properties are made to vary seasonally.

Discrete data infilling techniques

Rather than stand alone rainfall models, these exploit the cross correlation structure between adjacent rain gauge datasets in a network to infer a rainfall time series at a point. These techniques are used to remove erroneous outliers and improve the quality raingauge datasets. These include models proposed by Basson *et al.*(1997), Makhuvha *et al.* (1997a, 1997b) and implemented by Pegram (1997) and Smithers and Schulze (2000).

Continuous-time, multi-point rainfall distribution models.

Arguably the most successful of these are two-dimensional forms of the Neymann Scott shot-noise model. An excellent, concise overview of these types of models is given by Wheater et al. (2000). The most basic of these is that proposed by Cox and Isham (1988). Their model uses cylindrical pulses which are circular in space and distributed in space by a two dimensional Poisson process. They have exponentially distributed survival and arrival times, and maintain a constant rainfall intensity for the duration of their individual lifetimes. All of these cells are assumed to move at the same velocity. Rainfall intensity at a point is taken as the sum of all of the contributing cells at that point. A slightly modified version of this model is proposed by Northrop (1996) who uses elliptical shapes of variable eccentricity and orientation instead of the circular shapes of Cox and Isham (1988), in order to cater for the occurrence of line-storms. These models do not attempt to describe the continuous spatial properties of the rainfall fields, but rather the properties of a network of points within the rainfield. Their main advantage is that realistic point-wise statistics are preserved.

Hybrid state-pulse models.

These are models which employ a coarse process to describe the arrival of a weather state, coupled with a more detailed point process model to capture locale-specific rainfall statistics, particularly useful in areas where distinct rainfall generating mechanisms are observed. An example of this type of model is presented by Fowler et

al. (2000) who model the spatio-temporal rainfall distribution using a Markov-chain to switch between weather states in combination with a Neyman Scott Rectangular Pulse model (Cowpertwait et al. 1996a and b).

Like the single point process models, mentioned in the previous section, these multiple point process models represent special cases of the continuous space time models.

1.2.4 Continuous space models

With the advent of weather radar, the true space-time structure of the rainfall field was revealed and this sparked an interest in more realistic modelling techniques amongst both the hydrological and meteorological communities, namely the simulation of random fields. Until recently, many of these techniques have been limited by the available computing power but as processors have been improved, more elaborate space-time rainfall models have emerged. Good insights for the understanding of random fields are provided by VanMarke (1983) and Cressie (1991), the latter concentrating on the Kriging problem rather than the time series aspect.

Although the history of radar meteorology, the reflectivity vs. rainfall relationship and the radar calibration and quality control are not the focus of this study, they serve as useful background knowledge for the modelling process.

History of radar meteorology

A brief history of radar and radar meteorology is offered by Austin (1998). He explains how radar was developed first by the allied forces as a military tool in World War 2 used to predict German air raids on Britain. Rainfall was observed as a nuisance on these early images and techniques were developed to distinguish aircraft from rainfall. Only after the war, it was realised that rainfall intensity could actually be measured by radar and so was born the science of radar meteorology.

The Reflectivity vs. Rainfall relationship

Early research in radar meteorology was focussed on the relationship between radar reflectivity and rainfall rate leading up to the well known Marshall-Palmer relationship published in 1948. Research has continued in this area and numerous improvements to

the measurement of rainfall using radar have culminated in the modern weather radar. A good summary of the proposed models is given by Uijlenhoet (2001) who considers these relationships in detail and draws comparison between them with reference to work by Battan (1973). Most of these relationships are log-linear in nature and Uijlenhoet shows how the Marshal-Palmer (1948) approach is central amongst these.

Radar calibration and quality control

There exist a variety of algorithms which enhance and improve the rainfall estimate as measured by radar all of which use raingauge networks as the ground truth for calibration, some of which adjust radar estimates in real time using techniques such as Kalman filtering as demonstrated by Todini (2001) for example. Of particular interest to this study is the research on the MRL5 weather radar at Bethlehem, South Africa. These include improved algorithms to extract Cartesian fields from binned spherical fields (Mittermaier and Terblanche, 1997), algorithms to remove ground clutter from radar fields (Terblanche et al., 2001), algorithms to improve the rainfall estimate using the vertical reflectivity profile (Pegram and Mittermaier, 1998). In addition there are a number of validation experiments and comparisons to reliable networks of daily and telemetering raingauge networks by van Heerden and Steyn (1999) and Terbalnche et al. (2001) amongst others. An active research team in the SAWS:METSYS group in Bethlehem, South Africa, ensures that radar calibration is monitored and adjusted if necessary.

Marginal distribution of rainfall rates over an area

By far the most common estimate of the marginal distribution of instantaneous rainfall intensity at pixels on an image is the two-parameter lognormal distribution. This has been used for rainfields observed by satellite (Bell, (1987) and Wilson et al. (1991)) and radar (Crane, 1990) amongst others. Although other distributions may be appropriate, the quality of fit is usually excellent when using full precision data.

Spatially-clustered continuous-time, multi-point rainfall distribution models.

An expansion of the two-dimensional multiple point process models is to introduce clustering to the spatial distribution of the storm cells so that they are no longer entirely random. Models of this type were proposed by Cowpertwait (1995), Cowpertwait and O'Connel (1997) who categorised cells into convective and stratiform types, Northrop

(1997) who uses a bivariate displacement of a range of idealised storm shapes from a storm centre, and Wheater et al. (2000) who compare the use of a Gaussian bivariate clustering mechanism and a uniformly distributed cell centre over a random ellipse. These models produce continuous rainfall fields which are reasonably realistic when compared to real radar images.

Scale invariant rainfield models

Foufoula-Georgiou and Krajewski (1995) report how the trend in rainfall modelling has shifted towards models which are based on concepts of scale invariance which assumes that the stochastic structure of the rainfall field is independent of scale. This assumption has been shown to be true over a wide range of scales (Gupta and Waymire, 1990) and it lends itself to modelling through the use of multifractal algorithms. Models of this type have been proposed by Schertzer and Lovejoy (1987), Over and Gupta (1994) and Menabde et al. (1997) amongst others.

Over and Gupta (1994) use a multiplicative random cascade for the generation of their rainfields, but although statistically sound, the rainfield images produced by these models have the unfortunate property of appearing blocky. In spite of the blocky appearance of cascade models, a great deal of work has been done on the parameterisation and calibration of these models and to date, they are arguably good statistical descriptors of the rainfall field so long as it does not include gaps. Menabde et al. (1997) use a discrete random cascade to generate a field with the desired statistical structure which they then power law filter, thereby removing some of the blockiness and resulting in a more realistic looking rainfield.

An alternative to this approach is to simply generate a random field based on a suitable probability distribution and then power law filter the field to achieve the desired spatial correlation structure. This technique first became popular as a means of simulating fractional Brownian motion on which there is a considerable amount of literature (Mandelbrot, 1983 for example).

This is the approach that was used by Bell (1987) for the simulation of satellite rainfall intensity data. He found that the lognormal distribution was a good descriptor of the marginal distribution of rainfall intensities on a field and this was confirmed by Crane (1990). Bell's approach (1987) was simply to generate a Gaussian random field, power law filter it to achieve the desired spatial correlation structure and then exponentiate the field to obtain a spatially correlated, lognormally distributed field. These fields are convincing in appearance and possess a reasonable statistical structure. It is the same technique as the universal multifractal technique of Schertzer and Lovejoy (1987) for modelling rain and clouds.

The space-time continuum

In his work on fluid turbulence, Taylor (1938) hypothesised that the correlation of the local velocity at a point in time (i.e. at a fixed point in space) was related to the correlation of velocity at that point with that of its neighbours in space (i.e. at a fixed time) through some velocity vector. This hypothesis is discussed with reference to the space-time structure of rainfall intensity fields by Gupta and Waymire (1987) and by Crane (1990).

If Taylor's hypothesis is assumed to be true, the approach adopted by Bell (1987) for the simulation of two dimensional fields can be simply extended into three dimensional space-time by using a power law filter which defines the temporal correlation structure as a vector multiple of the average spatial correlation structure. The algorithm used to generate these three dimensional fields makes use of the Fast Fourier Transform thereby reducing the computational effort and making it possible to quickly generate a large number of fields.

In his paper, Bell (1987) points out that opposite edges of the Fourier filtered images behave as if they are joined. This property can be easily exploited to introduce an overall advection of the rainfield. This gives the Fourier filtered images an advantage when compared to generating correlated, lognormally distributed fields through the use of other algorithms such as the Turning Bands method (Matheron, 1973), and the Direct method. When extended into the third dimension, Bell's (1987) process produces a finite sequence of images which are correlated in both space and time and which incorporate

an advection characteristic. However, it would not be feasible to simulate long periods of data through the use of this process alone, since the computational effort (due to the need for a large, multidimensional Fourier Transform) and data storage requirements (storage of blank images) would soon become unworkable. This same approach is adopted by Wilson et al. (1991) who produced the first animated images of universal multifractals.

An alternative approach to these models is found in the works of Marsan et al. (1996) and Seed et al. (1999) who suggest that a one-layered image can be linked in time by a pair of high order autoregressive models (rather than a cascade) or ARMA(1,1) models, or a modified power-law filter to give the correct temporal correlation structure.

The *String of Beads* Model has borrowed ideas from all of the models discussed above and to arrive at a high resolution space-time model of rainfall which is capable of efficiently simulating many years of rainfall data as images which are very difficult to distinguish from real radar images. This is achieved without the need for super-computing equipment.

1.3 STRUCTURE OF THE THESIS

This document begins by discussing the radar and raingauge data used in the study and addressing issues of data collection, data quality and reliability. For the most part, data used in this study are of excellent quality and have been subject to rigorous pre-processing to remove holes and erroneous outliers.

After considering the data, Chapter 3 looks at the techniques used in the analysis of these data in order to extract as much information as possible, ultimately reducing them to uncorrelated white noise and a model structure.

The fourth chapter looks at the formulation of the *String of Beads* Model. These are the techniques used to reverse the analysis process of Chapter 3. That is, taking white noise and gradually imposing spatial and temporal structure and ultimately arriving at realistic sequences of simulated rainfall images.

Chapter 5 considers the calibration and verification of the model by simulating sequences of images and then measuring the statistics of the simulated images and comparing them to those of the observed data. This is a check of the integrity of the model. Since the modelling process is imperfect, this chapter also presents techniques by which the model output can be fine-tuned to achieve more realistic results.

Chapter 6 considers validation tests for the model. In essence, this involves extracting from the simulated sequences, statistics which are not used in the formulation of the model and testing them against the same for the observed sequence. The successful performance of the model in these tests help to validate the modelling process.

Chapter 7 presents the conclusions and recommendations for additional work.

Although no longer part of the modelling process, early attempts at modelling various components of the *String of Beads* Model are included in Appendix A of this document. Development of the model has taken a number of turns and these serve to illustrate other methods which have been tried and tested and found wanting for various reasons. They serve to illustrate the evolution of the model during the course of the research. Included in Appendix B are some standard tools which were collected for the Masters thesis and are included here for easy reference.

Chapter 2

Rainfall and Other Relevant Data

2.1 INTRODUCTION

There are two main types of rainfall data used in this study. The first is a set of daily raingauge data within a 100 kilometres of Bethlehem, SA which have been collected over the past eighty years, mostly by volunteers. The collective daily raingauge data set used in this study was provided by the Computing Centre for Water Research (CCWR) in Pietermaritzburg. The second type is the radar data set which is collected and archived by the South African Weather Services (SAWS) METSYS group in Bethlehem.

Daily raingauge data have been used since the birth of hydrology. Their limitations are well understood and they have become the trusted ground truth for rainfall measurement. At this point in time, they are the only quantitative, long-term rainfall measurements available.

Radar rainfall data presents instantaneous images of rainfall intensity over large areas thereby enabling the human mind to qualitatively assess the data much faster, and with far greater intuition, than can be done by a computer program. If poorly managed, the large volumes of radar data can be extremely cumbersome and lead to problems in computer disk storage space and processing. It is therefore essential to outline the processes that were used in this study to store, manipulate and display the radar data.

This chapter presents methods of collection of both types of rainfall data and addresses issues such as data quality, limitations and difficulties experienced when measuring rainfall using each technique. Data management and presentation as well as treatment of incomplete data sets are also discussed.

2.2 *STUDY AREA*

Figure 2.1 shows the South African Radar network as at 2002. The radar cover is shown by blue circles up to a range of 100km and green circles to a range of 150km. An area of great strategic interest in terms of water resources and flood control in South Africa is the catchment of the Vaal Dam (38 505 km^2) which provides most of the stored water for Johannesburg and Pretoria and the surrounding areas. The Vaal catchment is monitored by two radars, a C-band at Ermelo and an S-band near Bethlehem.

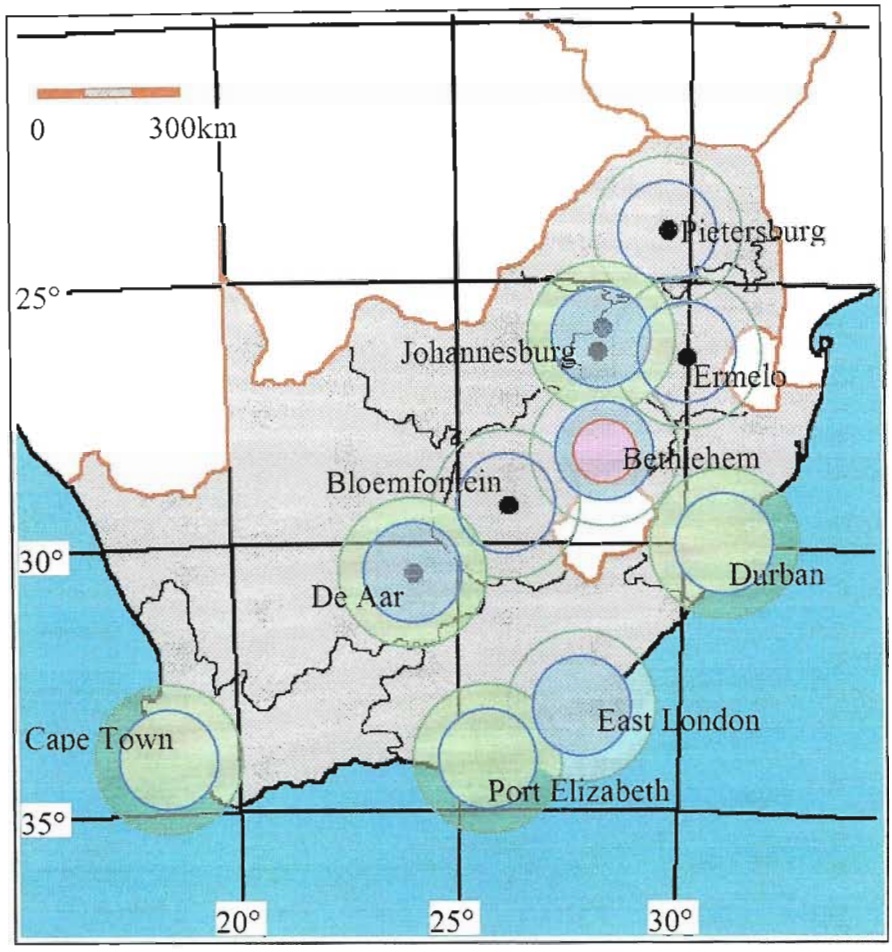

Figure 2.1 - The South African radar network

For the purposes of this study the data collected in a circular area of 13000km^2 is considered, centred on the MRL5 S-band radar near Bethlehem and shown as a red circle in Figure 2.1. The climate at Bethlehem is semiarid with an annual rainfall of approximately 650mm, most of which falls in the summer months between October and April. Bethlehem experiences both convective and stratiform rainfall which approaches mostly from the West and North West. Apart from the Maluti Mountains which are 50km South East of Bethlehem, the region is reasonably flat so orographic effects are minimal.

2.3 GAUGE DATA

The rainfall in the Vaal catchment is monitored by a network of daily raingauges maintained by local volunteers, and a good network of 45 logging tipping bucket raingauges on a sub catchment of approximately 4 600km^2 which is maintained by the South African Weather Services (SAWS). The daily raingauge data for this region extends over approximately 80 years and the network of logging tipping bucket raingauges has recorded data for approximately 10 years. Of primary concern for the purposes of this study is the network of daily raingauges. Figure 2.2 illustrates the positions of reliable daily raingauges which have records in excess of 10 years.

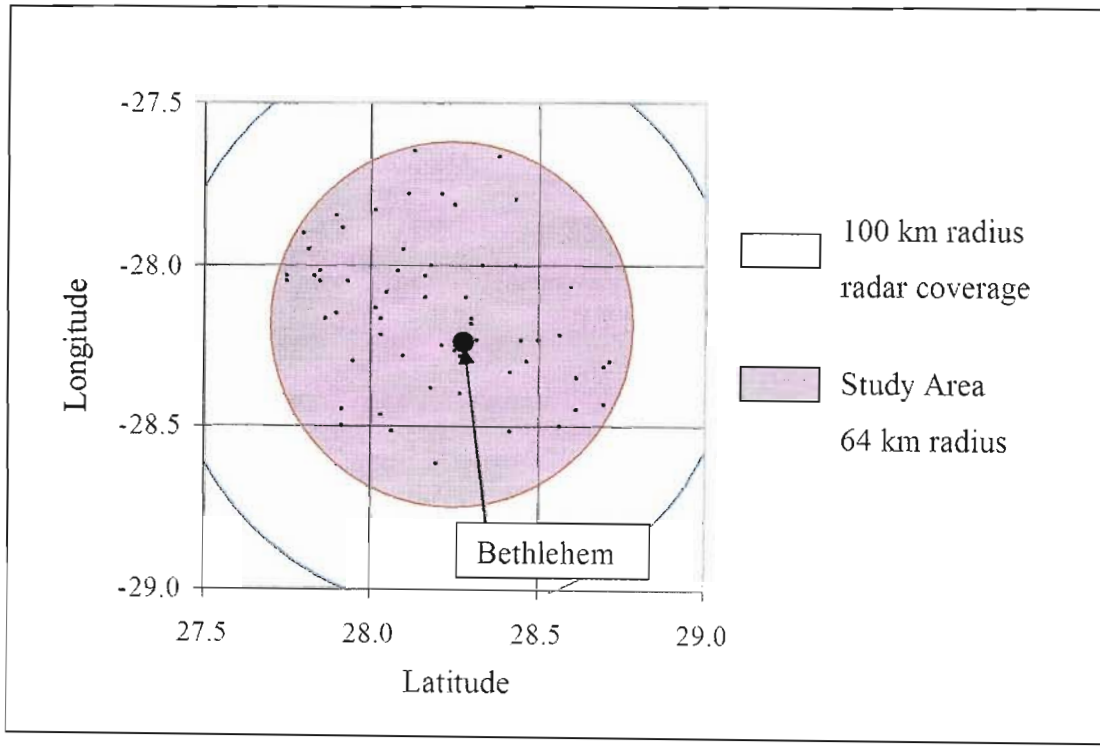

Figure 2.2 - Distribution of 54 daily raingauges over the study area which is 64 km radius centred on the MRL5 radar north-east of Bethlehem

2.3.1 *Gauge data pre-processing*

There are a number of difficulties associated with daily raingauge records and these include cases of

- missing data
- data out of phase caused by recording rain with an incorrect date
- rainfall for several days recorded as a single day's total
- erroneous data, such as gauges which have been irrigated

These types of errors can cause considerable difficulty when attempting to model the rainfall for a region. The daily raingauge data set used in this study was pre-processed by the Computing Centre for Water Research (CCWR) in order to eliminate as many of these errors as possible.

A detailed description of the quality control measures and infilling techniques used by the CCWR is given by Smithers and Schulze (2000). Their process employs:

- A set of routines (CLASSR) developed by Pegram (Basson *et al.*,1994) to detect outliers and define groups of stations for infilling of monthly rainfall totals.
- A modified version of the Expectation Maximisation Algorithm formalised by Dempster *et al.*(1977), (Makhuvha *et al.*, 1997a, 1997b), a recursive technique which uses a multiple linear regression relationship between the target station (missing data) and selected nearby control stations. This technique was used by Pegram (1997) to develop (PATCHR) routines (Basson *et al.*,1994) used to infill missing monthly rainfall totals.

Details of these techniques will not be discussed here, but the result of this treatment is a continuous record of good quality data for all 54 gauges for the past 80 years.

2.4 RADAR DATA

Radar measures rainfall by sending out an intense electromagnetic pulse, and then *listening* to the echoes of the pulse as minute fractions of it are reflected back to the antenna by the raindrops. The intensity of the echo increases with rainfall intensity. Being an electromagnetic wave, the pulse propagates in a reasonably straight line from the antenna and knowing the speed of light in air, the range of the raindrops can be computed by measuring the time between the pulse and the echo. The echo is a continuous analogue signal of *reflectivity* (Z) which is then digitised or binned by sampling the echo signal at discrete *small* time steps and averaging the signal over the sampling interval. The sampling interval will determine the length of the bin so that a sampling interval of 1μs would result in a bin length of 150m since light will travel approximately 300m in that time. Due to the wide range of reflectivity observed in the radar echo, it is usually measured in decibel units (dBZ). Radar measures the reflectivity in a cone perpendicular to the antenna. The apex angle of the cone is referred to the beam width and is dependent on the wavelength of the signal and the diameter of the

antenna – longer wavelength radars require a larger antenna in order to keep the beam width small. By rotating the antenna and sampling the reflectivity at all azimuths, a two dimensional field of reflectivity is measured and furthermore, if the elevation is increased after each revolution, a three dimensional volume scan is achieved.

By adjusting the way in which the azimuth and or elevation changes, or by adjusting the sampling interval, the radar can be configured to give different kinds of information. A common configuration in rainfall investigation is to keep the radar pointing vertically to record vertical profiles in a storm for the study of convective activity. Another possibility is to fix the azimuth and only vary the elevation to examine a slice through a storm cell. By far the most common configuration for weather radar are the base scan (fixed elevation as close as possible to the ground, all azimuths) and the volume scan modes (all elevations, all azimuths).

2.4.1 The MRL-5 S-Band radar near Bethlehem

The MRL-5 S-band radar was installed near Bethlehem by the Water Research Commission in 1994 and continuous rainfall data during the wet summer months (November to April) are available from this radar from January 1995. Data from the C-band radar at Ermelo are available from January 1998 and the remaining 8 radars in the network from mid 2000. For the purposes of this study, only data from the MRL-5 S-band radar were used.

A photograph of the MRL-5 Radar is given in Figure 2.3. The radar dome is mounted on top of a portable control centre. The raw data is transmitted via radio link to the offices of the Weather Bureau in Bethlehem which are roughly 20 km south east of the radar site. It is powered by an uninterruptible power supply (UPS) to ensure that the data set is not corrupted in the event of an electricity failure which is most likely to occur during a thunder storm.

Figure 2.3 – The MRL-5 S-band radar near Bethlehem

The MRL5 is a combined S-band (10 centimetre wavelength) and X-band (3 centimetre wavelength) weather radar. The S-band radar requires a larger, more expensive antenna but is less prone to errors due to attenuation and is therefore a better choice when measuring rainfall intensity at long range. The X-band radar data is not considered in this study. In S-band mode, the MRL-5 is configured to perform a full volume scan starting with a base scan at an elevation of 1.5° and incrementing in elevation in 18 increasing steps to its final elevation of 55°.

With an antenna 5 meters in diameter, the beam width of the S-band radar is 1.5°. Data is collected in 600m bins along the beam. Data received prior to 1999 records its first 600m bin echo at a range of 14km from the radar, resulting in a characteristic hole at the centre of the images. The controlling software of the MRL-5 has since been reconfigured so that the first bin is at a range of 600m and images now record rainfall everywhere within a 150km radius of the radar. The MRL-5 completes one volume scan every 4.5 minutes.

2.4.2 Methods of CAPPI extraction from volume scan data

Constant Altitude Plan Position Indicators (CAPPIs) are two dimensional fields of reflectivity or rainfall intensity which are extracted from the three dimensional volume

scan. The images considered in this study are extracted (by the SAWS) to a Cartesian grid space from the spherical volume scan data in one of two ways. The first is that used by Seed (1992) which involves simply projecting the nearest bin of radar data (shown as blue lines in Figure 2.4) onto a horizontal plane (shown as a red dotted line) at a chosen altitude which was taken as 2 km above ground level (AGL) for this study. The 2km level is chosen as a compromise between the lowest possible altitude with the least amount of ground clutter. Images used in this study have a spatial resolution of 1km.

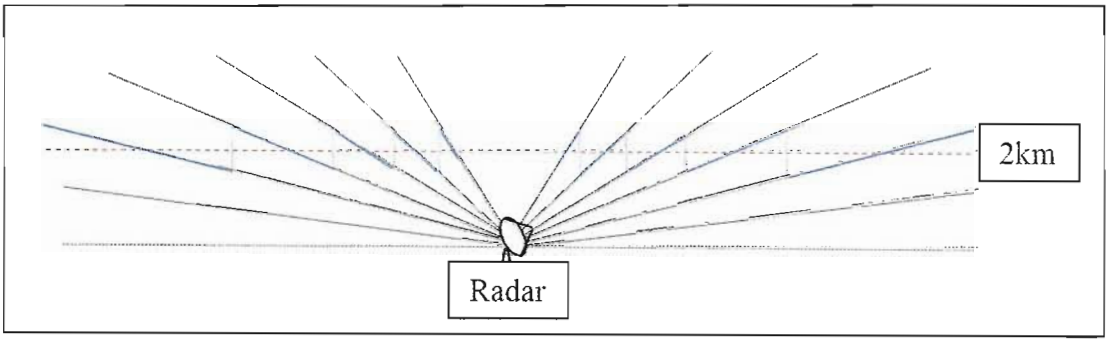

Figure 2.4 - Projection technique used to extract CAPPIs from radar volume scans

This method is a very fast and simple way to generate CAPPIs, however it has the disadvantage of having discontinuities in rainfall rate at the jump from one beam to the next. This effect is observed in the image of Figure 2.5.

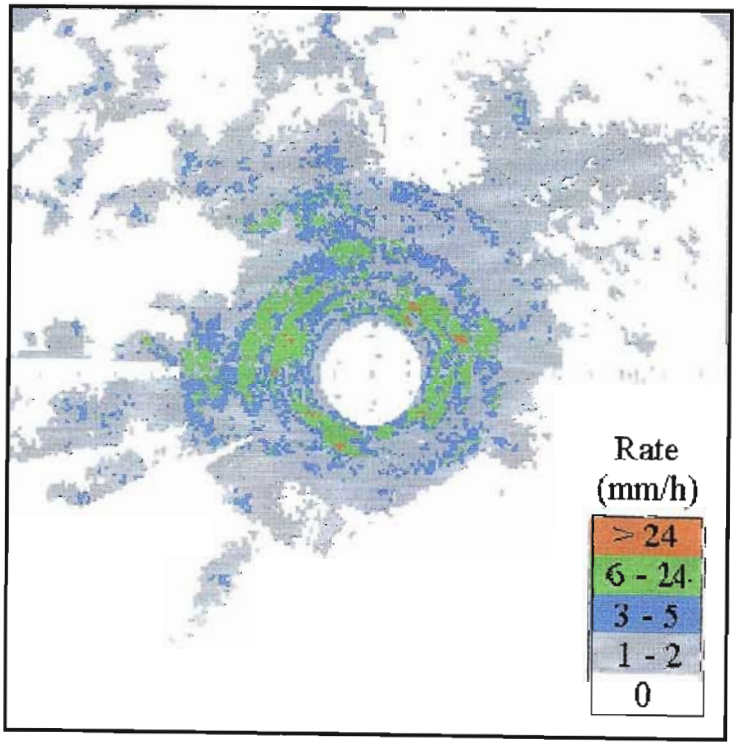

Figure 2.5 – An image extracted using the projection technique. Discontinuities are observed in the reflectivity field (Mittermaier and Terblanche, 1997).

A better method of extracting images from volume scan data is that suggested by Mittermaier and Terblanche (1997) which exploits the logarithmic arithmetic in the reflectivity – dBZ conversion algorithm . The algorithm is known as DISPLACE averaging and involves the interpolation between eight radar bins at two levels to achieve a weighted average reflectivity at a point ***(x, y, z)*** in Cartesian space.

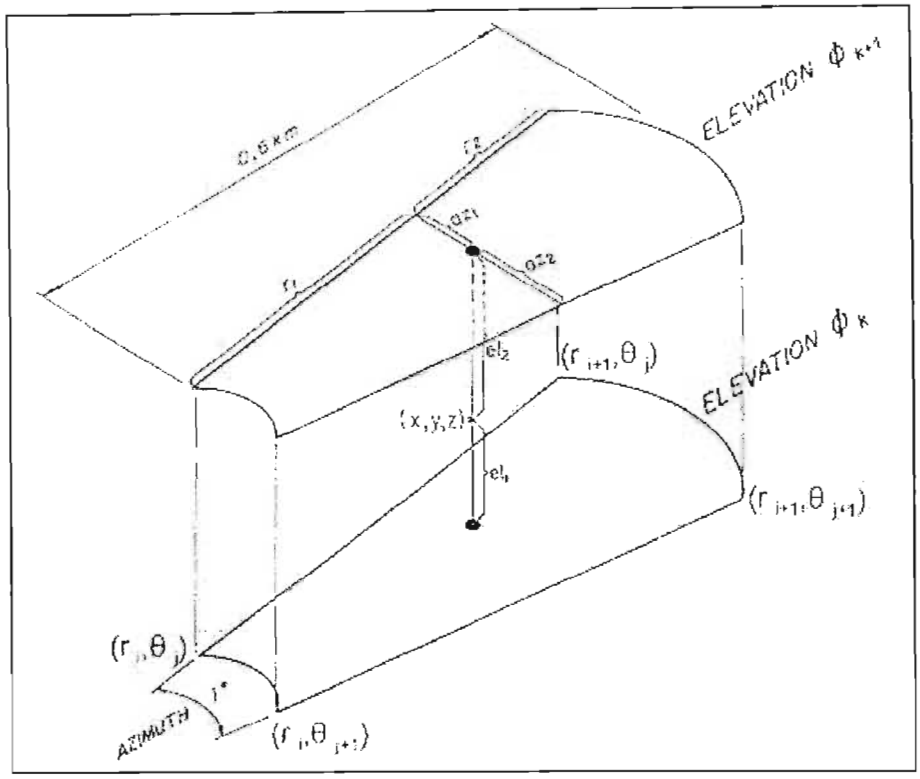

Figure 2.6 - Illustration of Cartesian point and the eight surrounding radar data points used in DISPLACE averaging to extract CAPPIs from volume scan data (Mittermaier and Terblanche, 1997).

The result of this process is a far smoother image which does not exhibit the concentric rings observed in images calculated using the projection technique. An example of the DISPLACE type of CAPPI, extracted from the same raw radar data as that used to generate the image of Figure 2.5, is given in Figure 2.7.

Most of the data analysed in this study were images generated using the DISPLACE averaging technique.

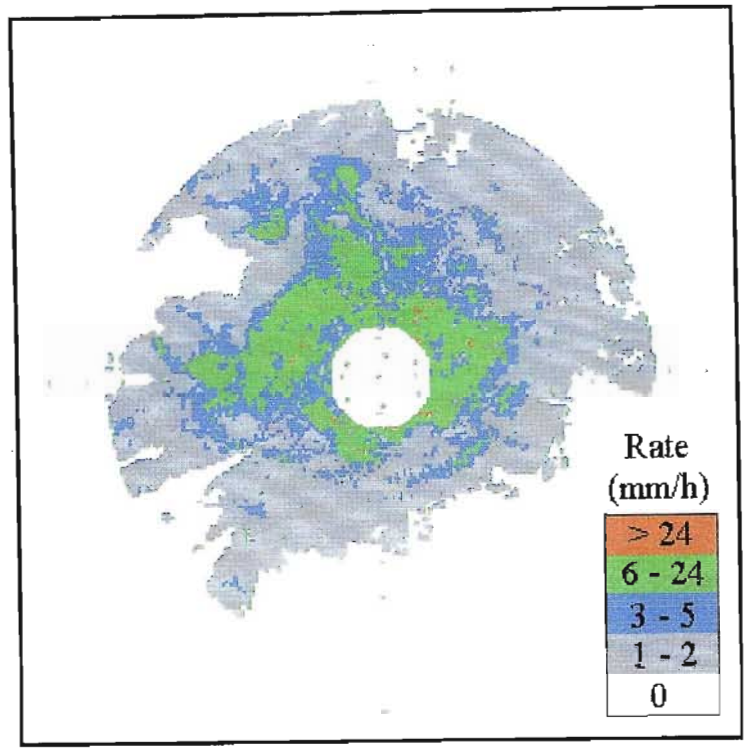

Figure 2.7 - CAPPI generated from volume scan data using the DISPLACE averaging technique (Mittermaier and Terblanche, 1997).

2.4.3 Quantitative limits of the radar data

The base scan of the Bethlehem radar is at an elevation of 1.5° above the horizontal so that the range at which the base scan exceeds an altitude of 2km above ground level is 67km. Data recorded beyond this range are useful for qualitative analysis of the weather system since these are usually within or above the melting layer. It is convenient in the analysis to use a radius of 64km as it is an integer power of two, a requirement for the radix-2 decimation in time Fast Fourier Transform algorithm employed in this study, whose use is discussed in Chapter 3. Using a beam width of 1,5 degrees at a range of 64km, the radar can resolve blocks of approximately 1,6km across so there is some smoothing in the sampling of the rainfall rate at this range.

This study has been done over a period of several years (1997-2002) during which time the configuration of the MRL-5 radar was changed. In the data-sets available at the start of the study, the radar did not record any reflections received within a 14km radius of itself although this "hole" has since been eliminated. In addition, algorithms to reduce the effects of ground clutter had not yet been instigated by the SAWS:METSYS. For these reasons, images in the first part of the study were masked as shown in Figure 2.8

so that only data within the $^3/_4$ doughnut shaped sample area of 9128 km^2, were considered.

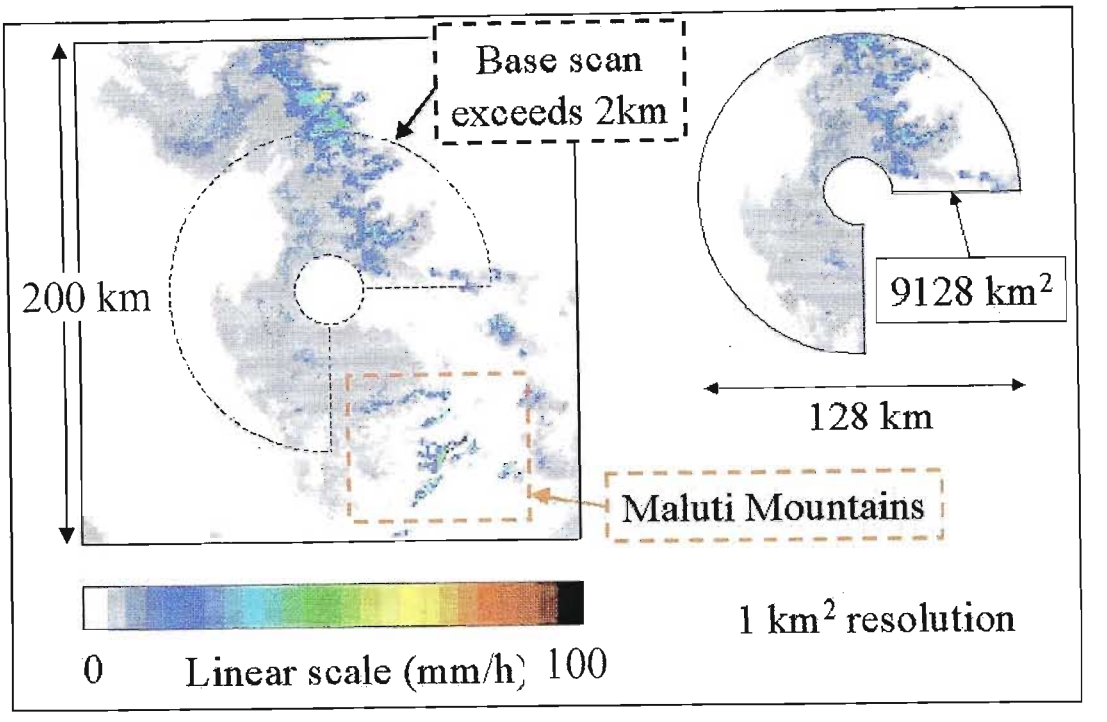

Figure 2.8 - Masking unreliable data

2.4.4 Sources of error in rainfall measurement using radar

A basic assumption made in the development of the model is that the data received from the radar is a reasonably accurate representation of what actually reaches the ground. There is a wide variety of systematic errors which can occur in the measurement process and these will be discussed briefly with reference to the conditions in Bethlehem.

Ground clutter and beam blocking

Ground clutter is caused by the radar beam or the side lobes of the radar beam colliding with the ground at some point and the result is a strong echo which could be interpreted as an intense rainfall cell on an image. In a sequence of images this is easily visually distinguished from rainfall due to the fact that it is stationary. The use of Doppler radar is one of the methods of eliminating the effects of ground clutter but this results in a much slower scan speed and consequently a much lower temporal resolution. Where the core of the radar beam strikes a large fixed object, the beam is said to be *blocked* and rainfall beyond that point will not be recorded by the radar. In mountainous areas

ground clutter presents a major problem in the measurement of rainfall using radar. Apart from the Maluti Mountains in Lesotho, the terrain surrounding the Bethlehem radar is quite flat and there is therefore very little ground clutter and the rainfields tend to be reasonably homogeneous. A CAPPI altitude of 2km above ground level ensures that the effects of ground clutter are minimal (although there are some low-grade effects which become obvious in accumumlation), whilst the rainfall passing through that level is thought to be a reasonable representation of what actually reaches the ground (ignoring the effects of evaporation and updraughts). For the purposes of this study the radar CAPPIs were modelled "as is", leaving the problem of image improvement to later studies.

Bright band

The so called *bright band* is a band of high reflectivity which, in the Bethlehem area, is usually at a reasonably constant altitude and corresponds to the melting layer in a cloud. When viewed by radar, snow and ice crystals have a reflectivity which is approximately 4 times lower than that of water. Consequently, the radar will tend to under estimate the rainfall on the ground when looking above the bright band. As snow and ice crystals pass through the melting layer en route to the ground, they begin to melt and become coated in a thin layer of water. This results in a very high reflectivity as the radar *sees* very large drops of rainfall. On an image this is seen as a ring of high reflectivity or rainfall intensity centred on the radar. When viewed with a vertically pointing radar this is seen as a band of high reflectivity. Research by Pegram and Mittermaier (1999), shows how it is possible to predict the rainfall reaching ground level using the vertical reflectivity profile. Due to the high summer temperatures experienced during the wet season in South Africa, the bright band is usually at or above the 2km level. A large proportion of the rainfall experienced in Bethlehem is convective in which case the bright band is not well defined and therefore does not present a significant problem.

Beam Attenuation

As the radar beam strikes any object, a small part of the beam is reflected back to the receiver, some is scattered and the remaining (slightly weaker) part continues in its original direction. The radar beam becomes weaker as its range from the radar increases. When passing through an intense storm cell, it is possible in extreme cases for the beam to become so attenuated that rain falling beyond the cell is significantly

under estimated. This is known as beam attenuation. Since the longer wavelength radars are more powerful and suffer from less Raleigh scattering, the effects of attenuation are kept to a minimum by using an S-Band radar.

Anomalous propagation

When calculating the three dimensional position corresponding to the reflectivity recorded by the radar, it is assumed that the radar beam has propagated in a straight line from the radar to the target and back again. This is not always the case. Being an electromagnetic wave, the radar beam is bent by a change in density of the medium through which it is travelling. Changes in air density are experienced with the approach of a weather front and in extreme cases, a temperature inversion can cause the radar beam to be bent to such an extent that it strikes the ground. Part of the reflected beam will follow the same path back to the receiver and will be observed as a point of high reflectivity. This is difficult to identify in an image as it is not observed in a fixed location. In these extreme cases the beam will be blocked at the point of intersection with the ground and will be difficult to distinguish from beam attenuation. More information is required in this case, either from ground based stations or from satellite images. These extreme cases of anomalous propagation are very rare in the data collected at Bethlehem because temperature inversions usually occur in the winter months which are the low rainfall months in the summer rainfall region of South Africa.

The Z-R relationship

Much research has been done on the relationship between the observed reflectivity ($\boldsymbol{Z}$) and the rainfall rate ($\boldsymbol{R}$) recorded on the ground. The Z-R relationship $\boldsymbol{Z = 200R^{1.6}}$ given by Marshall and Palmer (1948) is used to convert reflectivity to rainfall rate for the Bethlehem data. Various refinements of this relationship have been proposed, 69 of which are quoted by Battan (1973). It has been shown (Uijlenhoet, 1998) that the Marshall-Palmer relationship is very close to the mean of the 69 relationships referred to by Battan. An important point to note is that the measurement of rainfall intensity by *any* of the known methods is by no means infallible and at *best* a reasonable estimate will be achieved. The use of the Marshall-Palmer relationship seems to be a reasonable one and is adopted by the Bethlehem Precipitation Research Programme for data processing.

2.5 *SPATIAL DATA STORAGE*

A crucial aspect of this study has been the management of extremely large data sets and presenting the data in a meaningful and efficient way for quantitative and qualitative analysis. Part of the objective of this research was to perform the analysis and develop the model on personal computer in order to avoid the need for large, expensive computing resources. This constraint imposes significant limitations on the processing power, memory resources and hard disk storage space.

To give an idea of the scope of the data storage problem, the MRL-5 radar completes a full volume scan of 17 CAPPIs, at 1km spatial resolution, to a range of 200km from the radar, in approximately 4,5 minutes. Considering only the 2km AGL image, this amounts to 320 images per day, each of which is 400 x 400 pixels – approximately 19 billion pieces of data per year. If stored as 32bit floating point precision (4 bytes per pixel), this would require approximately 70 gigabytes of disk storage space per year – unreasonable even by modern standards and other methods must therefore be used.

2.5.1 *Meteorological Data Volume (MDV) file format*

The file format used to store the radar data by the SAWS is the MDV file format. This format is capable of storing multiple levels of multiple fields of data and is designed specifically for meteorological application. As an example, a single MDV file might contain

- radar reflectivity field at multiple levels
- visible satellite field
- infra red satellite field
- pressure field
- wind vector field at multiple levels

amongst others. On conversion to MDV format, each floating point field is shifted and scaled so that it has a range of 0 to 255 and the scale and shift terms for each field are recorded in the MDV file. Every floating point pixel in the scaled field is then converted to character precision and the field is then stored row-wise as a one dimensional array in the MDV file. There is some loss of information in this binning process, but the result is a saving in disk space of almost 75% since each pixel is then stored as an 8 bit character rather than a 32 bit floating point. On extraction of the MDV file, the precision

preserved in the field is dependent on the range of values in the field. The rainfall reflectivity fields observed in this study have a range between 12 and 58dBZ (0.2 up to 150.0 mm/h) so the precision is preserved to approximately 0.2 dBZ in the *worst* case. This equates in rainfall rate, to an uncertainty of approximately 0.004 mm/h (or 2%) at the lower end of the scale up to 5 mm/h (or 2,5%) at the upper end.

When MDV files are ZIP compressed, the storage space is typically reduced by 60% depending on the amount of rainfall on the image. Wetter images are less responsive to compression algorithms. Using *Zipped MDV* files, it is therefore possible to store one level of radar rainfall images, for one year, using approximately 7 gigabytes (40% of 25% of 70 gigabytes) of disk space. Most of the data analysed in this study were supplied on compact disk by the SAWS in this format.

2.5.2 Selective storage

At the start of this research in 1997, the entry level hard disk for personal computers could store under 2 gigabytes of data so that in order to store a single year's data would require several such disks. A sensible alternative was sought. Although it is possible to write the routines to extract the data and analyse them directly from compact disk, it would be an extremely long and complicated process. Firstly access from compact disk is extremely slow when compared to hard disk access. Secondly the data for a single year spans several disks so user interaction would be required for any large scale analysis.

The solution was one of selective storage. As stated in Section 2.4.3, only the heart of the radar image was required for quantitative analysis in this study. A single pass was made over the data, extracting each file and cropping the 128 x 128 pixel heart. In the early stages of the research, data precision was sacrificed and images were stored in character (8 bit) precision as rainfall rate to the nearest millimetre per hour (i.e. integer values between 0 and 255). This reduced the storage requirements to (128^2 x 320 x 365) 1,9 gigabytes *uncompressed*. In this format, the data can be easily and quickly accessed for analysis. In the latter half of the research, as hard disks became larger and less costly, the data extraction was repeated at floating point precision which requires four times the amount of space.

Further savings in disk space were achieved by only storing *wet* images. During the dry season (winter months) in Bethlehem, very little rain is observed and there is no need to store images recording zero rain. In the case of Bethlehem this cuts the typical required annual storage space by approximately 70% although the savings for a wet climate with indistinct seasonal variation will be considerably less dramatic; however in 2003 this problem is becoming less pressing.

2.5.3 Compression of radar data

Due to its erratic nature, radar data does not always respond well to compression algorithms and on occasions it can increase the file size, however a simple and effective means of compressing these data is to store only non zero pixels. The two dimensional rainfall fields are stored row-wise in a one dimensional array. Choosing a wet/dry threshold of 0.5mm/h (18dBZ), under prevailing South African conditions, a large proportion (usually more than 50%) of a typical 128km rainfield is dry at any instant in time. This is due to the fact that much of the rainfall experienced in the summer rainfall region of South Africa is in the form of isolated thundershowers. Setting all rainfall rates below 0.5mm/h to zero, the data array will usually contain a large proportion of zero values. The storage method adopted by Seed (1992) exploits this fact and stores radar images as a sequence of short *non-zero* arrays. Each non-zero array is preceded by a 32 bit integer *address* corresponding to the position of the first non-zero pixel in the one dimensional array, and a 16 bit integer *length* corresponding to the number of consecutive pixels before the next zero. In this way, radar data storage is typically reduced by approximately 60% depending on the extent of the rainfall recorded on the image.

2.5.4 Image File Nomenclature

A point of great importance when dealing with a large number of image files is that they should have sensible nomenclature so that they can be automatically accessed by computer code. A fixed number of characters in the filename is helpful for file referencing in a program and it is easier to sort and browse. The file nomenclature that was adopted for the data depending on the length of accumulation:

- Instantaneous images (16 character file name)

  First character - **S, C or X** (the type of radar)

  Second to fifth characters - the **year** (four characters for 2000 compliance)

  Sixth character - the **dash** character

  Seventh to twelfth characters - **minutes from the start of the year**

  (From 0 to 527040 minutes in a *leap* year.)

  Followed by the four character file extension (***.bmp*** or ***.flt***)

- Hourly images (14 character file name)

  First character - **H**

  Second to fifth characters - the **year**

  Sixth character - the **dash** character

  Seventh to tenth characters - **hours from the start of the year**

  (From 0 to 8784 hours)

  Followed by the four character file extension (***.bmp*** or ***.flt***)

- Daily images (13 character file name)

  First character - **D**

  Second to fifth characters - the **year**

  Sixth character - the **dash** character

  Seventh to ninth characters – **day** (From 1 to 366 days)

  Followed by the four character file extension (***.bmp*** or ***.flt***)

- Monthly images (12 character file name)

  First character - **M**

  Second to fifth characters - the **year**

  Sixth character - the **dash** character

  Seventh and eighth characters - **month**

  Followed by the four character file extension (***.bmp*** or ***.flt***)

- Annual Images (8 character file name)

  First character - **M**

  Second to fifth characters - the **year**

  Followed by the four character file extension (***.bmp*** or ***.flt***)

An example, an image filename could be:

S1996-034591.bmp

which would be a bitmap image (integer precision), recorded with an S-band radar in the 34591st minute of the year 1996. The same filename with a ***.flt*** extension would contain the same data, but in floating point precision.

2.6 A CAUTIONARY NOTE ON DATA QUALITY

Two datasets were used during the course of this study, Pre-1997 dataset and Post-1997 dataset. The Pre-1997 dataset comprised integer precision data, extracted using the projection technique of Section 2.4.2 and stored using the compression algorithm of Section 2.5.3. The routines used to extract and store the radar data were developed in the context of radar rainfall experiments testing the feasibility of Cloud Seeding. Output from the routines was therefore carefully verified and validated against networks of daily and tipping-bucket raingauges. The result was a simple, reliable, tried and tested dataset which was used in this study until 1999.

With the development of more sophisticated techniques for extracting and storing the radar data (DISPLACE averaging, MDV format, ground clutter removal) new routines had to be written to incorporate these algorithms. The new routines were written by the SAWS and new data were made available for analysis. After a significant amount of analysis it became apparent that the post-1997 MDV radar data overestimated the rain-rate on average, by a factor of 2. After a considerable amount of investigation into the cause of this error, the advice given by the SAWS was to simply scale the entire dataset to achieve the rainfall total recorded by the daily raingauges in the region. The validity of scaling the data in this way is questionable since the marginal distribution of the radar data is severely skew, but it served as a necessary work-around in order to continue with the research. For the purposes of this study, the scaled data are assumed to accurately represent the observed rainfall for the region.

2.7 SPATIAL DATA VISUALISATION

Perhaps the most attractive property of two dimensional data is that it is easily expressed as an image whose qualities can be efficiently analysed and judged by a trained human eye. In the context of radar rainfall data, sensibly displayed data will enhance the process of making decisions for real time application of rainfall data such as flood control or cloud seeding.

The 256 colour Windows Bitmap was chosen as being a suitable image file format to display the data. There are several reasons for this choice the first being that it is an uncompressed, simple file format which makes it easy to manipulate during the analysis process. Since the 256 colour bitmap uses one byte per pixel, the character precision raw data can be stored as an image without any loss of precision. Almost any Windows based graphics software will display, edit and convert the bitmap into any other desired image format. This makes it easily portable and accessible to most computers. The indexed colour palette is also readily edited in most simple graphics programs to give a different viewpoint without any manipulation of the data. The binary structure of the bitmap will be explained by simple example in Section 2.7.1.

2.7.1 Structure of the windows bitmap

To consider the structure of the windows bitmap, it is best to view the file as an array of hexadecimal (base 16) integers as shown in Figure 2.9.

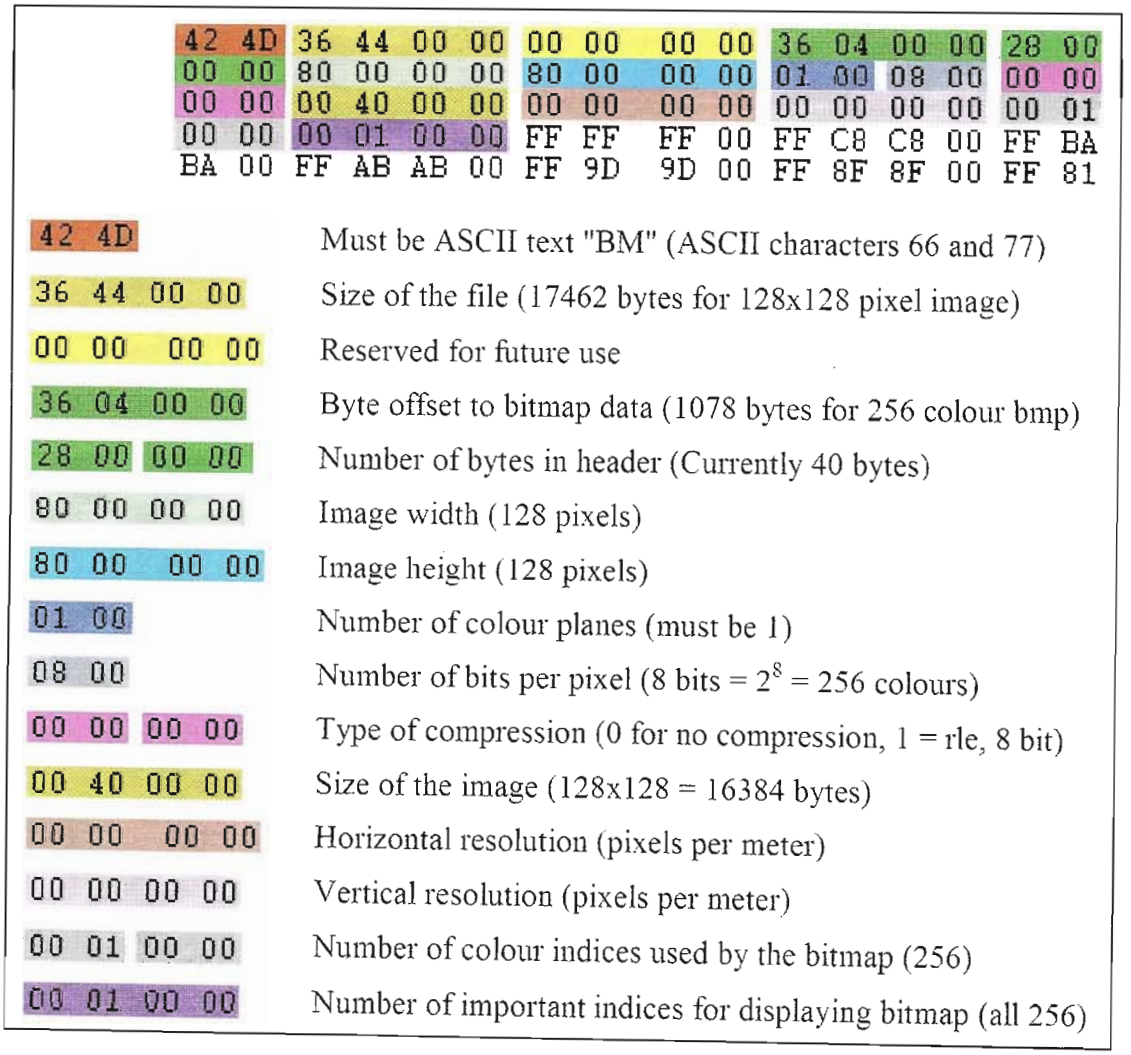

Figure 2.9 - Structure of the Windows Bitmap file header

Each pair of Hex numbers represents a single byte (8 bits) of data and has a range from Hex 00 up to Hex FF which equates to decimal 0 to 255. A 32 bit integer is stored as four consecutive bytes. The Windows operating system stores data in reverse byte order (least significant first) so that the decimal integer 306 263 for example, which converts to a hexadecimal integer of 4AC57, is stored as the 32 bit integer 57 AC 04 00.

The 256 colour bitmap structure comprises a 14 byte header, followed by 40 bytes of image information, then a 1024 byte colour palette (64 byte palette for 16 colour bitmap) and finally the data. The 1024 byte colour palette defines each colour index for the image at 24 bit colour depth. In the example presented in Figure 2.9, the first colour is defined immediately following the bitmap information as FF FF FF 00 which are the blue, green and red colour components of the colour respectively terminated by the 00 byte. This equates to the decimal 255 255 255 which is full blue, green and red intensities which combine to form the colour white. In this bitmap, any data which have value zero will be represented on the image as a white pixel. The data are arranged row-wise starting from the bottom left hand corner of the image. Each row is zero packed to a 32 bit (4 byte) word length. As an example, a 256 colour image which is 21 pixels wide, will be stored row-wise as 21 consecutive bytes of data followed by 3 zero bytes (since 24 is the first multiple of 4 which is greater than 21).

2.7.2 Selecting the colour palette to display a radar rainfield

Since the colour palette is a fixed size and format for a 256 colour bitmap, it is possible to change the way in which the image is displayed by simply redefining the colour palette in the binary file. As explained in Section 2.7.1, this is done by altering the mixture of primary colours (RGB) in each of the 256 indices. Manipulating the colour palette in binary is an extremely laborious process. A shortcut is to read a single image into any good graphics package and edit the palette using the tools provided in the graphics package. Once a satisfactory palette has been defined for the image, it should be saved as a new 256 colour bitmap. In order to change the palettes of a series of images, a very simple program can be written which opens the image with the newly defined palette and copies the first 1078 bytes into a buffer. Each image in the series is then opened in turn and the first 1078 bytes of the file is over written with the contents of the buffer. Figure 2.10 shows an image which is drawn using four different colour palettes.

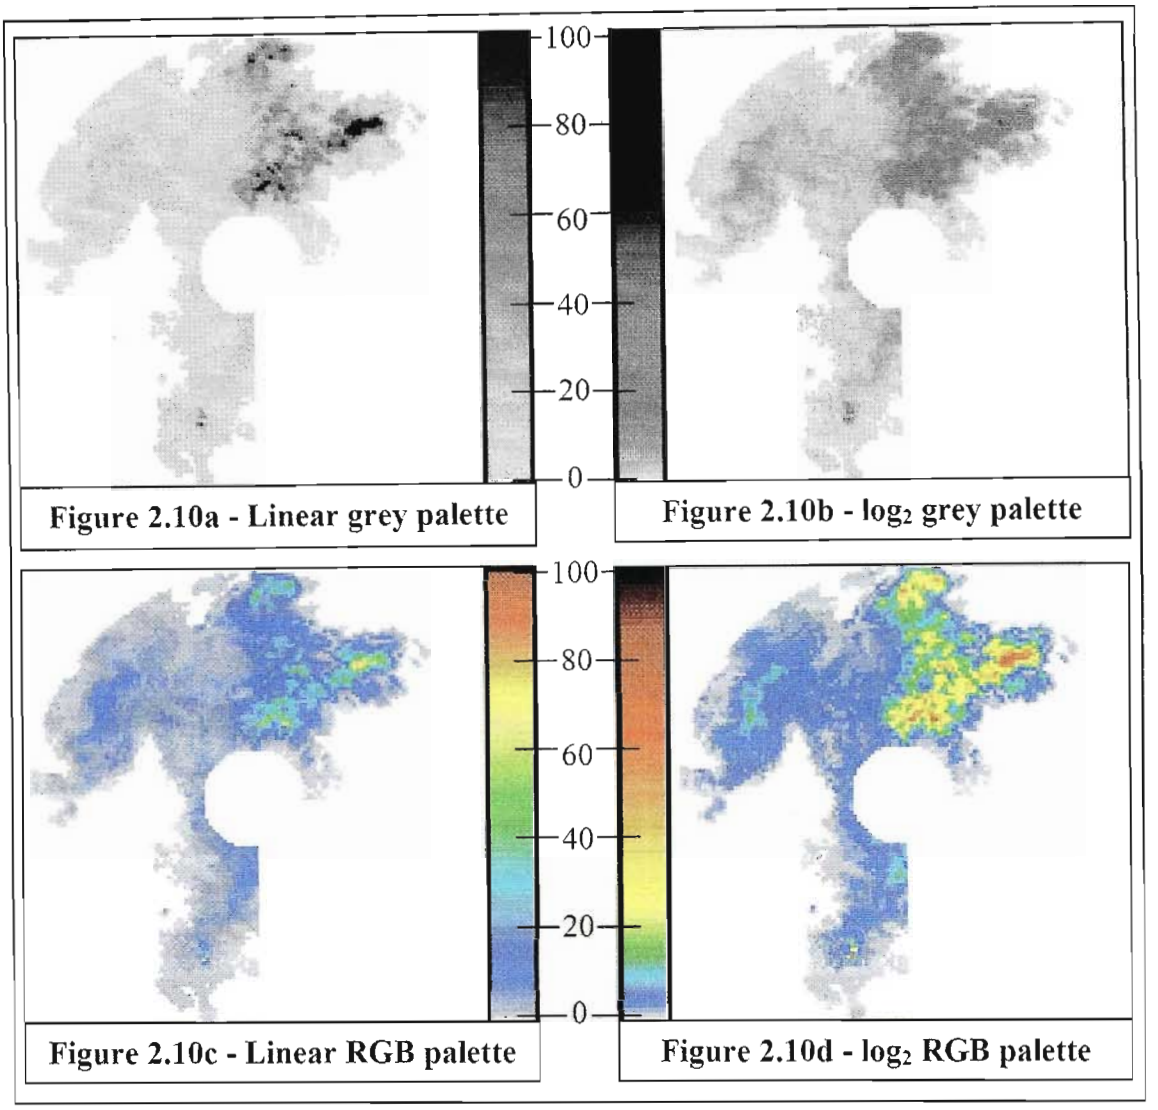

Figure 2.10 – Use of different colour palettes to display a radar rainfield

All palettes in Figure 2.10 have a range of 0-100mm/h, but the left hand images are drawn to a linear scale and the right hand images are drawn to a logarithmic scale. In linear greyscale (Figure 2.10a) it is very difficult to visualise the structure of the image because the human eye is not able to distinguish between the neighbouring intensities in the palette. Figure 2.10a serves to illustrate a poor definition of colour palette. The logarithmic greyscale (Figure 2.10b) provides a better picture of the image structure, but is still not able to convey the detail of the image structure, particularly at the higher rainfall intensities. Figure 2.10b was found to be the best choice of palette when limited to greyscale.

The use of colour proved to be far superior to greyscale in the communication of image structure. The linear Red-Green-Blue (RGB) palette of Figure 2.10c provides an image

which gives a clear idea of the distribution of rainfall intensity on a image and is particularly useful when considering the structure of less intense storm cells in the image. The logarithmic RGB palette better describes the structure of the more intense storm cells on the image although intuitively it gives the impression of a very intense weather system. When qualitatively comparing images, they should be viewed using both these palettes.

2.8 UPPER LEVEL WIND DATA

An important aspect of a space-time rainfall model is the direction in which the storm propagates. With a large radar data archive it would be possible to measure distribution of storm speed and direction from the radar data. However, since the MRL-5 radar has only been operational since 1994 very little information can be measured in this regard.

The advice offered by the meteorologists at the SAWS was to use the upper level wind data as an indicator of storm propagation. It was suggested that the 500 millibar wind data would be suitable for this purpose and this corresponds to a typical altitude of 5,8km. The 500 millibar wind speed and direction is measured using a weather balloon which is released from Bethlehem at midday every day. Daily weather balloon data for the past 20 years were obtained from the SAWS.

2.9 SUMMARY

This chapter has served to introduce the datasets used in this study. The extent and limitations of the datasets have been highlighted and the methods used to store and present the data have been outlined.

Chapter 3

Data Analysis – Statistics, Techniques and Results

3.1 INTRODUCTION

The purpose of this chapter is threefold. The first objective is to define the statistics used to describe the data. In most cases the statistical models used in the analysis are well known in the field of hydrology and should require only a brief discussion, however there are important, subtle details which must be clarified. The second objective is to describe the techniques used to measure the statistics. Again, many of the techniques are well known and will require only a short explanation, but in some cases techniques have been refined and devised which are specific to this study. The final objective is to present the results of the data analysis and to comment on the results based on the experience gained.

3.2 IMPORTANT PROPERTIES OF A SPACE -TIME RAINFALL MODEL

A first step in the analysis process is to identify the statistical properties of the data which are important in the context of their application. The "String of Beads" model is designed to simulate radar rainfall images for any region at 1km, five minute, spatial and temporal resolutions respectively, for use as a tool in hydrological design.

In truth, radar rainfall images present a slightly distorted space-time picture of the instantaneous rainfall rate in the rainfield. In the case of the CAPPI data obtained from the SAWS at Bethlehem, the rainfall rate at the outer pixels of the image are sampled first and the last of the pixels (at the centre of the image) are sampled approximately 4,5 minutes later. This phenomenon will be neglected and the radar rainfall images will be assumed to be a ***simultaneous, instantaneous*** representation of the rainfall rate over the entire field.

During the analysis process, radar data are considered in subsets which vary in spatial and temporal dimension and must be defined clearly in order to avoid confusion. Instantaneous radar images represent a subsets of the radar data set, each of which has two spatial dimensions and no temporal dimension. From a modelling perspective, four important properties of these instantaneous images are:

- The *Instantaneous Image Mean Flux* (IMF_i) or average rainfall rate in millimetres per hour over the instantaneous image.
- The *Instantaneous Wet Area Ratio* (WAR_i) defined as the proportion of the image experiencing a rainfall rate in excess of 1,0 mm/h, a *wet instantaneous image* is defined as an image with a WAR_i in excess of 1%.
- The *marginal distribution* of the rainfall rate over the image.
- The *spatial correlation structure* of the rainfield.

These will be referred to as *Instantaneous Image Scale Properties*, denoted by the subscript "*i*". The corresponding statistics will be referred to a the *Instantaneous Image Scale Statistics* and the model parameters which determine these properties will be the *Image Scale Parameters*. A *rainfall event* is defined as a sequence of consecutive *wet instantaneous images*.

The speed and direction with which the storm moves across the sample area is another important characteristic in a hydrological model since this can dramatically effect the time to peak of the runoff hydrograph and consequently the intensity of a flood. This is measured on average over the entire rainfield and is termed the *Mean Field Advection*.

By averaging the rainfall rate at a pixel in a sequence of images over an hour, the cumulative rainfall in millimetres, at that pixel, for that hour is obtained. These rainfall accumulations can be summed to produce hourly, daily, monthly and yearly images which will be referred to as *cumulative images* and whose properties are defined in a similar manner to those of the *instantaneous images*, only in terms of cumulative rainfall rather than rainfall rate. Of particular interest in this regard, are the daily images for which it is possible to measure:

- The *Daily Image Mean Flux* (IMF_d) or average rainfall rate in millimetres per day over the image.
- The *Daily Wet Area Ratio* (WAR_d) defined as the proportion of the image which receives rainfall in excess of 3,0 mm during the day, a *wet daily image* is defined as an image with a WAR_d in excess of 10%.
- The *marginal distribution* of the rainfall over the daily image.
- The *spatial correlation structure* of the daily rainfield.

These will be referred to as *Daily Image Scale Properties*, denoted by the subscript *"d"*, with corresponding statistics. Similar statistics are defined for hourly, monthly and yearly cumulative images so that IMF_m would refer to average rainfall rate in millimetres per month over the image, for example. A sequence of consecutive *wet daily images* forms a *wet run* and conversely, a sequence of consecutive *dry daily images* forms a *dry run*, so that daily rainfall is comprised of an alternating sequence of wet and dry runs. Daily classification is an idea taken from Pegram and Seed (1998) who define three daily states for a network of daily raingauges, a *dry day* (less than 3% of gauges report rain), a *scattered rain day* (more than 3% report rain, but less than 50% report more than 5mm) and a *general rain day* (more than 50 % report rainfall in excess of 5mm). Their classifications were based on the work of Court (1979).

Another crucial aspect of the model is its temporal structure. As for the spatial statistics, this too has a wide range of scales which must be considered. As an example, the time series of rainfall rate at a pixel in the image must be analysed in the radar data set and preserved in the modelling process. Equally important would be the time series of daily rainfall accumulations at the pixel, extracted from the daily images mentioned earlier in this section. Statistics of these time series are referred to as *Pixel Scale Statistics*. This is the closest equivalent to daily raingauge data which can be inferred from radar data.

Image Scale Statistics are measured using all of the data on individual images whereas *Pixel Scale Statistics* are measured using selected data from a sequence of images, but the images are not necessarily instantaneous images and the pixels are not necessarily 1 km pixels. In fact, there exists an entire range of space-time combinations which must be considered in the analysis. For this study, the spatial range is from 1km up to 128km and the temporal range from 5 minutes upwards. The spatial and temporal correlation

structure for these various combinations of pixel size and sampling interval is referred to as the *upscaling structure* of the data and this too must be reproduced in the simulated data.

The statistics discussed so far in this section are primary statistics in the sense that they can be measured directly from the data, but there is another set of statistics which are required for the definition of the model. These are the secondary statistics which are derived from the primary statistics and they are dependent on the definitions stated up to this point. Statistics which fall into this category are:

- the *Event Statistics* which describe the arrival and duration of the rainfall events, as well as the temporal behaviour of the image scale statistics during a rainfall event.
- the *Run Length Statistics* which describe the distribution of daily wet and dry runs.

The secondary statistics are used to bind the model together on the large scales by describing the behaviour of the large scale primary statistics. This provides an efficient means of measuring long term trends in the data such as seasonal and annual variation.

The statistical properties mentioned in this section form the principle components used to define and test the "String of Beads" model. The following sections of this chapter will consider each of these statistics in turn and present the methods used and the results obtained in the analysis process.

3.3 GAUGE ANALYSIS

Raingauge data represent the only long term quantitative record of the rainfall available in South Africa and has become the trusted *ground truth* amongst the hydrological community. Raingauge records for the Bethlehem region begin in the 1920s whereas the radar database begins in 1996. In order to incorporate seasonal and annual variation into the model it is necessary to make use of the gauge data and to develop a sensible means of relating radar data to gauge data. The major problem in this regard is one of spatial scale.

3.3.1 A Question of Scale

Raingauge data samples an area of 10^{-2} m^2 as opposed to radar image data with a pixel sampling volume of 10^9 m^3. If it is assumed that the rainfall rate does not vary with

altitude, radar image data could be considered to be sampled with a two dimensional pixel of area 10^6 m^2. In both cases, there is a severe mismatch in spatial scale. Due to the highly variant spatial nature of rainfall, comparing these two datasets is not likely to yield sensible results unless the comparison is made over a long period of time. Short term (hourly) gauge-radar comparisons have been attempted in Bethlehem (van Heerden and Steyn, 1999) with inconclusive results.

In the same study, van Heerden and Steyn considered annual accumulations using a network of raingauges with more success, but even these result in discrepancies in annual totals of approximately 20%. This problem is particularly noticeable in regions such as Bethlehem which experience a large proportion of the annual rainfall as isolated convective showers. Stratiform rainfall is widespread and less variable than convective rainfall, and consequently the gauge and radar measurements of stratiform rain are more likely to agree.

The infilled daily raingauge network in the study region has an average density of 1 gauge per 200 km^2 which equates to an average spacing between gauges of 14km. If the radar were somehow perfectly calibrated to the raingauges so that the rainfall accumulated on a 1 km pixel over a period exactly matched the rainfall recorded at a gauge on the ground, there would still be a mismatch in sampling scale. For the Bethlehem region this would be of the order of 200 times. This is the best case scenario.

To illustrate the effects of a mismatch of this magnitude, consider the following analysis in which the rainfall was sampled from the daily radar accumulations *at the positions of the 54 daily raingauges* shown in Figure 2.2. The average accumulated daily rainfall on the study area as sampled from the 54 pixels, versus the same for the entire image (10900 pixels), are compared in Figure 3.1 for 300 days.

The 54 pixel average rainfall for the study area had ranges of between 0.02mm and 0.40mm which correspond to the 0.1mm image average (i.e. for 10900 pixels), and between 17mm and 28mm corresponding to the 20mm image average (horizontal dotted lines in Figure 3.1 illustrate the observed range on the abscissa for a given point on the ordinate). This represents a range of uncertainty of 400% at the 0.1mm and 50% at the 20mm image averages respectively. An image average of 20mm is measured as a result

of a large system with continuous rainfall and an average of 0.1mm is as a result of an isolated convective event. This small analysis illustrates the limitations of measuring rainfall (particularly isolated convective rainfall) using raingauges. Comparing real daily raingauge data to radar cumulative pixel data will reveal an even more dire situation, van Heerden and Steyn (1999) show discrepancies between these two daily datasets (at a pixel) which are of the order of 60%.

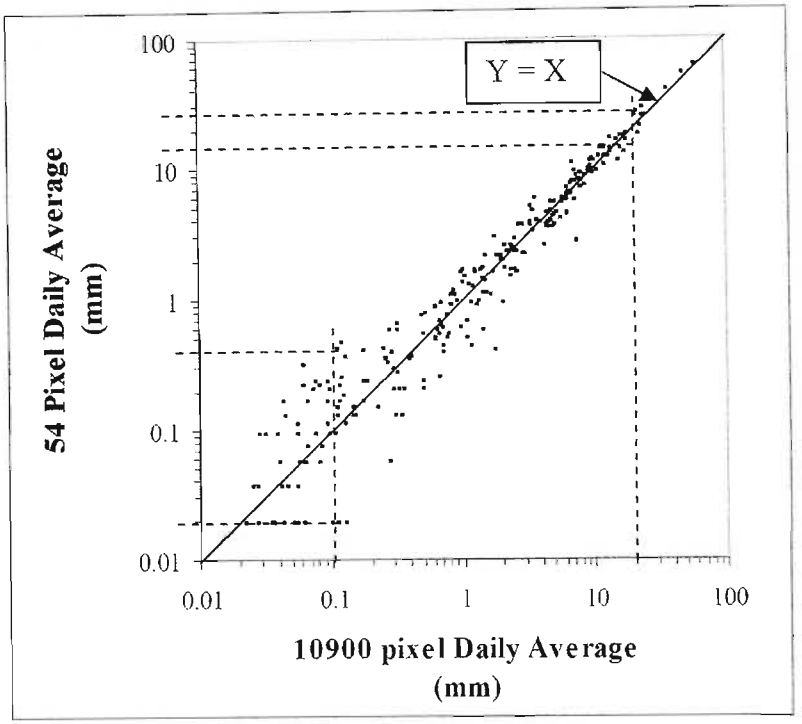

Figure 3.1 - An illustration of pixel sampling error measured from 300 days of daily radar rainfall accumulations during 1998 and 1999

Acknowledging these limitations, there are two encouraging results which are evident in Figure 3.1. The first is that the 54 pixel sample yields a reasonably unbiased estimate of the average daily rainfall (IMF_d) on the sample area. On a smaller scale, it should be reasonable to assume that a raingauge sample will give an unbiased estimate of the rainfall on a radar image pixel. The second is that this unbiased estimate is obtained in spite of the fact that the distribution of raingauges over the study area is not entirely uniform. With reference to Figure 2.2, there is an obvious scarcity of gauges in the North Eastern quadrant.

A linear regression (in normal space) of the IMF_d as estimated by the 54 pixel sample with that estimated by the whole image, the points in Figure 3.1, yields the relationship $y = 1.0296x - 0.1156$ with an R^2 value of 0.980. The gradient, slightly greater than unity

with a negative intercept indicates a slight overestimate of the average rainfall by the 54 pixel sample, for higher daily rainfall and a slight underestimate by the 54 pixel sample for lower daily rainfall.

The *Daily Wet Area Ratio* (WAR_d) is defined in Section 3.2 as the proportion of the image which receives rainfall in excess of 3,0 mm during the day. In a similar way to Figure 3.1, Figure 3.2 illustrates the uncertainty of estimating the WAR_d parameter by sampling at the same 54 pixels.

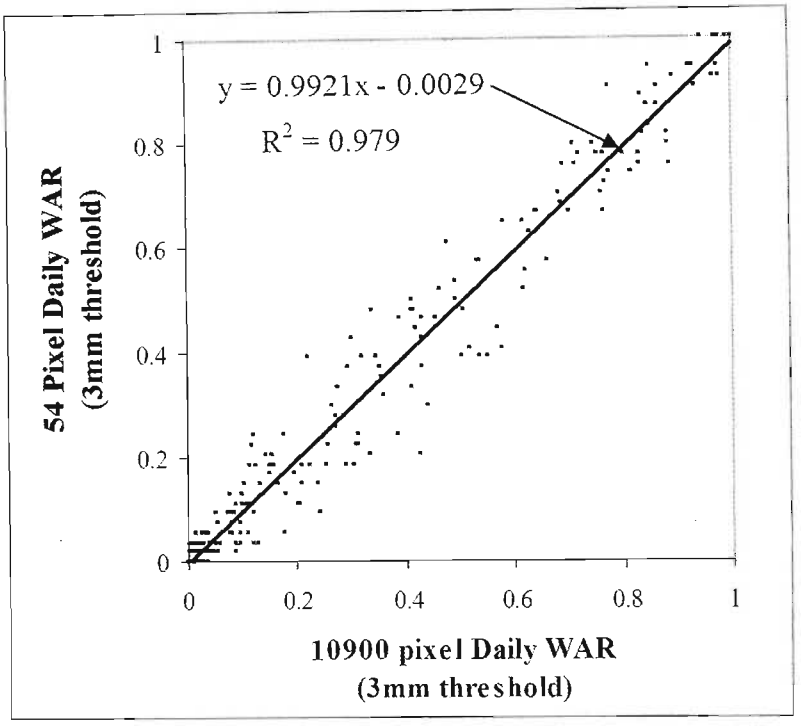

Figure 3.2 - The uncertainty of estimating the Daily Wet Area Ratio by using 54 of the 10900 pixels on the radar image for 300 images during 1998 and 1999

Again, this yields a reasonably unbiased estimate although in this case, both the gradient and the intercept indicate a small underestimation in WAR_d by the 54 pixel sample. This underestimation is of the order of 1% which is negligible considering the possible sources of error in rainfall measurement using gauges or radar. In the light of the results presented in Figure 3.1 and Figure 3.2, the *Daily Wet Area Ratio* and the *Average Daily Rainfall* estimates for the study area, obtained from the 54 raingauges, can be considered to be uncertain but unbiased. By the same argument, monthly and annual accumulations can also be shown to be unbiased and the uncertainty of the sample estimation decreases with longer accumulations.

3.3.2 *Daily Gauge Network Statistics*

Having established that the 54 gauge network provides an unbiased estimate of the rainfall on the study area, the statistics of the rainfall estimates can now be examined. Annual, monthly and daily statistics will be examined in turn. Figure 3.3 shows an estimate of the average annual total rainfall on the study area for the 50 years 1948 to 1997 inclusive.

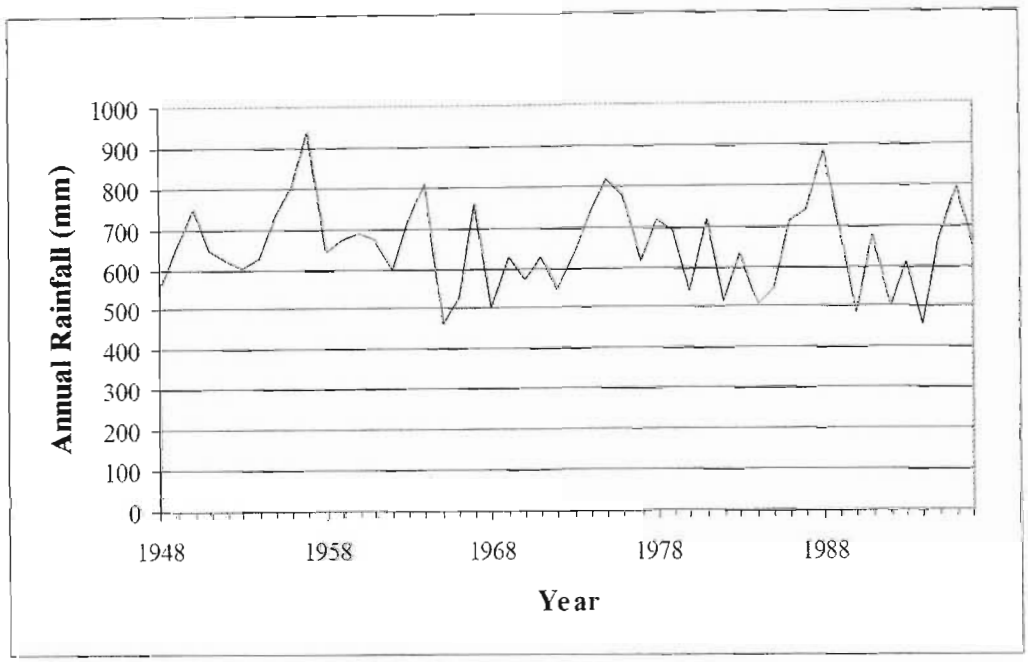

Figure 3.3 – Average annual total rainfall as sampled from the infilled 54 daily gauge network for the 50 year period 1948-1997

These annual totals have a mean of 654mm, a standard deviation of 108mm, a skewness of 0.3 and a kurtosis excess of -0.03. The skewness and kurtosis excess both reasonably close to zero indicate that a Gaussian distribution is a reasonable approximation for the marginal distribution of the total annual rainfall on the study area (although the total annual rainfall can obviously never be negative, zero is 6 standard deviations below the mean, and corresponds to a probability of less than 10^{-9}). Using this distribution it is possible to estimate the recurrence interval of wet and dry years, so a 1 in 50 year *wet year* would have a 2% probability of occurrence which corresponds to an annual rainfall total of 876mm and the corresponding *dry year* would be 432mm.

The serial correlation and power spectrum of this time series of annual totals are shown in Figure 3.4 and Figure 3.5 respectively. With only 50 data points it is not possible to draw any firm conclusions from these two analyses, but there is no *obvious* serial

correlation structure in the data. This is indicated by the fact that the serial correlation in Figure 3.4 dies off almost to zero for a lag of one year, and then fluctuates randomly about zero.

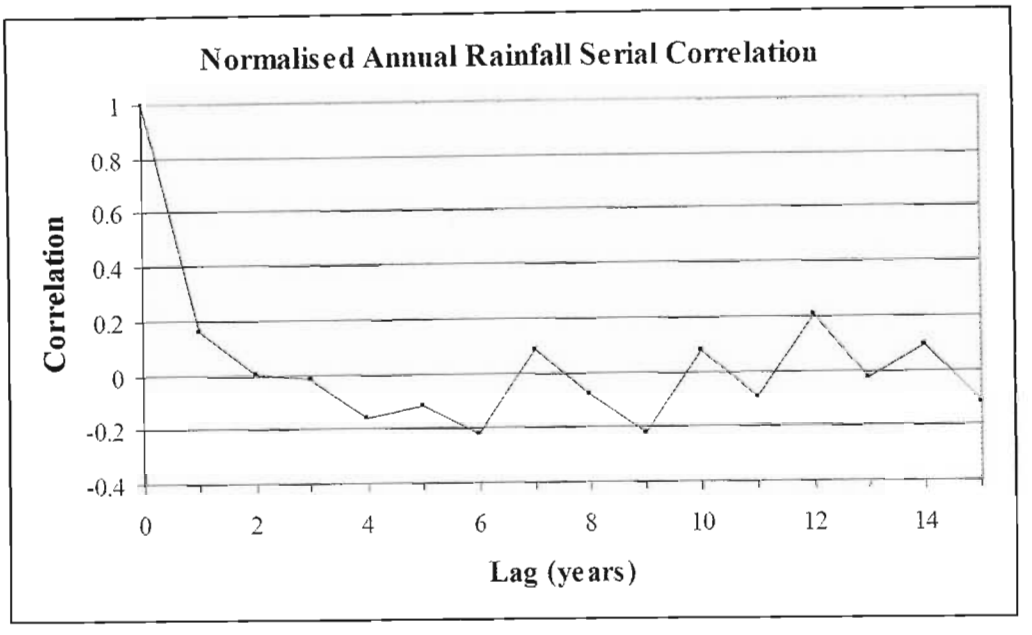

Figure 3.4 - Serial correlation of the time series of normalised, average annual rainfall total for 54 daily raingauges

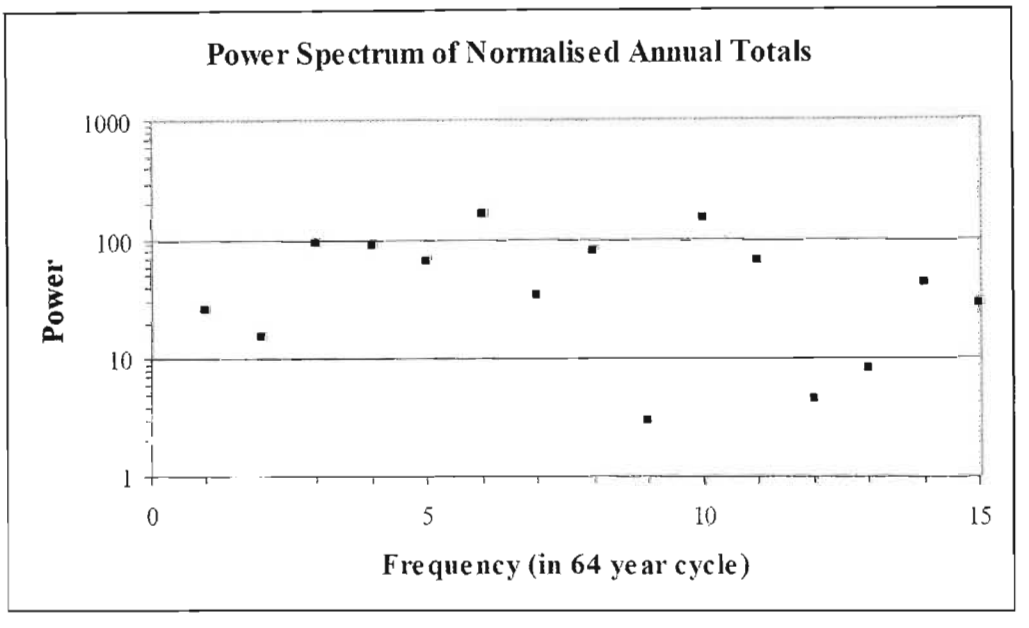

Figure 3.5 - Power spectrum of the time series of normalised, average annual rainfall total for 54 daily raingauges, sampled over 50 years.

The power spectrum also indicates no *obvious* cyclic trend in the data by the fact that the power appears to be randomly distributed amongst the frequencies. A cycle length of 64 is the lowest integer power of two which is longer than 50. There are two small

peaks at the 6th and 10th frequencies, corresponding to 10.6 years 6.4 years respectively, but these are not significant enough to conclude that there is any cyclic trend in annual rainfall totals. Consequently, annual rainfall totals appear to be approximately independently sampled from a near Gaussian distribution.

Another interesting aspect of the daily raingauge data is the average monthly distribution of rainfall. Figure 3.6 illustrates in cumulative form, how the annual rainfall is distributed amongst the months for the 50 year study period.

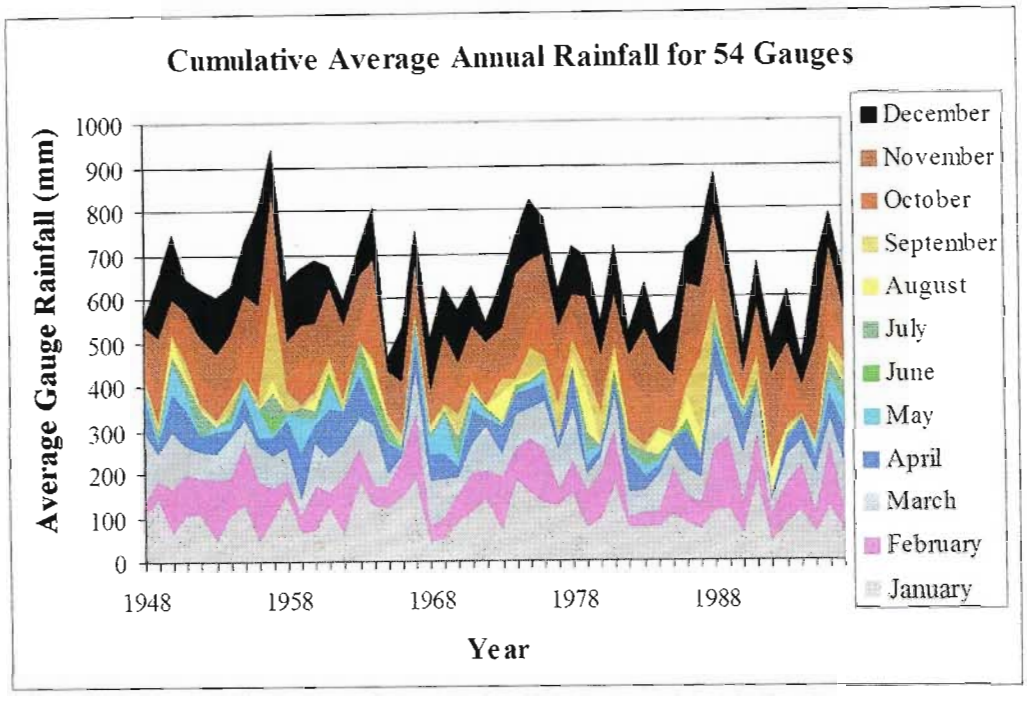

Figure 3.6 - Illustration of distribution of monthly average rainfall totals for 54 raingauges, sampled over 50 years

On average, 45% of the rain in Bethlehem falls in the three Summer months (December to February), 20% in the Autumn, only 5% in the Winter months (July to August) and the remaining 30% in the Spring. Figure 3.7 shows the monthly variation of the marginal distribution of rainfall totals, measured from the 54 daily gauges for the 50 year period. The high skewness and kurtosis measured in the dryer months (May to September) indicate that the two parameter Gaussian distribution is not well suited as an approximation of the marginal distribution of rainfall totals in these months. It should be noted however, that these higher order moments are particularly sensitive to outliers and their closeness to zero in the wetter months indicate that the Gaussian distribution is quite a reasonable approximation to the marginal distribution. With these limitations in mind, these statistics are extracted only to illustrate the seasonal trend in the data.

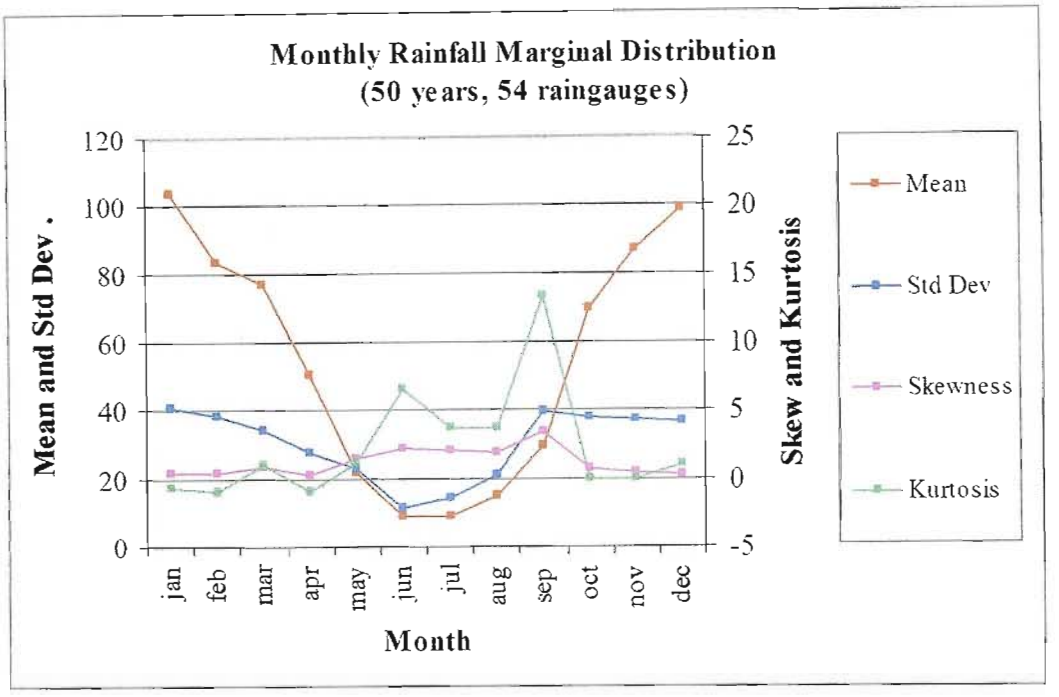

Figure 3.7 - Marginal distribution of monthly rainfall totals

Clearly there is significant seasonal variation in the rainfall at Bethlehem. Of the dryer months, September is outstanding as being prone to occasional large monthly totals. The mean of 29.7, standard deviation of 39.5, skewness of 3.4 and kurtosis excess of 13.3 for September are reduced to 22.6 , 18.7 , 1.4 and 1.9 respectively if the two extreme events of 1957 and 1987 are discounted.

Moving on to some daily statistics for the study area, Figure 3.8 shows the extreme analysis for the average daily rainfall on the study area. Five series are plotted in Figure 3.8, one for each season and one for the combined analysis considering all of the days, independent of the season. The dry winter season is clearly defined with a probability of observing *any* rain in the study area of 21%. By contrast, the probability of receiving any rain in the study area on a summers day stands at 88%. Autumn and spring probabilities are between those of summer and winter as expected. Spring shows a larger probability of high daily rainfall than Autumn over most of the range, but particularly between the 98.0 and 99,5 percentiles. This reinforces observations of monthly totals in Figure 3.6 and Figure 3.7 in which September appears to be particularly prone to occasional large rainfall events.

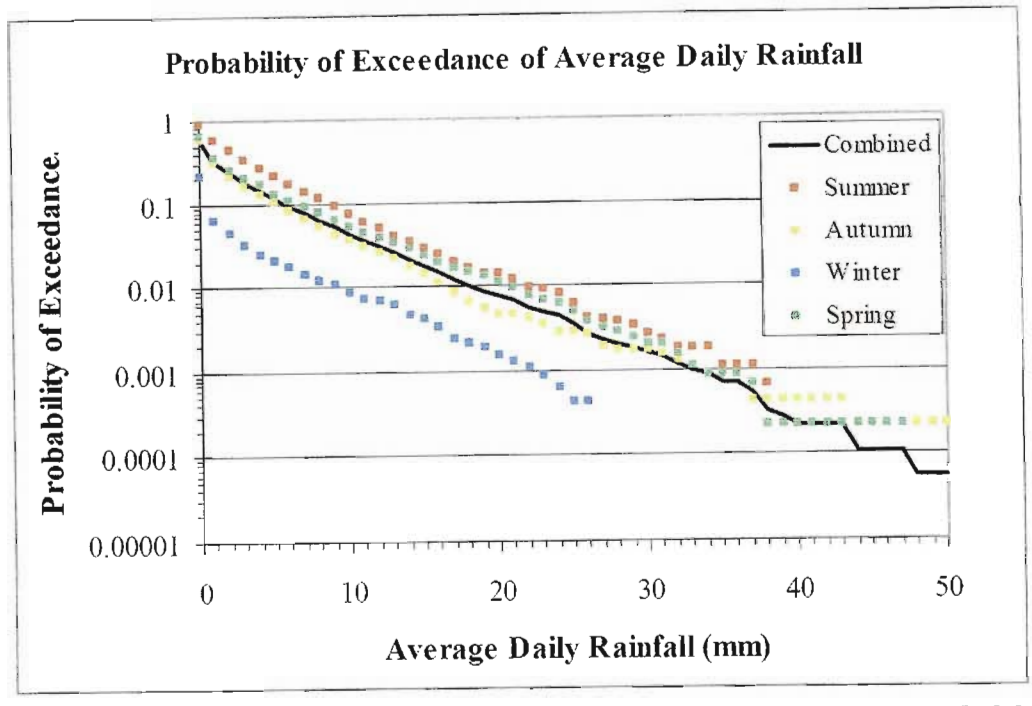

Figure 3.8 - Probability of exceedance of Average Daily Rainfall, sampled by 54 daily raingauges over 50 years

To complement the average daily rainfall, Figure 3.9 considers the probability of exceedence for the *Daily Wet Area Ratio* (WAR_d).

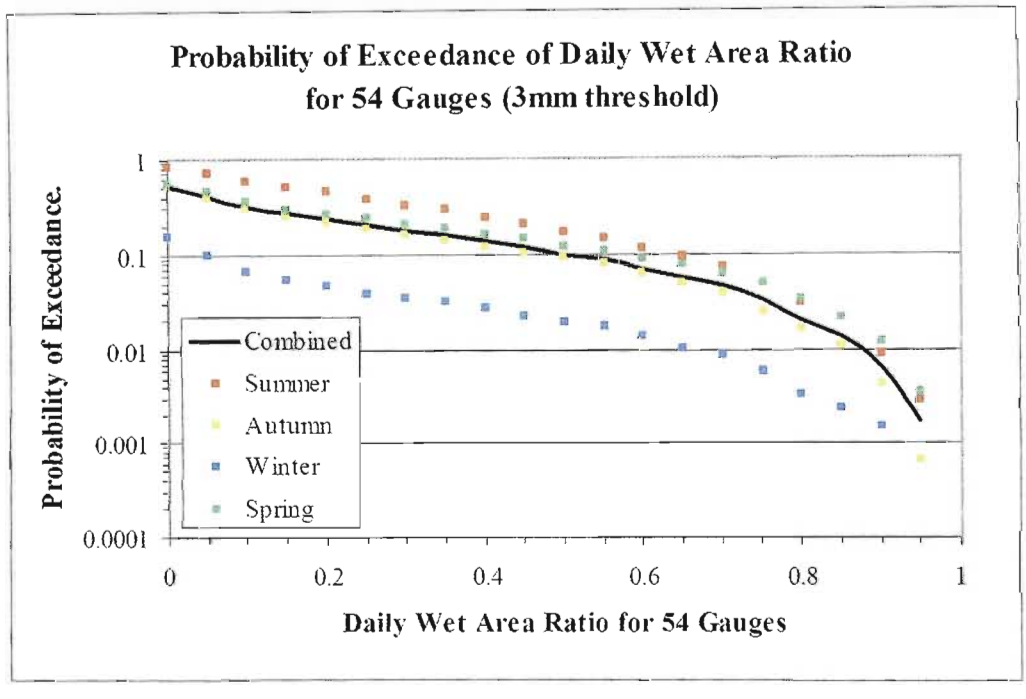

Figure 3.9 - Probability of exceedence of the daily wet area ratio (3mm threshold), estimated from 54 daily gauges sampled over 50 years.

Again the seasonal distinctions are clearly evident in this analysis. For low WAR_d the summer series is distinctly higher that the autumn and spring series and this is as a

result of more convective activity in the summer season. For higher WAR_d the spring and summer series are very similar and this is due the large systems which occur at these times of the year. A useful statistic in the validation of the "String of Beads" model is the probability of a *wet* or *dry day*. A *wet* day is defined in Section 3.2 as a day with a WAR_d in excess of 10%. In this regard, the probability of a wet day in summer is 58%, autumn is 31%, winter is 7% and spring is 35%.

Figure 3.10 shows the total number of dry days observed in the 50 year period on the study area, as measured by the network of 54 daily raingauges.

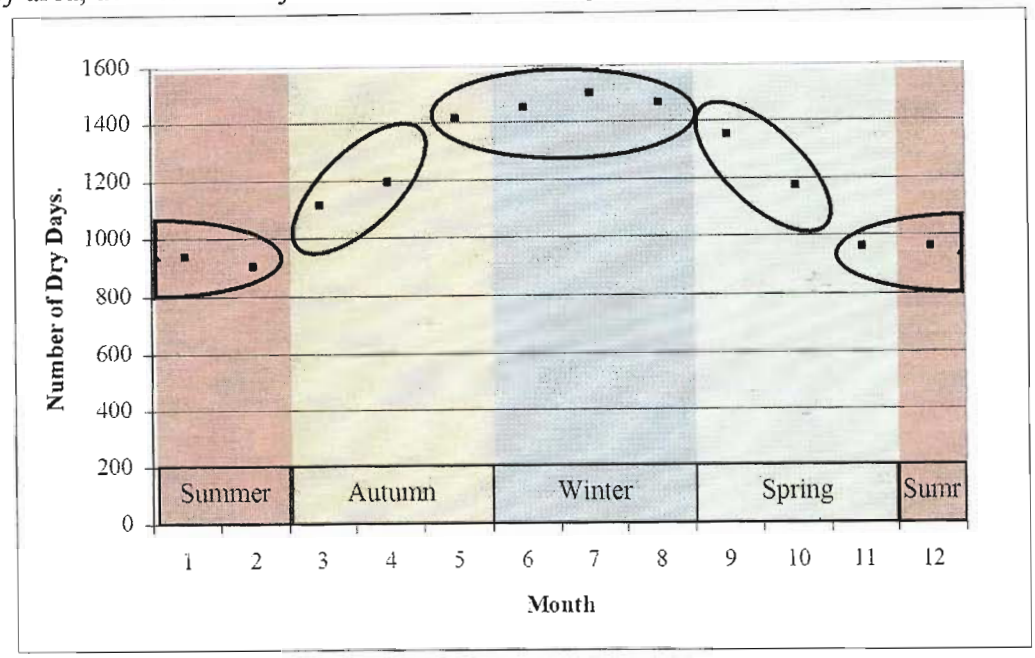

Figure 3.10 - Total number of dry days observed in each month in the 50 year period 1948 to 1997

This figure again clearly shows the seasonal trend in the rainfall experienced in Bethlehem, July being the driest month and January the wettest. The traditional calendar seasons are shown in Figure 3.10 as labelled bands which are convenient for the analysis which follows although it may be sensible to extend the summer months to include November and to extend the winter months to include May leaving March and April for Autumn and September and October for Spring as indicated by the elliptical shapes. A more detailed look at the daily wet-dry time series will reveal the distributions of the runs of consecutive wet days and dry days. Figure 3.11 through Figure 3.14 show the observed runs of wet days for each season, measured using the 3mm, 10% threshold for the 50 year period. These four figures are plotted to the same scale for easy comparison.

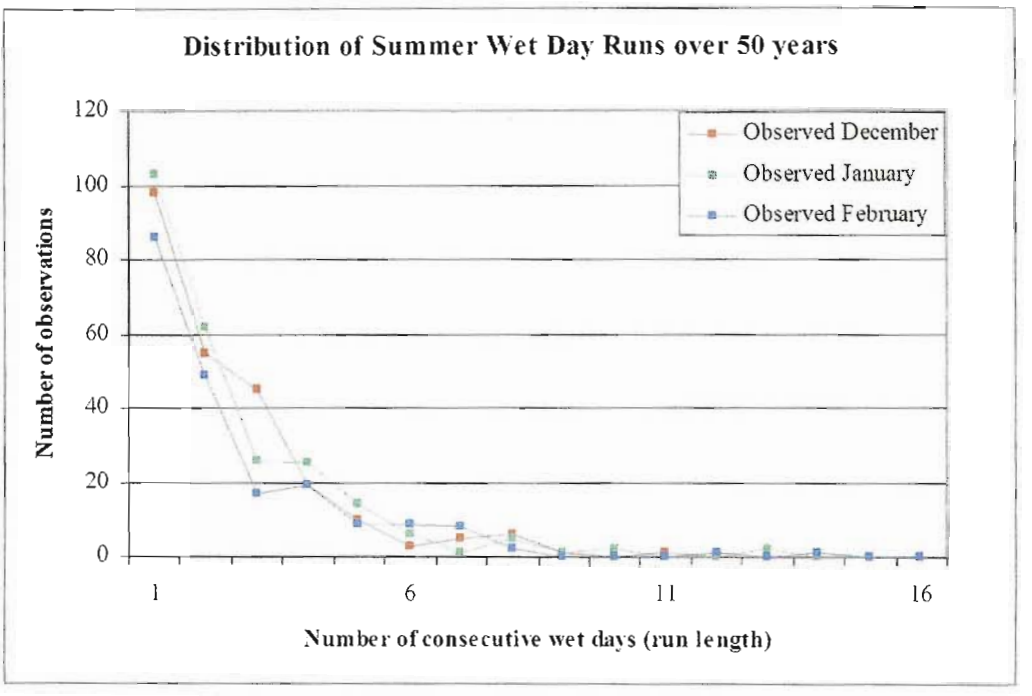

Figure 3.11 - Summer wet day runs observed by the 54 raingauges over the 50 year period

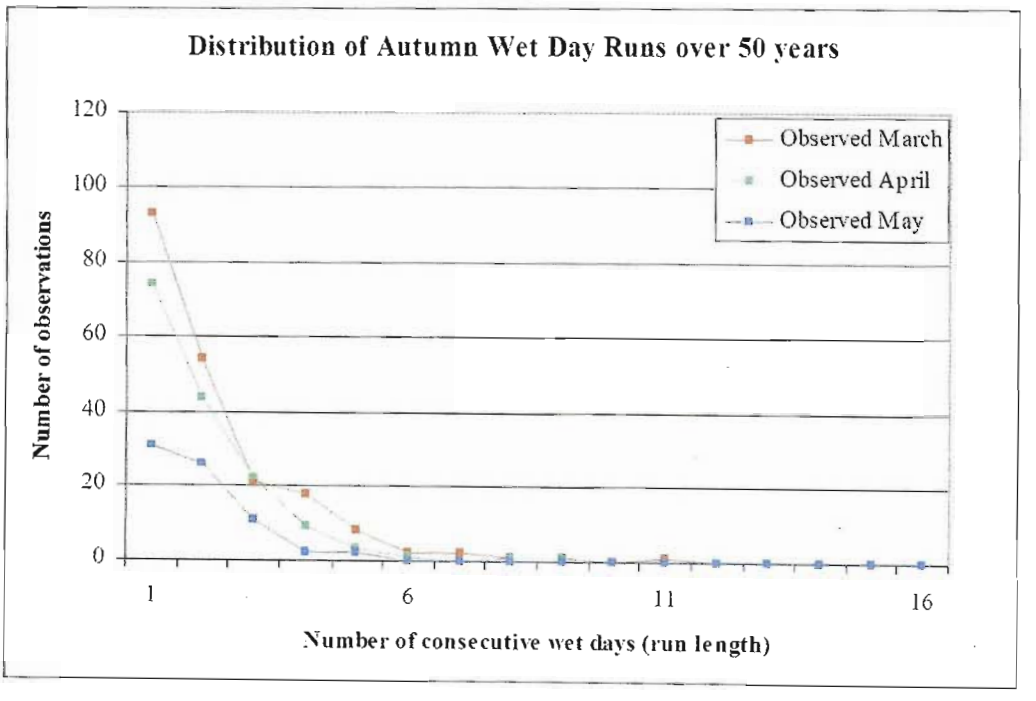

Figure 3.12 – Autumn wet day runs observed by the 54 raingauges over the 50 year period

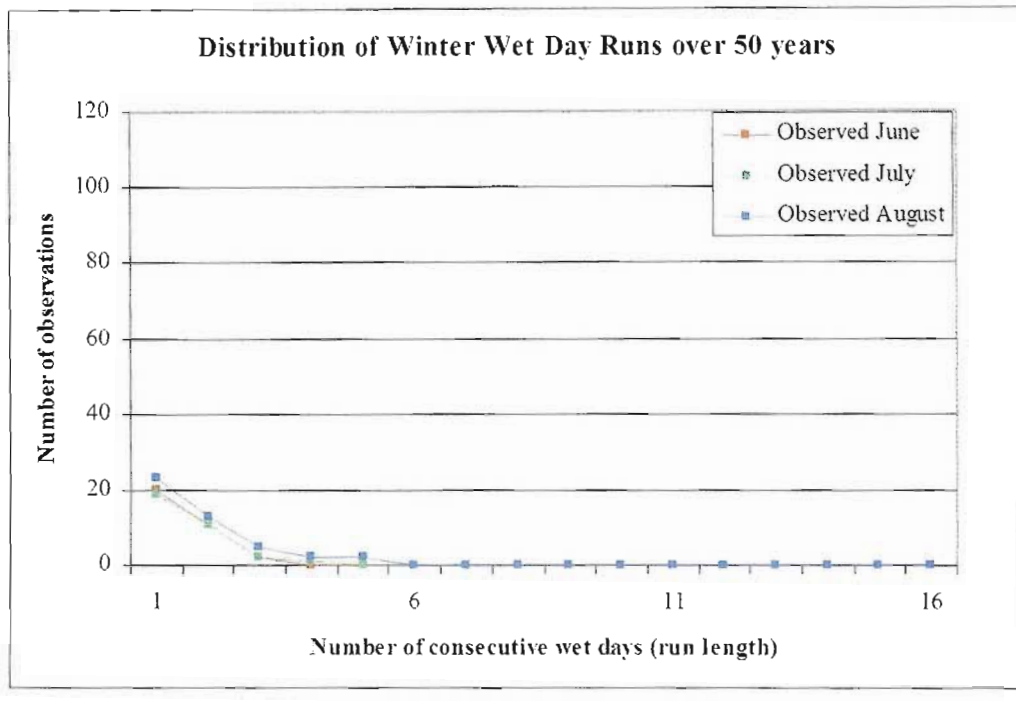

Figure 3.13 – Winter wet day runs observed by the 54 raingauges over the 50 year period

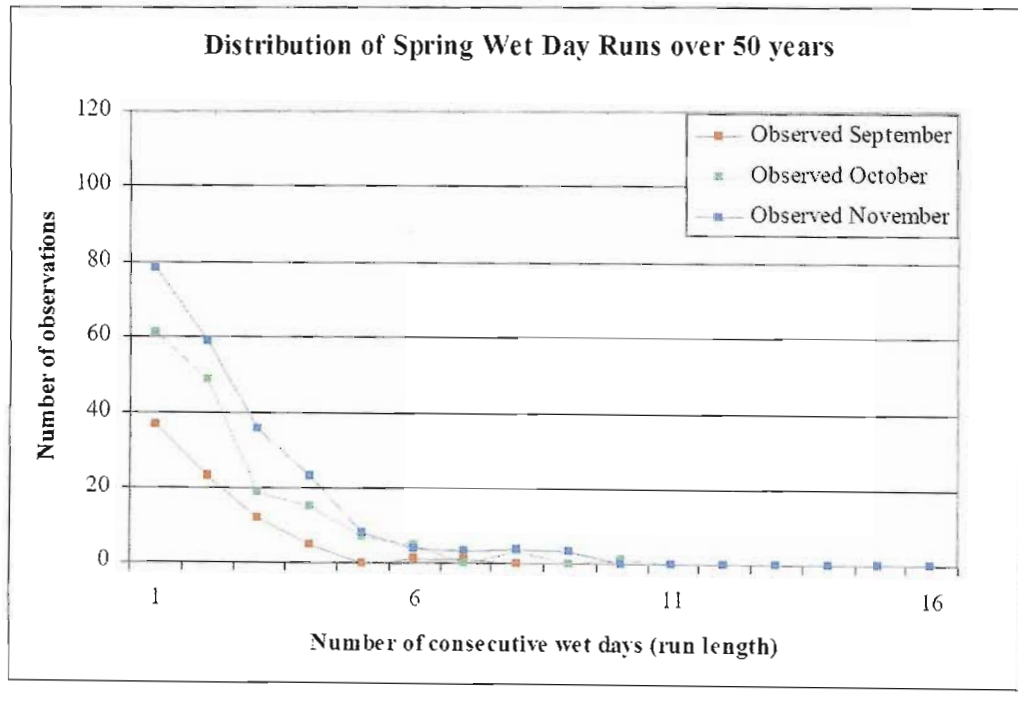

Figure 3.14 - Spring wet day runs observed by the 54 raingauges over the 50 year period

The three summer months appear to have very similar distributions of wet day runs and of the four seasons, most of the long wet runs (in excess of 6 days) are observed in these months. Autumn and spring have similar distributions of wet day runs, reducing almost linearly from around 50 occurrences of one day runs down to approximately 5 occurrences of six day runs in the 50 year period, with very few observations of wet runs in excess of seven days. Spring appears to be slightly more prone to the longer runs than autumn, and autumn shows more of the short (one or two day) wet runs. There are considerably fewer wet day runs in winter, as expected.

The analyses of the daily raingauge data presented in this section serve to illustrate the coarse resolution space-time behaviour of the rainfall in the region. Raingauge data are used as an averaged network sample which is representative of the daily, monthly and annual rainfall experienced on the entire 128km study area. Finer spatial and temporal resolution rainfall data (down to 1km, 5 minute) as measured by radar will be considered in the following sections of this chapter.

3.4 MARGINAL DISTRIBUTION OF RAINFALL ON A RADAR IMAGE

A major benefit of radar data is the fact that it gives a very detailed picture of the structure of the rainfield. For modelling purposes, a quantitative means of describing the rainfield is required and a first step is to consider the marginal distribution of the rainfall intensity at a pixel on an instantaneous radar image. This is a positive distribution which can vary over a wide range of shapes depending on the extent and type of rainfall on the image.

It has been shown that the *lognormal distribution* can be used effectively in describing the marginal distribution of rainfall rates over an area as estimated by both satellite data (Bell, 1987) and radar data (Crane, 1990). The formal definition and properties of the two parameter lognormal distribution are described by Aitchison and Brown (1957) and paraphrased In Appendix B for completeness and clarity of exposition.

For a lognormally distributed variate $X \sim \Lambda(\mu, \sigma^2)$, the two parameters of the lognormal distribution μ and σ^2, are related to the mean μ_x and standard deviation σ_x of the untransformed X-variate by the equations

$$\mu_x = exp(\mu + {}^1/_2 \sigma^2) \quad (3.1)$$

$$\sigma_x = exp(2\mu + \sigma^2)(exp(\sigma^2) - 1) \quad (3.2)$$

and the median is at $x = \exp(\mu)$. The maximum likelihood estimators of the parameters μ and σ^2, m and s^2 respectively, are given by:

$$m = \frac{1}{n}\sum_{i=1}^{n} \log x_i \quad (3.3)$$

and

$$s^2 = \frac{1}{n}\sum_{i=1}^{n} (\log x_i - m)^2 \quad (3.4)$$

Equations 3.3 and 3.4 can be used explicitly to find the maximum likelihood estimation of the parameters for continuous data sets (i.e. data sets in which each x_n is a real number given to its full precision). For integer precision data, such as the 8 bit precision radar data supplied by the SAWS, the multi-dimensional downhill simplex method of Nelder and Mead (1965) as given by Press et al. (1992) was used to maximise the log likelihood function L shown in Equation 3.5 with respect to the parameters μ and σ for each image.

$$L = \sum_i n_i \log\left\{N\left(\frac{y_i - \mu}{\sigma}\middle|(0,1)\right) - N\left(\frac{y_{i-1} - \mu}{\sigma}\middle|(0,1)\right)\right\} \quad (3.5)$$

where $i = 1, 2, \ldots, 255$

y_i = log (rainfall rate i mm/h)

n_i = number of pixels in the masked image recording a rainfall rate of i mm/h

μ = the mean of the logs of the rainfall rates in the image

σ = the standard deviation of the logs of the rainfall rates in the image.

To illustrate the range of distributions observed in the radar data set, Figure 3.15 shows five images (Figure 3.15a to e) varying from intense, scattered rainfall (usually as a result of convective activity) up to light, general rainfall (stratiform rainfall, the result of large frontal systems). These data were obtained in the earlier part of the study and are in integer precision, extracted from the volume scan using the projection technique discussed in Section 2.4.2.

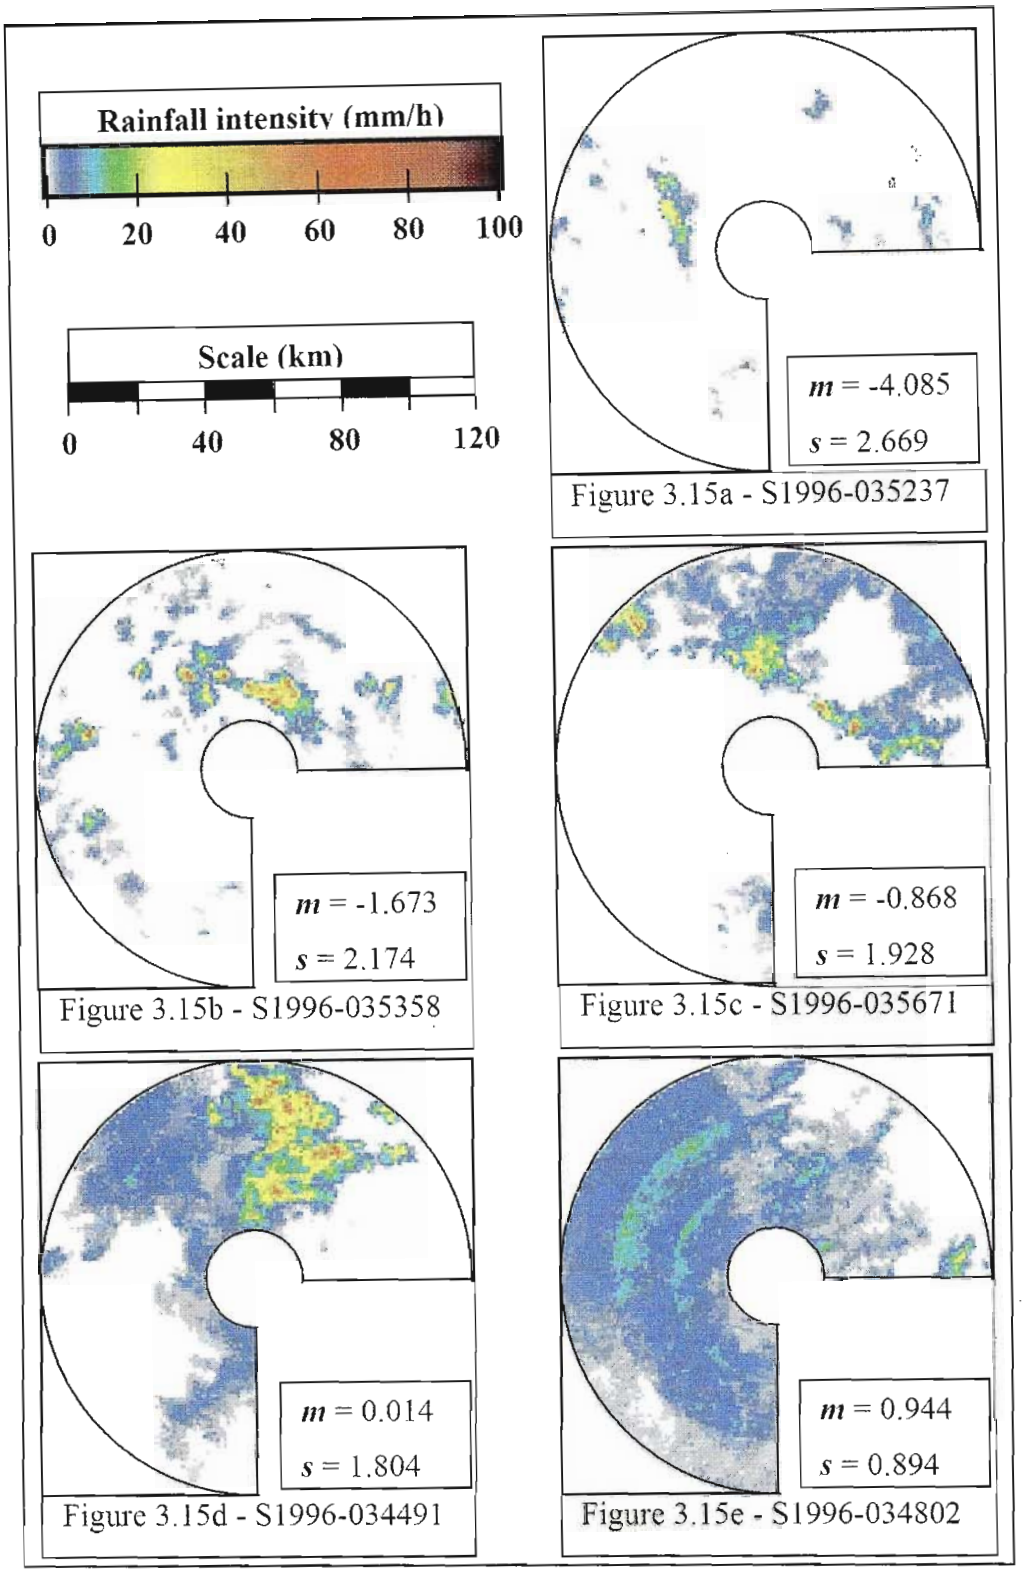

Figure 3.15 - Illustration of the types of rainfall experienced in Bethlehem

The concentric rings evident in Figure 3.15e also serve to illustrate the *Bright Band* phenomenon discussed in Section 2.4.4 which is revealed as a segmentally increasing rainfall rate with range (and therefore altitude) in steps.

The observed histogram of rainfall intensity and the expected histogram which corresponds to the maximum likelihood approximation of the lognormal distribution for each of the five images of Figure 3.15, are plotted in Figure 3.16 to Figure 3.20. To enable a critical examination of the fit, the histograms are plotted with a logarithmic abscissa and the final bin is lumped for the most intense 0.25% of the number of pixels on the image. The labels on the ordinate axis represent the lower limit of the bin (for grouped data) so that the first bin which is labelled 0mm/h includes all rainfall rates between 0 and 1mm/h (i.e. includes all values below 1 mm/h which convert to an integer of zero).

Figure 3.15a shows a *free convective system* with isolated convective cells which are less than 15 km across. Due to the large proportion of the image with rainfall between 0 and 1mm/h (bin 0), a low ***m*** is observed, however the high rainfall intensities experienced within the convective cells results in a high ***s*** parameter. The expected (calculated from the fitted ***m*** and ***s*** parameters) and observed histograms for this image are plotted in Figure 3.16.

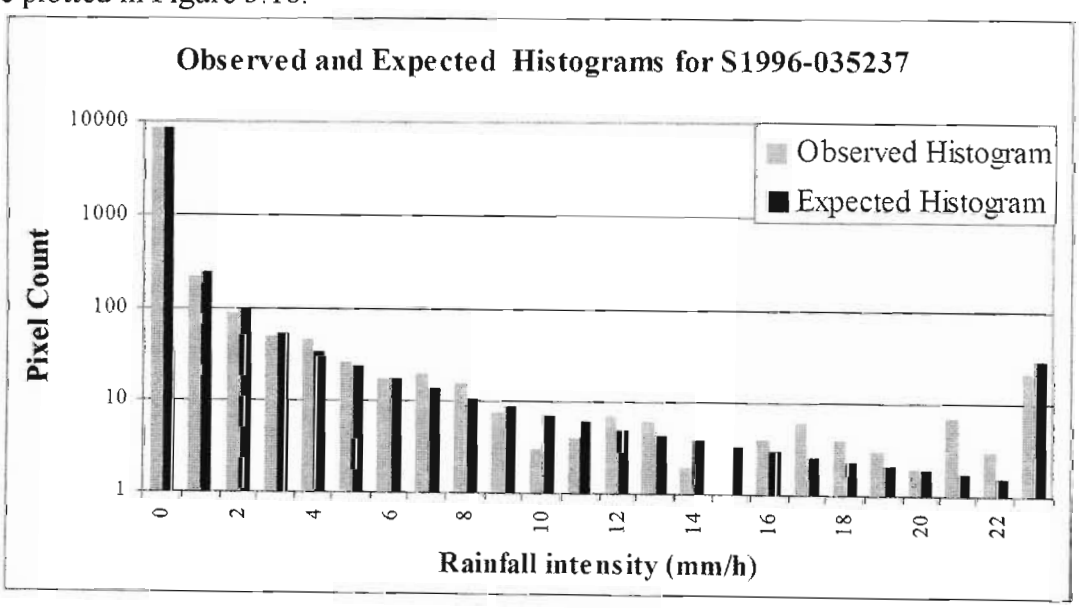

Figure 3.16 - Expected and observed histograms for image S1996-035237 depicted in Figure 3.15a

Figure 3.15b shows a *free convective system* with large isolated convective storm cells which are of the order of 20km in diameter. The result is an intermediate ***m*** parameter combined with a high ***s*** parameter. The relevant histograms are plotted in Figure 3.17.

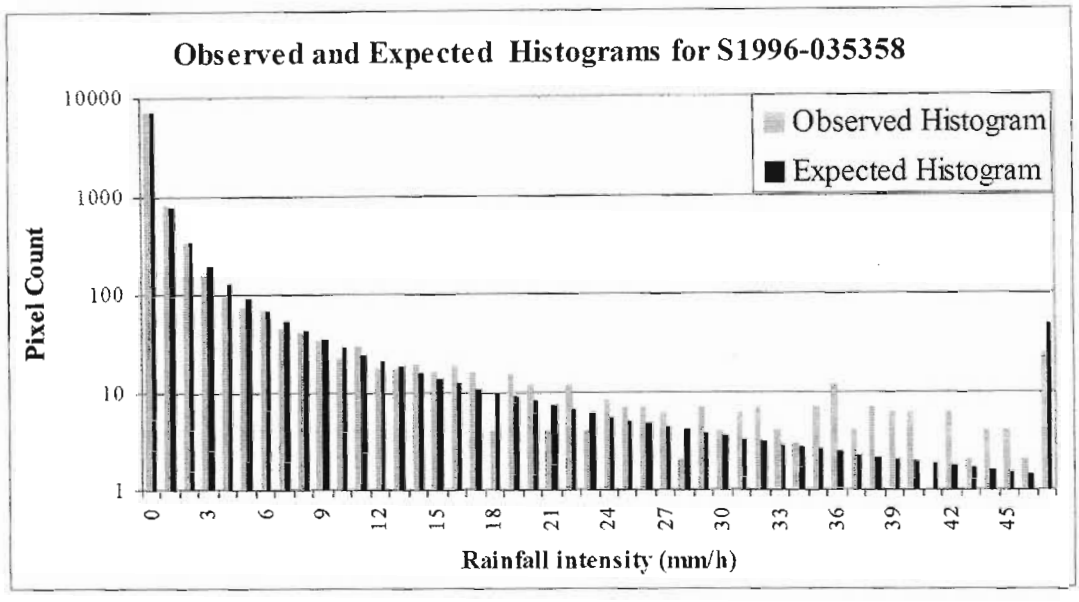

Figure 3.17 - Expected and observed histograms for image S1996-035358 depicted in Figure 3.15b

Figure 3.15c shows a weather system which is between a *free convective* and a series of *forced convective systems* with a characteristic high ***s*** parameter and a reasonably high ***m*** parameter. The most intense rainfall rate on this image is 85mm/h. The histograms for this image are plotted in Figure 3.18.

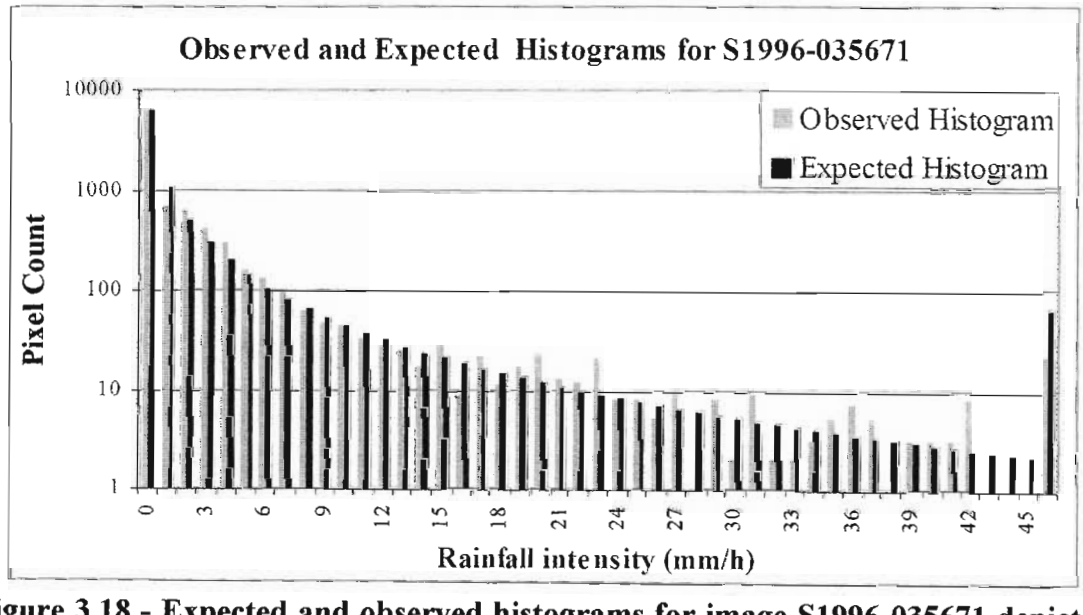

Figure 3.18 - Expected and observed histograms for image S1996-035671 depicted in Figure 3.15c

Figure 3.15d shows a *forced convective weather system*. The main convective cell is approximately 50 km across and the most intense rainfall experienced on the image is 75 mm/h. The expected and observed histograms for this image are plotted in Figure 3.19.

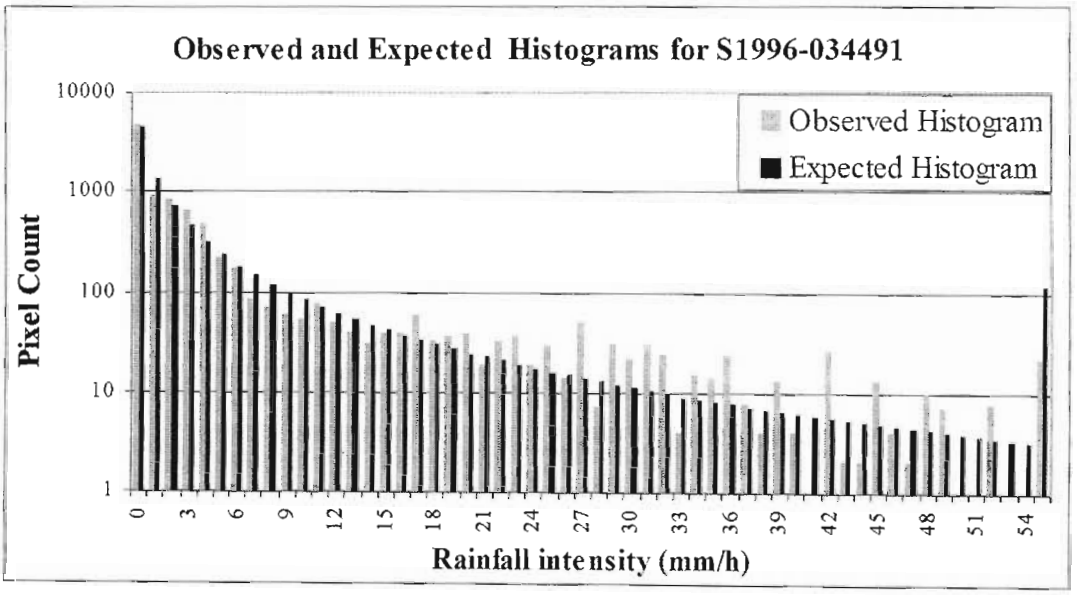

Figure 3.19 - Expected and observed histograms for image S1996-034491 depicted in Figure 3.15d

Figure 3.15e shows a typical *stratiform weather system* with widespread rainfall of low intensity. The maximum intensity experienced on this image is 30mm/h, however less than 0.25% of the sample area experiences intensities above 15mm/h. The expected and observed histograms for this image are plotted in Figure 3.20.

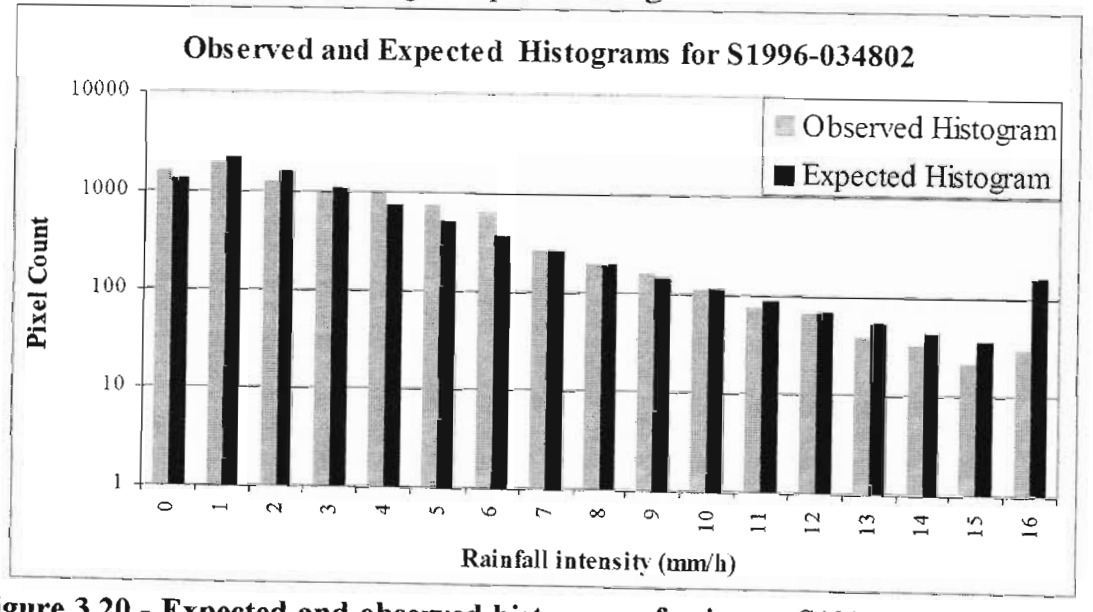

Figure 3.20 - Expected and observed histograms for image S1996-034802 depicted in Figure 3.15e

In all of the cases considered in Figure 3.16 through Figure 3.20, the lognormal distribution appears to approximate the observed distribution well. Although the observed and expected histograms sometimes differ markedly for a particular intensity of rainfall in the upper tail (higher rainfall intensities) these differences are of the order of 10 pixels out of a total of 9128.

3.4.1 Relationship between WAR_i , IMF_i , μ and σ

If the two parameter lognormal is assumed to be a suitable approximation to the marginal distribution of rainfall intensities on an image, then the Instantaneous Wetted Area Ratio (WAR_i) and Instantaneous Image Mean Flux (IMF_i) of the rainfall field uniquely define the μ and σ image scale statistics. IMF_i was defined in Section 3.2 as the average rainfall rate over the entire image and with reference to Equation 3.1, it is related to the μ and σ parameters as defined in Equation 3.6.

$$IMF_i = \exp[\mu + \sigma^2 / 2] \qquad (3.6)$$

In addition to this, WAR_i was defined in Section 3.2 as the percentage of the study area experiencing a rainfall rate in excess of 1mm/h. Assuming the lognormal distribution, WAR_i is given by:

$$WAR_i = 100 \times \int_1^{\infty} f_X(x)\ dx \qquad (3.7)$$

$$\text{where } f_X(x) = \frac{1}{\sqrt{2\pi}\sigma x} \exp[-\{(\ln x - \mu)/\sigma\}^2 / 2] \qquad (3.8)$$

Using the relationships in Equations 3.6, 3.7 and 3.8, it is possible to express μ and σ in terms of WAR_i and IMF_i which are both quick and easy to compute for an image.

Approximately 30 000 *wet* images were analysed to obtain the WAR_i and IMF_i and the corresponding μ and σ parameters using numerical solutions to Equations 3.6, 3.7 and 3.8. This more extensive analysis revealed that the conventional two parameter lognormal is only able to properly explain the statistics of the most intense (highest IMF_i) 30% of the analysed images. Figure 3.21 illustrates the range of WAR_i and IMF_i which can be properly explained by a two parameter lognormal.

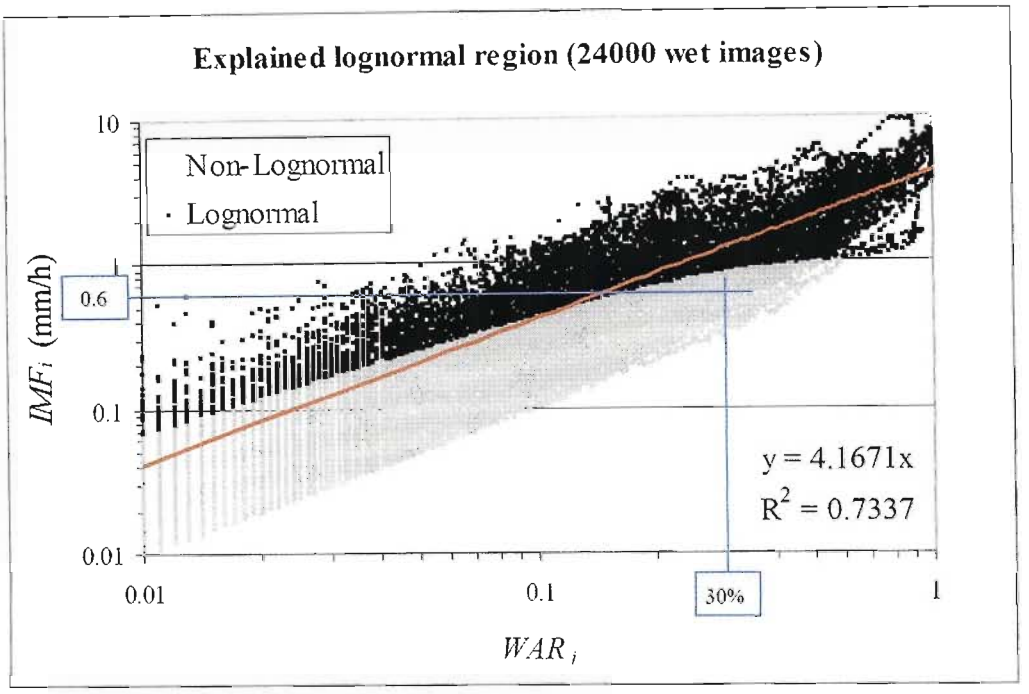

Figure 3.21 – WAR_i vs IMF_i for 24000 images illustrating lognormal and non-lognormal images: for an explanation of the lines, see discussion on Figure 3.22 which follows

The WAR_i and IMF_i are plotted on logarithmic scales in Figure 3.21 in order to give a better idea of the distribution of these parameters over the data set. [As an aside, a linear regression (in normal space) of these points, with a forced zero intercept, provides a rule of thumb: $IMF_i = 4{,}17.WAR_i$, $R^2 = 0{,}73$. So the IMF_i has a numerical value which is approximately four times the numerical value of the WAR_i]. The vertical lines of points are as a result of the discrete number of pixels on an image. Also evident from this figure is that there is a definite lower bound to IMF_i for a given WAR_i. In fact, for a given WAR_i the lowest possible value of IMF_i is obtained when every *wet* pixel is 1,0 mm/h and every *dry* pixel is 0,0 mm/h. In which case IMF_i is given by Equation 3.9.

$$IMF_i = \frac{N_w.1{,}0}{N_w + N_d} \qquad (3.9)$$

Where

N_w is the number of *wet* pixels (rainfall rate in excess of 1mm/h) on the image

N_d is the number of *dry* pixels (rainfall rate less than 1mm/h) on the image

By definition in Section 3.2, WAR_i is given by:

$$WAR_i = \frac{N_w}{N_w + N_d} \qquad (3.10)$$

so that Equations 3.9 and 3.10 show that the lowest value of IMF_i is given by the equality:

$$IMF_i = WAR_i \text{ mm/h}$$

Therefore the lower bound of the IMF_i is numerically equal to WAR_i as observed in Figure 3.21. The range of WAR_i for *wet* images is between 0,01 and 1,00.

Figure 3.21 also reveals that there is a clear distinction between the $WAR_i \sim IMF_i$ combinations which are consistent with a lognormally distributed variable and those that are not. In order to explain this distinction, consider the lognormally distributed variable R_i, distributed $\Lambda(\mu, \sigma)$. Defining L_i as $\log_e(R_i)$, L_i is normally distributed $N(\mu, \sigma)$ and can be transformed to a standardised normal $N(0,1)$ variable Z_i by shifting and scaling by μ and σ respectively.

The *wet* pixel threshold of R_i = 1.0 mm/h can be transformed to the corresponding $Z_i = Z_{WAR}$ in which case $P(Z_i > Z_{WAR}) = WAR_i$. Reversing this transformation, leads to the relationship given by Equation 3.11.

$$\exp(Z_{WAR}.\sigma + \mu) = 1{,}0 \qquad (3.11)$$

or:

$$-Z_{WAR}.\sigma = \mu \qquad (3.12)$$

Substituting for μ in Equation 3.6 yields:

$$Z_{WAR} \pm \sqrt{Z_{WAR}^{\ 2} + 2.\ln(IMF_i)} = \sigma \qquad (3.13)$$

Since σ must be real and positive

$$Z_{WAR}^{\ 2} + 2.\ln(IMF_i) \geq 0 \qquad (3.14)$$

resulting in the curve which separates the lognormal $WAR_i \sim IMF_i$ combinations from the non-lognormal ones in Figure 3.21 (note that WAR_i not Z_{WAR} which is plotted in Figure 3.21). This curve was evaluated numerically - it cannot be described explicitly since it requires an evaluation of the percentage points of the normal distribution. Furthermore, with reference to Equation 3.12, a real and positive σ also implies that the sign of Z_{WAR} is opposite to that of μ. To take a specific example, considering that $P(Z_i > Z_{WAR}) = WAR_i$, Z_{WAR} must be less than or equal to zero if WAR_i is greater than or equal to 50%, in which case μ must be greater than or equal to zero. With both μ and σ greater than zero, Equation 3.6 shows that the minimum IMF_i would be 1,0 mm/h. So for a WAR_i greater than 50%, IMF_i is greater than 1,0 mm/h which is consistent with the observations of Figure 3.21.

Since the simple two parameter lognormal is unable to properly describe the marginal distribution for a large percentage of the observed full precision images, a slightly different approach was adopted. For a WAR_i, IMF_i pair measured from a two dimensional field of data it is possible to generate a pseudo-random field with identical WAR_i and IMF_i by assuming an underlying censored lognormal distribution. This is best illustrated by example.

Consider the case in which a radar rainfall image is analysed and found to have WAR_i = 30% and IMF_i = 0.6 mm/h. Such an image is given in Figure 3.22, using the most recent (2003) clutter mask which leaves 10 939 data for analysis as opposed to the 9128 of the ¾ donut mask presented in Figure 2.8.

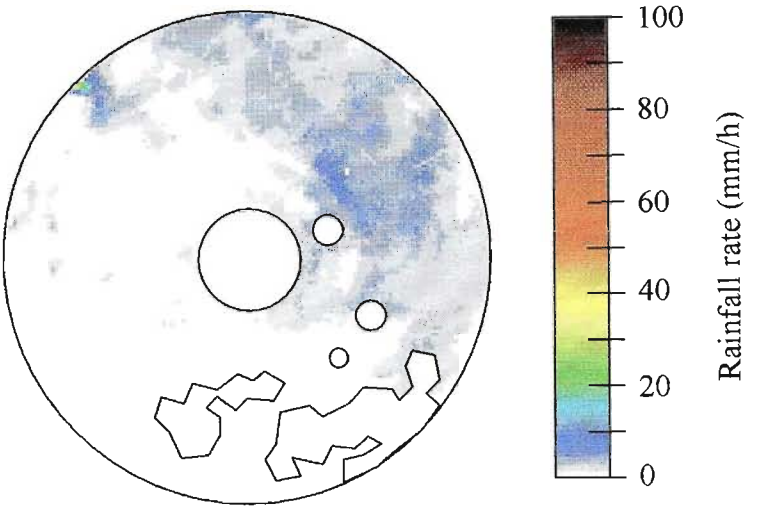

Figure 3.22 - Example of an image whose distribution cannot be described by the two parameter lognormal, WAR_i = 30% and IMF_i = 0.6 mm/h

With reference to Figure 3.21 the marginal distribution of the image shown in Figure 3.22 is not defined by the two parameter lognormal as seen by the intersection of the two lines in Figure 3.21. However, if it is assumed that the underlying distribution is described by a truncated lognormal, a field with similar properties (WAR_i, IMF_i) can be generated using the following algorithm:

1. Generate a standard-normally distributed, pseudo random field, $Z_j \sim N(0,1)$, where (j = 1, 2, …, n)
2. Calculate the percentile Z_{WAR} corresponding to the WAR_i such that $P\{Z_j > Z_{WAR}\} = WAR_i$.

3. Subtract Z_{WAR} from the field; Z_j - Z_{WAR} will then be normally distributed, $N(-Z_{WAR},1)$.
4. Compute $IMF'' = \sum_{j=1}^{n} \exp(Z_j - Z_{WAR})$, for all $Z_j > Z_{WAR}$.
5. Compute $\varepsilon = | IMF'' - IMF_i |$
6. Scale the Z_j - Z_{WAR} field by adjusting σ, and repeat steps 4 to 6 until ε is sufficiently small

This process is similar to fitting a *censored* lognormal distribution, ignoring values below 1mm/h. In the case where WAR_i = 100%, the distribution will be equivalent to a two parameter lognormal. Based on the content of the literature reviewed thus far, it appears to be a novel approach to fitting the censored lognormal to rainfall data. An example of this type of fit is given in Figure 3.23 in which the histograms of the image of Figure 3.22 and that of a simulated image with identical WAR_i and IMF_i are compared. In addition, Figure 3.23 presents the histogram of an *uncensored* two parameter lognormal field, simulated using the maximum likelihood estimators of μ and σ, as given by Equations 3.3 and 3.4, which were measured from the full precision image of Figure 3.22. This illustrates the limitation of the simple two parameter lognormal distribution as a descriptor of the marginal distribution of rainfall rate on a radar image.

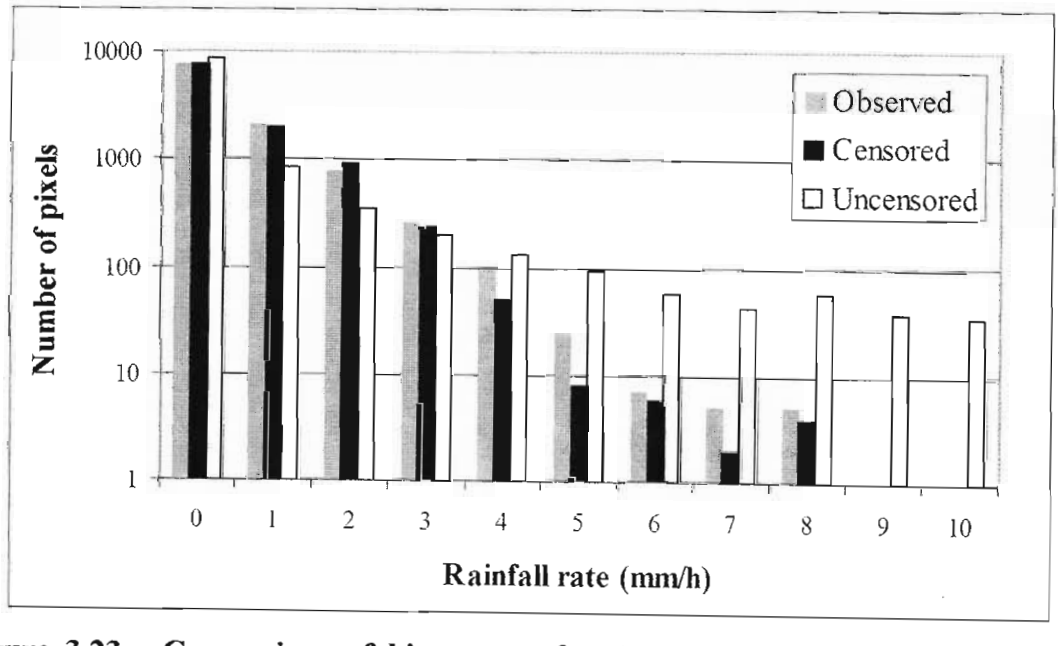

Figure 3.23 - Comparison of histograms for the image of Figure 3.22 and two equivalent simulated fields with censored lognormal and uncensored lognormal marginal distributions respectively.

Again the histograms are plotted on a logarithmic scale in order to better examine the quality of the fit. Clearly the *uncensored* lognormal is unsuitable for describing the marginal distribution on this image. The *censored* lognormal simulation of the field appears to be reasonable. Comparing the histograms using the chi-square statistic, defined in Equation 3.15 for two binned datasets R_i and S_i, reveals that even the censored lognormal is an unlikely descriptor of the distribution.

$$\chi^2 = \sum_i \frac{(R_i - S_i)^2}{R_i + S_i} \qquad (3.15)$$

The chi-square test appears to be very sensitive to differences in the tails of the distributions for large N. With reference to Figure 3.23, the differences in the tails of the observed and the censored distributions are of the order of tens of pixels out of a total of 10 900. Perhaps a more instructive means of illustrating the quality of fit of the censored lognormal is to consider the probability of exceedence of a particular rainfall rate for a range of radar images. Figure 3.24 does this for five images with WAR_i ranging between 5% and 80%, similar to those presented in Figure 3.15, only using full precision data.

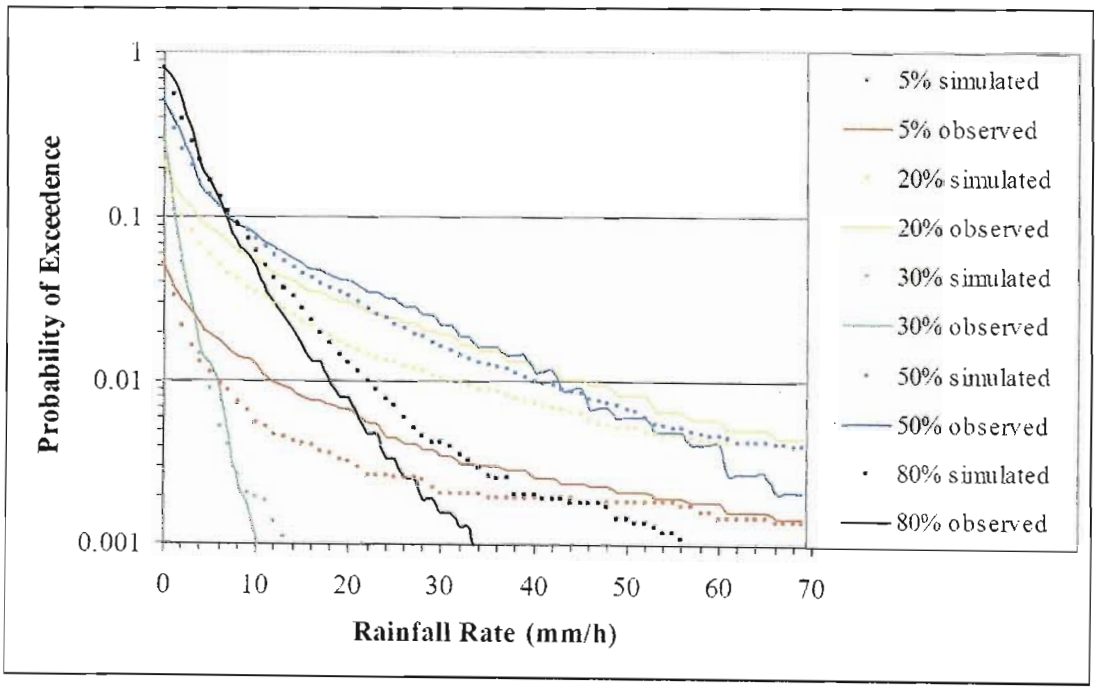

Figure 3.24 - Probability of exceedence of rainfall rate on 5 observed radar images

Figure 3.24 illustrates the type of fit which can be expected using the censored lognormal fitting technique discussed previously. In some cases, such as the WAR_i = 30% and 50%, the marginal distribution is closely reproduced by the censored lognormal simulation even beyond the 99 percentile. The WAR_i = 80% case is also

reasonably well simulated with the 99 percentile of the observed image at 18mm/h compared to the simulated of 22mm/h. The worst fits are obtained for images with low WAR_i, such as the 5% case. For low WAR_i, the simulated distribution underestimates the number of pixels in the middle order (10 to 30mm/h) and over estimates the number of higher order pixels (above 30mm/h), converging toward the observed distribution at the upper end of the distribution.

Further development of the "String of Beads" model may warrant revisiting this subject, but for the purposes of this study the censored lognormal will be assumed to be an adequate approximation of the marginal distribution of rainfall flux on a radar image as it can fit all of the images whose WAR_i and IMF_i are plotted on Figure 3.21.

3.5 *Spatial Structure*

With a reasonable means of describing the marginal distribution of rainfall flux on an image some additional information is required in order to quantify the spatial structure of the rainfall flux field. The assumption of scale invariance has been shown to be true over a wide range of scales (Gupta and Waymire, 1990) and this result is used and tested in this study. Scale invariance assumes that the stochastic spatial structure of the rainfall field is independent of scale. A property of this type of field is that it exhibits a power spectrum (which describes the spatial correlation structure of the field) that can be described by a power law relationship using a single parameter, β_{space}, which is the *negative exponent of the radially averaged, two dimensional power spectrum*. In order to estimate this parameter it is first necessary to transform the rain field into Fourier space also known as the frequency domain. The definition of the Fourier Transform is well documented and can be found in any good text on digital signal processing (Kuc (1988), for example). For ease of reference, its definitions in continuous and discrete form are included in Appendix B of this document.

3.5.1 *The Fast Fourier Transform*

Since rainfall data measured by radar are recorded in discrete space-time intervals, the discrete form of the Fourier transform is used, and specifically, the Fast Fourier Transform. The *Decimation in time Fast Fourier Transform* (Cooley and Tukey, 1965) is a method of computing the Discrete Fourier Transform using fewer calculations.

The data sets used in the development of the *String of Beads model* are extremely large and the computational efficiency of the algorithm used to compute the Fourier Transform plays a crucial part in both the analysis and the simulation process. It can be shown that for a sample of N points, the computational effort required for the Fourier Transform is proportional to N^2 calculations. The computational effort for the Fast Fourier Transform algorithm is proportional to $N.log_2 N$ calculations. To emphasise the importance of this algorithm, for large fields such as those encountered in this study, this can mean the difference between minutes of computation and months of computation on a modern personal computer.

The derivation of the *Decimation in Time Fast Fourier Transform algorithm* is covered in many texts (Kuc (1988), for example) and will not be repeated here. It takes advantage of certain symmetries in the transform by splitting the sequence h_k into two half sequences, even indexed $h_{k(even)}$ and odd indexed $h_{k(odd)}$, which are independently transformed and added together. It is for this reason that the sequence to be transformed must be an integer power of 2 in length (although the use of a radix-4 FFT algorithm or a combination radix-8,4,2 could possibly be used in order to minimise the constraint on the length of the dataset if required). A sequence of arbitrary length should be packed with a buffer of zeros to achieve this condition. This process is known as "zero padding". Efficient programming of these algorithms is not a trivial exercise. The routine given by Press et al. (1992) was used for this study.

3.5.2 Estimating the power spectrum and autocorrelation using the FFT

For a one dimensional sample of N complex valued points (where N is an integer power of two) in the time domain h_k (where $k = 0, 1, \ldots, N-1$), the one dimensional power spectrum is proportional to $| H_n |^2$, where the H_n ($n = -N/2, \ldots, N/2$) represent the corresponding sequence in the frequency domain (defined by Equation 9.12 in Appendix B). That is to say that in order to estimate the power spectrum of a sequence of N complex valued points h_k, the sequence is transformed into Fourier space via the FFT to get the transformed complex valued sequence H_n and the power spectral density $P(f_n)$ corresponding to each frequency (f_n) is obtained by squaring the absolute value of the H_n term. Consequently the power spectral density is real valued and positive.

The Discrete Fourier transform of a *real valued* sequence of numbers will exhibit conjugate symmetry about the Nyquist frequency in Fourier space and thus the power spectrum of a one dimensional sequence of real numbers will be symmetrical about the Nyquist frequency. It is therefore only necessary to consider one half of the power spectrum in this case.

The power spectrum is related to the *autocorrelation* function of the sequence. The correlation of two functions $g(t)$ and $h(t)$, written $Corr(g, h)$, is defined in Equation 3.16.

$$Corr(g,h) = \int_{-\infty}^{\infty} g(t+\tau).h(t).dt \qquad (3.16)$$

It is a function of τ which is known as the lag and it is defined in the time domain. An efficient means of computing the correlation is to make use of the "Correlation Theorem" which is defined for *real valued* functions g and h in Equation 3.17.

$$Corr(g,h) \Leftrightarrow G(f).H^*(f) \qquad (3.17)$$

where $G(f)$ and $H(f)$ are the Fourier transforms of $g(t)$ and $h(t)$ respectively and $H^*(f)$ represents the complex conjugate of $H(f)$. In other words, for *real valued* functions g and h, the product of the Fourier transform of g with the complex conjugate of the Fourier transform of h is equal to the Fourier transform of their correlation.

The *autocorrelation* of a function is simply the correlation of that function with itself and is a measure of the similarity between $g(t)$, the function evaluated at a point t and $g(t+\tau)$, the function evaluated at some lagged point $(t+\tau)$. Furthermore, it can be shown that:

$$Corr(g,g) \Leftrightarrow |G(f)|^2 \qquad (3.18)$$

which is know as the "Wiener-Khinchin Theorem". That is, the autocorrelation function is equal to the Fourier Transform of the power spectrum. An important point when estimating the correlation (or autocorrelation) using Fourier methods, is that it is implicitly assumed that the original data are cyclic in nature. In most circumstances, this is not the case and it is therefore necessary to pad the original sequence(s) with zeroes in order to avoid a biased estimate of the correlation function. The amount of zero padding required is equal to the maximum lag of interest, τ_{max}. For the case in which the correlation function is estimated for all lags in the sample, the sample must be padded with the same number of data as in the original sequence. In addition, to obtain the true

correlation function, not just the sum of products of the data, the data should be pre-treated by shifting them to have zero mean. Similarly, when estimating the power spectrum for time series data using Fourier methods, it is necessary to pad the data sample with zeros to avoid measuring erroneous power by assuming the data are cyclic when this is not actually the case. These properties of correlation functions and power spectra extend to the two dimensional case and beyond.

For the two dimensional case, the two dimensional power spectrum is itself a two dimensional image and an example of the power spectrum of a radar image is given in Figure 3.25. The oblique view of Figure 3.25 serves to illustrate the shape of the two dimensional power spectrum of a 128 x 128 pixel radar rainfall field, padded to 256 x 256 pixels – note the logarithmic vertical scale. The plan view of the spectrum illustrates its skew symmetry which is a consequence of the fact that the Fourier Transform of a real valued field exhibits conjugate symmetry about the Nyquist frequency (128 cycles per 256 km in this case). The term frequency is used loosely in this document and refers to the spatial frequency in wave cycles per kilometre (or wavenumber).

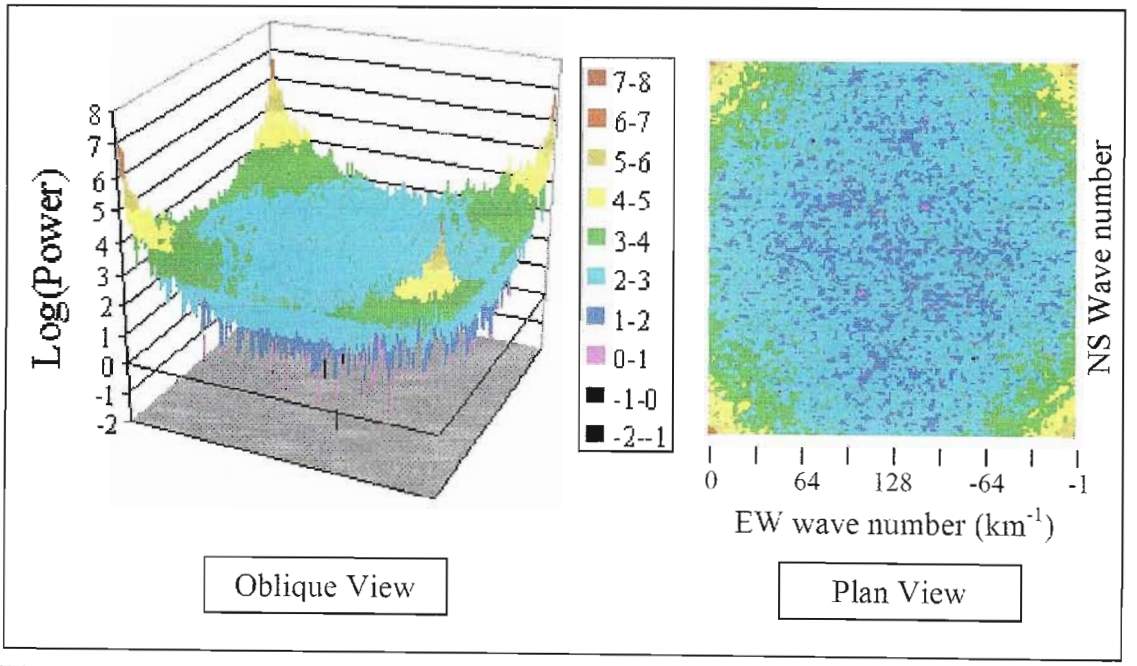

Figure 3.25 - An example of the power spectrum of a two dimensional field

By making the assumption that the rainfall field is isotropic, the two dimensional power spectrum can be conveniently reduced to a one dimensional representation by radially averaging it about the (0,0) wave number, out to the Nyquist frequency. An example of

a radially averaged, two dimensional (RA2D) power spectrum corresponding to that of Figure 3.25 is presented in Figure 3.26, with its fitted power law function.

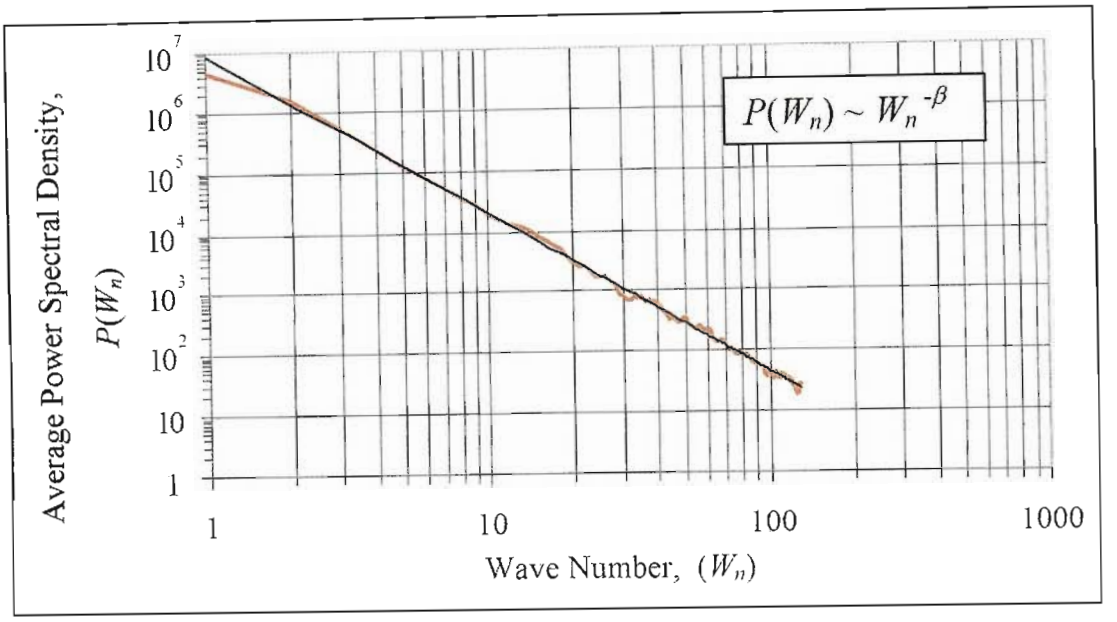

Figure 3.26 - The Radially Averaged, Two Dimensional Power Spectrum. A one dimensional representation of the two dimensional power spectrum in Figure 3.25.

Of great significance in this study is the fact that this RA2D power spectrum is reasonably linear when displayed in log-log space as is the case with most rainfields, indicating a power law relationship and thus the scale-invariant nature of the rain field observed by Gupta and Waymire (1990) amongst others and mentioned at the beginning of this section. Assuming isotropy and scale invariance in the rain field, the negative exponent, β_{space}, of the RA2D power spectrum is sufficient to describe the spatial correlation structure of the rainfall flux in the field.

Due to the severely skewed and highly variable marginal distribution of rainfall flux on a radar image, if the power spectra of two images are to be compared sensibly, the fields must first be transformed into Gaussian space and standardised before extracting their power spectra. If this is not done, the β_{space} of rain fields with relatively high IMF_i will tend to be underestimated (i.e. lower spatial correlation) when compared to that with a relatively low IMF_i. For full precision data it is possible to take logarithms of the pixels in the field to achieve a near Gaussian field from which the power spectrum can be computed. For character precision data, since it is not possible to take the logarithm of a dry (0 mm/h) pixel, it is necessary to assume a low rainfall rate (0.001 mm/h) in the

field, and then to take logarithms and extract the power spectrum. Unfortunately, this introduces a bias and frequently leads to poor estimates of the power spectrum, particularly for images with low WAR_i.

Another point to be considered when extracting the power spectrum from a field of data is the effect of the mountain mask used in the study. This was tested by simulating 200 images with various image scale statistics, and then measuring their power spectra both with and without the application of a mask. The result of this test is shown in Figure 3.27. It clearly illustrates that masking the field of data prior to measuring the power spectrum introduces a sampling error as would be expected, but most importantly, this error is unbiased. The differences in power spectral exponent between the masked and unmasked cases can be sensibly approximated by a Gaussian distribution. Based on these results it is reasonable to assume that measurement of a *masked* radar image produces an *unbiased* estimate of the power spectrum of the field.

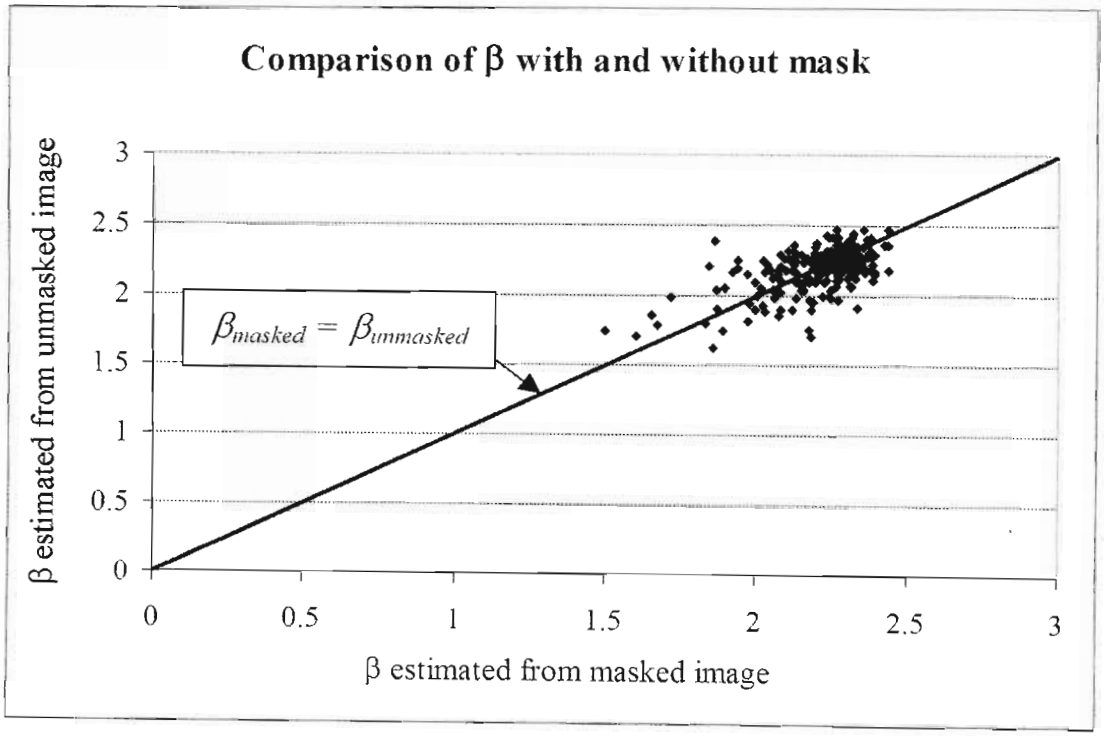

Figure 3.27 - The effect of masking the rainfield on the estimation of the power spectrum. Computed for 200 simulated images and using the $^3/_4$ doughnut mask shown in Figure 2.8

These Fourier techniques are an extremely efficient means of estimating correlation functions and power spectra, an important point when analysing large datasets such as radar rainfall data. For this reason, they were used almost exclusively in the estimation

of the parameters of the *String of Beads* model. However, an additional technique of describing the spatial structure of rain fields which was used as a means of validation in the *String of Beads* model is the Generalised Structure Function.

3.5.3 *Estimation of the Generalised Structure Function*

Unlike the two dimensional power spectrum and the lognormal distributions, the Generalised Structure Function is not used to determine the input parameters of the "String of Beads" model, and therefore serves as a useful validation test of the spatial structure of simulated images. It was developed for use in the field of fluid turbulence and has become a widely used technique for multifractal fields since the work of Anselmet et al. (1984). Its use is appropriate in the analysis of a non-stationary field of dimension D where $\beta_{space} > D$. When $D = 2$ dimensions (as is the case in rain fields), $\beta_{space} > 2$ suggests that ordinary correlation analysis is inappropriate. A concise description of this technique with specific reference to rainfall fields is given by Menabde et al. (1999).

Paraphrasing their work, the Generalised Structure Function $G_q(l)$ is defined on a non-negative scalar random field R(x) as being the ensemble average of the q^{th} power of the absolute differences of the R values a distance l apart. That is

$$G_q(l) = \left\langle \left| R(\vec{x} + \vec{l}) - R(\vec{x}) \right|^q \right\rangle \qquad (3.19)$$

In addition, the non-stationary random field is called *multiaffine* (Benzi et al., 1993) if $G_q(l)$ satisfies the condition:

$$G_q(l) \propto l^{\zeta(q)} \qquad (3.20)$$

where $l = |\vec{l}|$ and $\zeta(q)$ is some non-linear function of q. For the case in which $q = 2$, $\zeta(2)$ is related to β_{space}, the gradient of power spectrum, by Equation 3.21.

$$\beta_{space} = D + \zeta(2) \qquad (3.21)$$

where D is the dimension of the data ($D = 2$ for 2D radar data).

For the purposes of this study, the generalised structure function of Equation 3.19 was approximated for radar images in the two principle directions (north-south and east-west) and for 6 discrete values of l chosen to be 1, 2, 4, 8, 16, 32 and 64. Since both points $R(x)$ and $R(x + l)$ of Equation 3.19 are required to be within the unmasked region of the image, significant areas of the image do not contribute to the calculation of the

structure function for the larger values of l. This is illustrated for the case of the ¾ doughnut mask in Figure 3.28. The area shown in yellow represents the area of the image which contributes to the calculation of the structure function in both the east-west and the north-south directions. Areas shown in red and green represent regions of the image which contribute to the structure function in only the east-west and the north-south directions respectively. Regions shown in black do not contribute to the calculation of the structure function. The image on the left shows the case in which $l = 64$ and the image on the right when $l = 16$.

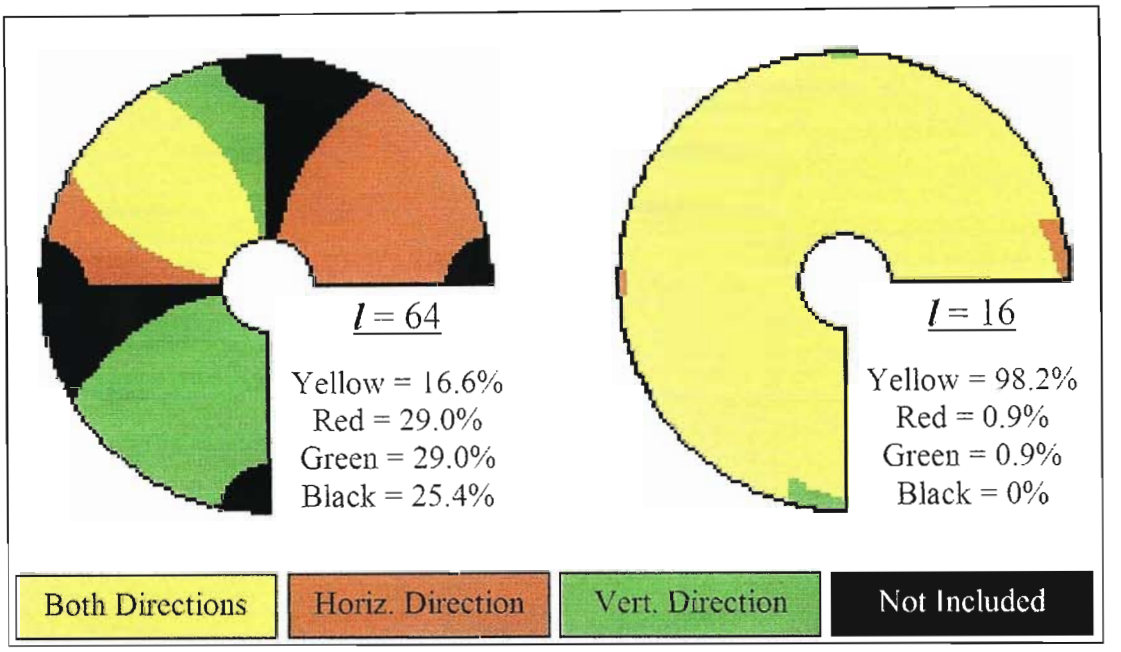

Figure 3.28 - Areas of the image contributing to GSF calculation for different values of l

Less than 75% of the image is used to calculate the structure function in the case where $l = 64$. In comparison, the entire image is used for the case when $l = 16$ and more than 98% of the image is used in both the north-south and east-west directions. Consequently the behaviour of the structure function at high values of l is dependent on the distribution of rainfall on the image and is somewhat erratic.

The Generalised Structure Function is defined for all $q \geq 0$ and an example has been plotted with logarithmic axes in Figure 3.29 for discrete values of l/l_0 up to 64 (where l_0 is the length scale of the field, 1km in this case), and discrete values of q up to 3.0.

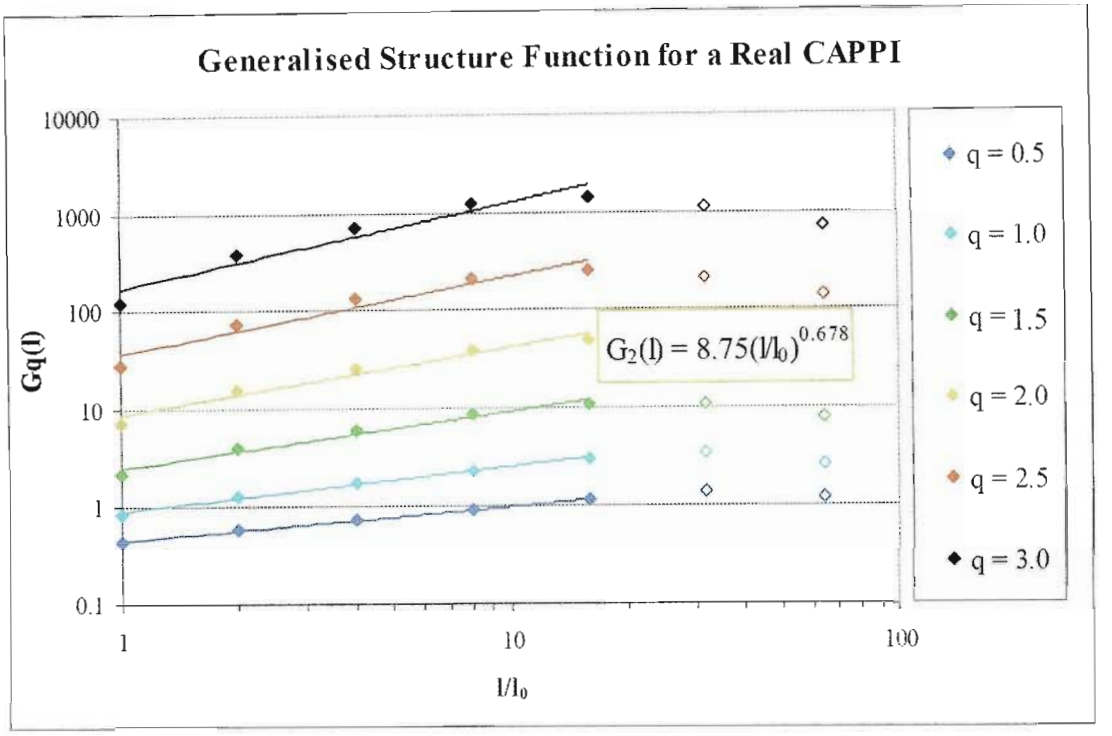

Figure 3.29 - Typical example of the Generalised Structure Function for a real radar rainfall image

For values of q less than 3.0 and l/l_0 less than 16 the function, $log\ G_q(l)$ versus $log\ (l/l_0)$ is approximately linear for a given q. In addition, the gradients $\zeta(q)$ of the linear approximations appear to be a function of increasing q. These observations in real rainfields suggest the *multiaffine* behaviour defined in Equation 3.20 and can be expressed in the form of Equation 3.22.

$$log\ G_q(l) \propto \zeta(q).log\ l \qquad (3.22)$$

The gradients $\zeta(q)$ of Equation 3.22 are obtained by fitting power law relationships by least squares in the linear space of $G_q(l) \sim l/l_0$. These can be plotted for discrete values of q and an example of such a plot is given in Figure 3.30 for the six values of q shown in Figure 3.29.

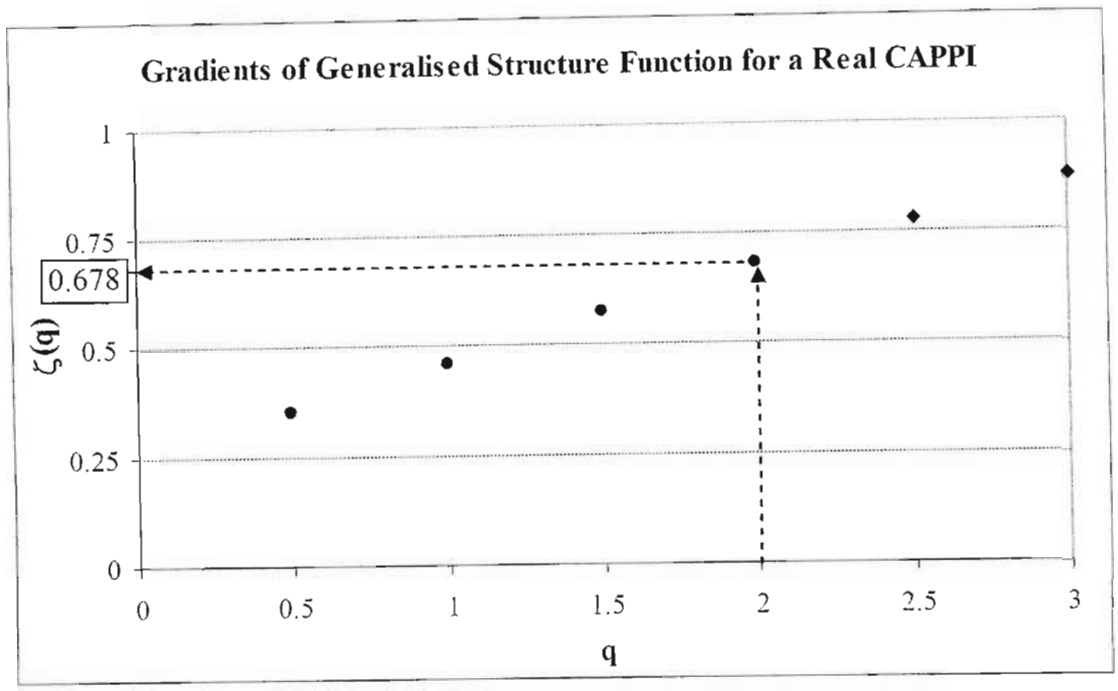

Figure 3.30 - Gradients of the GSF presented for a real image in Figure 3.29

According to the relationship defined in Equation 3.21 for a two dimensional field, the $\zeta(2) = 0.678$ shown in Figure 3.30 should correspond to a β_{space} of 2.678. The measured β_{space} for this particular image was found to be 2.692.

Since the *Generalised Structure Function* (GSF) only makes use of the data sampled from within the masked sample area for its calculation, the mask has no effect when comparing the GSF of two images. The precision of the data however, does effect the shape of the GSF.

The effect of data precision on the Generalised Structure Function

Much of the radar data analysed in this study were supplied at 8 bit precision and could assume any integer value of rainfall intensity between 0 and 255 mm/h. This presents a problem when calculating the Generalised Structure Function. The Generalised Structure Function $G_q(l)$ is defined in Equation 3.19 on a non-negative scalar random field R(x) and it suggests that for a continuous variable $R(x)$, the absolute difference between any two pixel intensities can assume any real, positive value z. That is

$$\left|R(x + l) - R(x)\right| = z, \qquad z \in \Re,\ z \geq 0. \tag{3.23}$$

In the case of *double precision data*, the z values of Equation 3.23 assume small *non-zero* values for small values of l, and this leads to a small $G_q(l)$. As q tends to zero, $G_q(l)$ tends to unity and the function gradient $\zeta(q)$ tends gradually to zero. That is to say that

$$\text{if } z \neq 0, \lim_{q \to 0} z^q = 1,$$

$$\therefore \lim_{q \to 0} \left\langle z^q \right\rangle = \lim_{q \to 0} G_q(l) = 1 \tag{3.24}$$

$$\therefore \lim_{q \to 0} \varsigma(q) = 0 \tag{3.25}$$

Conversely, as l increases, larger absolute differences will be encountered and the $G_q(l)$ will gradually increase for any given $q \neq 0$. A higher power of q will lead to a greater $G_q(l)$ and therefore a steeper function gradient $\zeta(q)$.

This is *not* the case when $R(x)$ can only assume *integer* values. The absolute difference between two integers must also be an integer value. This introduces zero values of z and the shape of the structure function is different from when it is calculated with full precision. Figure 3.31 and Figure 3.32 show the gradients of the structure function for an image. Figure 3.31 was calculated for a *double precision* image and Figure 3.32 for the same image stored in *integer precision*.

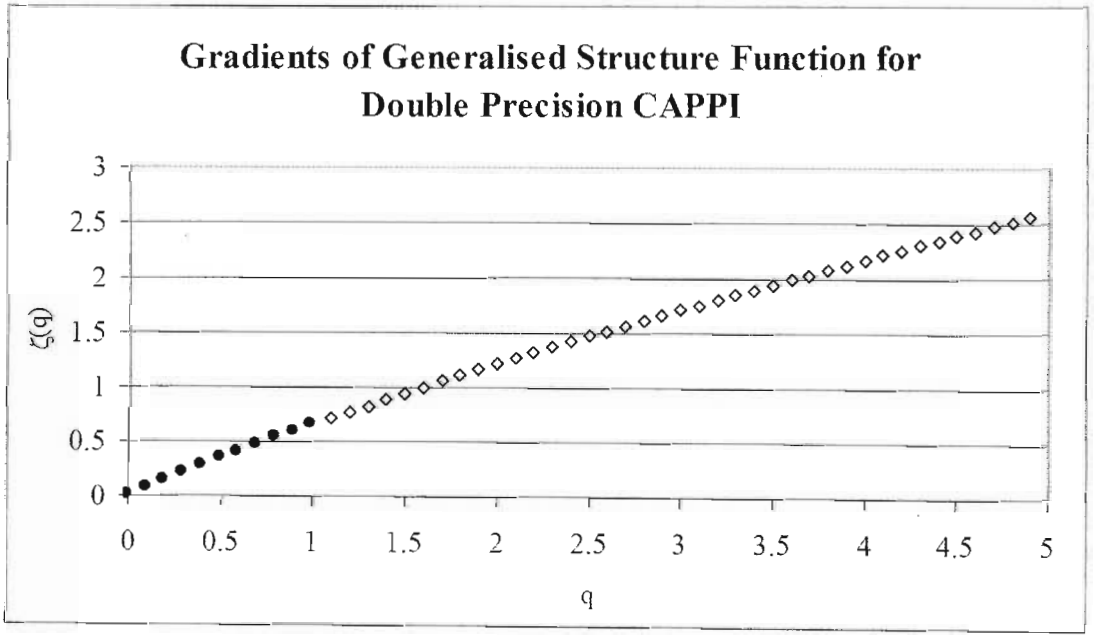

Figure 3.31 - Gradients of the GSF for a double precision image

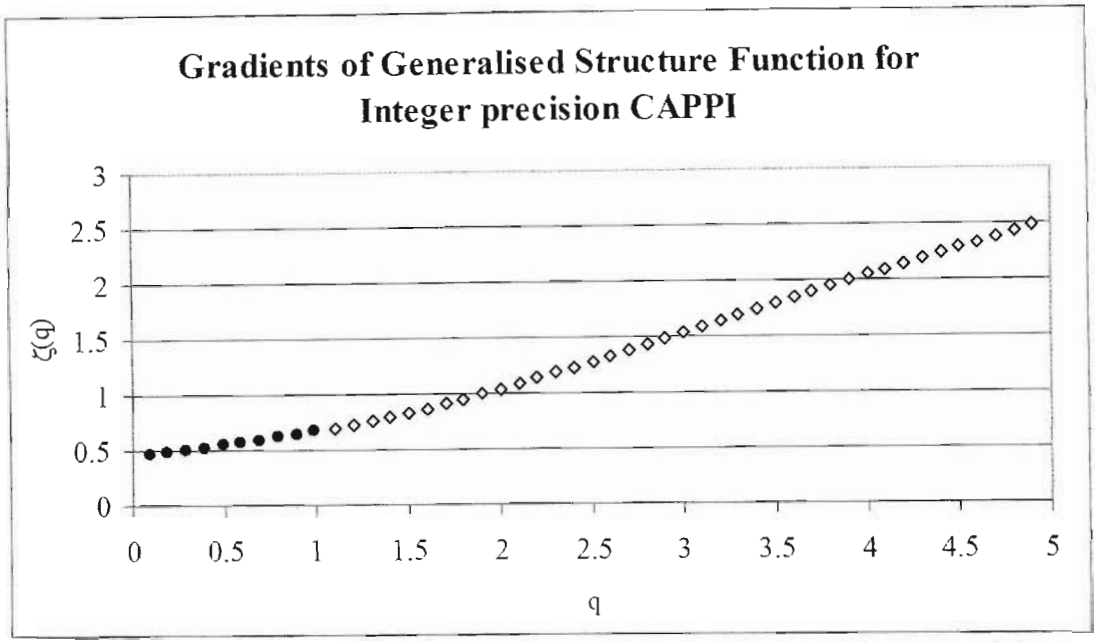

Figure 3.32 - Gradients of the GSF for an integer precision image to be compared with Figure 3.31

The explanation for the difference lies in the fact that the distribution of rainfall on an image is severely skewed towards the lower values and it stands to reason that the distribution of the *absolute differences* of rainfall intensities on the image will tend to be even more skewed towards the lower values. In addition, the sampled correlation of the intensity of a pixel with its closest neighbour is very high (usually above 0.99). The combination of these two conditions leads to the fact that for low values of l it is rare to observe an *absolute difference* of rainfall intensity in excess of 1mm/h. However, as l increases, the correlation between the two pixel intensities is reduced and there is a better chance of encountering a larger absolute difference between them. This is illustrated in Figure 3.33.

The horizontal bars in Figure 3.33 represent l values of 8 and 32 used to calculate the GSF in the simple one-dimensional example plotted in double precision as the dark grey series, and in integer precision as the light grey series. The vertical bars represent the corresponding z values for the double precision case. Even on the steepest part of the curve, the largest value of z that is experienced for an l value of 8 is approximately 2.5. On the same section of the curve a z value of 4.9 can be obtained when using an l value of 32.

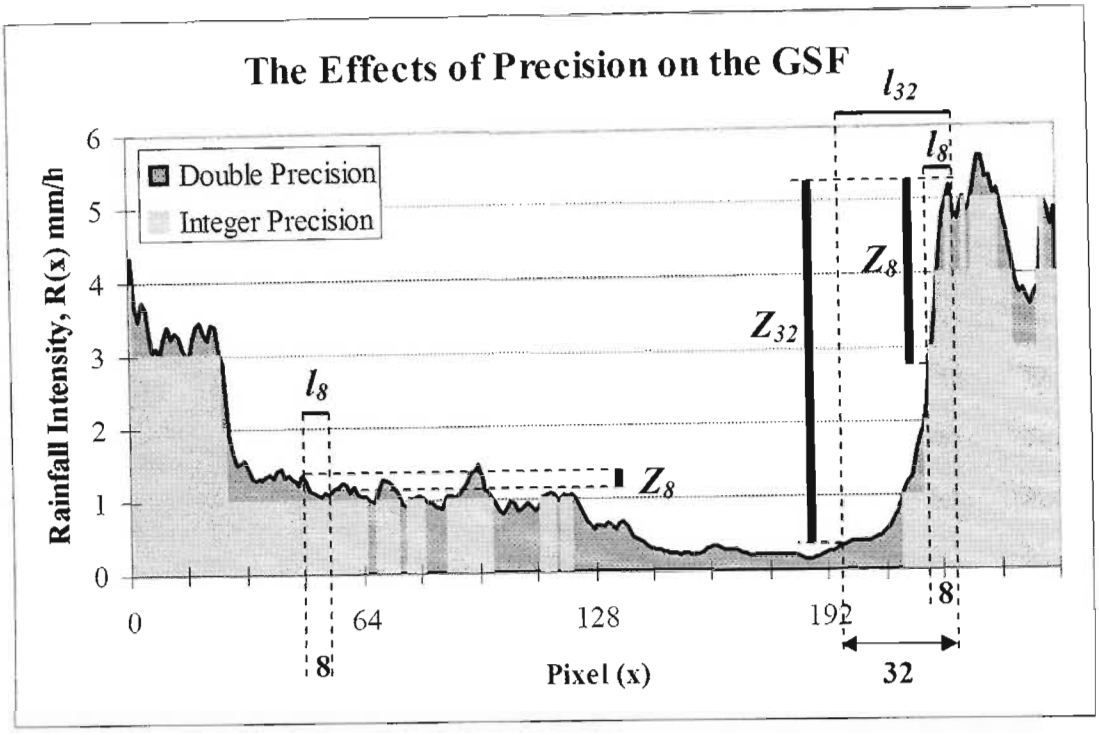

Figure 3.33 - Choice of larger *l* corresponds to a greater chance of observing a large absolute difference (*z*)

When working with integer values of rainfall intensity on the same series of data (plotted as the light grey series in Figure 3.33), a large proportion of the observed absolute differences (*z*) are zero. For a sample of *N* points containing ***m*** zero values

$$\lim_{q \to 0} \left\langle z^q \right\rangle = \lim_{q \to 0} G_q(l) = \frac{(N - m)}{N} \qquad (3.26)$$

since $0^q = 0$ and $1^q = 1$ for all q (defining $0^0 = 0$), and the ratio $(N - m) / N$ is strictly less than one. In the case where $m = 0$, Equation 3.26 reduces to Equation 3.24.

With reference to Equation 3.23, the smaller the value of *l*, the more z = 0 will be observed and consequently, the smaller will be the limit of Equation 3.26. Conversely, a larger value of *l* will produce fewer z = 0 and therefore a higher $G_q(l)$ will be calculated. This shows that when analysing data in which the absolute differences z *can* assume a value of zero, $G_q(l)$ will increase with increasing *l* even when $q = 0$. The only cases in which $\zeta(0)$ can be zero are if the rain field is either entirely zero or entirely non-zero. The result of this is that when calculating the structure function for integer precision data, the gradient $\zeta(0)$ will be finite and positive as shown in Figure 3.32. In the literature reviewed so far this appears to be a novel result.

3.6 *FIELD ADVECTION*

Measurement of the field advection is accomplished via a pattern search which finds the *shift vector*, which corresponds to the position of maximum correlation between two consecutive images. The algorithm is illustrated in Figure 3.34.

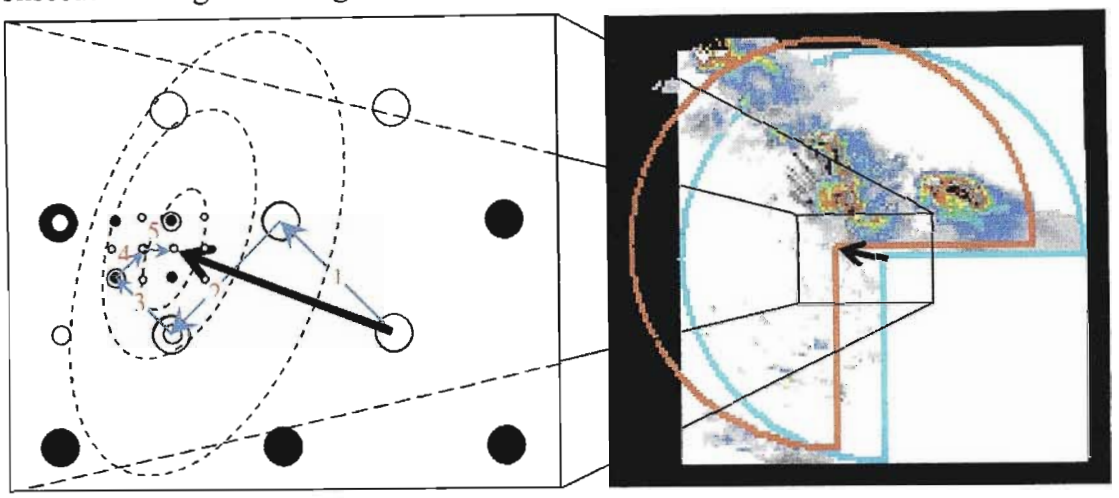

Figure 3.34 - Algorithm used to measure mean field advection

Consider two consecutive radar rainfall images. First, the pixel to pixel cross correlation is computed for the images with no spatial shift. The images are then diagonally shifted relative to each other to each corner of a 9 km square simplex and the correlation at each shift is computed. The correlations computed at each of the five shifts are then compared and the shift corresponding to the highest correlation is chosen as the centre for the new simplex. This is repeated until the highest correlation is at the centre of the simplex, at which stage the shifts corresponding to the two highest correlations in the simplex become the corners of a new simplex, half the size of the original (i.e. a 5 km simplex). The process continues until the size of the simplex is 1 km and then the correlations of the eight surrounding pixels are computed to find the maximum of the cross correlation between the images. It is a form of a 2-dimensional bisection search, on a square simplex, which is extremely fast on a correlation field which is usually unimodal. The best guess for the next consecutive pair of images is taken as the correlation shift of the previous pair. The shift vector (in kilometre pixels) divided by the time interval between the images (in hours) yields the mean velocity (in kilometres per hour) of the field. This will be referred to as the *mean field advection vector*.

More detailed analysis could be performed in order to determine the advection of individual storm cells within the rainfall field through the use of a storm tracking algorithm such as that used in the TITAN software (Dixon and Wiener, 1993), however

this would introduce considerable complexity to the model which seems excessive for the purposes of simulation. The model is intended as a fast and effective stochastic simulator. Measurement of fine scale advection would be beneficial in the application of the model as a forecasting tool. This refinement is left as a topic for future research.

3.7 *IMAGE SCALE ANALYSIS*

Up to this point in the chapter, most of the discussion has been concerned with the extraction of the image scale statistics, or primary statistics, from the data. As mentioned at the end of Section 3.2, the secondary statistics continue to describe the primary statistics and serve to bind the model together. This section and subsequent sections of this chapter are concerned with secondary statistics of the data. It is convenient to separate the analysis of the *instantaneous images* and *cumulative images* (which were defined in Section 3.2)

3.7.1 *Instantaneous image scale analysis*

It is now possible to present the results of analysis of the entire data set – over 98500 instantaneous radar images recorded during the 1995/1996, 1998/1999 and 1999/2000 rain seasons. Data from the 1995/1996 rain season were used in the initial part of the study and were stored in 8 bit precision (integers between 0 and 255mm/h). Data from 1998 onwards were extracted using the DISPLACE averaging technique discussed in Section 2.4.2 and stored in MDV format as reflectivity fields at higher precision.

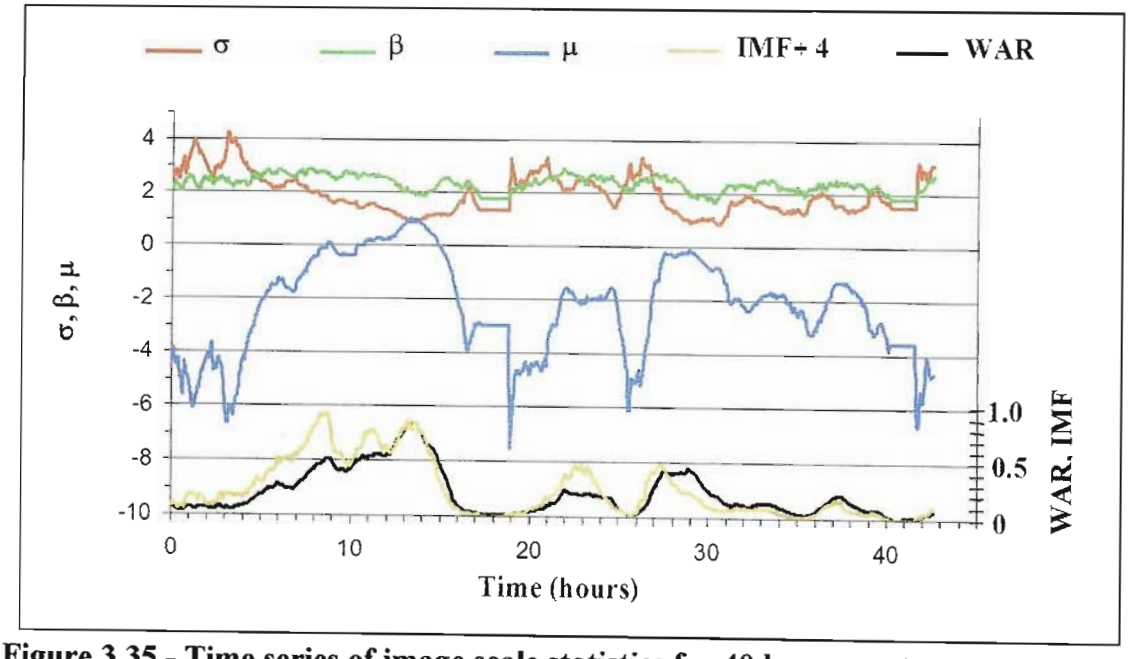

Figure 3.35 - Time series of image scale statistics for 40 hour event

Figure 3.35 contains time series of the image scale parameters for a 40 hour rainfall event which was recorded in February of 1996. These statistics were measured in the early part of the research from integer precision data.

This serves as a qualitative illustration of the type of relationships which can be expected between the image scale parameters discussed thus far. Bearing in mind that the MDV data set only became available later in the research, the rainfall event or *Bead* modelling process which will be defined in chapter 4, was originally formulated and tested on these integer precision data. This rainfall event resulted in a large flood event at the Vaal dam. There are several points to note in Figure 3.35.

- The μ and σ parameters are negatively correlated in log space so that a sequence of images with increasing μ is likely to have a decreasing σ. Conceptually, the more general rainfall is likely to be less variant in log space, and conversely the more isolated (usually convective) rainfall is likely to be more variant.
- The ranges of the various parameters are significant. Obviously the WAR_i ranges between 0 and 100%. The IMF_i ranges between 0 and 4mm/h - the upper limit of this range corresponding to a period of unusually wide-spread, intense rainfall. β ranges between 2.0 and 3.0.
- β appears to be reasonably independent of any of the other image scale parameters.
- Correlation between the image scale parameters of consecutive images is very high.
- The WAR_i time series exhibits a high correlation with the IMF_i time series. μ and σ are related to WAR_i and IMF_i as described in Section 3.4.1.
- The scaling of IMF_i by ¼ to plot it on the same axis as WAR_i is in line with the observation in Figure 3.21 that IMF_i is approximately 4 x WAR_i.

Of the 98500 floating point precision images analysed, approximately 24000 of them were *wet* images (WAR_i in excess of 1%). Since a rainfall event is defined in Section 3.2 as a sequence of consecutive *wet* images, the secondary statistics presented here will be confined to these 24000 images. Any image with WAR_i less than 1% is considered dry.

To begin, Figure 3.36 and Figure 3.37 respectively show on a logarithmic scale, the probability of exceedence of the WAR_i and IMF_i statistics for 24000 *wet* radar images analysed in this study.

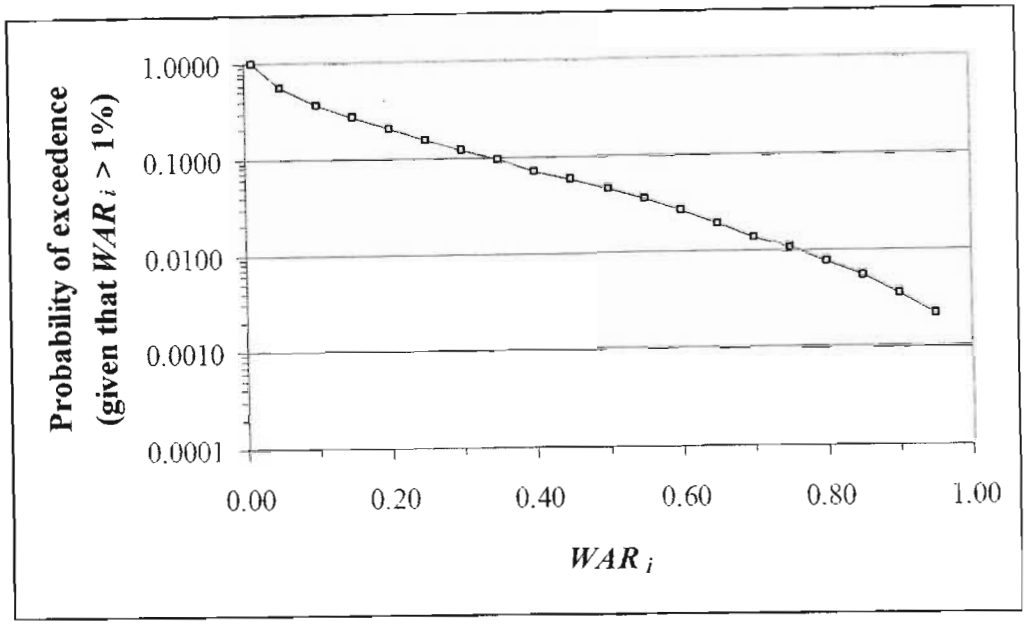

Figure 3.36 - Probability of exceedence of WAR_i as measured from 24000 wet images

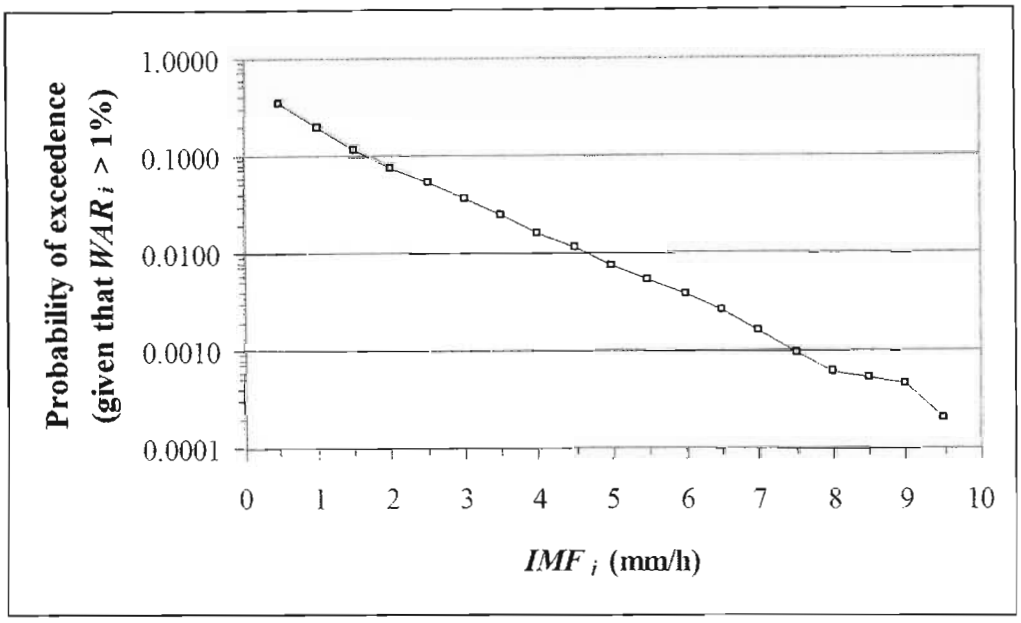

Figure 3.37 - Probability of exceedence of IMF_i as measured from 24000 wet images

The marginal distributions of both of these statistics are severely skewed towards the lower end of their respective ranges and very nearly exponential. 90% of the analysed *wet* images have WAR_i below 35% and IMF_i below 1.65 mm/h. The range of IMF_i is between 0 and 10 mm/h, but based on observations of the data set, images with IMF_i in excess of 6 mm/h usually show signs of bright band contamination and need to be treated with care. Figure 3.38 presents a histogram to illustrate the probability density function of β_{space}.

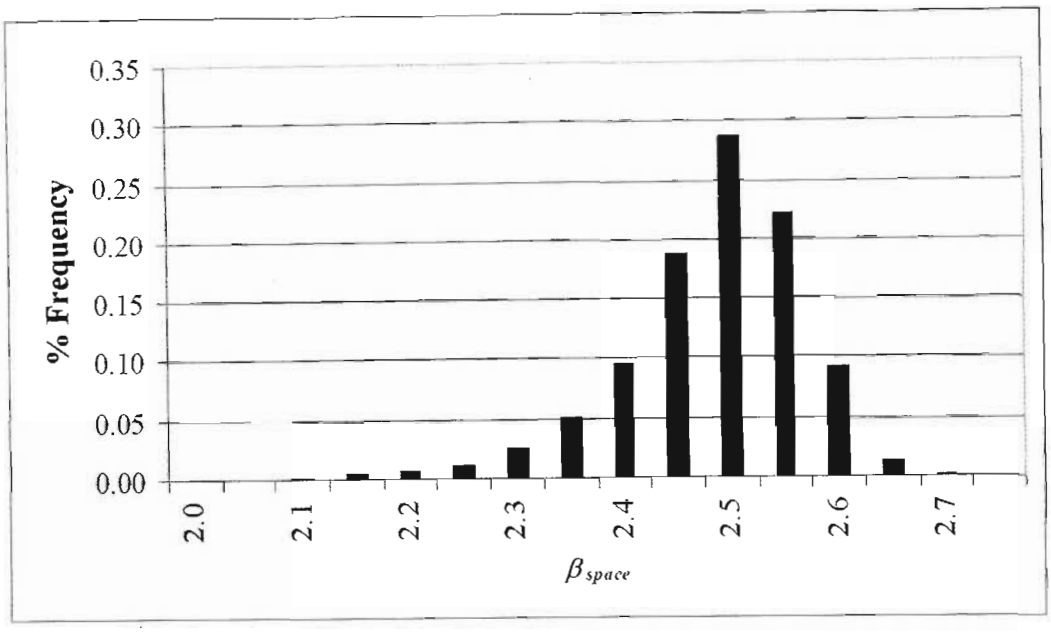

Figure 3.38 - Histogram to illustrate the probability density function of β_{space} for the 24000 wet images

90% of the observed *wet* data had β_{space} between 2.35 and 2.60. To a person who is unfamiliar with this type of analysis β_{space} is a slightly abstract parameter, but it is difficult to discern between the spatial correlation structure of an image with a β_{space} of 2.35 and that of an image with a β_{space} of 2.60. It is possible that much of the variation observed in the β_{space} parameter could be attributed to sampling error. As the qualitative observations of Figure 3.35 suggested, the β_{space} parameter is almost completely independent of the WAR_i and IMF_i parameters. This is confirmed in Figure 3.39 and Figure 3.40 which show β_{space} as a function of WAR_i and IMF_i respectively. Note that WAR_i and IMF_i are both plotted on a logarithmic scale in these figures.

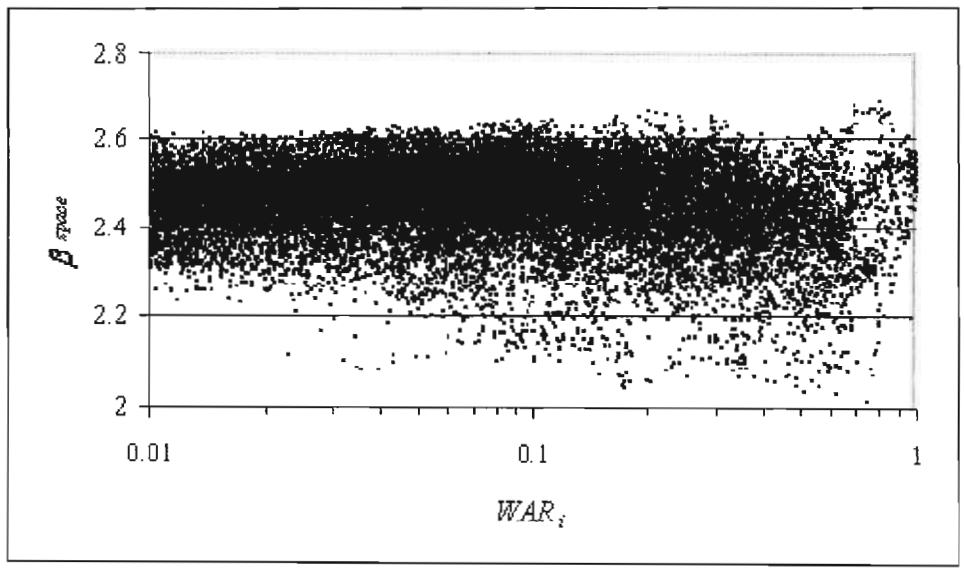

Figure 3.39 – β_{space} parameter as a function of WAR_i parameter for 24000 images

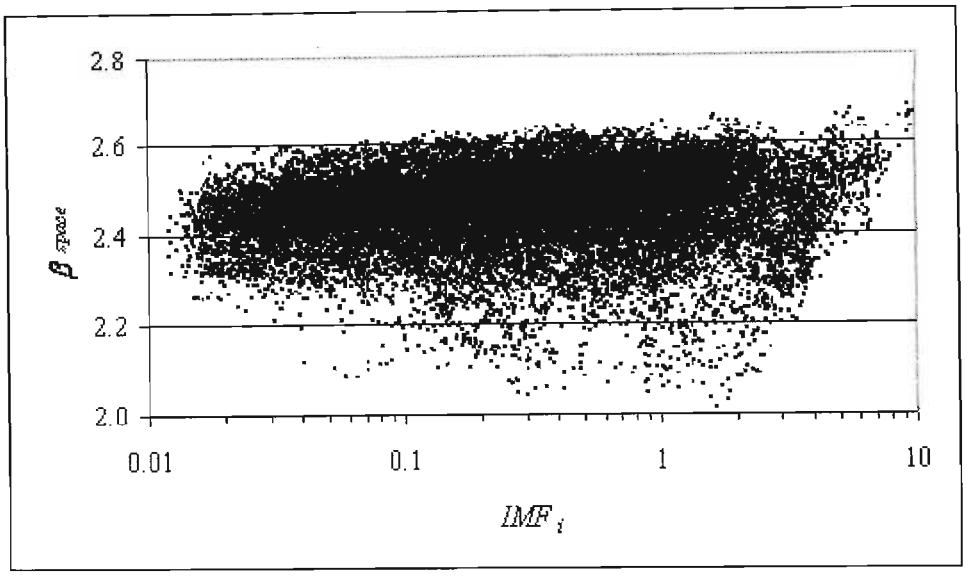

Figure 3.40 – β_{space} parameter as a function of IMF_i parameter for 24000 images

The average β_{space} for all of the images is 2.45 and regression lines fitted through the data of Figure 3.39 and Figure 3.40 reveal gradients of -0.10 in the case of WAR_i and 0.005 in the case of IMF_i. In the case of WAR_i, a gradient of -0.10 suggests some dependent structure in the β_{space} versus WAR_i relationship, however the *negative* gradient is contrary to qualitative observations of the rainfields. Although small, this negative would suggest that β_{space} is decreasing with increasing WAR_i, in other words the more general the rainfall, the more variable the field. This is not the case however, more general rainfall is almost always less variable than scattered rainfall. The negative gradient is as a result of a thresholding process in which rainfall below 0.2 mm/h (12dBz) is taken as noise and set to a constant (small) value. This is done for MDV data storage purposes. Consequently, fields with a low WAR_i appear to be constant over a large proportion of the field and measurement of the spatial correlation is biased upwards resulting in a slightly higher β_{space} at low WAR_i.

The last of the instantaneous image scale statistics to be measured from the space-time radar data is the mean field advection as discussed in Section 3.6. Representation of a vector time series in two dimensions poses a slightly difficult problem, however Figure 3.41 illustrates the cumulative mean field advection vector for the 40 hour event whose image scale statistics are presented in Figure 3.35.

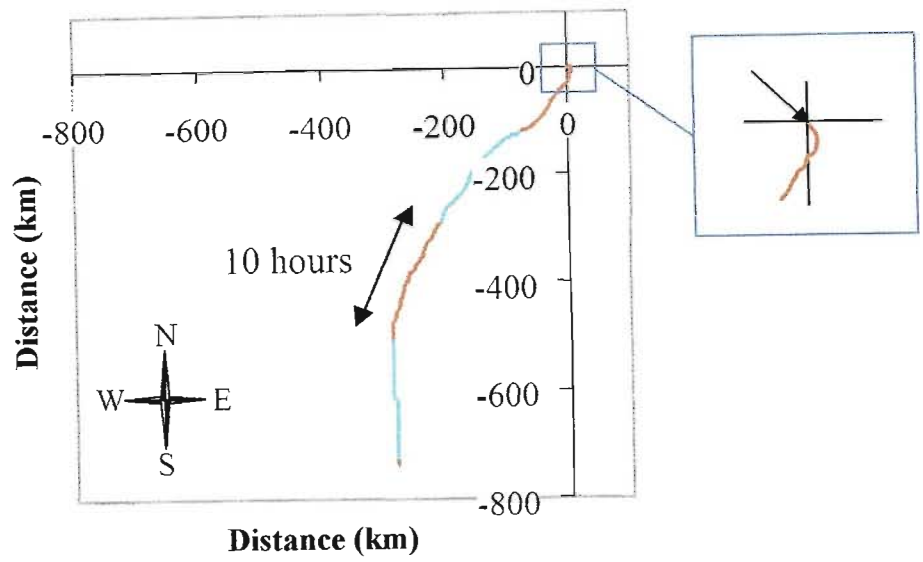

Figure 3.41 - Cumulative mean field advection vector for the 40 hour event of Figure 3.35

The cumulative mean field advection vector of Figure 3.41 was computed by summing the consecutive *shift vectors* (defined in Section 3.6) sampled at each five minute time step, starting at the beginning of the event, time Zero, position (0, 0). In this particular case, the storm began in a north-westerly direction (coming from the north-west) and then swung to a north-easterly direction after 3 hours where it remained until the 30th hour when it switched back to a northerly direction. The storm speed appears to have been reasonably constant at approximately 20 km/h which is typical of the advection speeds in this area.

With only a few years of radar data and consequently a limited number of rainfall events, it is difficult to draw firm conclusions as to the distribution of storm speed and direction. Additional information in this regard can be obtained from weather balloon data which have been in use for several decades. A good indicator of storm advection in the Bethlehem area is the 500 hecta-pascal wind speed and direction (Steyn, PCL and RT Bruintjes, 1990). Daily weather balloon data for Bethlehem (which include the 500Hpa wind data) commencing in January 1980 were obtained from the SAWS.

With 20 years of 500Hpa wind data, a reasonable estimate can be made of the monthly distribution of wind speed and direction. These data are generally presented in the form of a polar histogram, or wind rose, as shown in Figure 3.42.

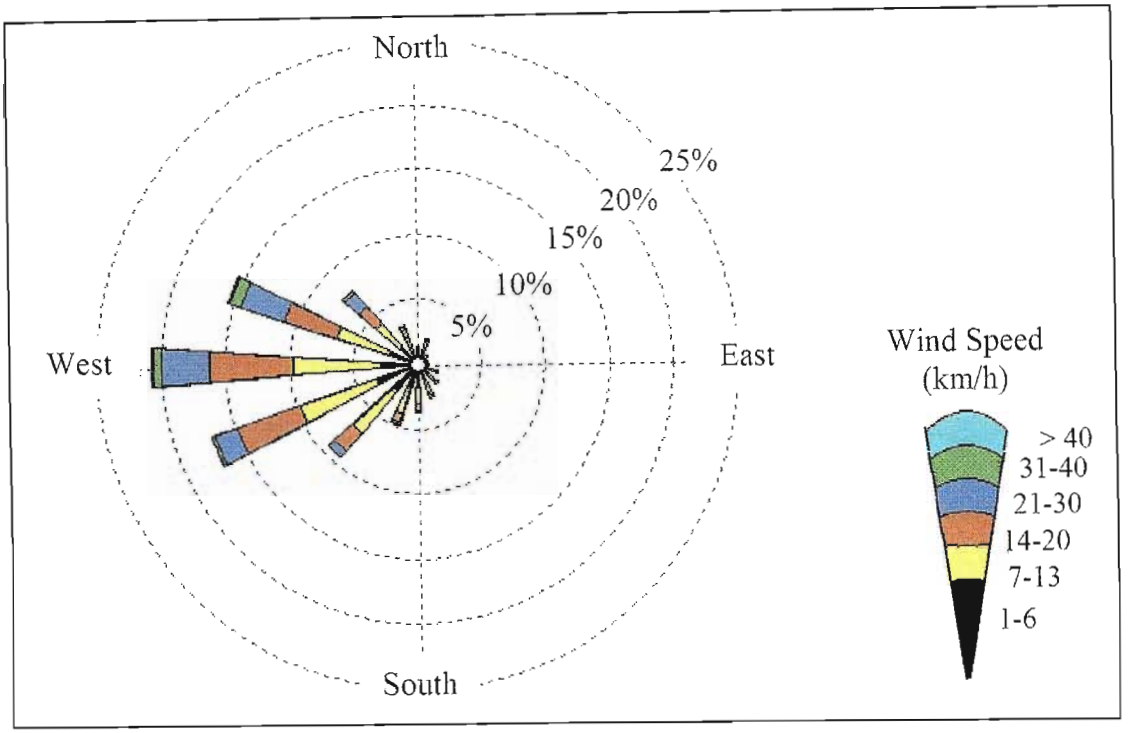

Figure 3.42 - Polar histogram of annual 500HPa wind for Bethlehem

Clearly most of the 500Hpa wind in any year approaches Bethlehem from a westerly direction at a speed between 7 and 20 km/h with 70% of the wind approaching from the western quarter-sector (between north west and south west). Analysis on a monthly basis reveals a similar pattern throughout the year, with a tendency towards the south in the summer months (prevalent direction of west-south-west) and a tendency towards the North (prevalent direction of west-north-west) in the winter months.

These data can be presented in the form of two cumulative frequency distributions - the first which shows the cumulative frequency of the wind direction (increasing clockwise, zero is north), and the second which shows the probability of exceedence of wind speed for each sector of the 16 point compass. The cumulative frequency curve for wind direction is given in Figure 3.43 for the 500HPa wind data as well as the observed rainfall events of the 1998/1999 datasets (non-overlapping datasets). Consistent with observations in Figure 3.42, most of the wind approaches between the west and west-south-west sectors, the point of maximum gradient on the curve of Figure 3.43. Furthermore, cumulative advection vectors of the observed rainfall events compare well to the 500HPa wind directions, supporting the advice of the SAWS.

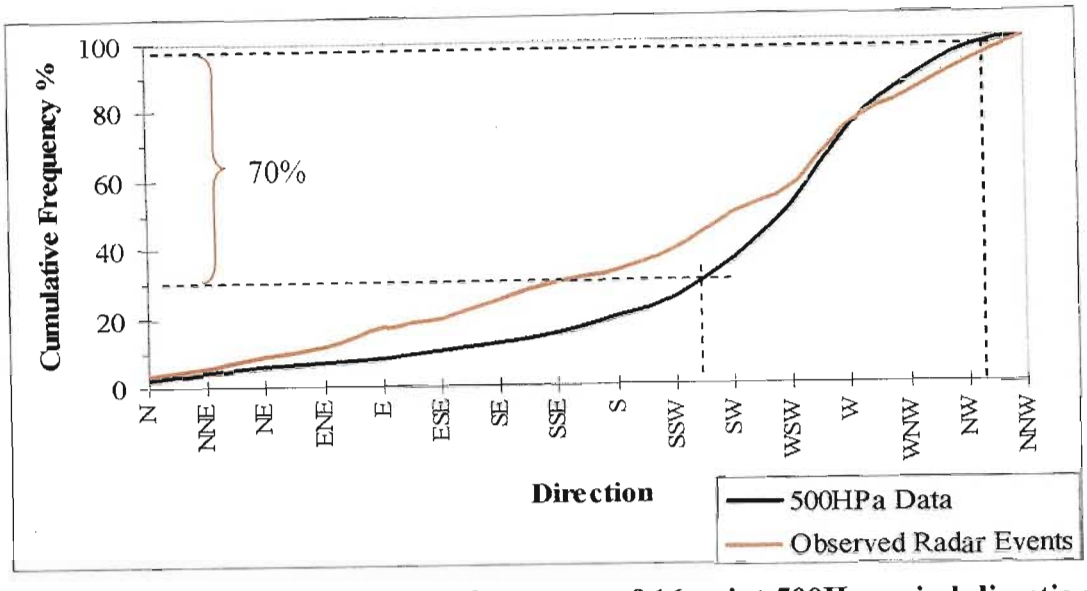

Figure 3.43 - Annual cumulative frequency of 16 point 500Hpa wind direction for the 500HPa data and observed radar events for the 1998/1999 season

Perhaps more revealing than Figure 3.42, Figure 3.44 presents the probability of exceedence of wind speeds in each of the 16 point directions. Winds approaching Bethlehem from the west-north-west tend to be the strongest and from the north-east tend to be weakest.

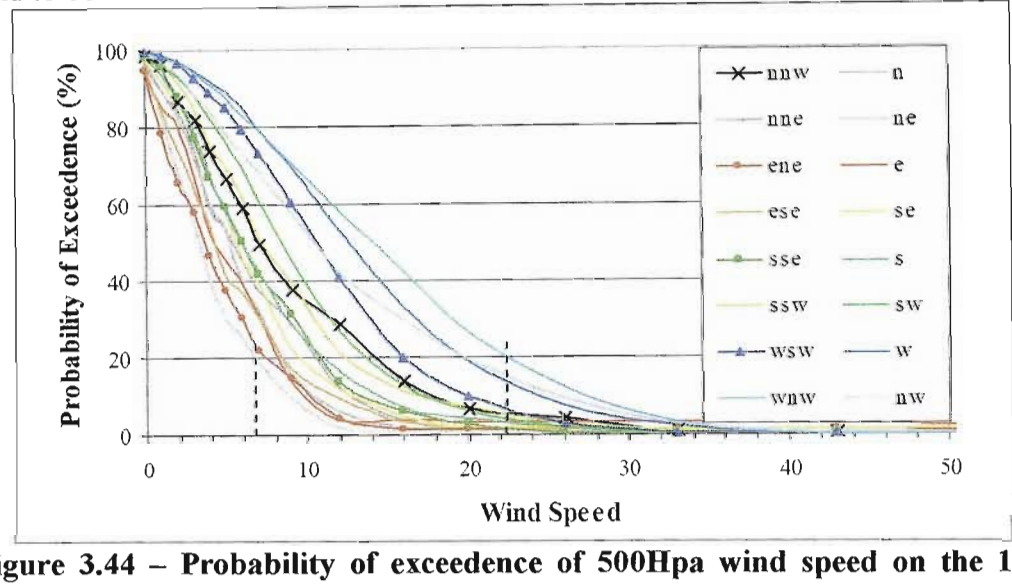

Figure 3.44 – Probability of exceedence of 500Hpa wind speed on the 16 point compass

As an example with reference to Figure 3.44, given that the wind is blowing from the north-east, there is a 20% probability that the wind will exceed 7 km/h (see dotted lines on Figure 3.44). By comparison, given that the wind is blowing west-north-west, the 20% probability of exceedence corresponds to 22 km/h, three times as strong as that from the north-east. Values for the observed radar events range between 0 and 43km/h which correspond well with those of the 500HPa wind speed shown in Figure 3.44.

These curves serve to illustrate the technique of analysis of the wind data. Although not included here, similar curves to those shown in Figure 3.43 and Figure 3.44 were extracted for each month of the year in order to capture the variation in wind speed and direction on a seasonal basis. In light of the analysis presented in Figure 3.42 through Figure 3.44, it is interesting to note that the 1996 event of Figure 3.35 and Figure 3.41 approached from a rather unusual north-easterly direction.

3.7.2 Cumulative image scale analysis

Moving on to the cumulative image scale analysis, a simple test of accumulating the radar images over a complete season serves as a good initial test to check the performance of the radar against the raingauge network. It was this test (performed late in the study when MDV data became available) which revealed that the data supplied in MDV format overstated the rainfall rate by a factor of two as was discussed in Section 2.6. The radar rainfall accumulation for the 1999 calendar year is shown in Figure 3.45.

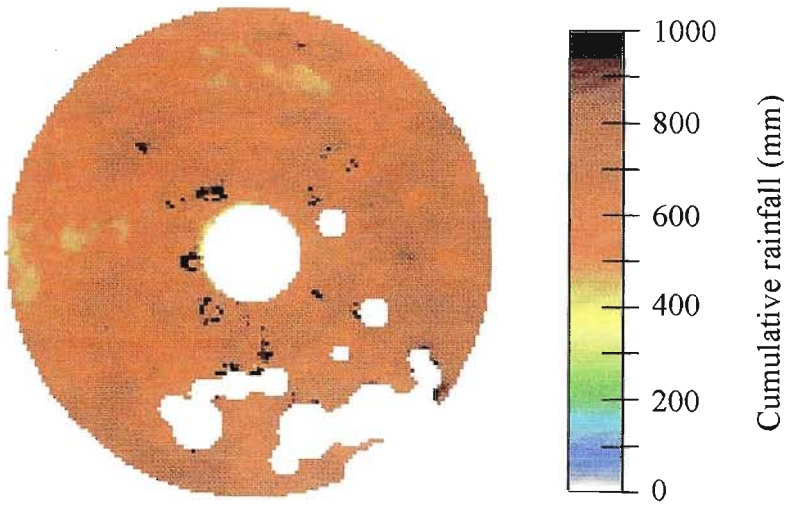

Figure 3.45 - Annual radar rainfall accumulation for the 1999 calendar year

The spatially averaged, accumulated rainfall during 1999 on the study area is 615 mm. With reference to Figure 3.3, which considers the distribution of annual mean rainfall as measured by raingauge network on the study area, the average rainfall recorded in 1999 is slightly lower than the mean of 654 mm estimated from the gauges. Distribution of cumulative rainfall on the image is approximately lognormal with wetter regions in the south west near the Maluti mountains. Note that the effects of ground clutter are exaggerated in cumulative images and although most of it is masked out for analysis, there is still evidence of it in parts of the image. Later studies will address the removal of this low-grade ground-clutter.

The 98000 instantaneous images were accumulated to 307 daily images from which the WAR_d and IMF_d statistics were measured. Figure 3.46 and Figure 3.47 show the probability of exceedence of the WAR_d and IMF_d statistics computed from these 307 daily radar accumulations and the comparison is made to the same statistics measured by the rain gauge network and plotted in Figure 3.8 and Figure 3.9 earlier in this chapter.

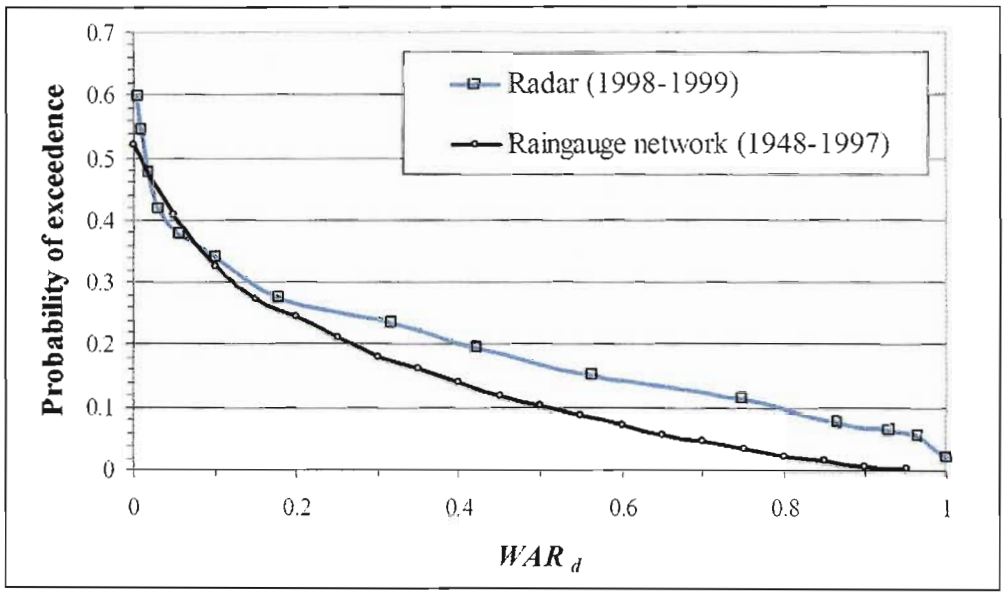

Figure 3.46 - Probability of exceedence of WAR_d as measured by radar and compared to the same measured by rain gauge network (note that these are sampled from non-overlapping periods)

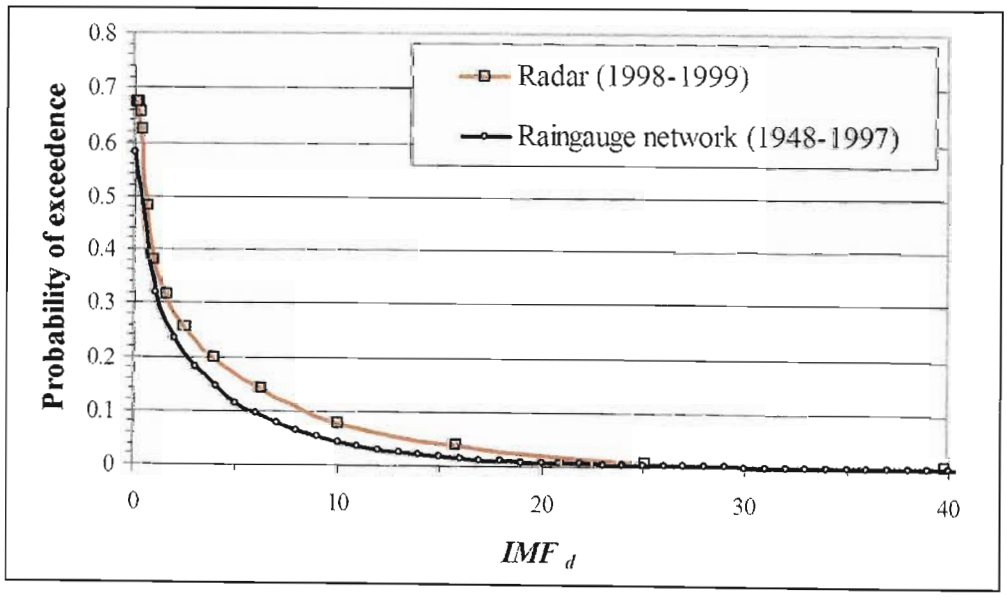

Figure 3.47 - Probability of exceedence of IMF_d as measured by radar and compared to the same measured by rain gauge network (note that these are sampled from non-overlapping periods)

Figure 3.46 and Figure 3.47 suggest that the radar accumulations yield a higher estimate of the cumulative rainfall on the study area. This is indicated by the fact that the probability of exceedence of WAR_d and IMF_d is significantly higher in the case of radar data when considering the mid and upper ranges of these parameters. Overlapping datasets would facilitate a more convincing comparison, but unfortunately infilled raingauge data beyond 1997 had not been processed at the time of analysis. Although 1998 and 1999 were relatively wet years, differences in exceedence of these magnitudes seems excessive and may warrant further investigation. Once the radar has been calibrated so that long term (annual) accumulations agree with those of the raingauge network, which of the two datasets better estimates the rainfall on the catchment poses a difficult question, but conventional thinking is inclined to favour the raingauge network.

3.7.3 An empirical relationship between instantaneous and daily statistics

A useful result which can be used to draw comparisons between the statistics of daily raingauge data and those of the instantaneous radar data, is the relationship between WAR_d and WAR_i. Since the daily images from which the WAR_d are measured, are accumulations of the instantaneous images from which the WAR_i are measured, it stands to reason that there will be some relationship between WAR_d and $\overline{WAR_i}$, the daily average of the WAR_i. Figure 3.48 illustrates this relationship.

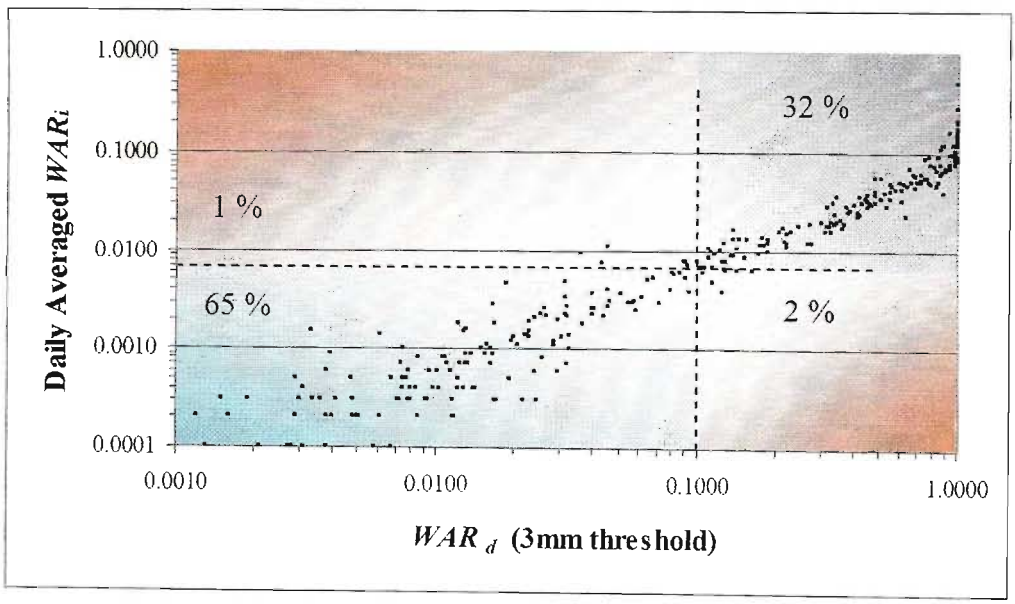

Figure 3.48 - Empirical relationship between $\overline{WAR_i}$ and WAR_d

Logarithmic axes are used in order to present a clearer illustration of the distribution of the WAR_d and $\overline{WAR_i}$ parameters, both of which are skewed towards the lower end of

their respective ranges. Obviously completely dry images with $WAR_d = 0$ or with $\overline{WAR_i} = 0$ cannot be represented on this plot. With reference to Section 3.2, a so called wet day has a WAR_d in excess of 10% and the increasing nature of the $\overline{WAR_i}$ versus WAR_d plot makes it possible to select a corresponding $\overline{WAR_i}$ to serve as an alternative means of classifying a wet or dry day. Based on the data analysed thus far, a $\overline{WAR_i}$ of 0.007 appears to be a suitable choice of threshold for this purpose. Using these thresholds, Figure 3.48 can be divided into three regions:

- blue region in which both classification methods suggest a dry day
- grey region in which both classification methods suggest a wet day
- red region in which the classification methods disagree

The 10% WAR_d threshold was chosen deliberately as the lowest convenient threshold where the scatter was not too great. The upper red region represents the case in which $\overline{WAR_i}$ suggested a wet day while WAR_d suggested dry and the lower red region represents the opposite scenario. The two methods of classification agree in 97% of the days considered. The 3% probability of disagreement is split between the upper and lower red regions, with a slightly higher probability of misclassifying in the lower red region where $\overline{WAR_i}$ suggests a dry day whilst WAR_d suggests a wet day.

Referring to Section 3.3.1 which illustrates the fact that the 54 raingauge network provides a reasonably unbiased estimate of WAR_d, combined with the results of this section, a crucial link between gauge and radar image scale statistics can be inferred, thereby bridging the gap in spatial and temporal scale between the two datasets. An immediate application of this result is to extract the daily wet/dry runs from the WAR_i statistics of a simulation in order to compare with those extracted in Section 3.3.2 for gauge data. This is a technique used in the validation of the *String of Beads* model to be discussed in Chapter 6.

Analysis thus far has been concerned with the *image scale* descriptive statistics of individual *instantaneous images* and of individual *cumulative images* constructed by aggregating the instantaneous images in time. The next stage of the analysis considers temporal aspects of the data. This includes:

- statistics which describe the duration of wet and dry spells

- time series analyses of the image scale statistics during the wet spells
- time series of the Pixel Scale Intensity (PSI), the rainfall intensity at a pixel in the rainfield.

3.8 *THE NORMALISING TRANSFORM*

In order to make use of Gaussian time series analysis and modelling techniques, a means of transforming skewed time series such as WAR_i and IMF_i, of undefined marginal distributions, into Gaussian time series (and back), is required. For this study, a normalising transform is defined using the u-statistics of the distribution. This technique is presented here as a crucial tool in the time series analysis of the image scale statistics.

Consider the set of n values, X_j, with a probability density function of arbitrary shape. The X_j can be transformed using the normalising transform into a corresponding set X_j^*, where the X_j^* are distributed as *N(0,1)*. This is done using Equation 3.27.

$$X_j^* = \Phi^{-1}\left(\frac{r_j}{n+1}\right) \qquad (3.27)$$

where

$j = 1, 2, \ldots, n$

r_j is the rank of X_j

$\Phi^{-1}(..)$ represents the inverse of the cumulative standard Normal distribution

This transform is appropriate only for large values of n.

3.9 *ALTERNATING WET AND DRY SPELLS*

Recalling that a *wet spell* or *rainfall event* is defined as a series of wet images, the Wet Spell Duration (WSD) is defined as the length of time between the start and end of the rainfall event, or the time for which the WAR_i exceeds 1%. The complement of the Wet Spell Duration is the Dry Spell Duration (DSD) between rainfall events, during which the WAR_i is less than 1%. Two key aspects need to be addressed with regards to the duration of wet and dry spells, the first is the marginal distribution and the second is the dependent structure between consecutive spells.

3.9.1 *Marginal distribution of the wet and dry spell durations*

During the one and a half seasons of available data, 405 rainfall events occurred with WSD in excess of 60 minutes. 1998 data are only available from October of that year. 1999 is the only complete year of data during which 293 of the observed 405 events occurred. Again, the marginal distribution of the Wet Spell Duration (WSD) is skewed towards the lower end of its range and the probability of exceedence of WSD (given that WSD is greater than 60 minutes) is shown on a logarithmic scale in Figure 3.49.

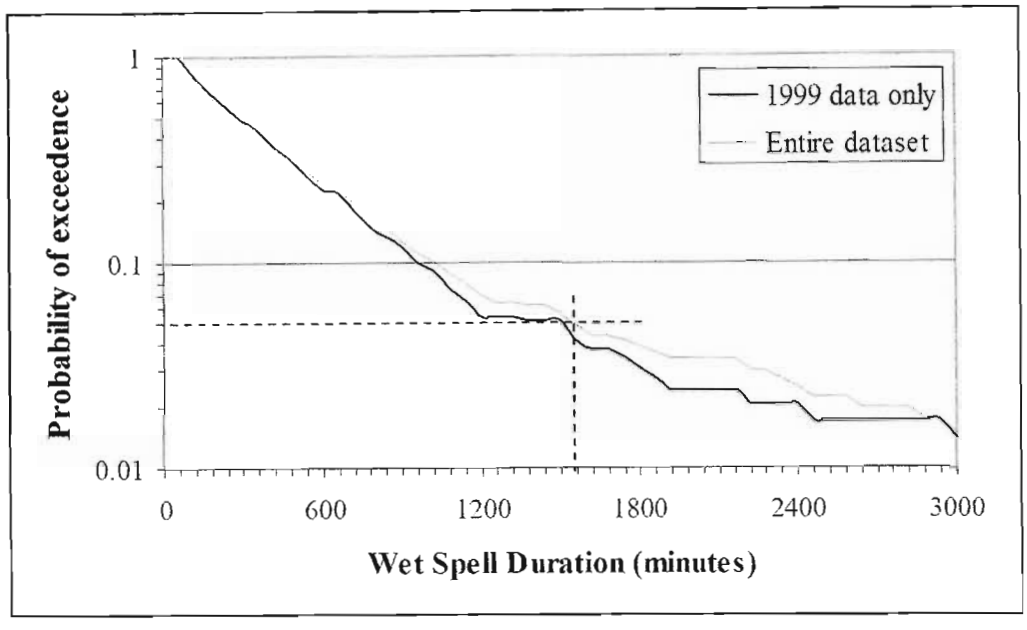

Figure 3.49 - Probability of exceedence of WSD given that WSD > 60 minutes

Two series are plotted in Figure 3.49, the first being that of a complete year (1999) and the second for the entire dataset which spans one and a half years (1998 and 1999). The sixty minute WSD minimum was chosen in order to eliminate beginning/end effects of the event, where the WAR_i can alternate on either side of the 1% WAR_i *wet* threshold. The two series are very similar and this is due to the fact that the vast majority of the rainfall at Bethlehem occurs in the summer season and since the 1998 data is sampled from the summer season, the distribution of WSD is likely to be similar to that of a complete year.

Figure 3.49 is plotted for all observed wet spells, irrespective of the month during which they occurred. Although monthly variation in WSD may be significant, with only one and a half years of data it is difficult to investigate this point with any certainty. It seems unlikely that this will be the case however, since the rainfall generating mechanisms in Bethlehem are reasonably consistent throughout the summer months and virtually non-

existent during the winter. The WSD 95% probability of exceedence stands at approximately one day. Following the work of Haberlandt (1998), this distribution can be approximated by a lognormal distribution and Figure 3.50 illustrates the observed histogram of the WSD and the probability density function of the fitted lognormal approximation to the distribution.

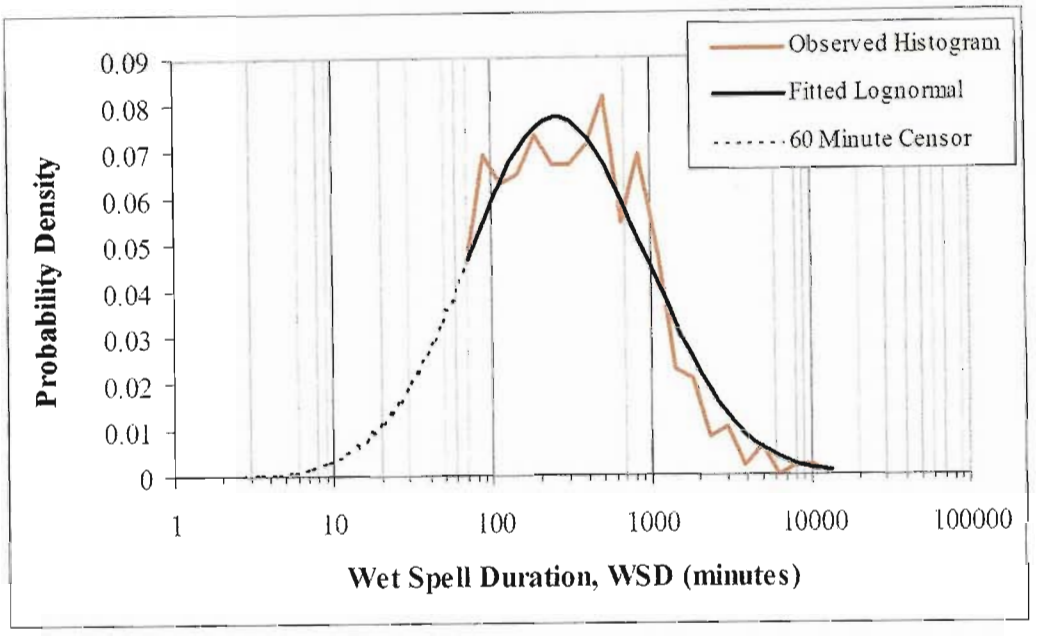

Figure 3.50 - Lognormal approximation of the marginal distribution of the WSD

A censored lognormal distribution is used in this case to account for the fact that only wet spells with duration in excess of 60 minutes were considered in this part of the study. Again the fit appears to be reasonable, although it will fail the Chi-square test.

Moving on to the dry spells, the probability of exceedence of the DSD is plotted on a logarithmic scale in Figure 3.51. As for the WSD above, two series are plotted, one for the entire dataset and the other for the 1999 calendar year. Again the two series are very similar and the same argument applies for the DSD as that given for the WSD. The dry spells plotted here are up to five days and only really represent the distribution of dry spells during the wet summer season. DSD during the winter season is likely to be comprised of one or two dry spells which are between 30 and 90 days long. Spring and autumn are likely to present a slightly different picture, but with only one complete year of data, plotting the probability of exceedence for only a handful of events in each month is a non-informative exercise which will not be attempted here.

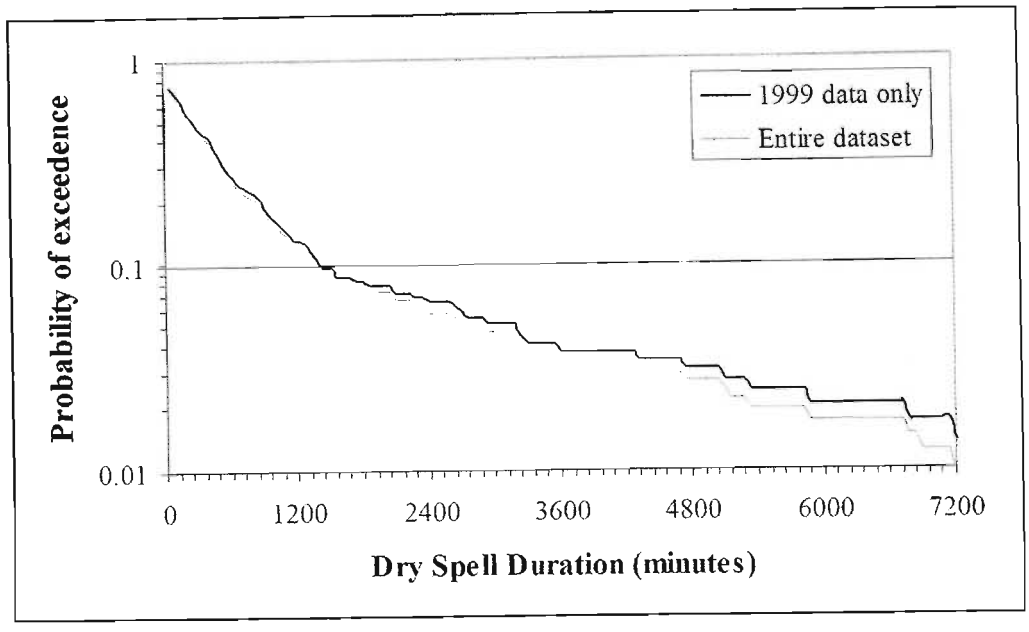

Figure 3.51 - Probability of exceedence of DSD

Harberlandt (1998) found the Gamma and lognormal distributions to be suitable approximations for marginal distribution of the DSD and Figure 3.52 illustrates the fitted approximations (least squares fitting of equal probability bins) of these two distributions for the Bethlehem dataset.

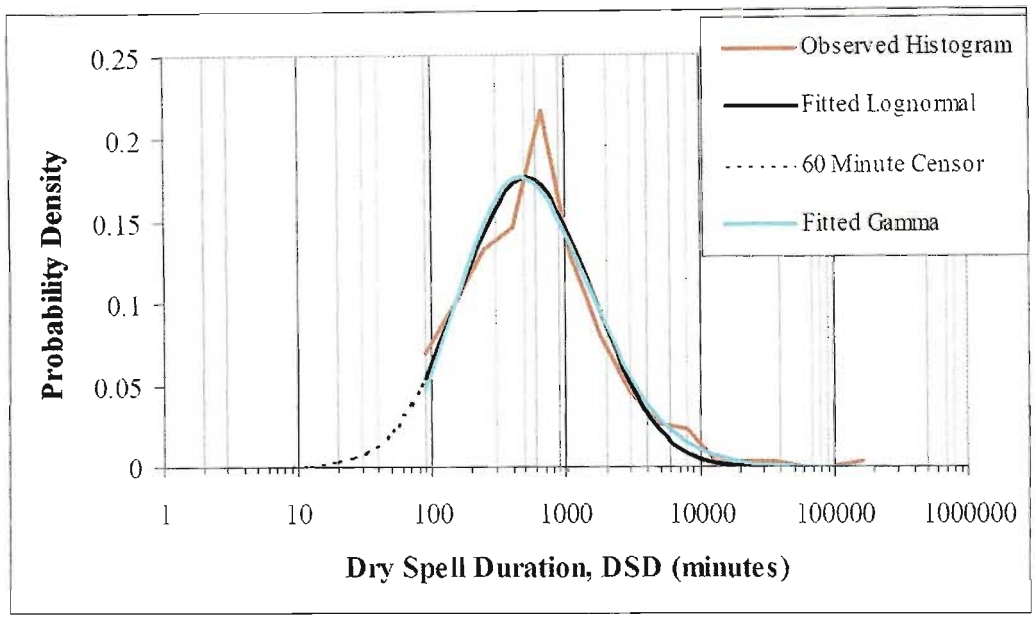

Figure 3.52 - Gamma and Lognormal approximations to the marginal Distribution of the DSD

Clearly there is very little difference between the two fitted distributions, both seem to give reasonable approximations to the observed histogram although neither is quite able to capture the peak. The Gamma is perhaps slightly better in the tail.

3.9.2 *Spell duration dependent structure*

This part of the analysis addresses the question: *Is there any dependent structure in the duration of the wet and dry spells?* Specific examples of this may be:

Is a long wet spell likely to be followed by another long wet spell?

or, equally interesting,

Is a long dry spell likely to be followed by a short wet spell?

amongst other possible permutations. In order to investigate the dependent structure of the wet and dry spell duration, the *normalised* duration of the $(k+1)^{th}$ event is plotted as a function of that of the k^{th} event. Computation of the normally distributed WSD^*_k from the WSD_k was discussed in Section 3.8. Figure 3.53 illustrates the (lack of) Wet Spell Duration dependent structure.

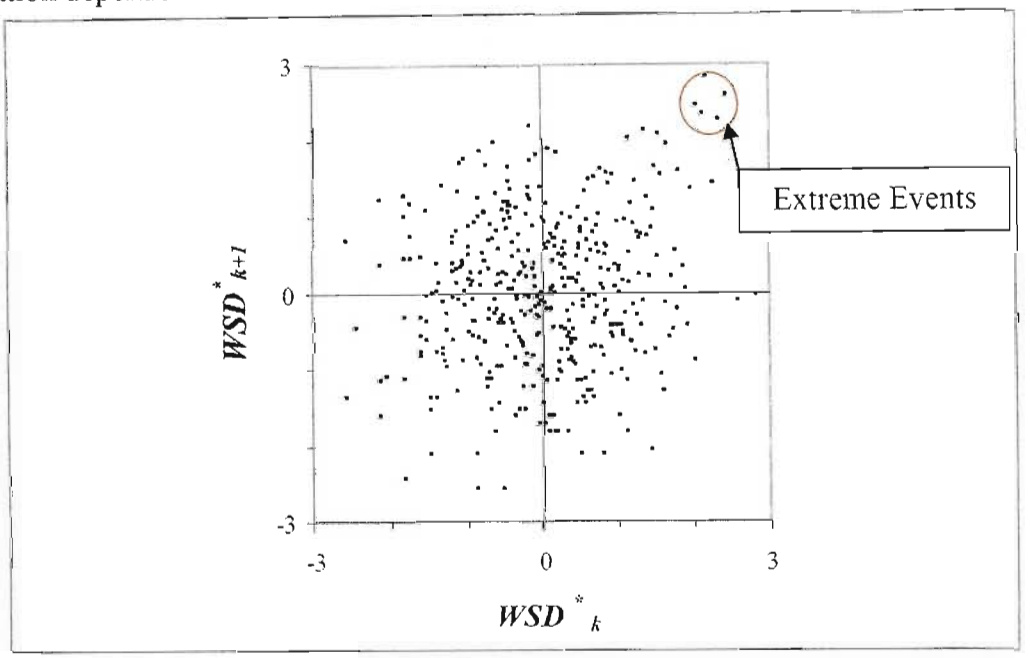

Figure 3.53 - Wet Spell Duration dependent structure

If it were present, dependence structure in the WSD would be revealed in Figure 3.53 as an elliptical or banded distribution of points. There is no obvious dependence structure in the wet spell duration although a linear regression of these points reveals a gradient of 0,14 which is mainly attributable to the extreme wet events marked on the plot. Closer inspection reveals that in three of these cases, what appear to be two consecutive long events, is in fact a single long event which just dips below the 1% WAR_i threshold for a short time (less than 20 minutes). Elimination of these points reduces the gradient of the regression line to below 0.1. Another indicator of the lack of dependence structure is the fact that the R^2 statistic is very nearly zero (0.027). In a similar way,

Figure 3.54 illustrates the (lack of) Dry Spell Duration dependence structure and Figure 3.55 illustrates the (lack of) Dry Spell ~ Wet Spell dependence structure.

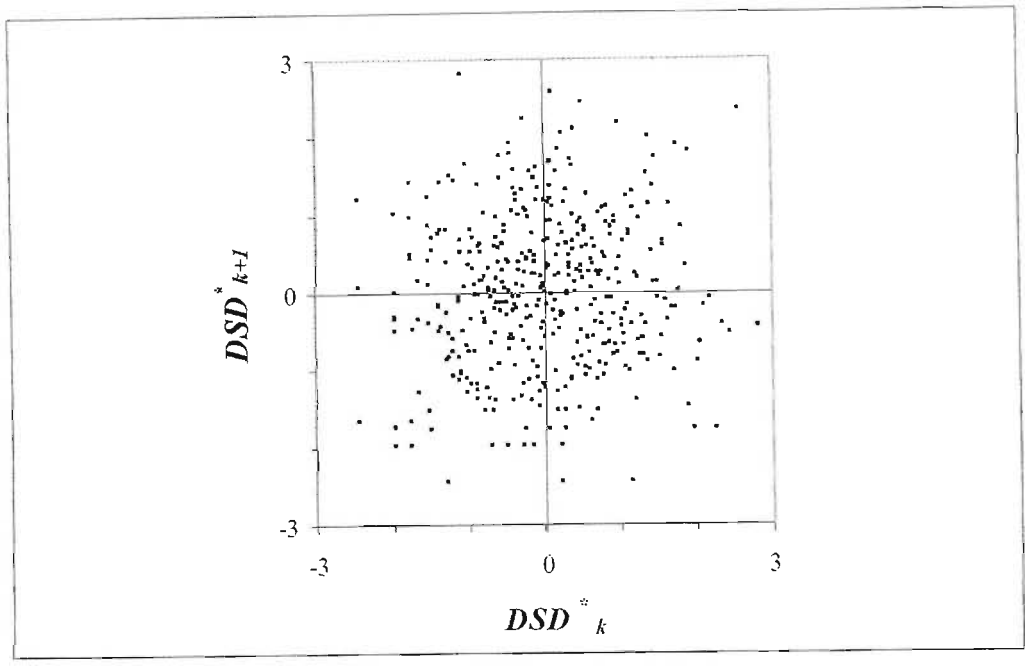

Figure 3.54 - Dry Spell Duration dependent structure

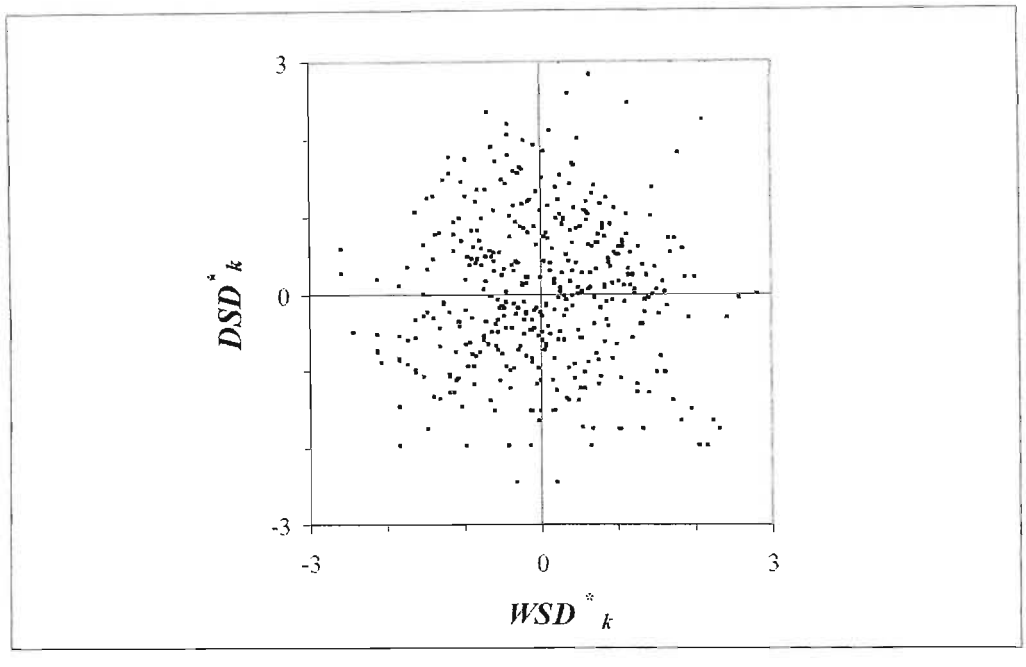

Figure 3.55 - Dry Spell ~ Wet Spell duration dependent structure

In both cases, the lack of any distinct structure in the bi-plot suggests that there is no dependence structure in either the Dry Spell Duration or the Wet Spell ~ Dry Spell Duration. This is confirmed by linear regressions, both of which yield almost zero gradients and intercepts and extremely low R^2 terms of 0,004 and 0,0006 respectively. Consequently, the wet and dry spells can be considered to be independent of each other.

3.10 IMAGE SCALE TIME SERIES

This part of the analysis is confined to the rainfall event, the *wet* time during which the WAR_i exceeds 1%. It is particularly concerned with the temporal behaviour of the *instantaneous* image scale statistics WAR_i and IMF_i, defined in previous sections of this chapter.

3.10.1 Time series transformation

To facilitate the use of Gaussian time series analysis, WAR_i and IMF_i are first transformed, as discussed in Section 3.8, to WAR_i* and IMF_i*, both of which are distributed as *N(0, 1)*. Figure 3.56 illustrates the transformation which links the WAR_i to the normalised WAR_i*.

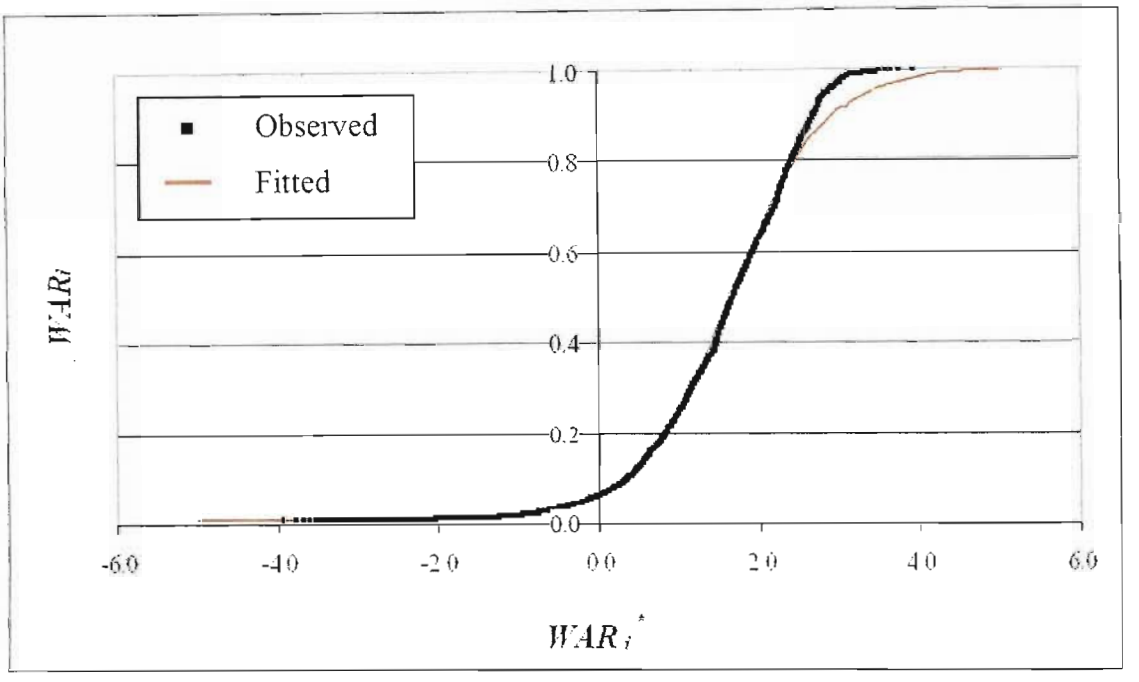

Figure 3.56 - The WAR_i normalising transform (24000 *wet* images)

The fitted curve shown in Figure 3.56 is a scaled and shifted sigmoid function which has the form of Equation 3.28.

$$y = \left(\frac{1}{1 + \exp(b - a.x)} \right).c + d \qquad (3.28)$$

Where a, b, c and d are the fitted constant values and x and y can be substituted for WAR_i* and WAR_i respectively. The function, y, is bounded between c and $(c+d)$. Similarly, the transformation which links the IMF_i to the normalised IMF_i* is plotted in Figure 3.57 and another sigmoid function is fitted in this case.

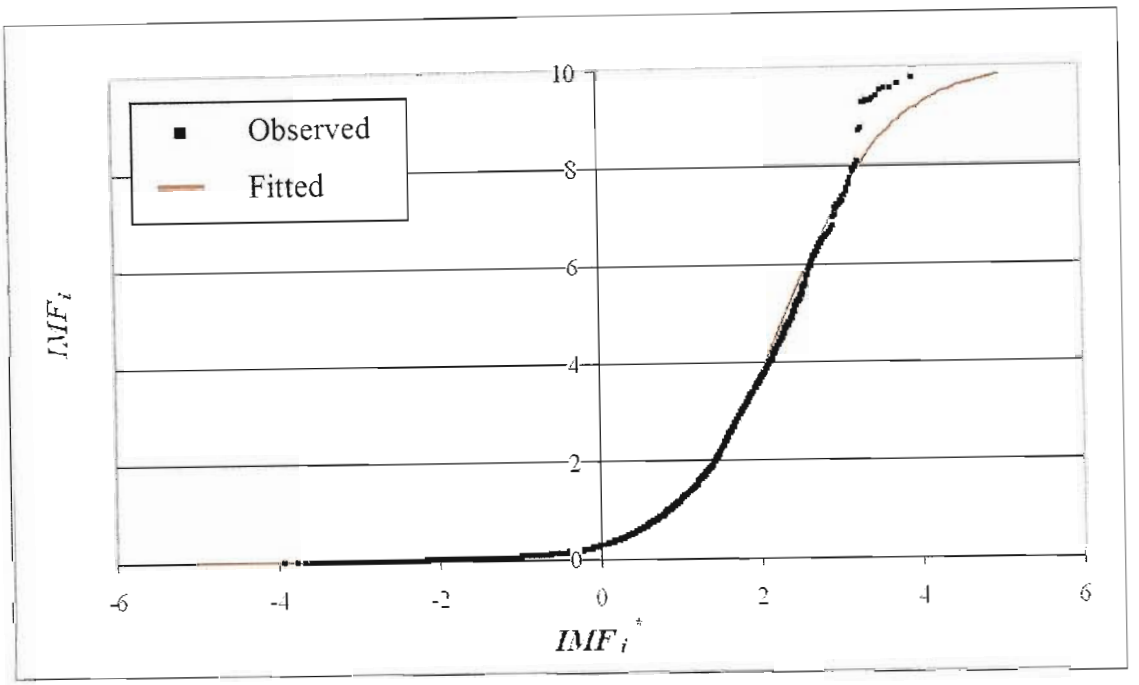

Figure 3.57 – The IMF_i normalising transform (24000 *wet* images)

In both cases the sigmoid function is a good descriptor of the transform in spite of the fact that it does not appear to fit well at the upper end of the range. In the case of the WAR_i transform, values above 0.8 not well described by the sigmoid, but with reference to Figure 3.36, images with WAR_i in excess of 0.8 represent less than 1% of the dataset. In the case of the IMF_i, with reference to Figure 3.37, the fit is even better and the fitted sigmoid accurately represents the transform over more than 99,9% of the observed range. The constant values were fitted by least squares for the 24000 wet images and are found to be:

	WAR_i transform	IMF_i transform
a	1.70	1.50
b	2.80	3.50
c	0.99	9.99
d	0.01	0.01

Note that $d = 0.01$ exactly in both cases, consistent with the definition of a wet image and the discussions following Figure 3.21.

Other functions for describing these transforms include a non-parametric approach of interpolating between the observed values, or a compound function comprised of two separate sigmoid functions, one for the lower half of the transform and the other for the upper half, meeting in the middle. An important computing consideration in this regard

is the complexity of the function, since it must be evaluated hundreds of thousands of times for each year of data. A poor choice of function would therefore consume a large amount of computing time. Use of the transform is illustrated in Figure 3.58.

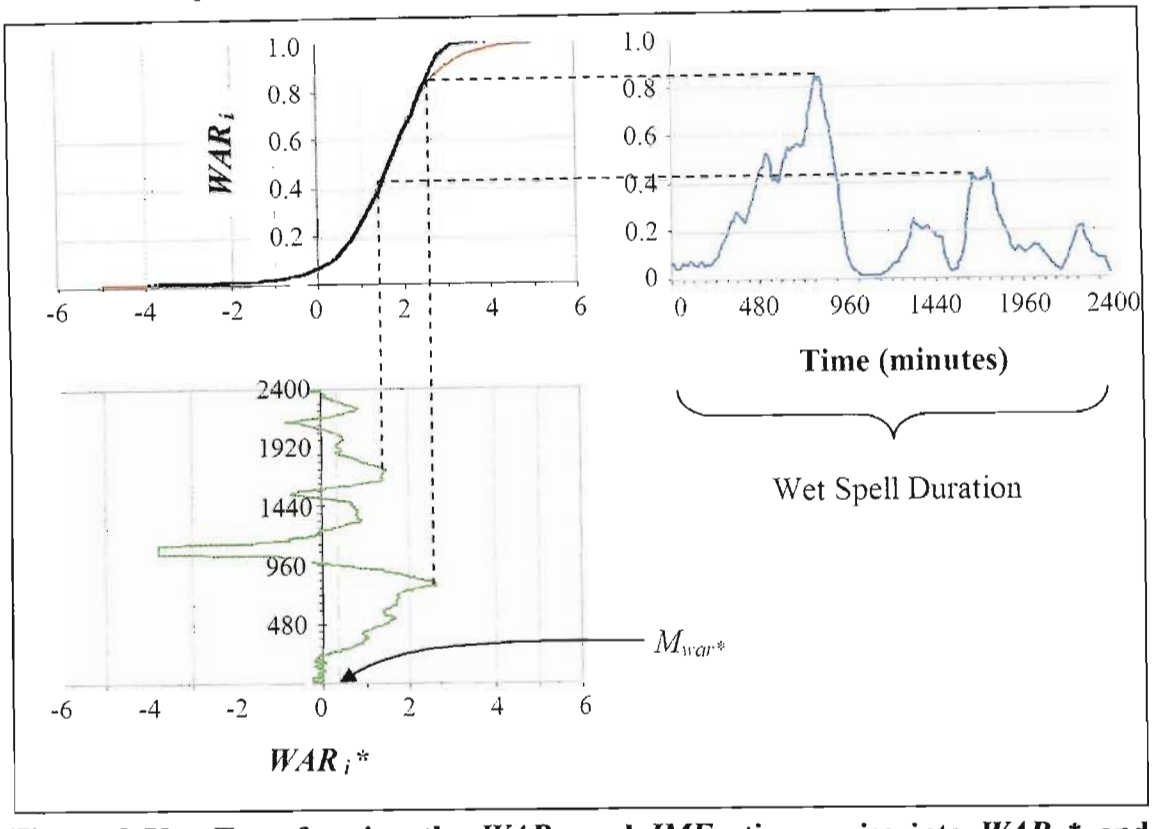

Figure 3.58 - Transforming the WAR_i and IMF_i time series into WAR_i* and IMF_i* respectively

In the example given in Figure 3.58, the WAR_i time series (upper right) of the 1996 event of Figure 3.35 is transformed into the corresponding WAR_i* time series (lower left). With these transforms in place, it is now possible to analyse time series of WAR_i* and IMF_i* using conventional Gaussian time series analysis techniques.

A point worth noting in the use of the normalising transform is its effect on the distribution of any errors in the original time series. It is an unfortunate side-effect of this method of manipulation of time series that the distributions of any errors in the in the original sequence will be skewed in its representation as a transformed sequence by the use of the normalising transform. This applies to the upper and lower parts of the observed range for the transforms used in this study. As an example, with reference to Figure 3.58, small errors in the upper ($WAR_i > 0.9$) and lower ($WAR_i < 0.1$) range of the WAR_i time series, would result in relatively large errors in the corresponding transformed time series due to the relatively flat gradient of the transform beyond these

limits. The same applies to the IMF_i transform. This is an artefact which cannot easily be accounted for. However, the transform is applied to the measured statistics of the data rather than the data themselves, namely the WAR_i and IMF_i statistics, which are estimated from several thousand pixels on each image. Errors in estimation of these statistics are expected to be very small and consequently the distortion of the error terms due to the normalising transform should not significantly affect the modelling process.

3.10.2 Event scale statistics

Event scale statistics describe the distribution of WAR_i* and IMF_i* during an event. The WAR_i* for this particular event are distributed as shown in the histogram of Figure 3.59, as well as the corresponding histogram for the most likely Normal approximation to the distribution, $N(M_{war*}, S_{war*})$, where M_{war*} and S_{war*} respectively represent the most likely approximations of the mean and standard deviation of the WAR_i* for the event. M_{war*} is also indicated in Figure 3.58. A similar approach can be adopted for the IMF_i time series to yield M_{imf*} and S_{imf*}, the corresponding mean and standard deviation of the IMF_i* for the event.

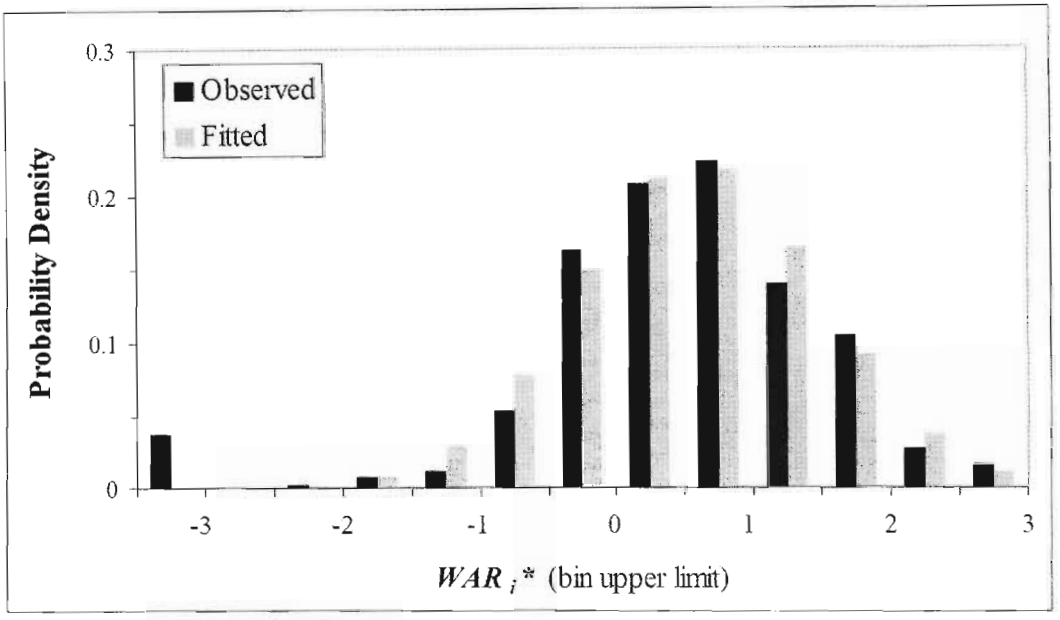

Figure 3.59 - Histogram of WAR_i* for the event of Figure 3.58

The distribution of WAR_i* is well approximated by the Normal distribution, indicating that the use of Gaussian time series analysis techniques is a reasonable approach. Observations below –3 occurred where the WAR_i fell below 0.01 for a short spell and were set to 0,0101 in order to investigate long lag correlations when there was a shortage of data in the early part of the study. This particular event has a D_w of 2400

minutes – longer than the average, with reference to Figure 3.50 – resulting in M_{war*} and S_{war*} statistics of 0,55 and 0,88 respectively. Although the entire set of observed WAR_i (24000 wet images) were transformed onto WAR_i* such that WAR_i* is normally distributed $N(0,1)$, the distribution of WAR_i* sampled from an isolated event is unlikely to be $N(0,1)$. Longer duration events are usually the result of large systems and consequently higher WAR_i. Conversely short duration events are usually associated with isolated convective activity with lower WAR_i. To investigate this phenomenon further, Figure 3.60 considers the M_{war*} as a function of D_w for 140 events. The 140 events were chosen as the best of the 405 observed in the dataset because they contained no radar interruptions and therefore had continuous time series of both WAR_i and IMF_i. The Wet Spell Duration in hours (D_w) is plotted on a logarithmic scale.

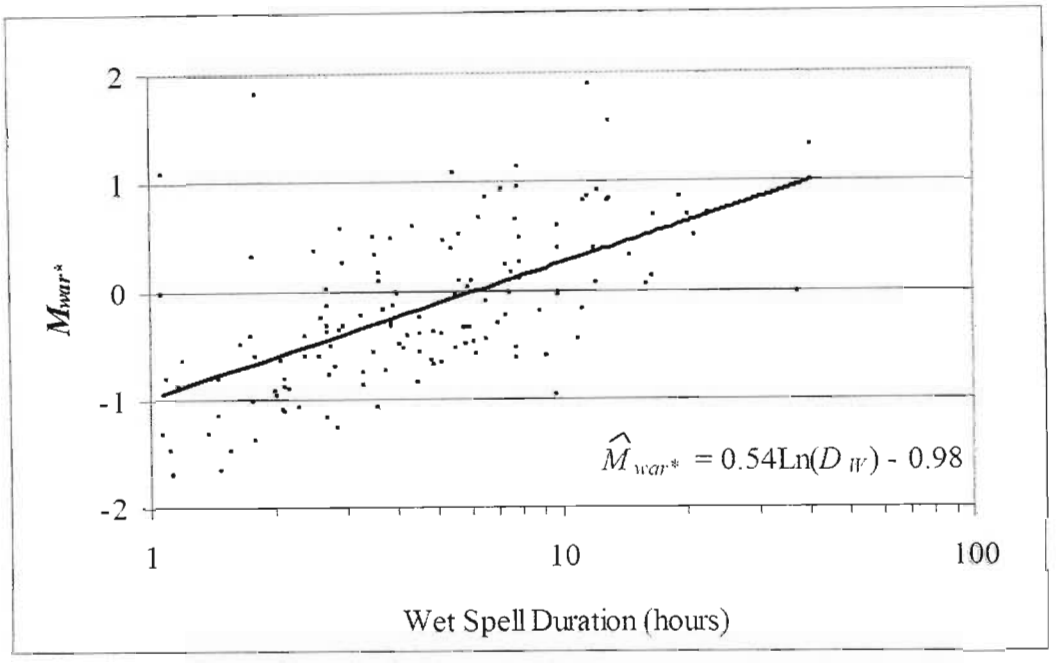

Figure 3.60 – M_{war*} as a function of the wet spell duration, D_W

Clearly there is an increasing trend in the M_{war*} with increasing WSD, and the fitted log-linear function appears to capture this trend reasonably well, with points randomly distributed on either side of the trendline. Defining $\hat{M}_{war*}$ as the log-linear approximation of M_{war*} for any given D_W, and ε_{war*} as the residual value computed as $\varepsilon_{war*} = M_{war*} - \hat{M}_{war*}$, an interesting analysis is to consider the distribution of ε_{war*}. The observed density function of ε_{war*} is plotted in histogram form in Figure 3.61, along with the most likely Normal approximation to the distribution.

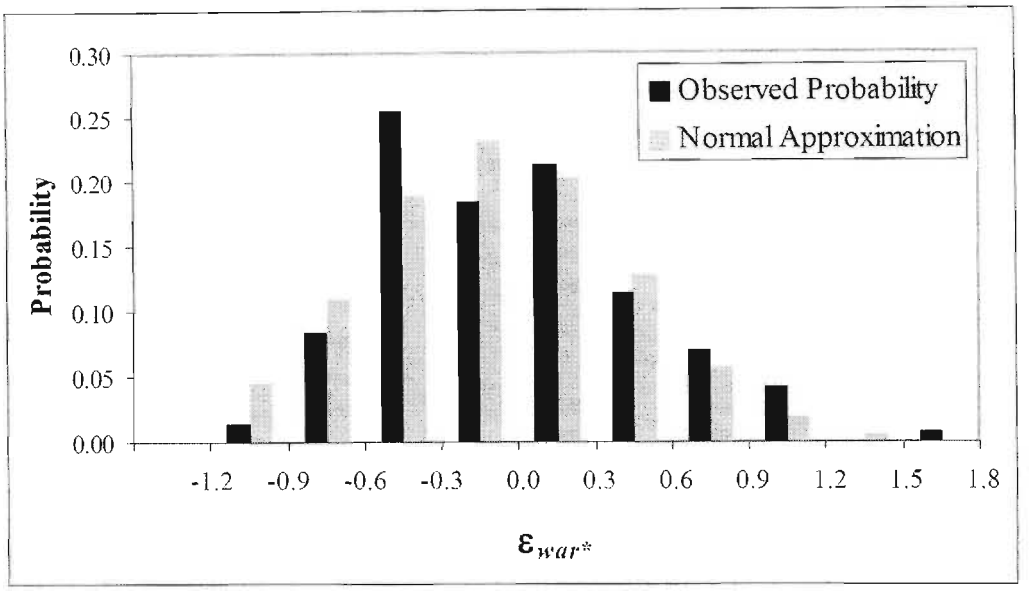

Figure 3.61 - Histogram of $\hat{M}_{war*}$ residuals and most likely Normal approximation

The distribution of the ε_{war*} is slightly skewed, but otherwise reasonably well approximated by the Normal distribution with mean of -0,12 and standard deviation of 0,51. Moving on to the S_{war*}, this too is dependent on D_w and consequently the M_{war*}. It is convenient from a modelling perspective to look at the relationship between S_{war*} and M_{war*} rather than S_{war*} and D_w. Figure 3.62 considers this relationship.

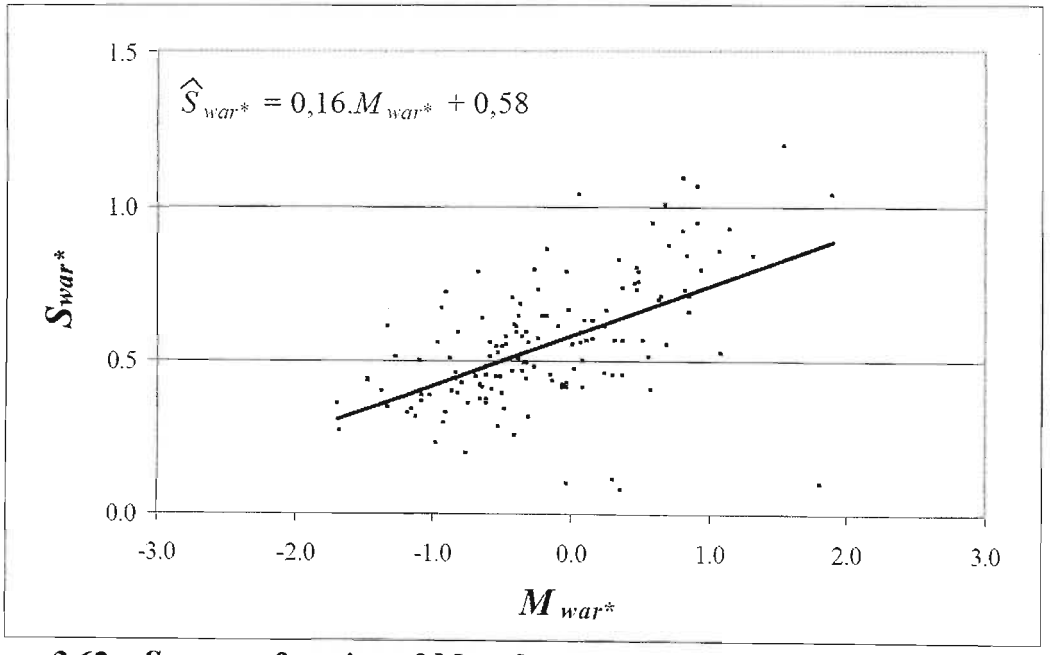

Figure 3.62 – S_{war*} as a function of M_{war*} for 140 observed events

The relationship between these two parameters is not particularly well defined, however there is an increasing trend in the S_{war*} with increasing M_{war*} as indicated by the regression line. Defining $\hat{S}_{war*}$ as the linear approximation of S_{war*} for any given M_{war*}, and δ_{war*} as the residual value computed as $\delta_{war*} = S_{war*} - \hat{S}_{war*}$, Figure 3.63 illustrates

in histogram form, the observed distribution of the δ_{war*} as well as the corresponding most likely Normal approximation, with mean of 1,00 and a standard deviation of 0,25. Again, the Gaussian is a reasonable approximation of the marginal distribution of the residuals δ_{war*}.

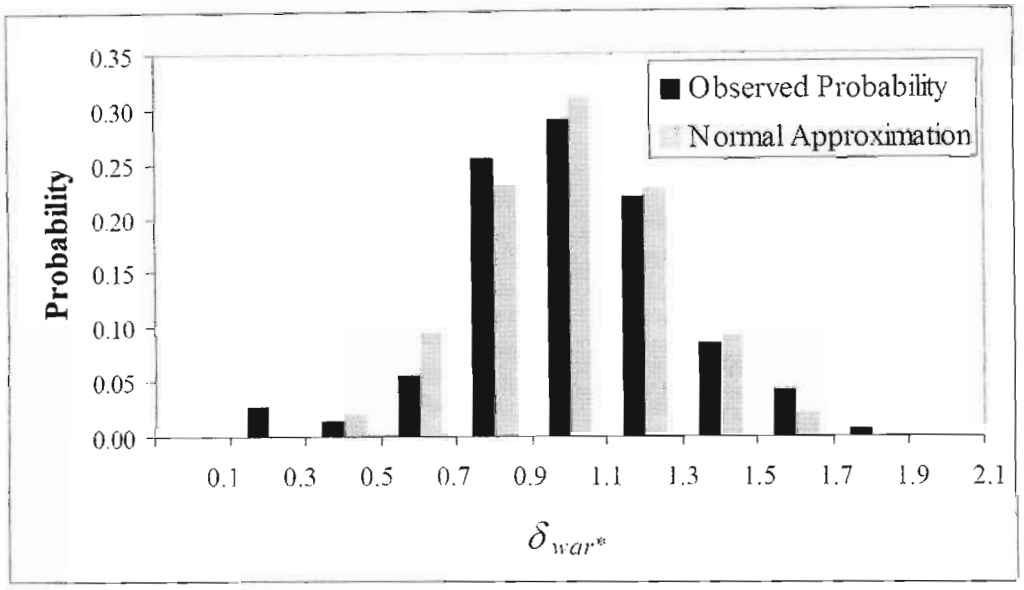

Figure 3.63 – Histogram of $\hat{S}_{war*}$ residuals and most likely Normal approximation

Following the same theme, the relationship between M_{war*} and M_{imf*} is illustrated in Figure 3.64 and the corresponding distribution of residuals $\varepsilon_{imf*} = \hat{M}_{imf*} - M_{imf*}$, is plotted in histogram form in Figure 3.65 with its most likely Normal approximation, a mean of 0,01 and a standard deviation of 0,35.

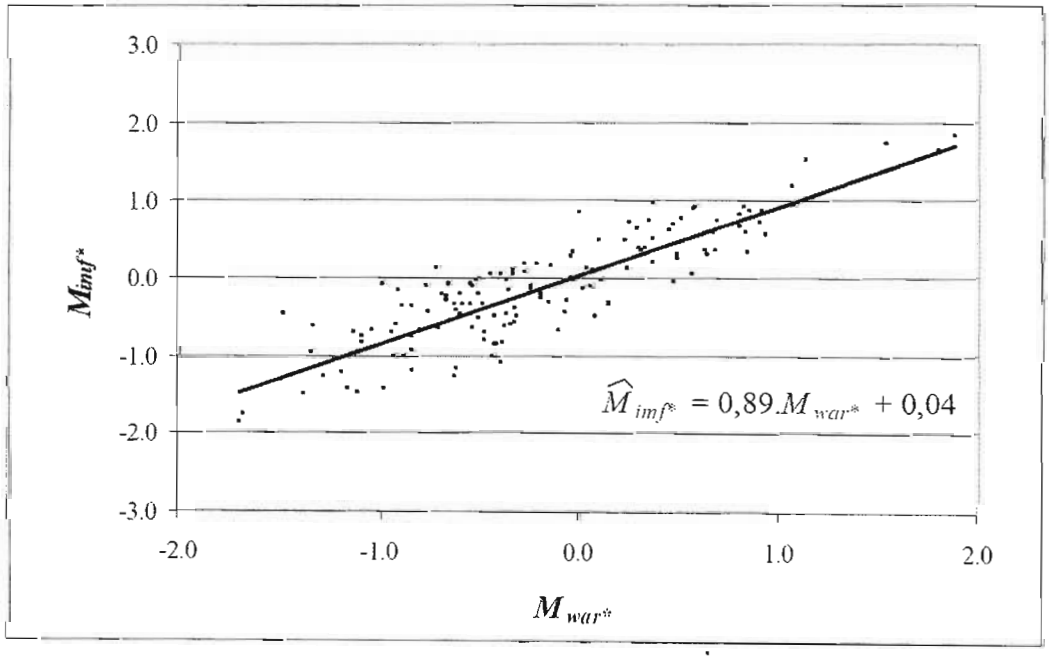

Figure 3.64 - The relationship between M_{war*} and M_{imf*}

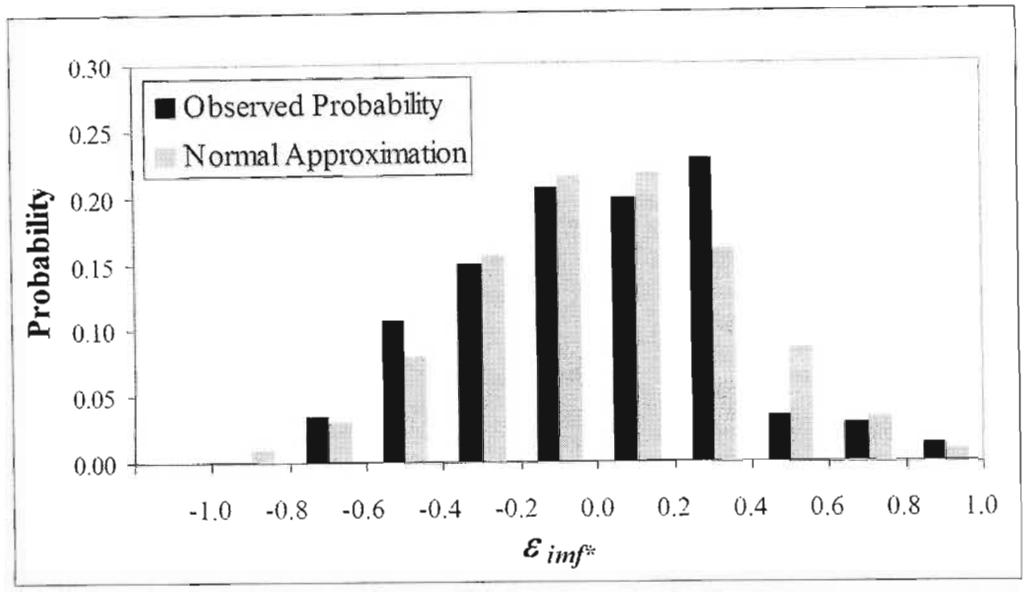

Figure 3.65 – Histogram of $\hat{M}_{imf*}$ residuals and most likely Normal approximation

Finally, the relationship between S_{war*} and S_{imf*} is plotted in Figure 3.66. The $\hat{S}_{imf*}$ regression line is computed for any given S_{war*} and the distribution of residuals, $\delta_{imf*} = \hat{S}_{war*} - S_{war*}$, is plotted in histogram form in Figure 3.67 with its most likely Normal approximation, a mean of 0,00 and a standard deviation of 0,11.

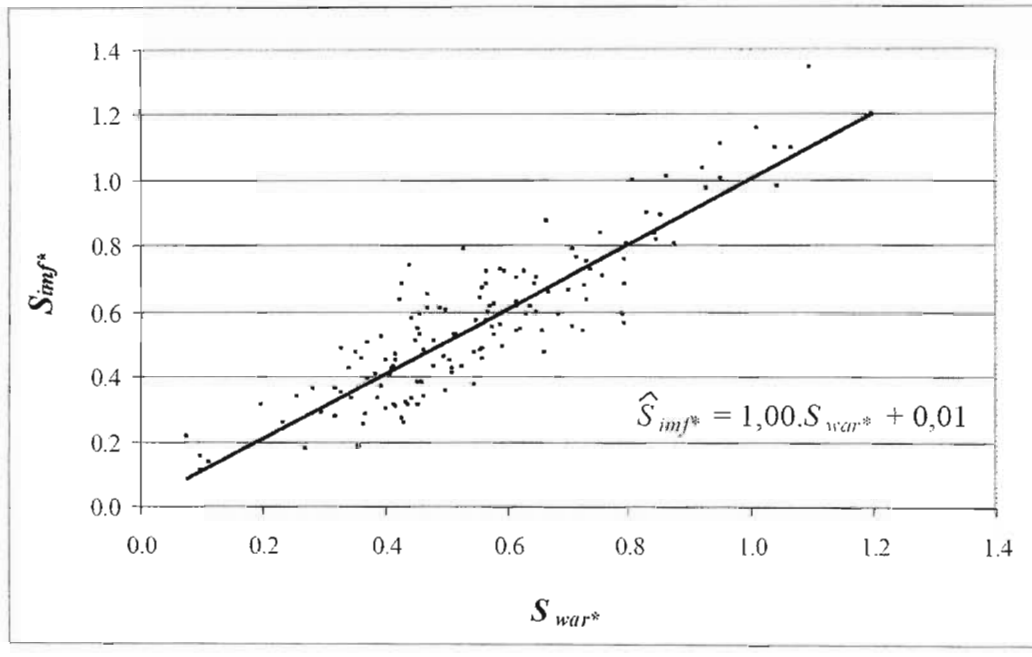

Figure 3.66 - The relationship between S_{war*} and S_{imf*}

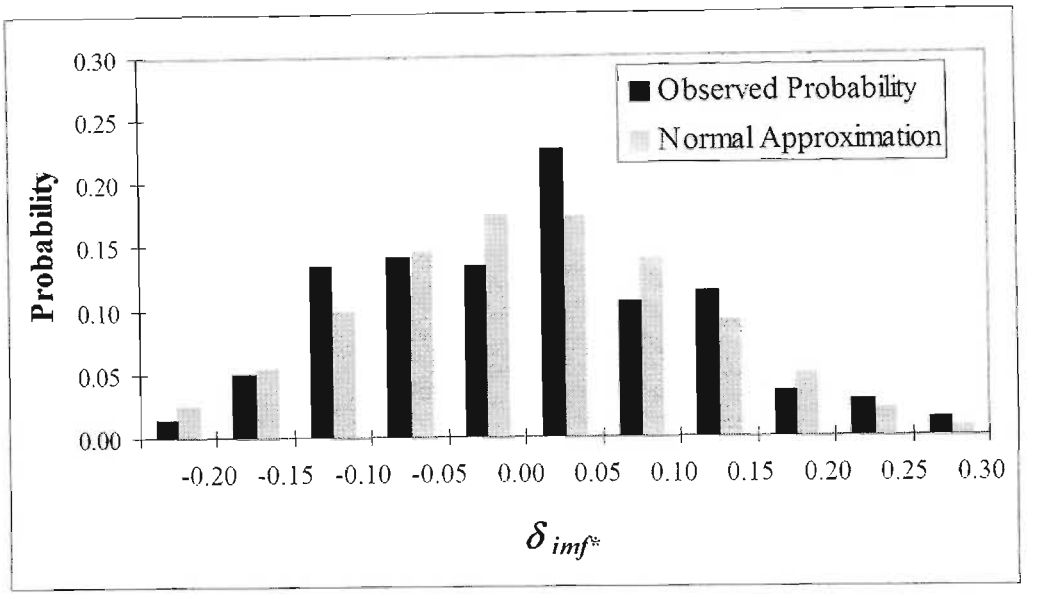

Figure 3.67 – Histogram of $\hat{S}_{imf^*}$ residuals and most likely Normal approximation

The normal distribution is used effectively to approximate the residual histograms of Figure 3.61, Figure 3.63, Figure 3.65 and Figure 3.67. Other distributions or even a non-parametric option could be considered, however the Normal distribution offers a simple approach which fits the residuals well, is well understood and therefore a good starting point with this limited number of events.

The important thing in simulation is to re-create with the correct inter-correlation between the variables. The technique adopted here can be seen as a *cascade of conditional distributions* describing the multivariable distribution of a complex process.

3.10.3 Correlation in time

What remains in the analysis of the image scale parameters, is to consider the temporal correlation (defined in Equation 3.16) of the WAR_i^* and IMF_i^* time series. Four cases are considered in this analysis, all up to a lag (τ) of 120 minutes (30 time steps):

WAR_i^* autocorrelation, $Corr(WAR^*_{i(t)}, WAR^*_{i(t+\tau)})$

IMF_i^* autocorrelation, $Corr(IMF^*_{i(t)}, IMF^*_{i(t+\tau)})$

WAR_i^*~IMF_i^* cross correlation, $Corr(WAR^*_{i(t)}, IMF^*_{i(t+\tau)})$

IMF_i^*~WAR_i^* cross correlation, $Corr(IMF^*_{i(t)}, WAR^*_{i(t+\tau)})$

For discrete data, the correlation can only be sampled for discrete values of τ (integer multiples of Δt, the sampling interval of the data). For a sequence of N discrete values and a given lag $\tau = a.\Delta t$ (for some integer $a < N$) the correlation can be sampled $(N - a)$ times, so for larger N a better estimate of the correlation function is obtained. These correlation functions for the 1996 event of Figure 3.35 are plotted in Figure 3.68.

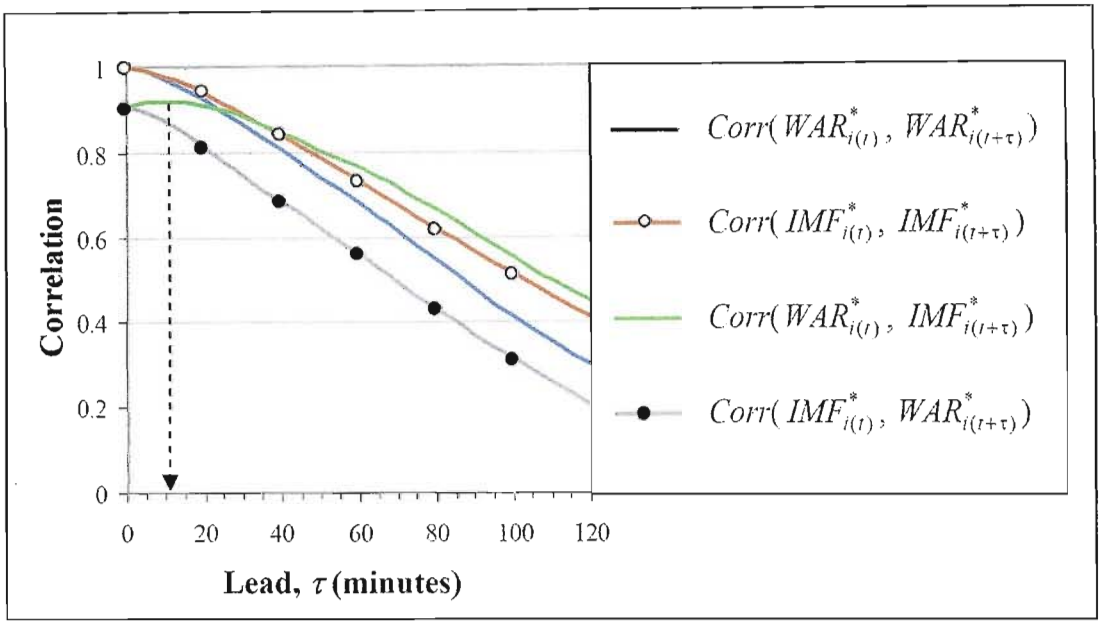

Figure 3.68 - Correlation of WAR_i* and IMF_i* time series (40 hour event of 1996)

Figure 3.68 is plotted for positive values of τ, negative values can be inferred from these. An interesting feature of Figure 3.68 is that the IMF_i appears to lead the WAR_i by approximately 11 minutes as indicated by the peak of the cross-correlation, marked by the arrow. This phenomenon can also be observed in Figure 3.35 and is evident in most of the events considered in this study so far. It suggests that the storm grows (or decays) in intensity before it expands (or contracts) in area – information which could be helpful in real time assessment of the growth or decay state of the storm.

Being relatively long, the 1996 event is a good candidate for measurement of the correlation functions, however similar functions can be extracted for any of the observed events. The correlation functions for twenty-one other events were measured using the same method and the results are presented in Figure 3.69.

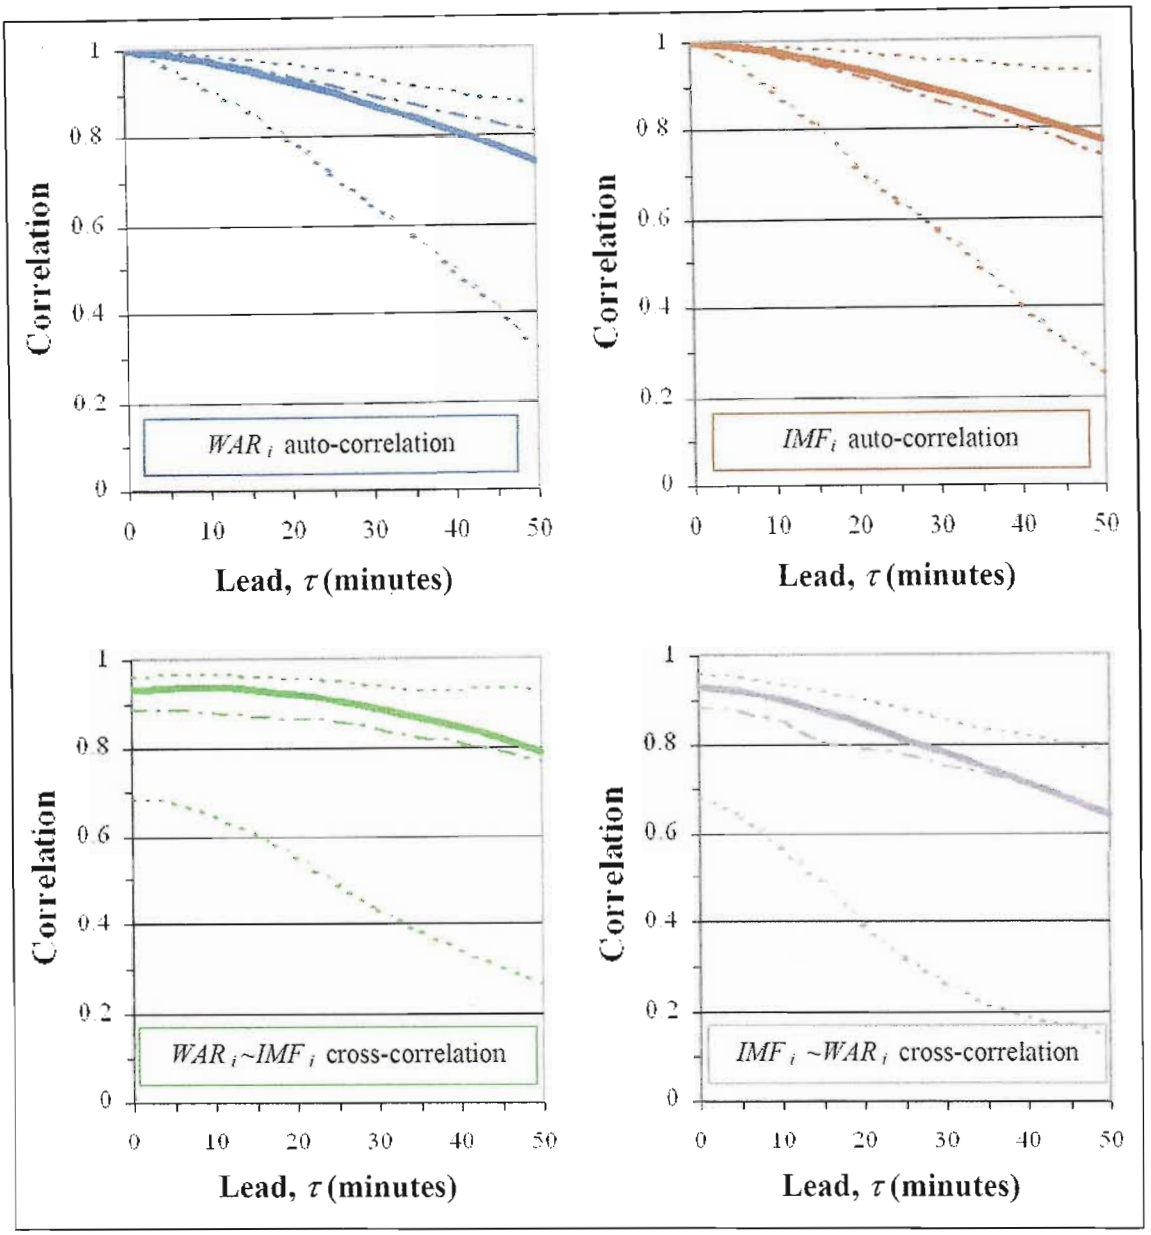

Figure 3.69 - Correlation envelopes for twenty-one analysed events

Analyses of twenty-one events are presented in Figure 3.69 as four correlation envelopes where the upper (dashed) line represents the maximum measured correlation for a particular lag, the lower (dashed) line represents the minimum, the middle (dot-dashed) line represents the median and the solid line represents the corresponding correlation measured for the 1996 event. In all four cases the correlation of the 1996 event is close to the median value indicating that it exhibits typical behaviour of an observed rainfall event in this respect.

3.11 *PIXEL SCALE TIME SERIES*

The analyses presented in Sections 3.11 and 3.12 were carried out at a time when only the 1996 integer precision data were available. To date they have not been repeated on the full precision dataset, instead the focus of the study shifted towards the long term (seasonal) temporal behaviour of rainfall and the inclusion of the raingauge dataset in the modelling process. Many of these findings have already been presented and published (Pegram and Clothier, 2001b).

This section is concerned with the time series of the rainfall intensity at a pixel in the field – the time series of the *Pixel Scale Intensity* (PSI). Explicit analysis of the PSI time series was attempted, but was fraught with difficulty and was ultimately abandoned in favour of an iterative calibration exercise whereby the structure of the time series could be inferred rather than measured directly from the data. The explicit methodology is described here for the sake of completeness. There are four major factors which complicated this time series analysis:

- Advecting field
- Masked (windowed) field
- Skew marginal distribution
- Integer precision data

The first of these problems is extracting the time series at a pixel from an advecting field. Measurement of the mean field advection by optimising the spatial cross-correlation between consecutive images was described in Section 3.6 and once the cumulative advection vector (Figure 3.41) has been measured for an event, the images can be translated into a Lagrangian reference frame by subtracting the cumulative advection vector. When viewed as an animated sequence in this Lagrangian reference frame, storms appear to remain stationary (whilst growing and decaying in the usual manner) and the ¾ doughnut mask appears to move with a velocity opposite to the cumulative advection vector. With the images aligned in this manner, runs of PSIs sampled at a point in the Lagrangian space can be extracted and analysed thereby removing the effect of the mean field advection. The only obstacle encountered while working in the Lagrangian reference frame is that the then moving mask limits the length of the runs which can be extracted, thereby making it difficult to measure the correlation for long lags.

Unfortunately this method cannot account for the fact that parts of the 128km field are likely to be advecting in slightly different directions at any point in time. Storm advection is driven at a whole range of scales and subject to many contributing factors and as a result the space-time rainfield is often quite turbulent in nature. The consequence of this phenomenon with regards to the time series analysis, is that the correlation is under-estimated. There is no easy way around this problem and for the purposes of this study this effect was neglected.

Analysis of the runs of PSIs extracted from the Lagrangian reference frame revealed another severely skewed marginal distribution. On its own, the skew marginal distribution can be addressed using a transform such as the Normalising transform discussed in Section 3.8 and used in Section 3.10.1 for the $\dot{W}AR_i$ and IMF_i time series. However, when combined with the fact that the data were only available in integer precision at the time of analysis, the correlation structure becomes very difficult to measure.

Figure 3.70 illustrates why this is the case, by presenting two time series with skewed marginal distributions similar to that observed for the PSI, the first continuous and the second an integerised version of the first (Figure 3.70, upper right). Applying the normalising transform of Section 3.8 yields an identical transform for the continuous and the integer time series (Figure 3.70, upper left). The problem arises when computing the correlation from the normalised, integer time series (Figure 3.70, lower left). Due to the skew nature of the time series, the transform is steep at the lower end of the range and consequently changes in the original time series (at the lower end of its range) are magnified considerably in the corresponding normalised time series. Changes in the integer time series are relatively large when compared to the changes in the continuous time series and the result is that the correlation as estimated from the normalised integer time series, is severely biased downwards as shown in Figure 3.70, lower right.

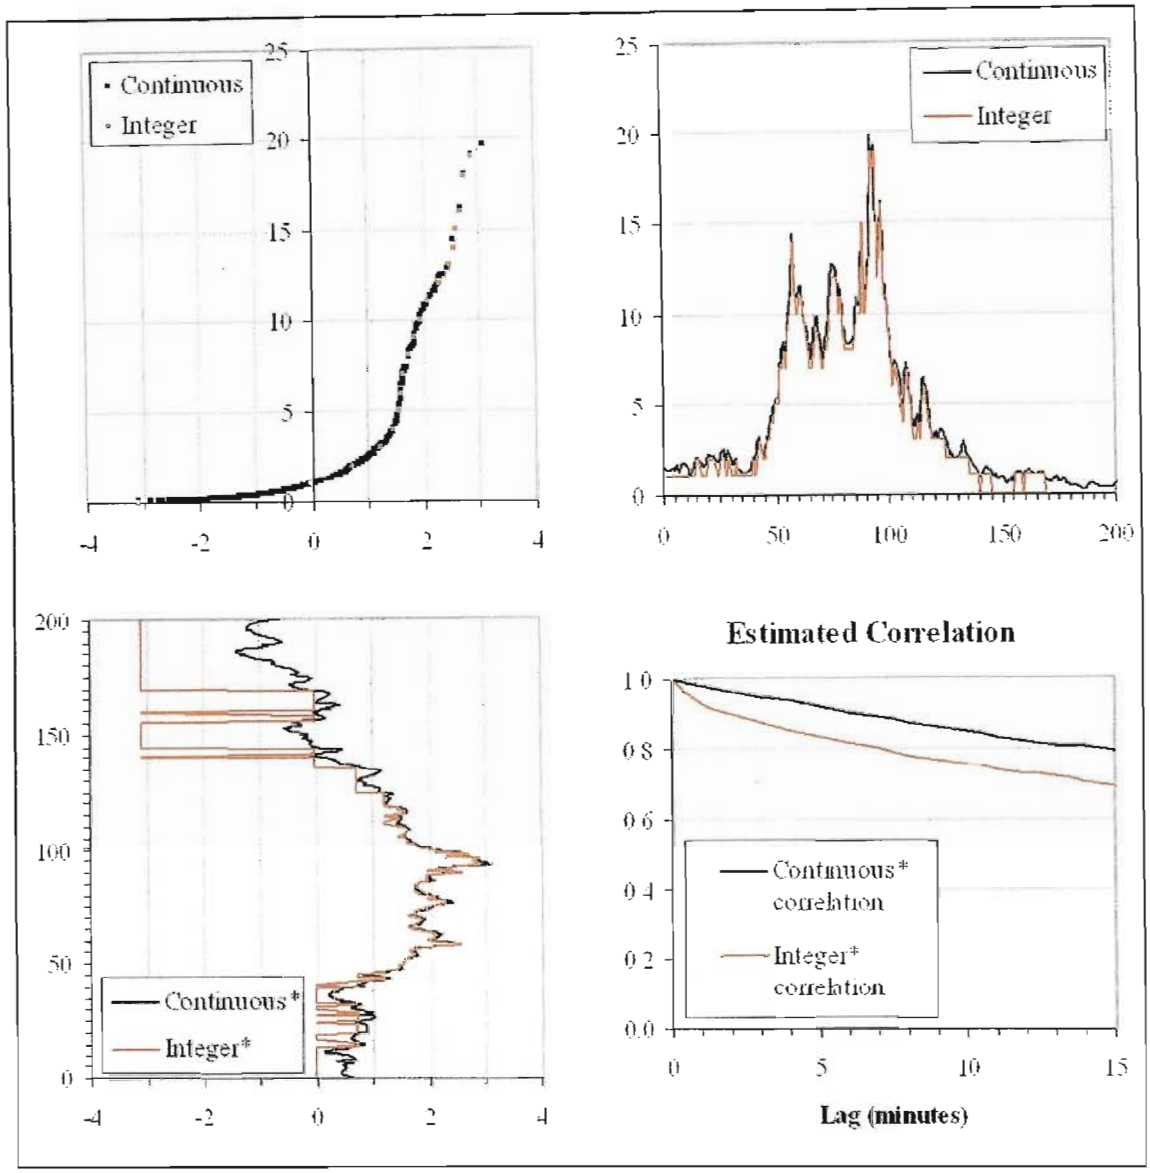

Figure 3.70 - Difficulty in estimating the correlation from normalised, integer data

The combined effect of these 4 factors meant that the correlation structure of the PSI was almost completely lost when attempting to apply conventional Gaussian time series analysis to the data available at the time. For the sake of completeness, Figure 3.71 shows the PSI auto-correlation function measured in Lagrangian space from normalised, integer precision data. With a lag of only five minutes (one image) the PSI correlation falls below 0,2 – a result which is contrary to qualitative observations of the data in which PSIs appear to have high correlation for short lags. This part of the analysis may be worth repeating now that full precision data are available, but this will not be attempted here, instead the PSI auto-correlation is inferred by iterative calibration and the description of this technique is deferred until Chapter 4.

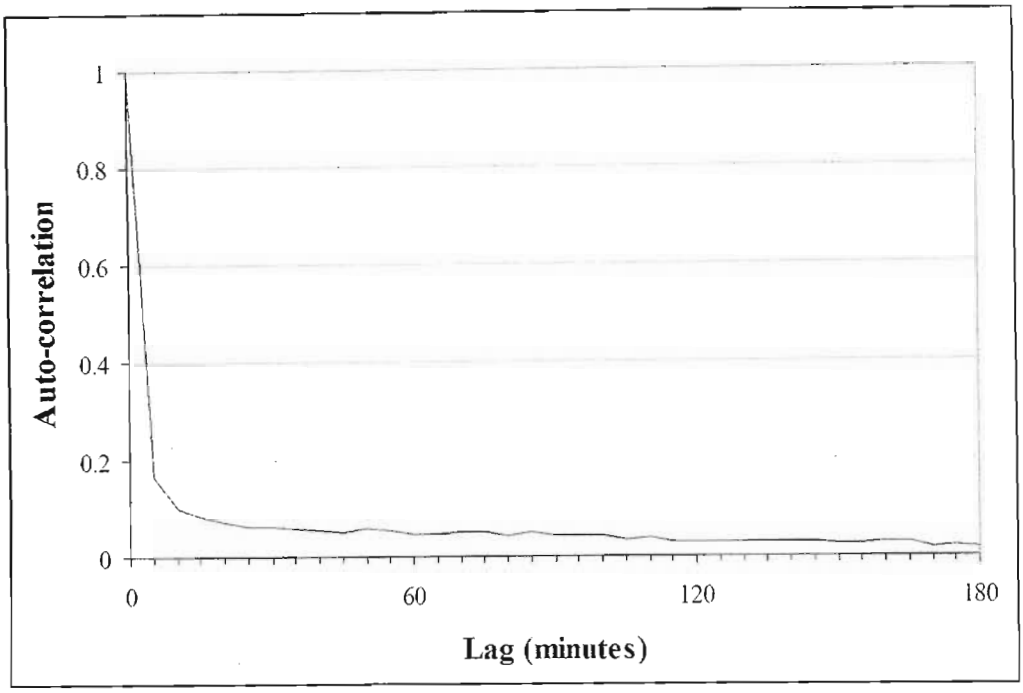

Figure 3.71 – Pixel Scale Intensity (PSI) auto-correlation function measured in Lagrangian space from normalised, integer precision data

3.12 SPACE-TIME UPSCALING OF RAINFIELDS

Data Upscaling in the context of the "String of Beads" model entails the statistical properties of the data when they are aggregated into larger space-time blocks. It is the opposite to downscaling. At the time of this analysis only the 1996 integer precision data stored in 128x128 km fields at 1 km, 5 minute resolution were available. Selecting the 40 hour event (512 image) of Figure 3.35, the first stage in the analysis was to average the rainfall rate in the individual images into 2, 4, 8, 16, 32, 64 and 128 km pixels. Figure 3.72 illustrates this process for a typical image.

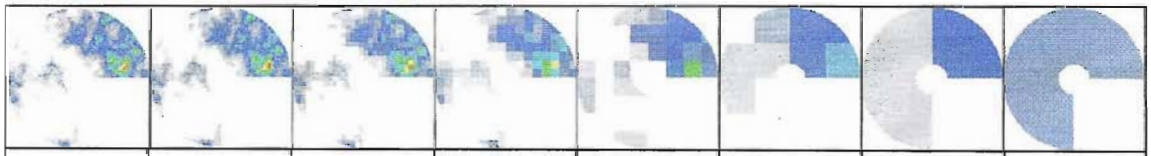

Figure 3.72 - Illustration of spatial upscaling for a typical 1 km 5 minute image

The images in Figure 3.72 are examples of 1km-5min, 2km-5min, ...,128km-5min space-time images and since they are all sampled from the 40 hour event at a temporal resolution of 5 minutes, they each represent one of a sequence of 512 images at their respective spatial resolutions. They are spatially up-scaled images.

Just as the individual images can be spatially up-scaled by averaging the PSI over larger spatial areas, it is possible to temporally up-scale the data by averaging the PSI over longer time periods. This was done for 15, 30, 60, 120, 240 and 480 minute intervals (averages of 3, 6, 12, 24, 48 and 64 consecutive images respectively) for each of the spatial scales illustrated in Figure 3.72. With only 512 images in the original sequence, the number of images in the temporally up-scaled sequences were 170 fifteen minute aggregations, 85 thirty minute, 42 sixty minute, 21 two hourly, 10 four hourly and 5 eight hourly aggregations for each spatial scale.

Two major aspects were measured from these sequences of space-time aggregations:

- The spatial correlation structure of each image in each sequence, using the radially averaged two-dimensional power spectrum
- The temporal correlation structure, using the spatially averaged, autocorrelation function of each sequence

It should be made clear at this stage that this part of the analysis was done with a view to validation of the "String of Beads" model and not to provide measurements which can be compared to other analyses. Estimation of power spectra from images which have only a few pixels such as the 8 and 16km examples of Figure 3.72 is an inaccurate process, but there are few other options - measurement of the spatial structure is limited by the resolution of the original data. Similarly in the temporal domain, estimation of autocorrelation functions from short sequences of aggregated images is also undesirable, but at the time of analysis the radar database was limited and there were very few long sequences to choose from. The idea behind this analysis was to measure these statistics using well known techniques and then to simulate a sequence of images with identical characteristics and apply these same analysis techniques to facilitate a fair comparison between the observed and the simulated data.

Computation of the radially averaged two-dimensional power spectrum was described in Section 3.5.2, however, with reference to Figure 3.72, as the pixel resolution is reduced with increasing spatial scale, estimation of the power spectrum becomes meaningless due to lack of pixels on the image. For this reason, power spectra were computed only for images with a spatial resolution of 1, 2, 4, 8 and 16km. The means and standard deviations of the spectral exponents for various temporal resolutions were computed and these are plotted in Figure 3.73.

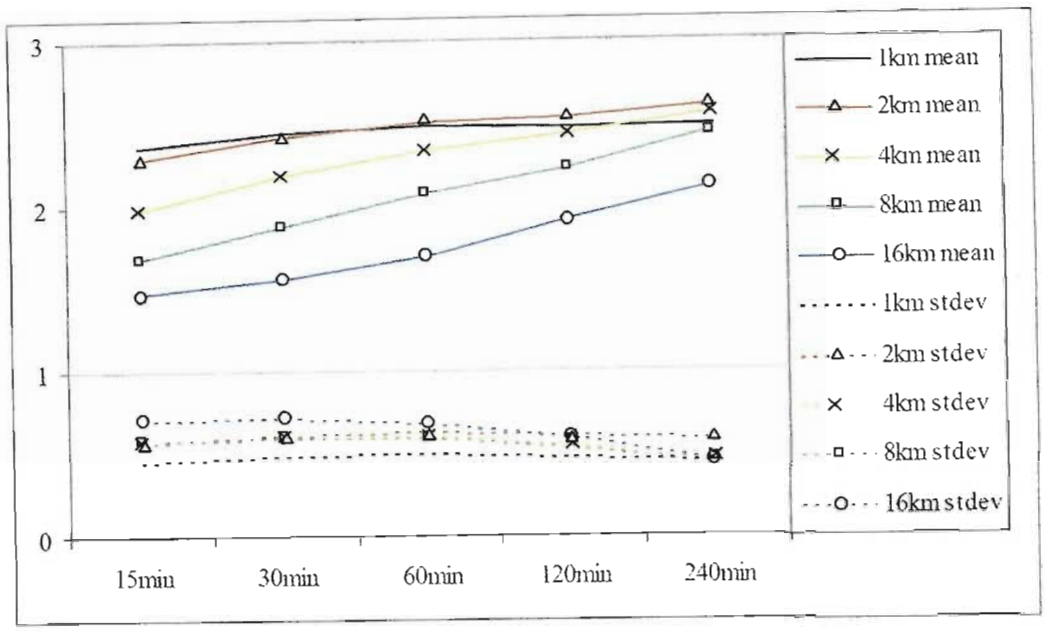

Figure 3.73 - Mean and standard deviation of spatial spectra exponents for a range of space-time aggregations

The main points of interest noted in Figure 3.73 are the increasing trend in mean of the spectral exponent with longer duration aggregations and the slightly decreasing nature of the standard deviation. Averaging data over long periods has a smoothing effect which results in a higher correlation and therefore a higher spectral exponent. The effects of the change in statistical parameters due to spatial aggregation has been extensively discussed in the field of geostatistics and is known as the support effect. A good explanation of this effect is given by Cressie (1991) who explains how the variance of the average of many values taken within an area (the support) is considerably reduced; by a factor of $^1/n$ if the field is uncorrelated and to a lesser extent in a correlated field.

Moving on to the temporal correlation structure, the correlation in time for each of the space-time aggregated sequences was computed at a each (spatially aggregated) pixel in space and then averaged over all other pixels in the space. So, for example, at 1km spatial resolution, there would be 9128 pixel sequences from which the autocorrelation function can be computed and these would then be averaged for each temporal resolution. At the other end of the scale, the 128km spatial resolution, the autocorrelation function can be computed only once for each temporal resolution. Figure 3.74 shows the autocorrelation functions for the 1, 8, 32 and 128 km spatial scales at 15, 60 and 240 minute temporal scales.

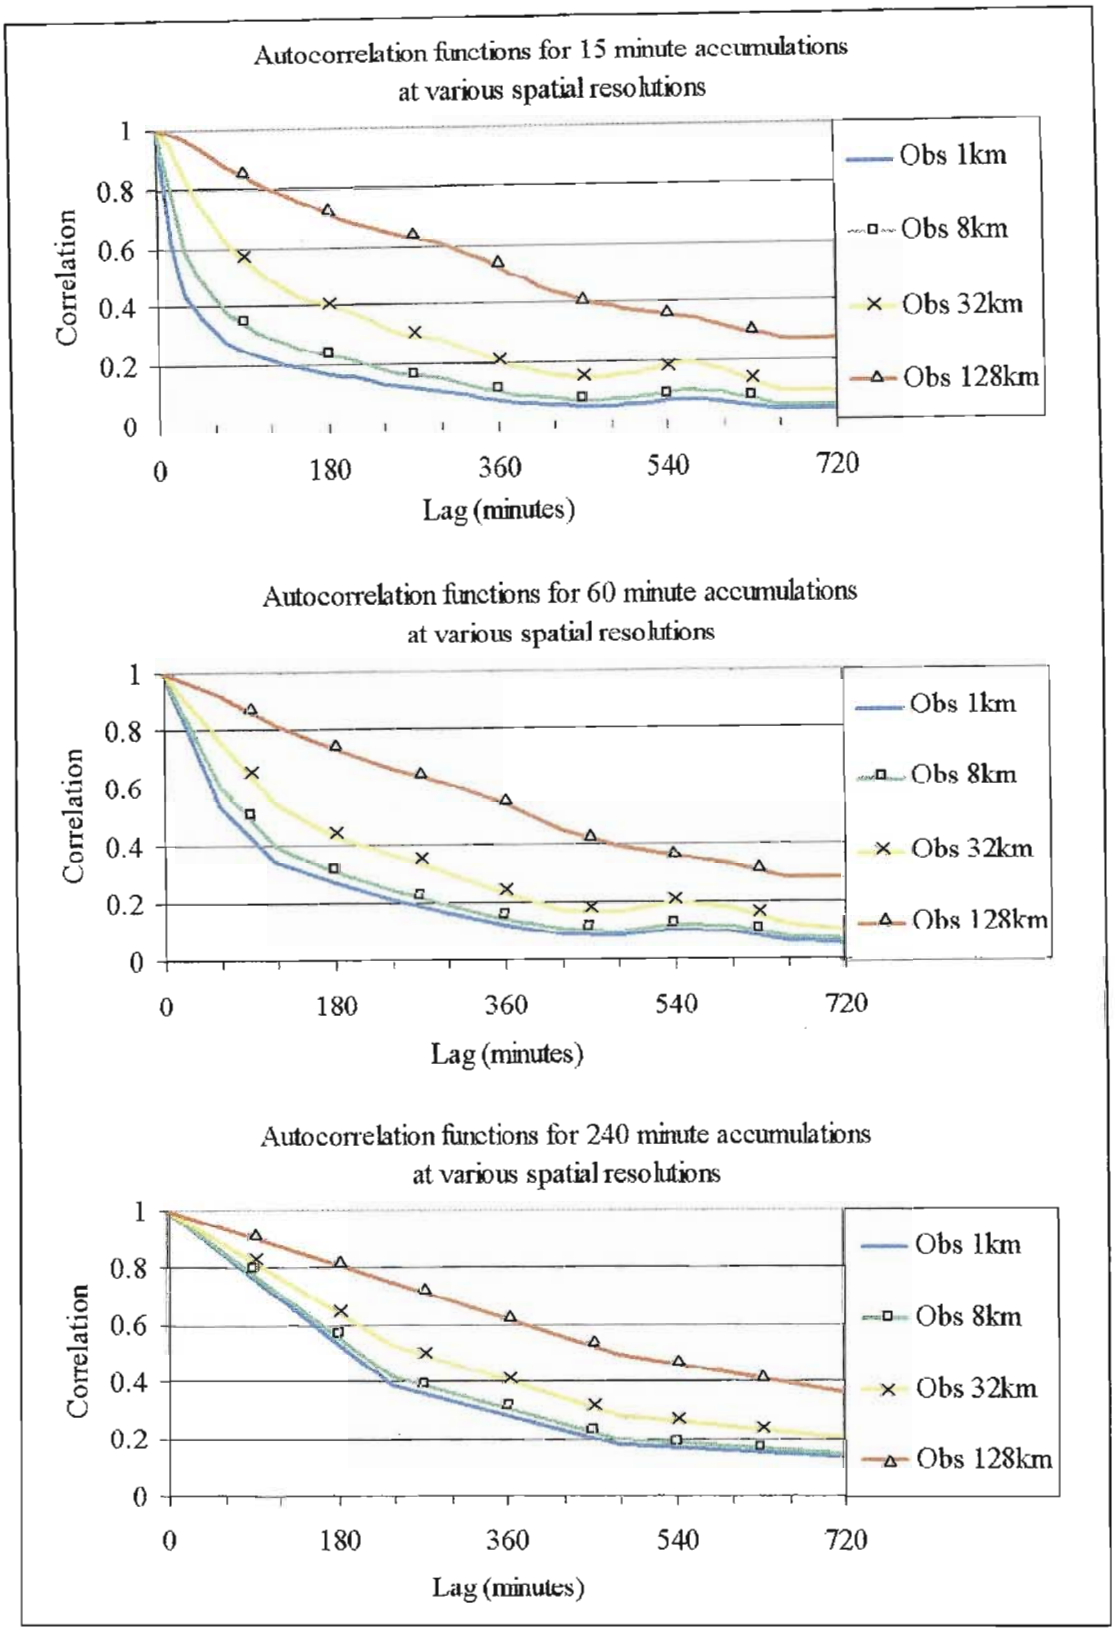

Figure 3.74 – Spatially averaged temporal correlation functions for the 1, 8, 32 and 128 km spatial scales at 15, 60 and 240 minute temporal scales for observed data

The functions plotted in Figure 3.74 exploit the assumption that the temporal correlation between consecutive pixels is independent of the position in the space-time event. In other words, provided that the pixels are in the same position relative to each other in space-time, the sampled correlation between them can be averaged over all feasible positions in the dataset. This procedure exploits assumed ergodicity.

Comparison between these spatially averaged temporal correlation functions reveal some reassuring trends. In all cases, larger spatial scales (larger pixels) yield temporal correlation functions with a longer temporal dependence. In other words, for a given lag at a given temporal scale, the temporal correlation increases with increasing spatial scale as would be expected. The large scale changes occur more slowly. Similarly for larger spatial scales, for a given lag at a given spatial scale, the temporal correlation increases with increasing temporal scale. In the three examples of Figure 3.74, with the limited dataset available at the time of analysis, the temporal correlation functions for the 15, 60 and 240 minute temporal scales were sampled from sequences of 170, 42 and 10 aggregated images respectively. Consequently, the error in estimation of the autocorrelation increases with decreasing temporal resolution since there are fewer opportunities from which to measure the data. Nevertheless, the correlation functions faithfully increase with decreasing temporal resolution.

3.13 SUMMARY

This chapter has served to demonstrate the theory and techniques behind the analyses performed on the datasets which were discussed in Chapter 2. Where it is impossible to present the variety observed in the 98500 images analysed in this study, typical examples have been given to illustrate the range in results which can be expected. The intention is really to communicate trends in the data rather than detailed numerical results. Analyses have been presented in an order convenient for the purposes of discussing the formulation of the model which will be addressed in Chapter 4. In many cases there are more sophisticated methods of analysis than were used here, these were sacrificed in the interests of clarity, simplicity and computational speed, the objective being to capture the general trends of the data reasonably well over all of the observed scales, deferring refinements for a time when the model is operational.

Chapter 4

The "String of Beads" Model Formulation

4.1 INTRODUCTION

This chapter serves to present the structure of the model and to outline the theory behind the modelling techniques used. The chapter is structured so that the conventional theory used in the development of the model is first presented, followed by a detailed description of the processes and additional theory which bind together to form the *String of Beads* model. Various operational modes of the model are introduced and computing considerations in this regard are discussed.

4.2 PSEUDO-RANDOM NUMBER GENERATION

The object of the *String of Beads model* is to generate a realistic pseudo-random rainfall scenario in space-time and consequently, one of the key ingredients in the model is a good *pseudo-random number generator*. Two algorithms are used in the *String of Beads* Model for this purpose; the first is *an efficient and portable pseudo-random number generator* (Wichmann and Hill, 1982) to generate uniformly distributed pseudo-random numbers; the second is to transform a sequence of uniformly distributed random numbers into a Gaussian white noise sequence using, *the percentage points of the normal distribution* (Beasley and Springer, 1985). Details of these algorithms and a quantitative measure of the quality of their output is given in Appendix B.

4.3 AUTO-REGRESSIVE PROCESSES

There are two main autoregressive (AR) processes used in the String of Beads model. The first is a univariate AR process applied at the pixel scale to simulate the temporal relationship between consecutive pixels in the rainfield (pertinent analyses presented in Section 3.11). The second is a bivariate AR process applied at the image scale to simulate the temporal relationship between the WAR_i and IMF_i image scale statistics (pertinent analyses presented in Section 3.10). The theory of univariate autoregressive processes can be found in any good text on time series modelling such as Box and

Jenkins (1970). However, slightly less common, is the theory of multivariate autoregressive processes and since this is an essential part of the String of Beads model, this theory presented by Wei (1990) is outlined in Appendix B. The optimal order of the autoregressive model can be chosen by minimising the Corrected Akaike Information Criterion (Hurvich and Tsai, 1989), also presented in Appendix B.

4.4 POWER-LAW FILTERING OF GAUSSIAN NOISE

The method of simulating sequences of images presented in this section are based on the works of Bell (1987), Schertzer and Lovejoy (1987), Brenier (1990) and Wilson et al. (1991). Bell (1987) used the technique of power-law filtering of a Gaussian noise field to simulate rainfields in three dimensional space-time (two space and one time) as viewed by satellite. This idea was used by Brenier (1990) and Wilson et al. (1991) to produce a sequence of images in time which were assembled into an animated image and published on video. Image sequence simulation in the String of Beads model is done using a modified version of their techniques and this will be illustrated first in the simple one-dimensional case and then extended into two and three dimensions.

4.4.1 One-dimensional power-law filtering

Three parameters WAR_i, IMF_i and β_{space} are required to describe the spatial distribution of rainfall on an image. β_{space} is the gradient of the radially averaged two-dimensional power spectrum. Any *pure* white noise process is uncorrelated and therefore has a β_{space} of zero. By power-law filtering a field of uncorrelated Gaussian variates it is possible to increase the correlation of the data from zero to any desired level specified by β_{space}. This is done by transforming the sequence into Fourier space via the Fast Fourier Transform, multiplying (term by term) the transformed sequence as a function of frequency by a filter function, and then back transforming into natural space. In the one-dimensional case, the power-law filter is defined in one-dimensional, *complex* valued Fourier space, as a distance function which decays exponentially (with exponent $\frac{1}{2}\beta_{space}$) from the first frequency to the Nyquist critical frequency and is then reflected about that point. The coefficients of a typical filter are plotted in Figure 4.1.

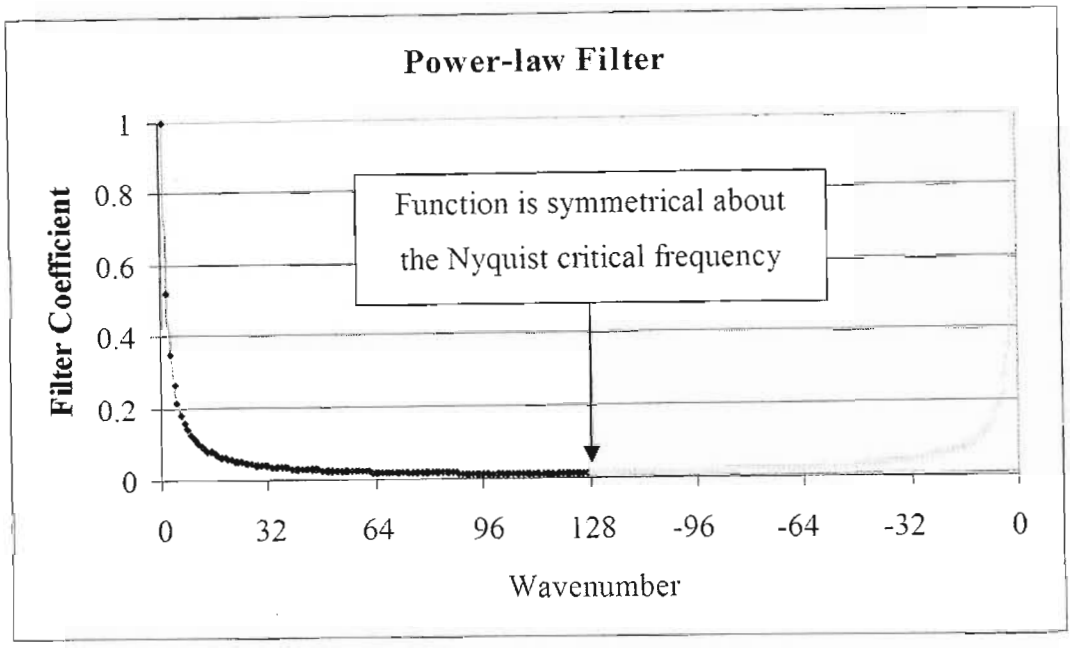

Figure 4.1 - Plot of coefficients for a typical power-law filter

To illustrate the power-law filtering process, consider the one-dimensional sequence of 256 standard, normally distributed, real valued random numbers plotted in Figure 4.2.

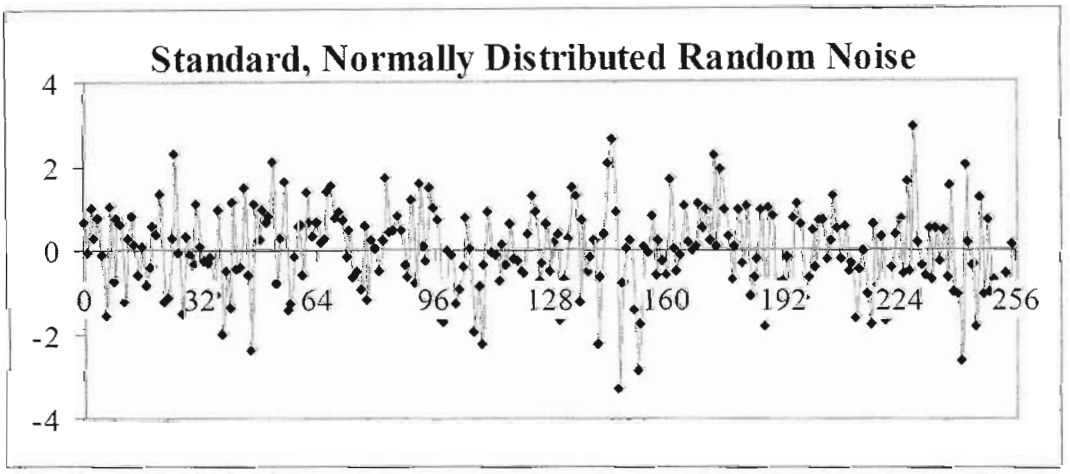

Figure 4.2 - Sequence of standard, normally distributed random numbers

This *real valued* sequence can be transformed into Fourier space via the FFT to produce a sequence of 256 *complex valued* numbers comprised of 128 pairs of random complex conjugates. Next, each *complex valued* member of the transformed sequence is multiplied by its corresponding *scalar* filter function coefficient at that point. The *real* and *imaginary* components of the product of the transformed sequence and its scalar multiplier, are plotted separately in Figure 4.3. The conjugate symmetry of this power-law filtered sequence is illustrated by the fact that the real component of the sequence is

symmetrical about the Nyquist frequency (128), whilst the imaginary component exhibits skew symmetry about the Nyquist frequency.

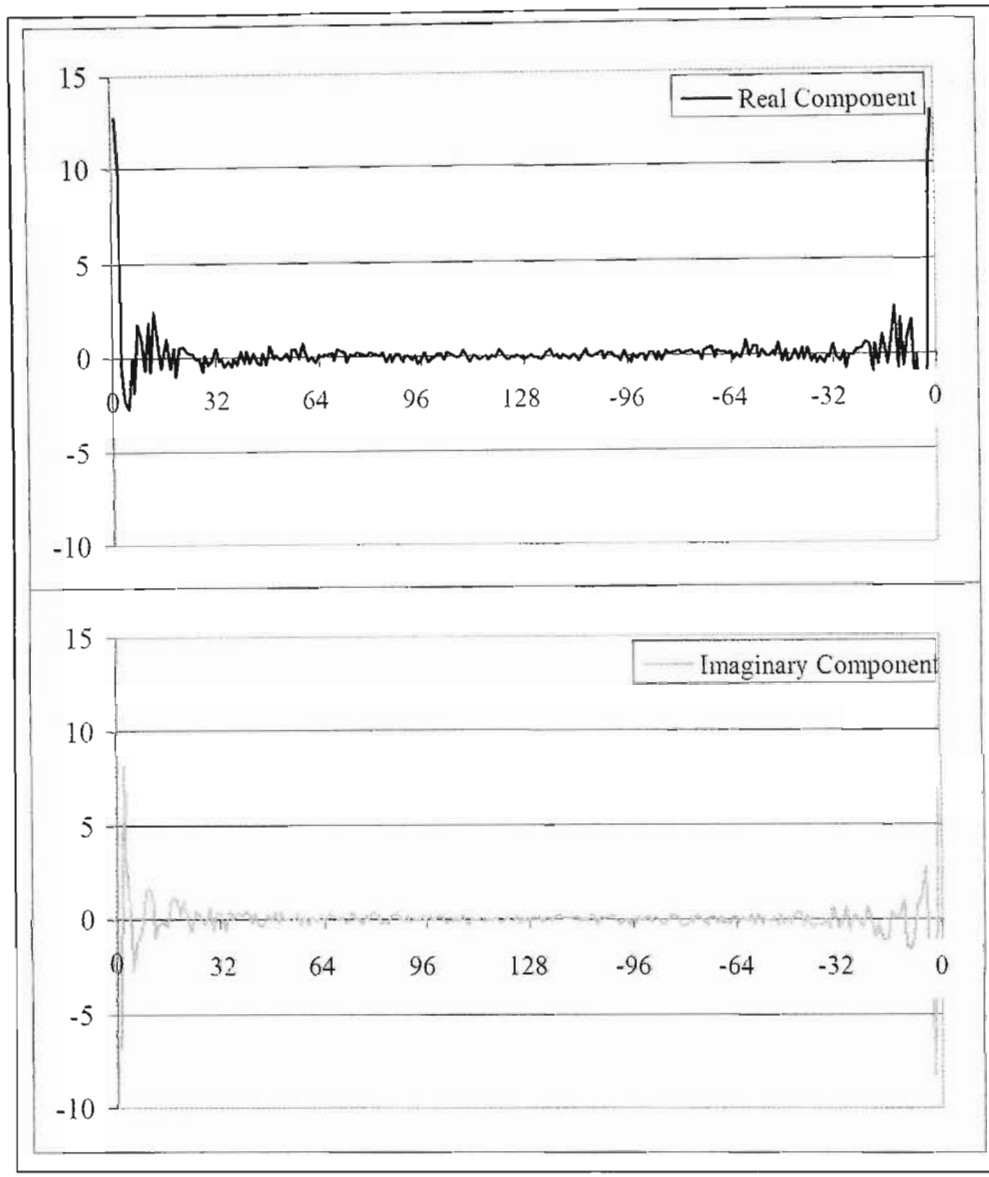

Figure 4.3 - Real and imaginary components of the product of the filter of Figure 4.1 and the complex valued, Fourier transformed sequence of random noise

Using the *Inverse Fast Fourier Transform*, the sequence of complex numbers in Fourier space shown in Figure 4.3 can then be back-transformed into natural space to yield the corresponding sequence of real valued, normally distributed random numbers which now have a correlation structure defined by β_{space}. The sequence can then be scaled and shifted to achieve the required marginal distribution with lognormal parameters μ and σ. For the purposes of this example, μ and σ values of 0 and 1 respectively were chosen and the resulting sequence is plotted in Figure 4.4.

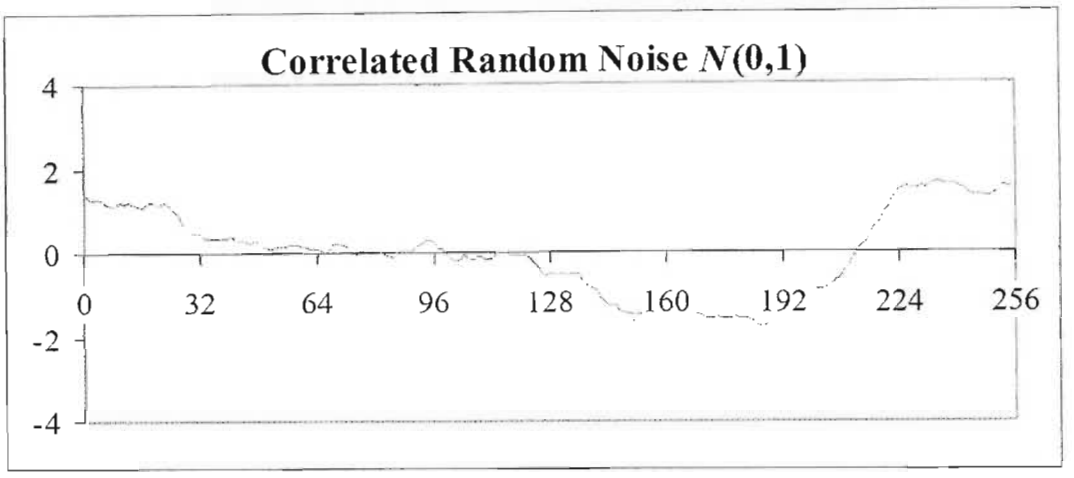

Figure 4.4 - Power-law filtered (Exponent β = 1.9) sequence of the standard, normally distributed random noise given in Figure 4.2

The sequence of numbers plotted in Figure 4.4 has the same marginal distribution as that of Figure 4.2, however the spatial correlation has been introduced by power-law filtering the sequence and thereby increasing the β_{space} parameter from 0 in Figure 4.2, to 1.9 in Figure 4.4. Since this sequence is normally distributed $N(\mu,\sigma)$, the conversion to a lognormal distribution $\Lambda(\mu,\sigma)$ is achieved by exponentiating each element in the sequence. The result is a *real* valued sequence of *positive* numbers, plotted in Figure 4.5, with marginal distribution $\Lambda(\mu,\sigma)$ and a spatial correlation defined by β_{space}.

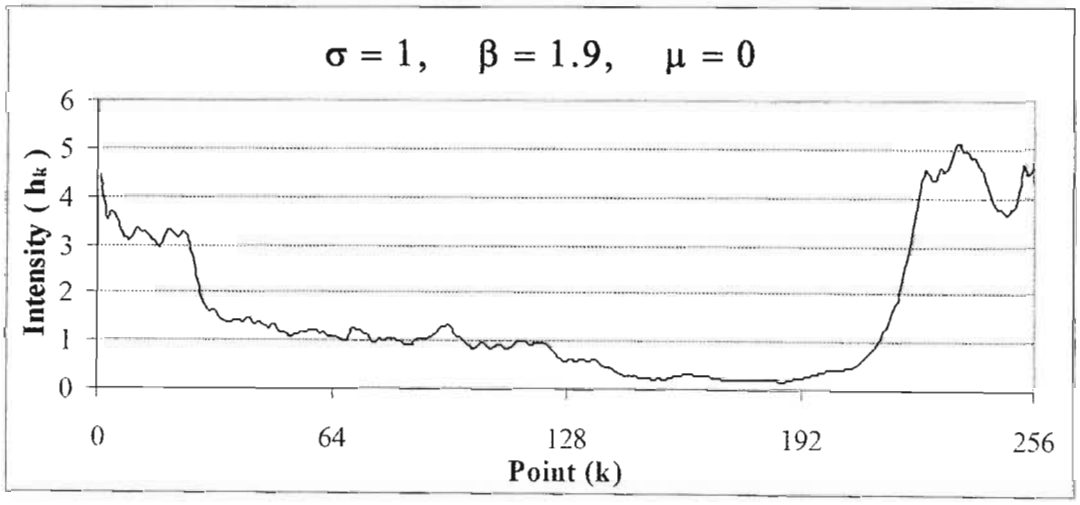

Figure 4.5 - Lognormally distributed, power-law filtered white noise

By varying the β_{space} parameter, it is possible to alter the spatial correlation in the filtered sequence. A higher exponent in the filter results in a smoother sequence. This is due to the fact that there is more power in the lower frequency Fourier coefficients and therefore a much higher correlation between any point in the sequence and its

immediate neighbours. Figure 4.6 is a plot of a sequence which was generated using the set of white noise plotted in Figure 4.2, but with a β_{space} parameter of 3.0 as opposed to the value of 1.9 used to produce the sequence plotted in Figure 4.5. The two sequences have been scaled and shifted so that they have identical marginal distributions.

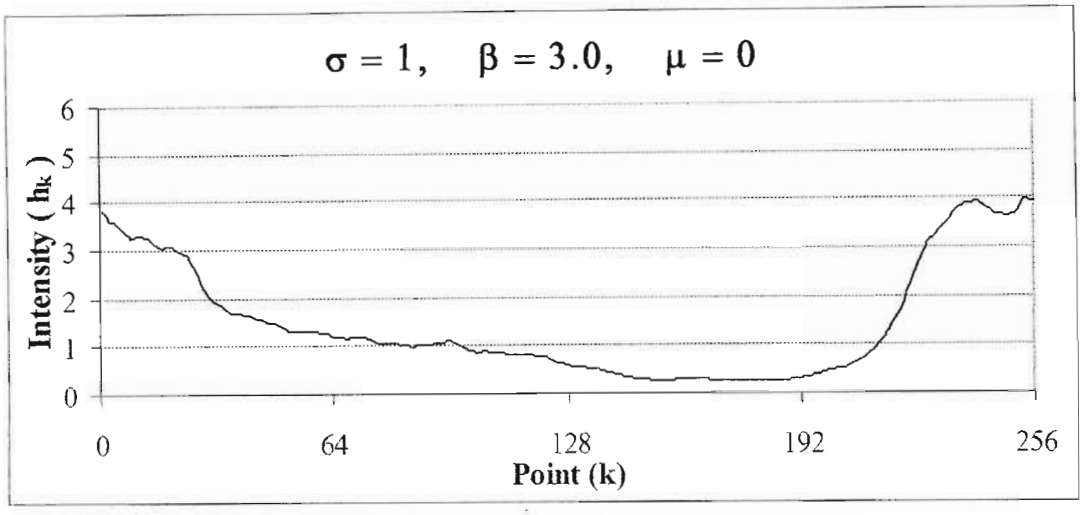

Figure 4.6 - Lognormally distributed, power-law filtered white noise, the greater β_{space} (=3,0) parameter produces a smoother sequence

Figure 4.5 and Figure 4.6 are plotted to the same scale and illustrate how the higher exponent power-law filter smoothes out the extreme points in the sequence. The periodic nature of the Fourier filtered sequence can be seen by the fact that both generated sequences finish with the intensity with which they began.

The sample power spectra of these two sequences are plotted in Figure 4.7 and Figure 4.8 and least squares approximations of their gradients have been calculated. The gradients of 1.928 and 3.028 are very close to the intended values of 1.900 and 3.000. The additional 0.028 is due to the fact that the original noise process is not absolutely pure and prior to any filtering had a spectral gradient of 0.028.

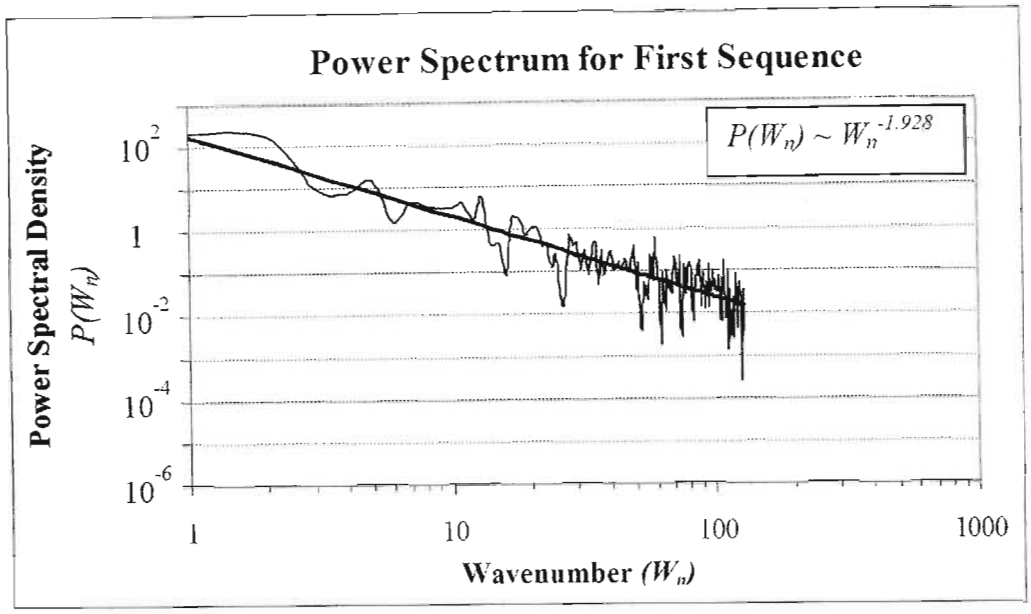

Figure 4.7 - Power spectrum for sequence plotted in Figure 4.5

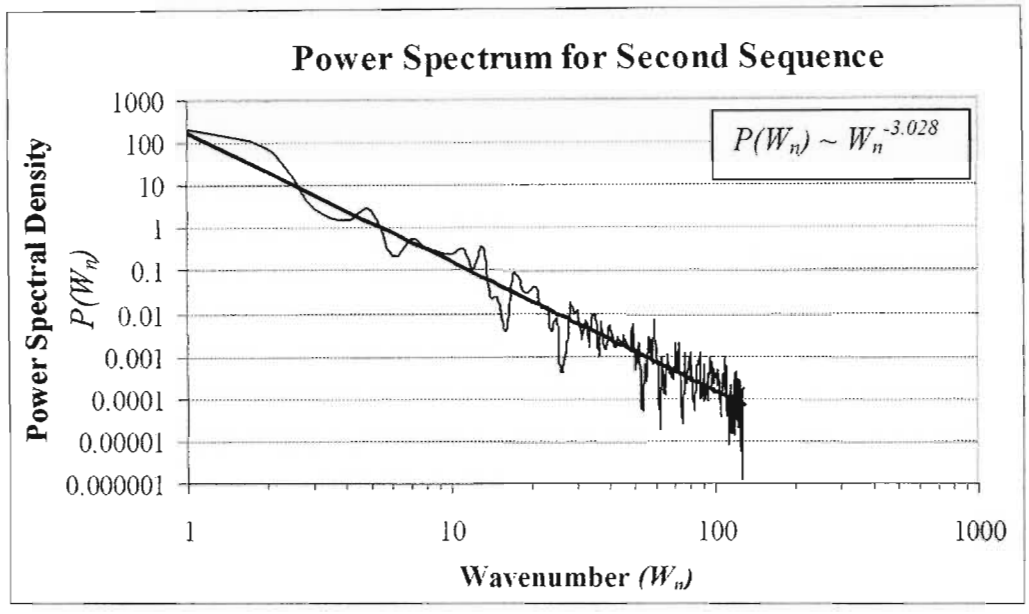

Figure 4.8 - Power spectrum for sequence plotted in Figure 4.6

All of the properties and methods presented here for the one-dimensional case, readily extend into two space and beyond, however it becomes increasingly difficult to visualise the shape of the power spectrum, filter function and random field as the dimension of the problem increases.

4.4.2 *Two dimensional power-law filtering*

Following the format adopted for the one dimensional case in Section 4.4.1, consider two two-dimensional fields generated using the same technique presented for the one-dimensional case, only using the two-dimensional Fast Fourier Transform on a two-dimensional field. This is the universal multifractal technique of Schertzer and Lovejoy (1987). With reference to Figure 4.9, the field on the left was generated using a β_{space} exponent of 2.0 in both x and y directions and the field on the right was generated using a β_{space} exponent of 3.0 in both directions. The two fields were generated using the same set of white noise and have the same marginal distribution.

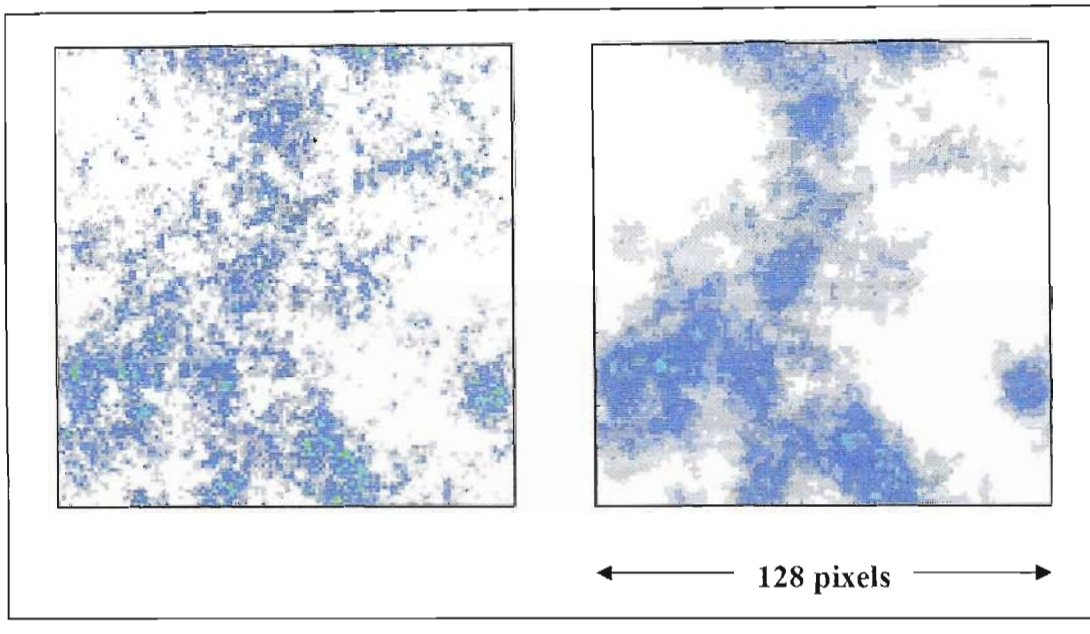

Figure 4.9 - Comparison of two two-dimensional power-law filtered fields

As was observed in the one-dimensional case, the higher β_{space} exponent in the filter produces a smoother image with a lower maximum intensity. The Fourier periodicity is evident in both the x and y directions of both images. The radially averaged two-dimensional power spectra of these two fields are plotted in Figure 4.10 and Figure 4.11 and least squares approximations of their gradients have been calculated. The gradients of 2.049 and 2.939 are very close to the intended values of 2.000 and 3.000. The difference between the calculated values and the intended values are again due to the fact that the original noise process is not exactly uncorrelated.

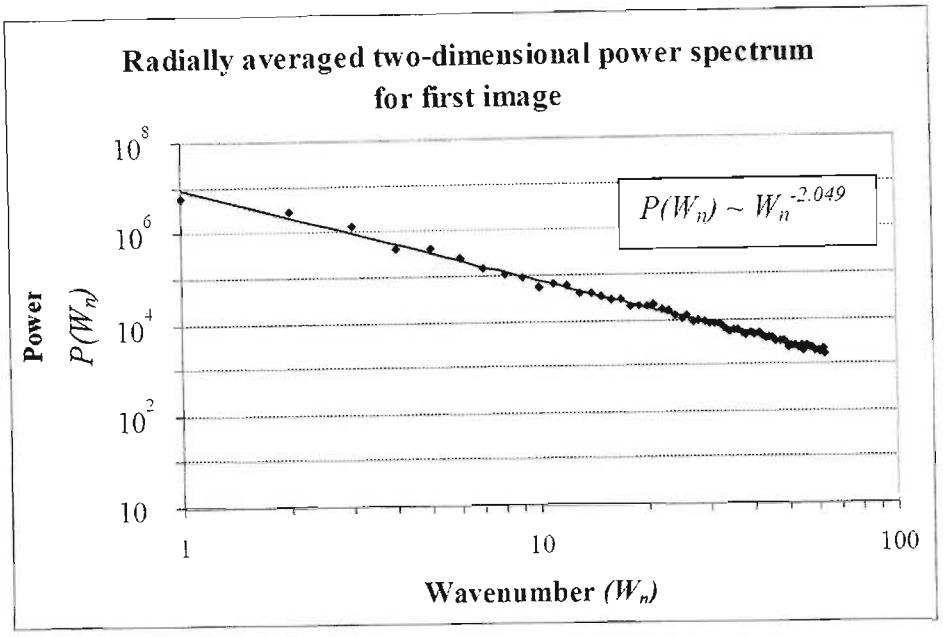

Figure 4.10 - Averaged power spectrum for left hand image of Figure 4.9

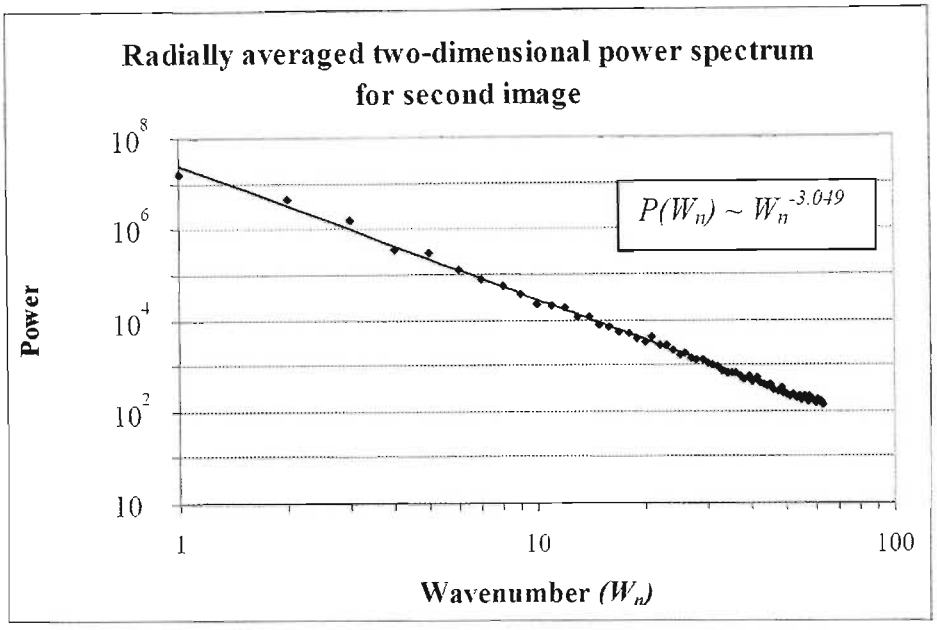

Figure 4.11 - Averaged power spectrum for right hand image of Figure 4.9

The two spectra shown in Figure 4.10 and Figure 4.11 are plotted up to a wavenumber of 64 since the original images were a size of 128 x 128 pixels. Due to the averaging processes about the Nyquist frequency *and* about the apex, the power spectra are of a much smoother nature than those seen in the one dimensional cases of Section 4.4.1, however, the basic properties observed in the one-dimensional case are preserved in the two-dimensional case and beyond. That is, the simulated fields are scaling, non-stationary random fields, with a lognormal distribution $\Lambda(\mu, \sigma)$ and a correlation structure defined by a single parameter β_{space}.

4.5 SBM STRUCTURE IN SIMULATION

There are three main stages of rainfall simulation in the *String of Beads* model and these correspond to the *event scale*, the *image scale* and the *pixel scale* statistics presented in Chapter 3. The *event scale* simulation concerns the *event arrival* and *event duration* as well as the *event advection vector*. The *image scale* simulation concerns the one-dimensional time series of the WAR_i, IMF_i and β_{space} image scale statistics. The *pixel scale* simulation concerns the spatial distribution of rainfall on the simulated images. Each of these stages will be discussed in turn in this section.

4.5.1 SBM First stage – event scale simulation

Two independent sets of the one-dimensional event scale statistics are pseudo-randomly generated in this first stage of the *String of Beads* model, the event cumulative advection vector and the event arrival, duration and intensity statistics. Once calibrated for a particular region, model input at this stage is simply the start and finish dates of the required simulation, together with three seeds for the random number generator which ensure that the simulation can be repeated if necessary.

Event arrival, duration and intensity

The analysis of Section 3.9.2 revealed that there is no significant dependence structure in the durations of consecutive wet and dry spells. Consequently, wet and dry spells can be sampled alternately from mutually independent processes. With only a short record of high temporal resolution data, the seasonal variation in the wet-dry process must somehow be inferred by making use of the only long term dataset available – the daily raingauge dataset.

Five main attempts were made in the search for a model which described the wet/dry process by making use of the daily raingauge dataset and these were

- Three state daily Markov chain
- Alternating renewal process fitted by direct measurement
- Alternating renewal process fitted by daily runlength optimisation
- Markovian alternating renewal process inferred from daily transition probabilities
- The scattered event alternating renewal process

The first four of these methods proved unsuccessful for a variety of reasons and were ultimately abandoned. In order to avoid confusion, the details of these attempts are included in Appendix A, Section 8.4 and will not be discussed here. The fifth and final attempt yields promising results in the tests conducted thus far and will be described in the following paragraphs. It introduces the idea of an *expected event* which is essentially a single event that represents an average of the 405 observed radar events. The process is modelled as a pure (exponential) alternating renewal process, but the generation of the simulated sequences is novel, hence the qualifier "scattered".

The results of the analyses in presented in Chapter 3 provide the key ingredients in the formulation of a model for the wet-dry process. In particular, we have:

- 50 years of daily raingauge data from a network of 54 gauges distributed over the catchment. These were analysed in Section 3.3.2 to obtain the monthly mean rainfall, IMF_m, illustrated in Figure 3.7.
- Marginal distribution of the wet spell durations (D_w) for 405 rainfall events observed by the MRL5 radar at Bethlehem during the 1998 and 1999 rainfall seasons, conditional on the 1% WAR_i threshold. This is presented in Figure 3.50.
- Marginal distribution of IMF_e, the total rainfall for the event, averaged over each pixel in the image, for the same 405 rainfall events. In other words, the spatially averaged depth of rainfall observed in each of the 405 events.

Harberlandt (1998) found that the lognormal distribution and gamma distributions were well suited to describing the distribution of wet spell durations in Germany and the analyses of Section 3.9.1 revealed that this is also true of the dry spell durations in South Africa. Figure 3.52 also reveals that there is very little difference between the fitted gamma and lognormal approximations of these distributions. For the sake of simplicity, the lognormal is chosen for the SBM. That is:

$$D_w \sim \Lambda(\mu_w, \sigma_w)$$

where

D_w is the duration of the wet spells

$\Lambda(\mu, \sigma)$ represents the lognormal distribution

μ_w represents the mean of the logs of D_w

σ_w represents the standard deviation of the logs of D_w

If it is assumed, as a first approximation, that the nature of a rainfall event is independent of the season in which it occurs, in other words that the marginal distributions of D_w and IMF_e are independent of season (justified by the values of E[D_w] appearing in Table 8.1), it stands to reason that the expected number of rainfall events in any month, is given by

$$E[Ne_m] = \frac{E[IMF_m]}{E[IMF_e]}, \qquad \text{where } January \le m \le December \qquad (4.1)$$

That is, the expected number of rainfall events in any month is given by the quotient of the expected depth of rainfall in that month and the expected depth of spatially averaged rainfall in any randomly selected event. This simple, but powerful result provides the link between daily raingauge data and radar data.

Extending the idea slightly, by considering the record of monthly rainfall depths over a Y-year period, IMF_{m_j} (where $January \le m \le December$ and $1 \le j \le Y$), an estimate of the number of events which occurred in each month, Ne_{m_j}, can be made using the relationship

$$Ne_{m_j} = \frac{IMF_{m_j}}{E[IMF_e]} \qquad (4.2)$$

Me_m and Se_m, introduced as estimates of the mean and standard deviation of the number of events in any month, can be computed as

$$E[Ne_m] \approx Me_m = \frac{1}{Y}\sum_j Ne_{m_j} \qquad (4.3)$$

and

$$Se_m = \sqrt{\frac{1}{Y}\sum_j \left(Ne_{m_j} - Me_m\right)^2} \qquad (4.4)$$

where: $January \le m \le December$ and $1 \le j \le Y$

Me_m and Se_m become parameters of the SBM, which are used to generate a normally distributed random number of events $Ne_m \sim N(Me_m, Se_m)$ for any month m. For simulation purposes, Ne_m is rounded to the nearest integer value.

The hypothesis behind this approach is that extreme rainfall records in any month are the result of extreme natured rainfall events rather than an extreme number or rainfall events in the month. This was observed by Pegram and Seed (1998) who found that the number of rainfall events did not vary significantly between years but that wet years experienced more wide-spread events of greater depth.

Thus far, two models have been discussed as components of the wet-dry process, one for the marginal distribution of the wet spell duration for a randomly selected event assumed independent of season $D_w \sim \Lambda(\mu_w, \sigma_w)$, and the other for the marginal distribution of the number of events in any month $Ne_m \sim N(Me_m, Se_m)$. In order to complete the description of the wet-dry process, a third component is required to describe the event arrival within a month.

Following the ideas of Cowpertwait (1994), amongst many others, an exponentially distributed event arrival time is assumed for this purpose. That is, the time measured from the start of an event to the start of the following event in the series is an exponential variate.

It can be shown that for a set of uniformly distributed numbers, X_i, arranged in rank order, the distribution of the differences between consecutive numbers in the set $\Delta_{Xi} = (X_{i+1} - X_i)$ is exponential. This useful result is exploited in order to assemble the String process.

Figure 4.12 illustrates the method by which the wet-dry process is constructed for a typical month in the simulation, given that the number of events (Ne_m) has been determined by the methods discussed in previous paragraphs of this section. In the example of Figure 4.12, Ne_m is taken to be 8 events in the month.

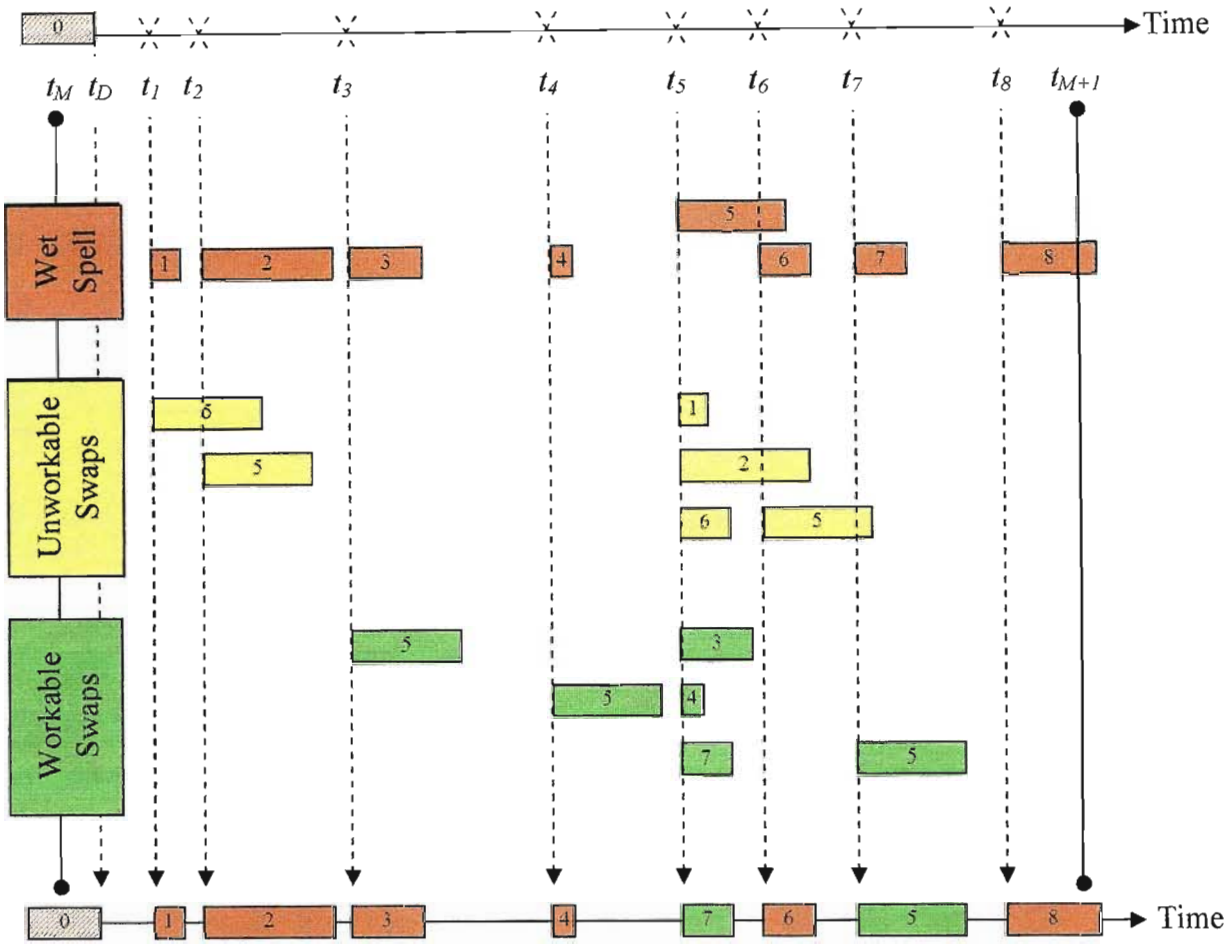

Figure 4.12 - Assembly of the wet-dry process. The scattered event renewal process

The first step is to compute Ne_m number uniformly distributed numbers X_i between the t_M , the time (in minutes from the start of the simulation) at the start of the month, and t_{M+1}, the time at the start of the following month. If the last event of the previous month extends into the current month, the range of X_i is limited to $t_D \leq X_i \leq t_{M+1}$, where t_D represents the end time of the last event of the previous month. The X_i are then arranged in rank order along the timeline as event arrival times which will have an exponentially distributed arrival time. That is $(t_{n+1} - t_n)$ is an exponential variate.

Next, Ne_m number wet spell durations is computed, $D_w \sim \Lambda(\mu_w, \sigma_w)$ by transformation of a uniformly distributed random number through the appropriate lognormal cumulative distribution function. These events are placed on the event arrival timeline as shown in red in Figure 4.12 between t_D and t_{M+1} with each event starting at its corresponding event arrival time.

The integrity of the wet-dry process is then tested, to see if there are any occurrences of event overlap. An overlapping event is defined as an event which finishes after the start of the next event. Event 5 in Figure 4.12 is an example of an overlapping event. Correction of this overlap is achieved by identifying all of the events in the set which are suitable candidates to be swapped with the overlapping event and will not result in an overlapping event. Events coloured yellow shown pair-wise at the same level in Figure 4.12 are unworkable swaps with Event 5 because they would result in additional other overlapping events, while pairs of events coloured green represent workable swaps and would eliminate the event overlap in this case. Once the workable swap candidates have been identified, one suitable candidate is randomly selected from the set. In the example above, event 7 was chosen, although events 3 and 4 are equally suitable. The last event in the sequence is not considered amongst the swappable candidates and is allowed overlap into the following month.

The only difficulty with this algorithm occurs when no workable swaps can be found, due to two of the event arrivals being too close together. In this case, the arrival timeline is simply re-sampled and the same events are distributed on the new arrival timeline. After all overlaps in the month are eliminated, the wet-dry process is fixed for that particular month and the process is repeated for the following month until the end of the simulation is exceeded.

The inter event arrival times are exponential and the event durations (D_w) are lognormal, however the marginal distribution of D_w is such that an exponential would also be a reasonable approximation of the distribution. The distribution of the dry spell durations D_d is accounted for without the need for measurement from the time series data. Based on the literature reviewed, this appears to be a novel approach to modelling the wet-dry binary process.

Model output at this point in the first stage is in the form of a binary process of alternating wet and dry periods – the *String* – which describes the event arrival and event duration process. From this, it is possible to infer pseudo-random statistics pertaining to the average intensity of the events, namely the M_{war*}, S_{war*}, M_{imf*} and S_{imf*} statistics for each event, by making use of the results of the analyses presented in Section 3.10.2. Figure 4.13 illustrates the algorithm for simulating these statistics.

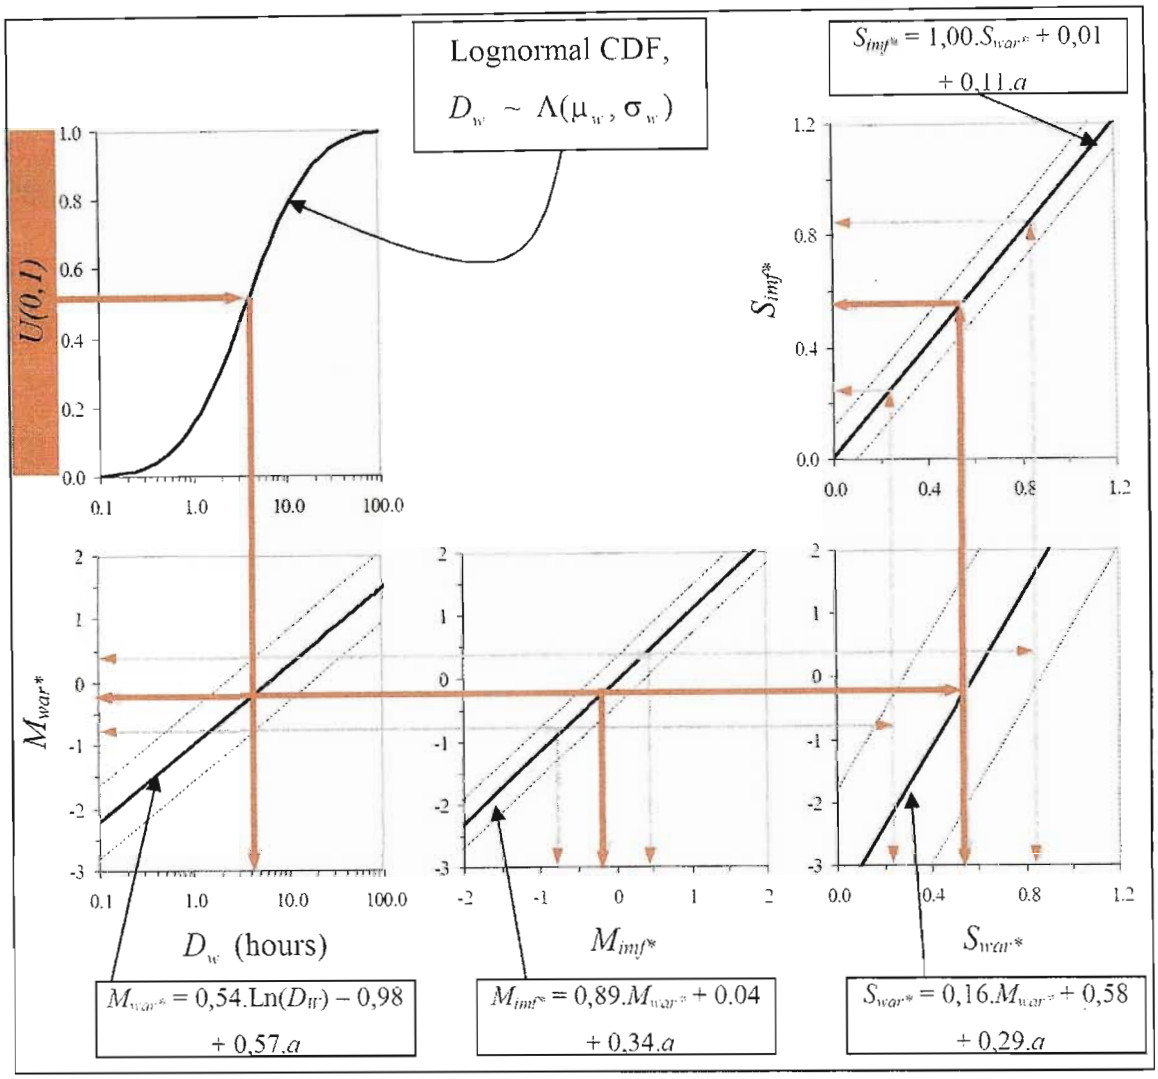

Figure 4.13 - Algorithm used to generate event scale, wet spell duration and intensity statistics, showing cumulative bands of uncertainty

The technique of independently modelling the event arrival, duration and intensity is similar in many ways to the Neymann-Scott Shot noise model described by Cowpertwait (1994).

The arrival date and time of a wet spell of duration, $D_w \sim \Lambda(\mu_w, \sigma_w)$ in hours, is announced by the wet-dry process. From this it is possible to compute an appropriate M_{war*} from the relationship given in Equation 4.5.

$$M_{war*} = 0{,}54.\mathrm{Ln}(D_w) - 0{,}98 + 0{,}57.a \qquad (4.5)$$

This relationship was devised from the analysis of D_w which was presented in Figure 3.60 and Figure 3.61. Figure 3.60 presents the regression relationship between M_{war*} and D_w while Figure 3.61 presents the residual values of this regression. These are

shown to be reasonably normal in distribution, with zero mean as expected from the regression and a standard deviation of 0,57. The a term in Equation 4.5 is a standard normal variate.

Dependent on the M_{war*} the terms M_{imf*}, S_{war*} and indirectly S_{imf*} are simulated via the Equations 4.6 , 4.7 and 4.8 respectively.

$$M_{imf*} = 0{,}89.M_{war*} + 0.04 + 0{,}34.a \quad (4.6)$$

$$S_{war*} = 0{,}16.M_{war*} + 0{,}58 + 0{,}29.a \quad (4.7)$$

$$S_{imf*} = 1{,}00.S_{war*} + 0{,}01 + 0{,}11.a \quad (4.8)$$

Again, the a terms in these equations are independent standard normal variates. These three relationships were extracted from the analyses presented in Figure 3.62 through Figure 3.67.

The bands of uncertainty (indicated by a single standard deviation) are represented in Figure 4.13 as dotted lines on either side of these relationships. The arrows illustrate the progress of the algorithm for a generated, pseudo random D_w, as well as the corresponding likely range of pseudo-random outcomes of the M_{war*}, M_{imf*}, S_{war*} and S_{imf*} terms, based on their respective (cumulative) bands of uncertainty which are computed by combining the consecutive shock terms of the dependent processes.

Extreme events

It is appropriate at this point to mention the possible inclusion of extreme rainfall events in the model, an important consideration in the context of design. Extreme events are either as a result of unusual duration or unusual intensity (or a combination of the two). Generation of event duration and intensity by the methods described above will result in a distribution of rainfall events similar in nature to that of the observed events on which the model is calibrated. Since radar rainfall records have only really become available relatively recently when compared to other methods, it is likely that other datasets (such as telemetering raingauge datasets will need to be examined in order to identify the distributions of duration and intensity and the frequency of occurrence of extreme events. Once these extreme events have been identified and their event scale statistics have been measured, inclusion into the modelling process is relatively simple a slight modification of the scattered event renewal process described above.

Similar relationships to those presented in equations 4.5 through 4.8 for the normal case would need to be determined for the extreme case and a suitable model for the distribution of the wet spell durations of extreme events would need to be determined. Once the number of events (Ne_m) has been determined in the scattered event renewal process, and knowing the probability of occurrence of an extreme event in any month, a random switch could be used to select either the extreme event distribution or the normal event distribution from which to calculate the event duration, D_w and the intensity statistics M_{war*} , M_{imf*}, S_{war*} and S_{imf*}. With this switch in place, simulation would continue as normal with the occasional occurrence of an extreme event in the string process. This is a refinement to the model which is left as a topic for future research.

Event advection

Advection of the rainfall images is a very important component of the spatio-temporal modelling process. Experience has shown that it is crucial to model this correctly in order to achieve realistic rainfall accumulations. With only one complete year of observed radar data comprising 140 analysed events of unbroken data, it is not possible to extract the seasonal marginal distribution of event speed and direction to any useful degree of accuracy, particularly in the drier months when there are very few rainfall events. Instead, it has been shown that the 500HPa wind speed and direction is usually a good indicator of storm speed and direction (Steyn, PCL and RT Bruintjes, 1990). In its current state, the model does not account for the fact that there may be some relationship between the duration of an event and its advection vector. This (perhaps naïve) assumption can be refined as more radar data are made available for analysis.

Acknowledging these limitations, simulation of the *event advection vector* is achieved via a simple non-parametric approach. Analyses of the 500HPa wind direction and speed were presented in Figure 3.43 and Figure 3.44 respectively. The distribution of 500HPa wind direction is shown in Figure 3.43 as a cumulative probability distribution function with directions binned into sectors according to the 16 point compass, beginning at north, and increasing (clockwise) through east, south and west and finishing at north-north-west.

Although presented in Figure 3.43 for the combined 20 years of daily data, similar curves can be extracted on a monthly basis to account for monthly variation in the distribution of wind speed and direction. Figure 4.14 illustrates the algorithm used to generate a Pseudo-random cumulative advection vector for an event. This vector is comprised of two components, the event advection direction, B_{adv}, and the event advection speed, C_{adv}.

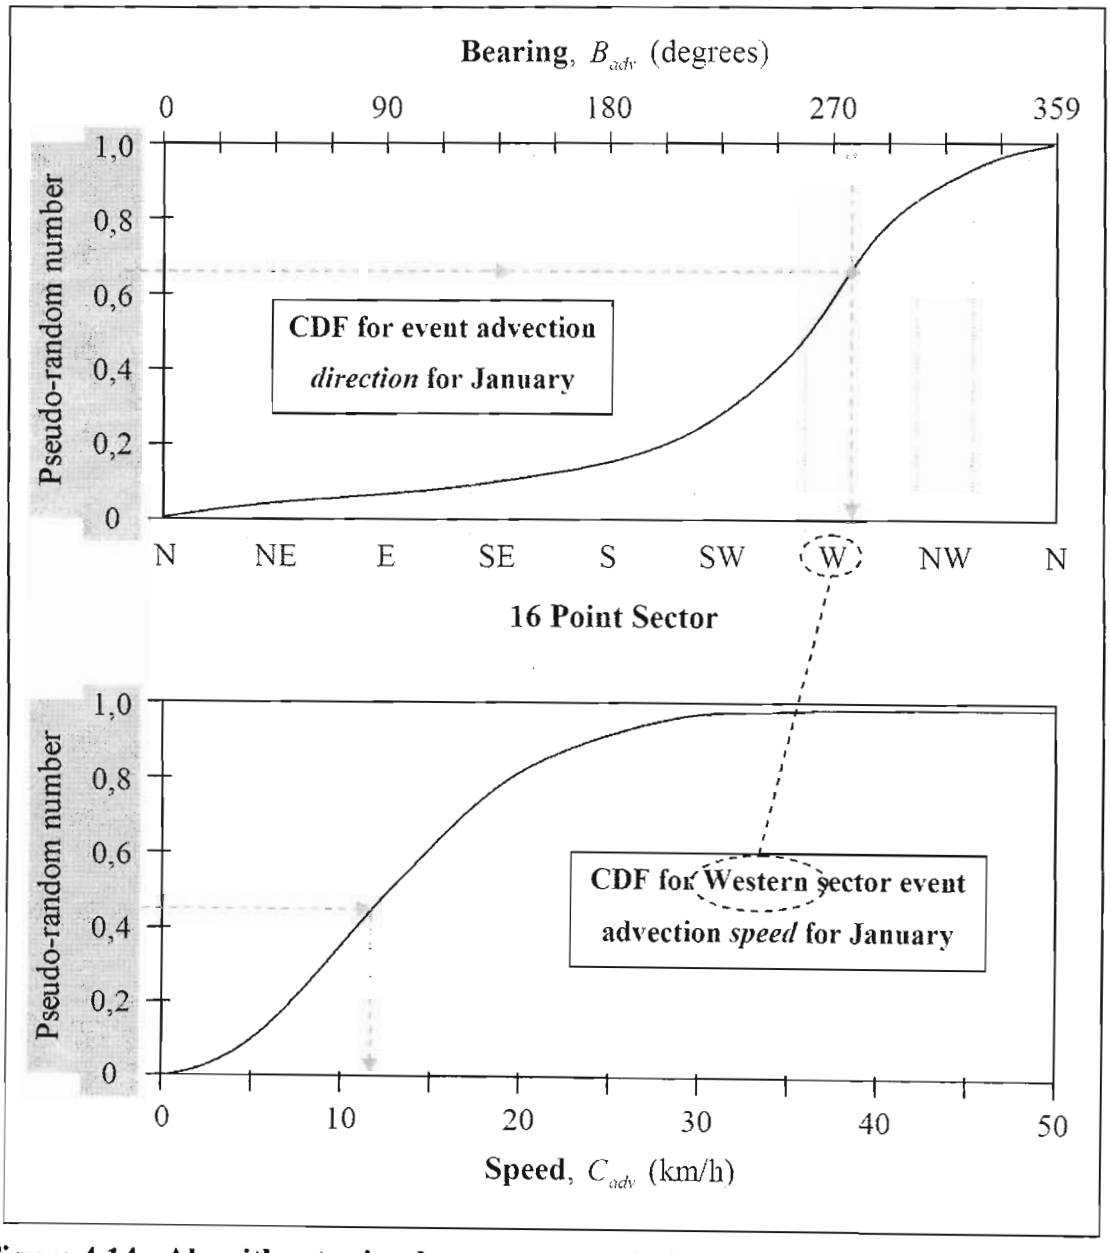

Figure 4.14 - Algorithm to simulate event cumulative advection vector

To simulate the first component (direction), depending on the month in which the event starts, a suitable curve which describes the Cumulative Density Function (CDF) for the event advection direction, B_{adv}, is selected. In the example given in Figure 4.14, the event starts in January. Next, a uniformly distributed random number is generated on the interval (0,1) using the Wichmann and Hill (1982) generator discussed in Appendix B, Section 9.4. This pseudo-random number corresponds, on the selected CDF, to a bearing on the range (0, 360) degrees and consequently to one of the 16 compass sectors [N, NNE, NE, ..., NNW]. In the example of Figure 4.14, a random number of 0,66 was generated which corresponds in the target distribution for January, to a bearing of 277 degrees which falls in the Western sector.

For the second component, selecting the CDF for the Western sector event advection *speed*, C_{adv}, for January, a second uniformly distributed pseudo-random number is generated and the corresponding advection speed is computed from the target distribution. In the example of Figure 4.14, the second random number was 0,45, which corresponds to an advection *speed* for the Western sector, in January, of 12 km/h. The simulated advection speed and direction are assumed constant during the whole rainfall event of duration D_w.

In its current configuration, the "String of Beads" model uses twelve (monthly) CDFs for the event advection direction, each binned into sixteen compass sectors (whose limits are denoted by vertical lines in Figure 4.14), each with its own CDF for event advection speed – a total of 12 x 16 curves. A matter of model calibration, the CDFs for *event advection direction* and *event advection speed* are locale specific and must be measured from available data, appropriate to the region for which rainfall is to be simulated.

First stage simulation output

To summarise, the output from the first stage of simulation using the *String of Beads* model is illustrated in Figure 4.15.

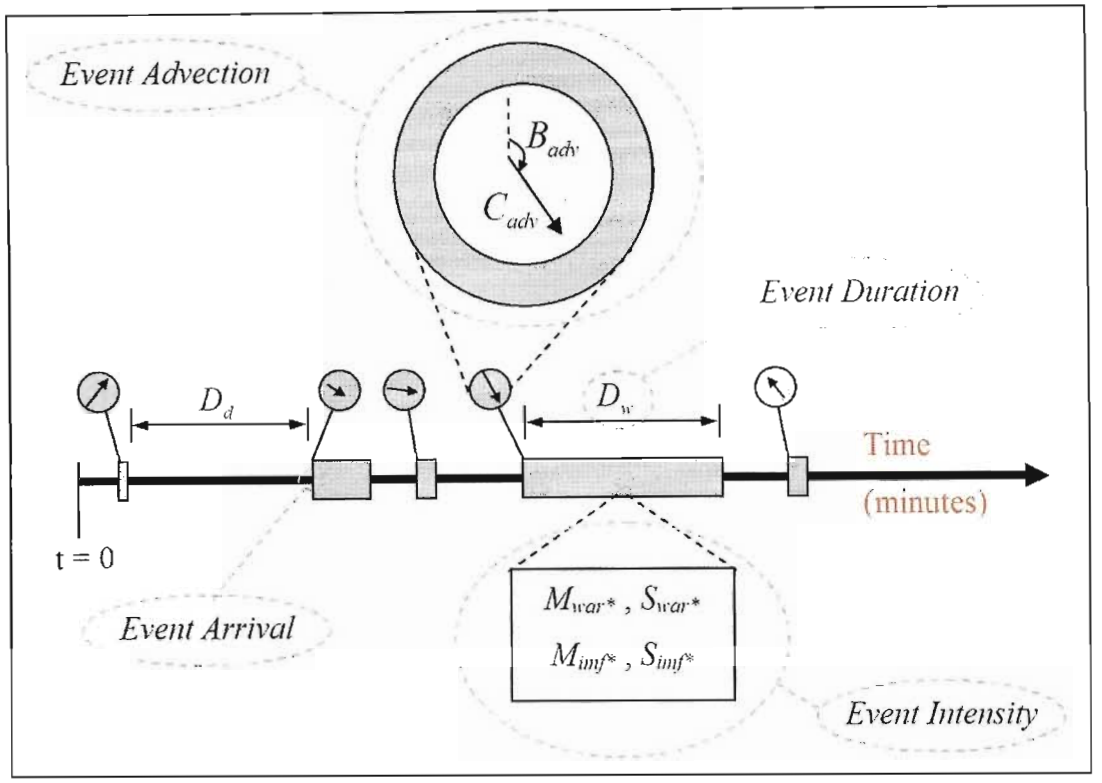

Figure 4.15 – Output of the *String of Beads* model at the end of the first stage of simulation

The seasonally-varying *scattered event alternating renewal process* simulates the wet-dry process, thereby providing information regarding the arrival of the events in the simulation and their respective durations. Each event duration, is then used in conjunction with the algorithm presented in Figure 4.13 to provide information regarding the respective event intensities. In a separate pseudo-random process, depending on the month of arrival, an event advection direction and event advection speed are simulated using the algorithm presented in Figure 4.14.

4.5.2 *SBM Second stage – image scale simulation*

The second stage of simulation uses the event scale parameters, output by the first stage, to generate pseudo-random time series of image scale parameters for each event. Three image scale parameters are required in order to simulate a single image and these are β_{space}, WAR_i and IMF_i. β_{space} is assumed to be a constant 2,5 based on the analysis presented in Section 3.7.1 and in particular, those presented in Figure 3.39 and Figure 3.40. The main process used in this stage of simulation is a bivariate autoregressive process used to model the WAR_i and IMF_i time series. The basic theory for the autoregressive process is outlined in Appendix B, Section 9.5 and pertinent analysis in this regard was presented in Section 3.10.3.

The ε_i were extracted using Equation 9.31 for a range of p and the optimum order of the autoregressive model was found to be a 5^{th} order process. This is illustrated by Figure 4.16 in which the Akaike Information criterion, AIC_c, (Hurvich and Tsai, 1989) is plotted for increasing order, p.

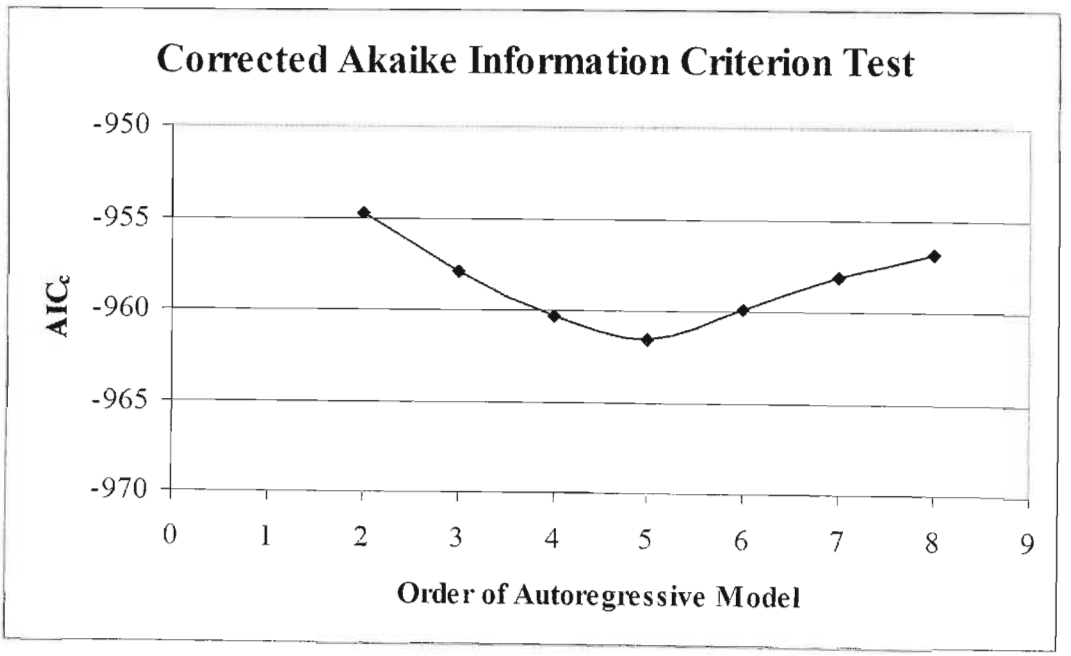

Figure 4.16 - Optimisation of the order of AR model

Since the bivariate autoregressive theory was developed for use in the simulation of Gaussian time series, the WAR_i and IMF_i time series are first simulated in Gaussian space as WAR_i* and IMF_i* as shown in Figure 4.17.

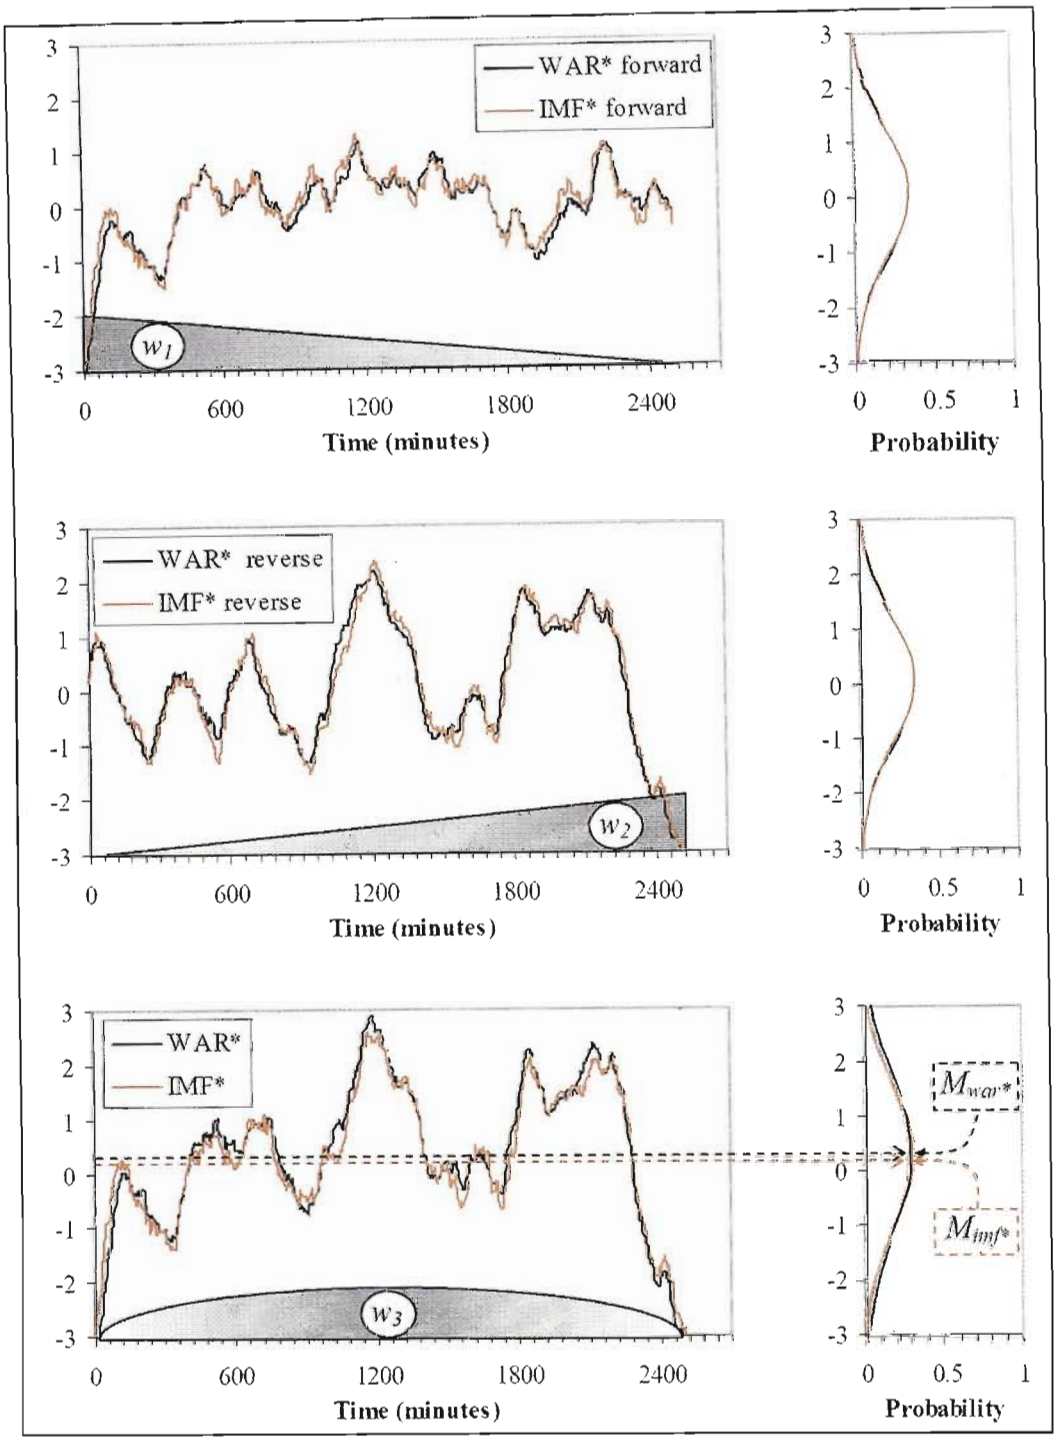

Figure 4.17 - Algorithm used to simulate image scale time series for an event. WAR_i^* and IMF_i^* *are trivariate Gaussian variables at this stage*

The parameters of the bivariate autoregressive process are computed using the Yule-Walker relationship given in Equation 9.25, in combination with the serial and cross correlation of the WAR_i* and IMF_i* time series, measured in Section 3.10.3 and presented in Figure 3.69. Owing to the fact that the autoregressive process is pseudo-random, there is no way of controlling the finishing state of the time series simulation with the generation of a single sequence. This presents a problem since a rainfall event

is defined as a period during which the WAR_i exceeds 1%, so the WAR_i must start and finish at 1%. The solution is to simulate two bivariate AR processes for each event, one in the forward time direction ($FWAR_i^*$, $FIMF_i^*$) and one in the reverse direction ($RWAR_i^*$, $RIMF_i^*$) as illustrated in Figure 4.17.

The two time series are then linearly combined using respective weighting functions w_1, which *decreases* linearly from 1 at the start of the event to 0 at the end, and w_2, which *increases* linearly from 0 at the start of the event to 1 at the end. A side effect of combining the two time series, is a reduction in variance of the time series closer to the middle of the event. The correction is to introduce a third weighting function, w_3, which is applied to the combined process and defined by Equation 4.9. Derivation of this weighting function is given in Section 9.6 of Appendix B.

$$w_3 = \frac{1}{\sqrt{\left(1 - \frac{2t}{D_w}\right) + \left(\frac{2t}{D_w}\right)^2}} \qquad (4.9)$$

where t is the time from the start of the event.

Symbolically, the linear combination of the time series is represented in Equation 4.10.

$$WAR_i^*(t) = \left(w_1(t).FWAR_i^*(t) + w_2(t).RWAR_i^*(t)\right)w_3(t) \qquad (4.10)$$

The starting points for the forward and reverse simulated (WAR_i^*, IMF_i^*) time series are chosen so that when transformed back into real space via the target distribution, WAR_i is 1% with a corresponding, appropriate IMF_i. Based on the literature reviewed, this appears to be a novel approach towards conditioning the autoregressive model on a specific finishing value.

Being Gaussian, the marginal distributions of the forward and reverse time series combine to form a third Gaussian time series. The simulated time series of WAR_i^* and IMF_i^* are then scaled and shifted to achieve their respective $N(M_{war*}, S_{war*})$ and $N(M_{imf*}, S_{imf*})$ marginal distributions, passed to the second stage of simulation by the first. Finally, the (Gaussian) WAR_i* and IMF_i* time series are transformed into real space (WAR_i and IMF_i) using the target distribution curves presented in Figure 3.56 for the WAR_i and Figure 3.57 for the IMF_i. This transformation process is illustrated (in reverse) in Figure 3.58.

Second stage simulation output

At this, the end of the second stage of simulation, the output of the *String of Beads* model is a continuous time series of image scale statistics, at five minute temporal resolution. The WAR_i and IMF_i are assumed to be zero during the dry spells and they start and finish each wet spell with a WAR_i of 1%.

From a hydrologist's point of view, this can already provide useful information for the study area as a whole. The IMF_i for a single image represents the average rainfall rate (in mm/h) over the entire image, so the average rainfall accumulated over the image, $\overline{R}_e$ (mm), for an event of duration D_w, simulated at a temporal resolution of δt minutes, can be approximated by Equation 4.11.

$$\overline{R}_e = \frac{\delta t}{60} \sum_{k=1}^{n} IMF_i(k) \qquad (k = 1, 2, \ldots, n) \qquad (4.11)$$

where n is the number of images in the event given by:

$$n = \frac{D_w}{\delta t} \qquad (4.12)$$

So for five minute resolution data, the average rainfall accumulated over the image for any period (including both wet and dry spells), can be *quickly* computed by summing the IMF_i during that period, and dividing by twelve to convert the values to millimetres per hour. This is an important relationship in the context of the *String of Beads* model calibration, which will be discussed in Chapter 5.

4.5.3 SBM Third stage – pixel scale simulation

The third and final stage of simulation accepts the event advection and image scale parameters output by the first and second stages and simulates 128km, two-dimensional images at the pixel scale, consistent with those parameters. The algorithm for this stage is illustrated in Figure 4.18 and requires careful explanation because it is a multiple stage process.

There are three main underlying processes that need to be combined in this third stage, the *pixel scale temporal process*, the *field advection process* and the *pixel scale spatial process*. The first is designed to manage the temporal behaviour at the pixel scale between consecutive images. The second is to introduce the advection component of the event. The third is to impose the required spatial structure on individual images.

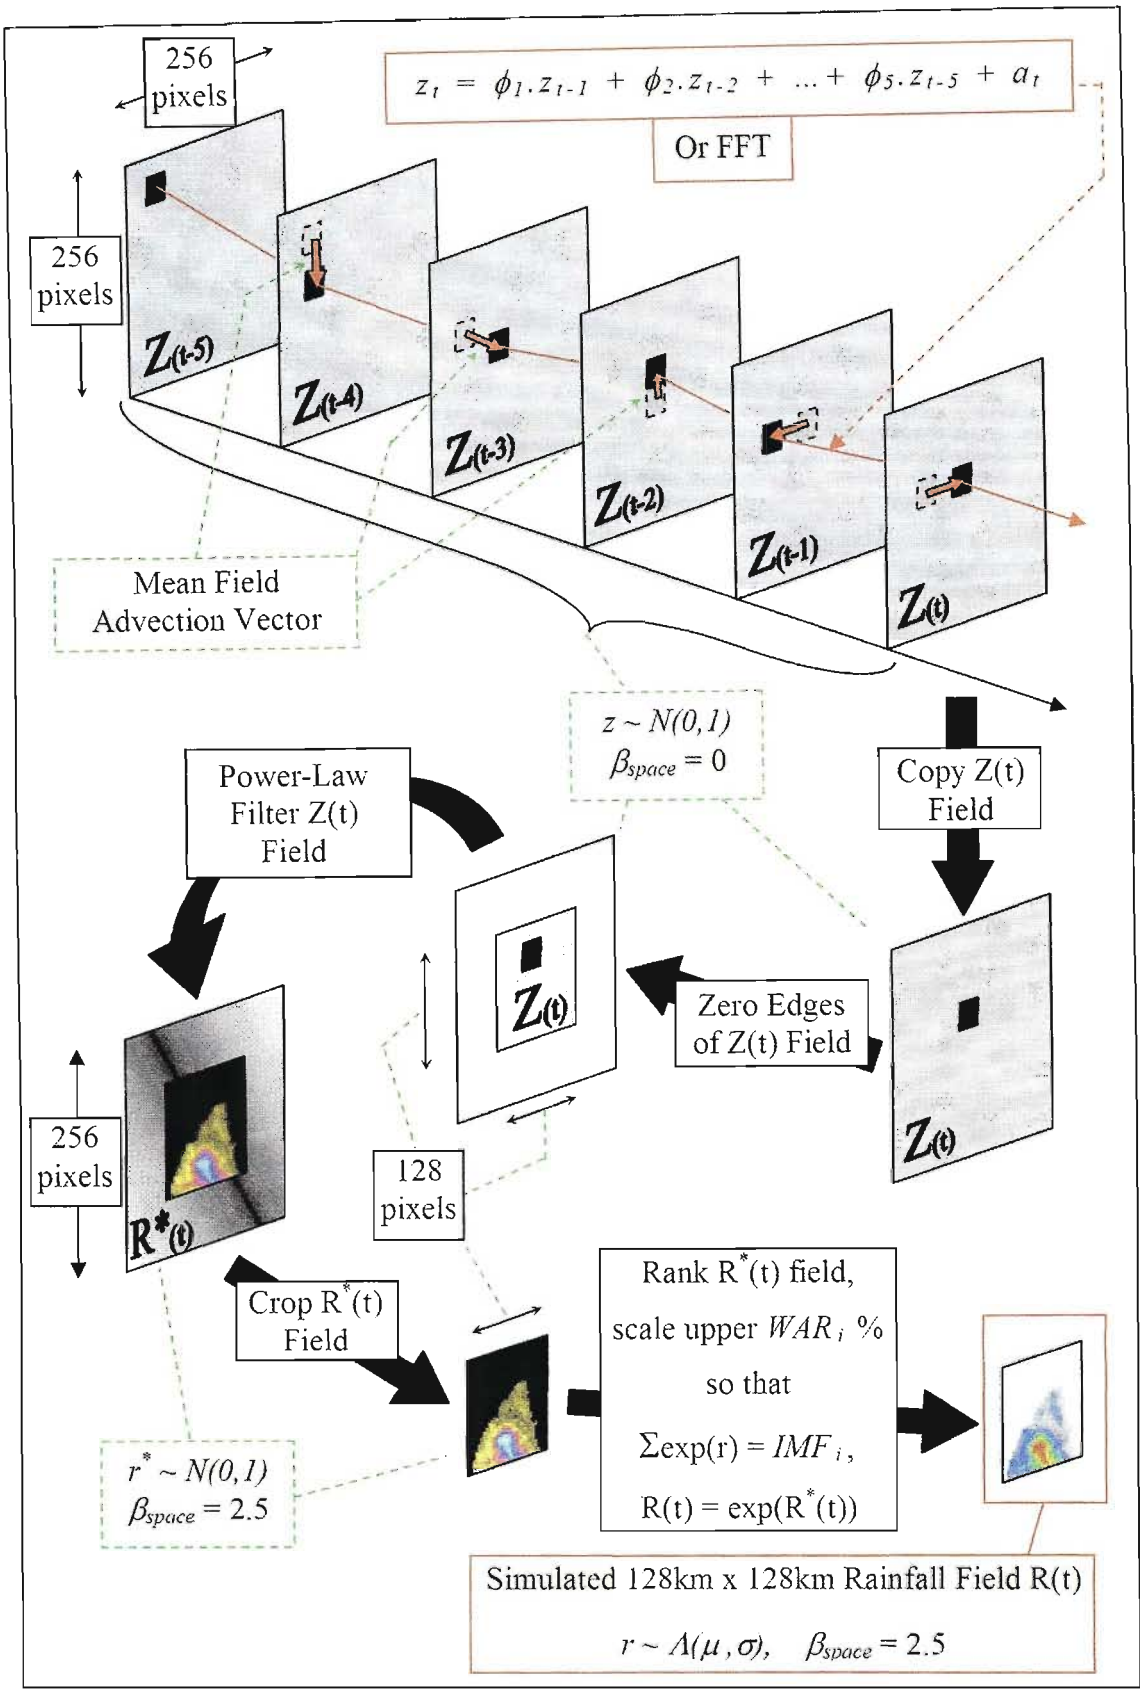

Figure 4.18 - Stack algorithm used to simulate pseudo-random images from image scale statistics using the autoregressive approach

Pixel scale temporal process

The *pixel scale temporal process* generates temporally correlated, spatially uncorrelated, pseudo-random fields of pixels. When viewed as individual images, these fields appear to be entirely random with no measurable spatial correlation – two dimensional fields of white noise. When viewed as a sequence however, these spatially uncorrelated fields exhibit a well-defined temporal correlation structure even at image scale. Development of an individual pixel along the timeline is illustrated in the upper half of Figure 4.18 As an advection shifted pixel in the spatially uncorrelated $Z_{(t-5)}$, $Z_{(t-4)}$, ..., $Z_{(t)}$ arrays.

Two different methods are proposed for this part of the third stage: one in the time domain which makes use of a univariate autoregressive process (resulting in the *Autoregressive Bead*), the other in the frequency domain which makes use of a one-dimensional power-law filter process (resulting in the *Frequency Bead*). Both processes achieve satisfactory results on the whole, but each has advantages over the other and a comparison is given in Table 4.1. Relevant theory is presented in Sections 4.3 and 4.4 and relevant analysis was presented in Sections 3.11 and 3.12.

Table 4.1 -Comparison of the autoregressive and power-law filtered forms of the pixel scale temporal process.

Comparison	Autoregressive Bead	Frequency Bead
Computational speed	Relatively fast (2 hr/year on 1GHz processor)	Relatively slow (9 hr/year on 1GHz processor)
Memory requirements	Relatively small and small by modern standards (3Mb)	Relatively large, but reasonable by modern standards (17Mb)
Forecasting abilities	Short term	None
Short term temporal correlation structure, upscaling structure	Described by 5 parameters, reasonably good short-term upscaling, visually stagnant	Described by 1 parameter, not quite as good as short-term upscaling, visually excellent
Cumulative spatial structure	Reasonable at the event scale but not as good for daily, monthly and annual accumulations	Good for all accumulations

The most significant benefit of the autoregressive approach is the fact that it affords some short term forecasting abilities, while the biggest benefit of the power-law filtered approach is that it provides sequences which accumulate sensibly over the long term. Comparisons of the sequences will be presented in Chapter 6 - Output and Validation. For these reasons, the autoregressive approach should be used for operation in forecasting mode, while the power-law filtered approach should be used in simulation mode. Each method is described in turn in the following paragraphs.

The ***univariate autoregressive*** approach for image generation used in this stage of pixel scale simulation is a fifth order process, so in order to accommodate the noise fields, a stack of six two-dimensional arrays is allocated in the computer memory, $Z_{(t-5)}$, $Z_{(t-4)}$, ..., $Z_{(t)}$, each 256 x 256 pixels in size (See Figure 4.18). To start the process off at the beginning of an event, the arrays $Z_{(t-5)}$, ..., $Z_{(t-1)}$ are filled with normally distributed pseudo-random noise. The $Z_{(t)}$ array is then computed at each of its 65536 pixels, $z(x,y)_t$, independently using Equation 4.13.

$$z(x,y)_t = \phi_1 . z(x,y)_{t-1} + \phi_2 . z(x,y)_{t-2} + \ldots + \phi_5 . z(x,y)_{t-5} + a_t \quad (4.13)$$

where

$0 \le x < 256$

$0 \le y < 256$

a_t is a shock term

$\phi_1 \ldots \phi_5$ are autoregressive parameters

To homogenise the temporal structure of the $Z_{(t-5)}$, ..., $Z_{(t-1)}$ arrays, the autoregressive process is allowed to run for several iterations before image simulation begins.

After each time step is complete, the $Z_{(t)}$ array is passed to the field advection process. (The autoregressive parameters ϕ_n, describe the temporal relationship between consecutive pixels in a *normalized Lagrangian* reference frame, that is the reference frame centred on the storm in which the radar appears to move. Measurement of these parameters from the dataset is a matter of model calibration which will be discussed in Chapter 5).

The ***one-dimensional power-law filter*** approach to image generation also makes use of a stack of two-dimensional noise fields, each 256 x 256 pixels in size, but instead of just 6 fields, 64 are used, $P_{(t-63)}$, $P_{(t-62)}$, ..., $P_{(t)}$. (Any less than 64 and patterns start to be seen in the upscaling correlation structure, to be discussed in Chapter 5). Like the autoregressive approach, to start the process off at the beginning of an event, the arrays $P_{(t-63)}$, ..., $P_{(t)}$ are filled with normally distributed pseudo-random noise. In contrast to the autoregressive approach however, these fields are, and remain, temporally independent of each other and require no homogenisation. At each timestep, the oldest field $P_{(t-63)}$, is first discarded and all remaining 63 fields are demoted by one timestep to make way for the newest field $P_{(t)}$ which is generated as an independent array of N(0,1) noise.

The pixel p at a position (x,y) in space, denoted $p(x,y)$, is then sampled from each of the 64 noise fields , $P_{(t-63)}$, $P_{(t-62)}$, ..., $P_{(t)}$ and copied into a one-dimensional, 64 value array $p(x,y)_{(t-63)}$, $p(x,y)_{(t-62)}$, ..., $p(x,y)_{(t)}$ which is then transformed into Fourier space via the FFT and power-law filtered with exponent β_{time} to achieve a temporally correlated sequence of pixels, $z(x,y)_{(t-63)}$, $z(x,y)_{(t-62)}$, ..., $z(x,y)_{(t)}$. Only the $z(x,y)_{(t)}$ value is then copied into the 256 x 256 pixel $Z_{(t)}$ array. As for the autoregressive process, the $Z_{(t)}$ array is computed at each of its 65536 pixels, $z(x,y)_t$, in this way before being passed to the field advection process. The β_{time} parameter alone describes the temporal relationship between consecutive pixels in a *normalized Lagrangian* reference frame. Measurement of this parameter from the dataset is a matter of model calibration which will be discussed in Chapter 5.

Field advection process

At each time step, t, the $Z_{(t)}$ array is received from either the *univariate autoregressive process* or the *one-dimensional power-law filter process* described above. The field is spatially shifted in space according to the mean field advection vector output by the first stage of simulation, in order to introduce the advection component of the rainfall event. The spatial shift is illustrated in Figure 4.18 by red arrows on the spatially uncorrelated $Z_{(t-5)}$, $Z_{(t-4)}$, ..., $Z_{(t)}$ arrays. Pixels which are pushed off the edge of the image are allowed to rejoin at the opposing edge. Although shown as a constantly changing vector in

Figure 3.41, in its current state, the advection during the event is modelled as a constant direction, B_{adv}, and speed, C_{adv}, however a more sophisticated approach could be adopted if required. Relevant analysis was presented in Sections 3.6 and 3.7. The spatially advected, uncorrelated $Z_{(t)}$ array is then passed to the *pixel scale spatial process*.

Pixel scale spatial process

The *pixel scale spatial process* accepts the temporally correlated, spatially uncorrelated, advection-shifted fields provided by the *field advection process* and imposes a realistic spatial correlation structure on the fields whilst maintaining an underlying temporal correlation structure and field advection. Theory pertinent to this stage of simulation was presented in Sections 4.4.2 and relevant analysis was presented in Section 3.5. The process is depicted in the central (reverse flow) section of Figure 4.18.

The shifted $Z_{(t)}$ array (comprised of z_t ~N(0,1)) is cropped and copied into a separate working array where the edges of the array are set to zero, leaving central 128km part of the field unchanged. The zero padding around the edge of the field eliminates Fourier wrapping, an artefact of power-law filtering which makes opposite edges of the image behave as if they were continuous in space.

The next step is to power-law filter the two-dimensional 256 x 256 pixel noise field, according to the description given in Section 4.4. The result is a (Gaussian) field, $R^*_{(t)}$, in which the 128km core has a grainy, but spatially correlated nature, as observed in radar rainfall fields. The edges of the field, zero before the filtering process, form a smooth function which gradually links the opposing edges of the 128km core. Cropping the edges leaves a 128km image with a spatial power spectrum characterized by a power law function with exponent $\beta_{space} = 2{,}5$.

The final process shown in the bottom of Figure 4.18 is to scale and shift the $R^*_{(t)}$ field to achieve the desired WAR_i and IMF_i image scale statistics in the simulated rainfall field. The $R_{(t)}$ field is obtained by *exponentiating*, pixelwise, the appropriately scaled and shifted $R^*_{(t)}$ field. Following the argument presented in Section 3.4 (but substituting

the $R^*_{(t)}$ field for the Z_j field, R_{WAR} for Z_{WAR} and r_t^* for Z_j) the percentile R_{WAR} corresponding to the WAR_i is first determined such that

$$P\{r_t^* > R_{WAR}\} = WAR_i \tag{4.14}$$

where r_t^* is a pixel in the $R^*_{(t)}$ field. The $R^*_{(t)}$ field is then *shifted* by subtracting R_{WAR} from all r_t^* in the field so that

$$\frac{number(r_t^* \geq 0)}{number(r_t^*)} = WAR_i \tag{4.15}$$

Following the argument of Section 3.4.1, scaling of the $R^*_{(t)}$ field has to be done iteratively by first computing IMF'' using Equation 4.16,

$$IMF'' = \sum_{r_t^* \geq 0} \exp(\lambda . r_t^*) \tag{4.16}$$

where λ is a scalar term to be determined iteratively. The objective function for the iterative solution, ε, is then given by Equation 4.17.

$$\varepsilon = | IMF'' - IMF_i | \tag{4.17}$$

where ε is the absolute value of the difference between the simulated IMF'' and the IMF_i passed to the third stage of simulation by the second stage. The value of λ is adjusted with each iteration until ε is small enough. The advection characteristics of the original random field are preserved throughout this process.

The current form of the third stage of simulation represents a significant evolution from the initial form of the *String of Beads* model as presented by Pegram and Clothier (1999) and Pegram and Clothier (2001a). The initial form of the *String of Beads* model employed Taylor's Hypothesis (1938) as discussed by Gupta and Waymire (1987), and generated a three dimensional (2 space and 1 time) space-time hypercube of consecutive radar rainfall images by power-law filtering a three-dimensional noise field. Although effective, this process is extremely expensive in terms of computer memory and processing power and has been dramatically reduced through the introduction of the autoregressive and the modified power law filtered methods of the third stage of simulation discussed earlier in this section.

Third stage simulation output

The third stage of the *String of Beads* model outputs simulated, 128km rainfall images, at five-minute, one-kilometre resolution. The algorithm can be configured to store images at any spatial scale between 1 and 128 km and at any temporal scale between 5 minutes and the duration of the entire simulation by accumulating the five-minute, one-kilometre images into appropriate space-time blocks as illustrated in Section 3.12 for the image upscaling analysis. Images can be stored as floating point precision arrays, (floatmaps) or as bitmap images as discussed in Section 2.7.

4.5.4 Summary of the SBM structure

The structure of the *String of Beads* model and the required input and output at each stage is summarised in Table 4.2. The model can operate in one of three modes:

1. *String of Beads simulation*
2. *String* calibration
3. *Bead* calibration

String of Beads simulation mode uses all three stages namely the *event scale simulation*, the *image scale simulation* and the *pixel scale simulation*. Once calibrated, operation in simulation mode requires only user input in the form of the duration of the simulation, seeds for the pseudo-random number generator and program configuration options for the type of output.

String calibration mode uses only the first and second stages namely the *event scale simulation*, the *image scale simulation*. It requires the same input as for the simulation mode, as well as long term daily raingauge data to measure seasonal variation in the string.

Bead calibration mode uses only the third stage, the *image scale simulation*, and requires the full set of image scale statistics from an observed event. Calibration modes will be discussed in Chapter 5.

Table 4.2 - Summary of the structure of the *String of Beads* model

	First Stage – Event Scale	Second Stage – Image Scale	Third Stage – Pixel Scale
User input	***String of Beads* Simulation:** • Simulation start date • Simulation stop date • Random seeds • Output options ***String* calibration:** • All of the above (excluding output options) • Raingauge data ***Bead* calibration:** • None	None	***String of Beads* Simulation:** • None ***Bead* Calibration:** • Event Advection and WAR_i and IMF_i time series from observed event
Algorithm Input	User input	First Stage Output	***String of Beads* Simulation:** • Event Advection and Arrival parameters from first stage • WAR_i and IMF_i time series from second stage ***Bead* Calibration:** • User input
Process	• Seasonally varying scattered event renewal process for event arrival and duration. • Target distribution process for event advection	• Combination of linearly weighted forward and reverse bivariate autoregressive process • Target distribution process to convert Gaussian parameters into real parameters	• Univariate autoregressive pixel scale process **OR** One-dimensional power-law filtering process • Two-dimensional power-law filtering process
Output	• Event Duration and Event Arrival parameters (String Process), D_w • Event Advection parameters, B_{adv} and C_{adv}. • Event Intensity parameters, M_{war^*}, S_{war^*}, M_{imf^*} and S_{imf^*}	***String of Beads* Simulation:** • Continuous WAR_i and IMF_i time series for the entire simulation ***String* Calibration:** • Daily, monthly and yearly image-average rainfall totals	Two-dimensional radar rainfall images at desired space-time resolution specified in the first stage user input

For the *String of Beads simulation* and *String calibration* modes, output from the first stage is the full set of simulated Event scale statistics which are passed on to the second stage. *Bead calibration* mode does not use the first and second stages.

Output from the second stage for both the *String of Beads simulation* mode and the *String calibration* mode is a full set of simulated image scale statistics, at five minute resolution, for the entire duration of the simulation. In the case of the *String of Beads simulation* mode, these are passed on to the third stage for pixel scale simulation. In the case of *String calibration* mode the third stage is not used. Instead image averaged daily, monthly and yearly simulated rainfall accumulations are output at the end of the second stage using the result derived at the end of Section 4.5.2.

For the *String of Beads simulation* and *Bead calibration* modes, output from the third stage is a sequence of simulated radar rainfall images. The spatial and temporal resolution of the output for the *String of Beads simulation* mode depends on the output options specified by the user at the beginning of the first stage. For the *Bead calibration* mode the full range of spatial and temporal images are output.

4.6 COMPUTING CONSIDERATIONS

Until recently (probably 1995), the processes simulated in the *String of Beads* model would have been confined to the realm of the supercomputer. Personal computers were inadequate in terms of processing power, memory and data storage capacity. In recent years this has changed and the *String of Beads* model has been designed to run on a personal computer. For annual simulations, the minimum practical computing requirements would be a 500MHz processor, 128MB of memory and 10GB of available storage. Although simulation is reasonably quick on a machine of this specification, the performance of the model can be vastly improved by exploiting its modular nature by distributing the processing tasks amongst a network of personal computers. This is known as Massive Parallel Processing and the concepts have been tried and tested during the course of this study.

The first and second stages of the model are relatively un-intensive when compared to the third stage. For a one-year simulation in South African conditions, execution of the

first stage (the univariate event-scale specification), on this minimum specification machine, requires less than 1 seconds. The second stage (univariate image scale stochastic process) would be complete in under 10 seconds and the third stage (the full set of 5-minute images) would take 5 to 10 hours.

To be useful as a design tool in hydrology, the model would be used to run a 50 year simulation of only the first and second stages and then to select from the image scale statistics, extreme wet or dry days, months or years, depending on the whether the critical design criteria involve flood or drought conditions. The selected extreme periods can then be passed to the third stage of simulation in order to simulate the pixel scale detail. It is the processing of the third stage which can be easily distributed amongst network computers. Third stage simulation of a single, typical event runs for approximately 3 minutes on the minimum specification machine. For a given number of events, the processing time is inversely proportional to the number of processors used in the simulation. Distribution of processing tasks can either be automatically assigned over the network using operating system dependent commands, or alternately, the selected image scale statistics can be output to file and manually run through the third stage on separate computers.

4.7 SUMMARY

The *String of Beads* model can be separated into three stages. The first stage simulates statistics at the *event* scale including event duration, event arrival, event advection and event intensity statistics. The second stage uses the event scale statistics output by the first stage to simulate time series of statistics at the *image* scale, namely the WAR_i, IMF_i and β_{space}. The third stage uses the image scale statistics output by the second stage to simulate radar rainfall images at the 1km, 5minute pixel scale. The model is used in one of three modes , *String of Beads simulation*, *String Calibration* or *Bead calibration* mode (model calibration will be discussed in Chapter 5). The model can run on a single personal computer, or to reduce processing time, the pixel scale (third stage) simulation processes can be distributed amongst available machines on a network.

Chapter 5

Model Calibration and Verification

5.1 INTRODUCTION

There are two stages of calibration in the *String of Beads* model, one for the *Bead* process and the other for the *String* process each of which will be discussed in turn in this chapter. Wherever possible and sensible, verification of the model output, after calibration, will be presented. Validation of the model is deferred until Chapter 6.

5.2 THE 'BEAD' CALIBRATION PROCESS

Two bead processes have been developed for use in the *String of Beads* Model and were introduced in Section 4.5.3, the *autoregressive bead* used mainly in short term simulations and forecasting and the *frequency filtered bead* used for long term simulations. These processes are calibrated independently, but the calibration process is common to both. There are two stages to the bead calibration process. The first is an explicit measurement process to determine the event scale statistics and their inter-relationship with each other. The second is an iterative upscaling exercise to determine the underlying pixel scale temporal correlation structure.

5.2.1 Explicit measurement of event scale statistics

Explicit measurement of the event scale statistics, namely the D_w, M_{war*}, S_{war*}, M_{imf*} and S_{imf*} and their inter-relationships and the field advection vectors, was illustrated in the analyses presented in Section 3.10.2. The relationships between these parameters are presented in Equations 4.5 through 4.8, repeated here:

$$M_{war*} = 0{,}54.\mathrm{Ln}(D_w) - 0{,}98 + 0{,}57.a \tag{4.5}$$

$$M_{imf*} = 0{,}89.M_{war*} + 0.04 + 0{,}34.a \tag{4.6}$$

$$S_{war*} = 0{,}16.M_{war*} + 0{,}58 + 0{,}29.a \tag{4.7}$$

$$S_{imf*} = 1{,}00.S_{war*} + 0{,}01 + 0{,}11.a \tag{4.8}$$

The a terms in these equations are independent standard normal variates.

These relationships have been determined independent of season, however there is scope, given enough radar data, to perform this analysis on a seasonal basis and thereby introduce seasonally varying event scale parameters. For the purposes of this study, as a first approximation, it has been assumed that the event scale statistics are independent of season; seasonal variation is left as a topic of future research.

The Cumulative Distribution Functions (CDFs) for the 500Hpa wind speed and direction are determined as shown in Section 3.7.1. Seasonal variation in wind speed and direction is accounted for in this process. With sufficient radar data, it would be possible to measure this directly from the data using the technique described in Section 3.6.

5.2.2 Iterative upscaling calibration exercise

Although the process is almost identical for both types of bead, it is important to bear in mind the end objective. For the autoregressive bead, the object of the exercise is to determine the $\phi_1 \ldots \phi_5$ parameters of the univariate autoregressive process defined in Appendix B, Section 9.5. For the frequency filtered bead, it is the β_{time} parameter which must be determined.

Depending on which type of bead is used, the $\phi_1 \ldots \phi_5$ parameters *or* the β_{time} parameter, defines the temporal correlation structure of the model at the pixel scale – that is, the correlation of the Pixel Scale Intensity (PSI). Computation of these parameters by measuring the pixel scale correlation from the original radar data is theoretically possible, but practically infeasible as demonstrated by the attempts made in Section 3.11. With reference to Section 2.6, this module of the *String of Beads* model was developed on the Pre-1997, integer precision radar dataset.

The technique uses the spatial upscaling analysis of a rainfall event, presented in Section 3.12, in conjunction with the third stage of simulation, presented in Section 4.5.3, to simulate events which are statistically identical to the observed sequence of radar-rainfall measurements at the image scale, using a choice of $\phi_1 \ldots \phi_5$ parameters and β_{time} parameters, until a satisfactory upscaling behaviour can be verified. In detail, the upscaling calibration process is:

1. For a chosen sequence of observed radar images, measure the WAR_i, IMF_i and β_{space} image scale statistics (Sections 3.4 and 3.5), as well as the mean field advection vector for each time step in the sequence (Section 3.6).
2. For the same observed sequence of radar images, quantify the space-time upscaling properties of the observed event (Section 3.12).
3. Choose a set of $\phi_1 \ldots \phi_5$ parameters *or* a β_{time} parameter
4. Using the chosen parameters, run the third stage of simulation (Sections 4.6.3), using the image scale statistics measured in step 1 as input parameters, as though they were passed from the second stage of simulation.
5. Convert the output of the simulation in step 3, into a format identical to that of the observed sequence (same precision, mask and spatial resolution).
6. Using the same routines as in step 2, quantify the space-time upscaling properties of the simulated event (Section 3.12).
7. Compare the space-time upscaling properties of the observed and simulated events, measuring the sum of the squares of the differences between the autocorrelation functions for the observed and simulated sequences, for a range of space-time accumulations. Examples of these autocorrelation functions are given for an observed sequence in Figure 3.74.
8. Iterate from step 3, adjusting the model PSI $\phi_1 \ldots \phi_5$ parameters or the β_{time} parameter, to obtain the least squares solution.

The idea behind this bead calibration technique is to eliminate all of the variables in the *Bead simulation* process, except for the underlying parameters of the process which describe the temporal correlation structure of the (spatially uncorrelated) PSIs in the normalised, Lagrangian reference frame as described in Section 4.5.3. A fundamental assumption made in the third stage of the *String of Beads* model is that the temporal structure of the PSI is independent of weather type. To date, this assumption has not been tested.

Although conceptually simple, this is a computationally intricate and an extremely processor intensive process which has not been fully automated at the time of writing

this document. Instead the parameters were manually manipulated to obtain what appeared to be a reasonable solution and are given for the autoregressive bead as:

$$\phi_1 = 1{,}6 \qquad \phi_2 = -0.4 \qquad \phi_3 = -0.2$$

$$\phi_4 = 0.0 \qquad \phi_5 = 0.0 \qquad a = 0{,}1$$

And for the frequency-filtered bead as:

$$\beta_{time} = 4{,}0$$

These parameters are applicable to the 5 minute time scale. For simulations at a different time scale, appropriate parameters would need to be chosen. The issue of different time scales is left as a topic for future research.

Ideally, several such events should be simulated so that a spread of upscaling results should be examined in order to compare the upscaling behaviour between observed and simulated events. This is a topic which could be revisited when more data are available in order to refine the model performance. Verification of the output from both of the 'bead' processes is presented in the following section.

5.2.3 Upscaling verification of the bead process

Using these relatively optimal parameters for simulation of both the frequency filtered and the autoregressive beads, the comparative downscaling analysis for the 1996 observed sequence (whose image scale statistics are presented in Figure 3.35) and the simulated sequences, are presented in the analyses which follow.

Figure 5.1 compares the spatially averaged temporal correlation functions for a range of spatial and temporal scales, extracted from an observed 1996 event and one of its autoregressive simulated equivalents. Figure 5.2 compares the same observed event with one of its frequency filtered equivalents. Apart from the effects of re-sampling using the donut mask, the three sequences (observed and two simulated) are identical in their WAR_i, IMF_i and advection vector image scale statistics.

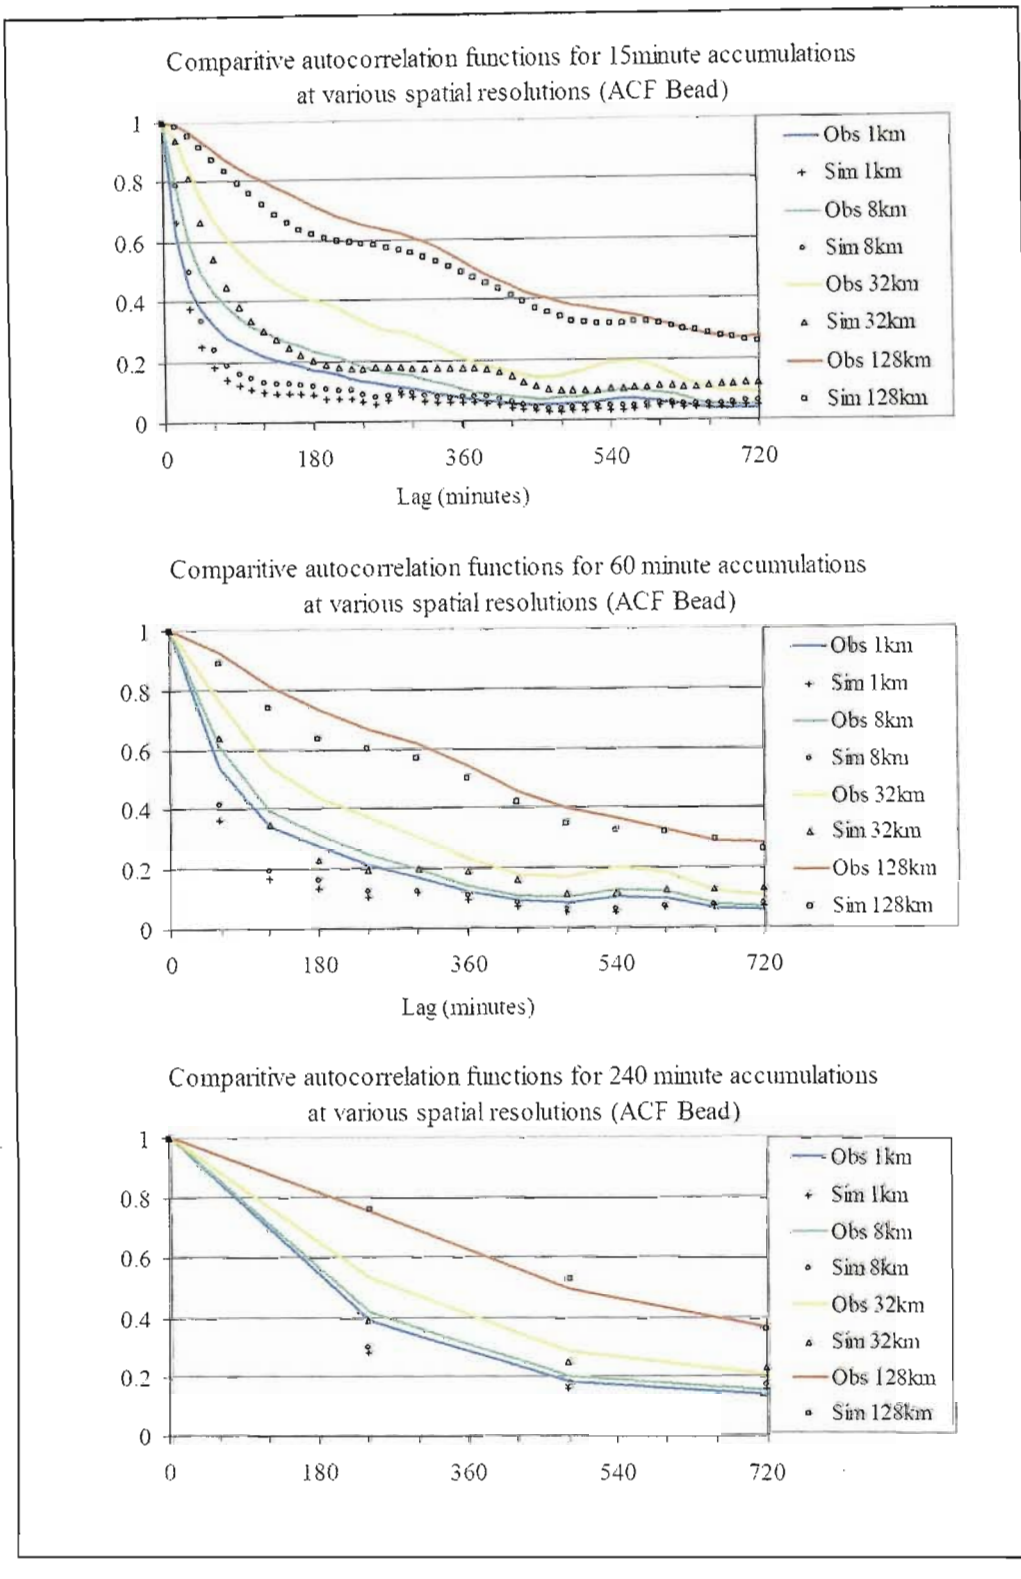

Figure 5.1 - Spatially averaged, temporal correlation functions for the 1 km, 8 km, 32 km and 128 km spatial scales at 15 min, 60 min and 240 min temporal scales for autoregressive simulated bead and observed events

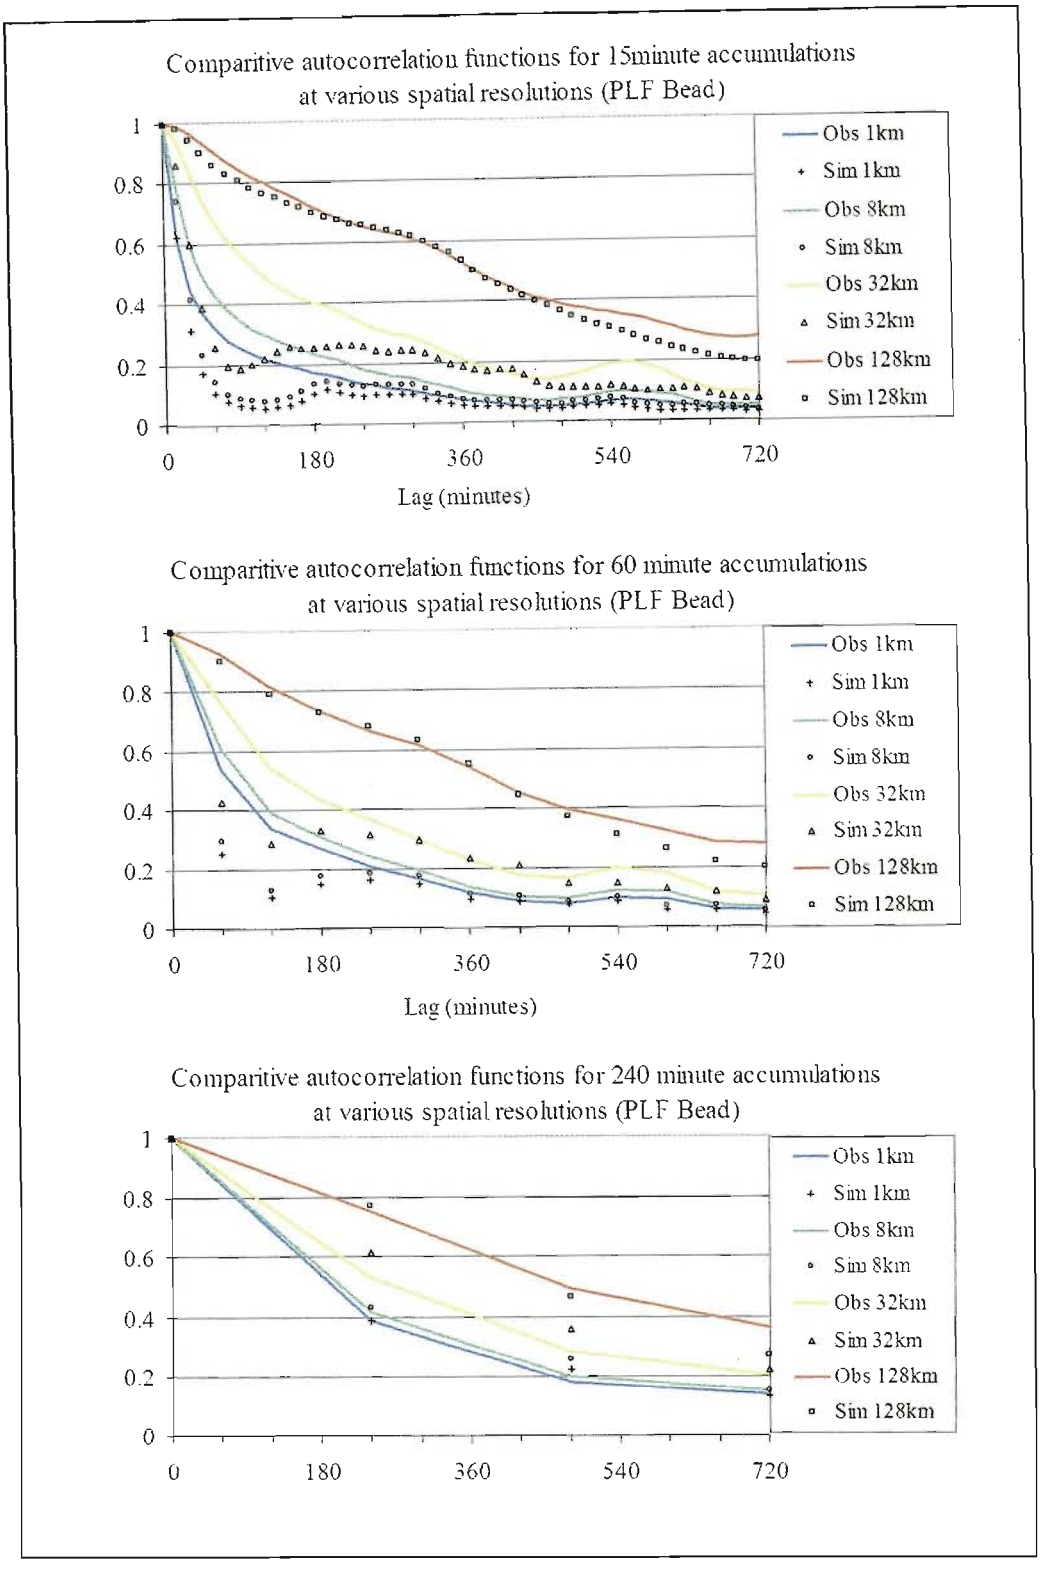

Figure 5.2 - Spatially averaged, temporal correlation functions for the 1 km, 8 km , 32 km and 128 km spatial scales at 15 min, 60 min and 240 min temporal scales for frequency filtered simulated bead and observed events

Since the third stage of simulation in the *String of Beads* model is driven by the image scale statistics, it is not surprising that the observed and simulated sequences exhibit

very similar characteristics at the 128km scale for all observed temporal scales, as illustrated in Figure 5.1 and Figure 5.2.

In both Figure 5.1 and Figure 5.2 finer spatial resolution aggregations are not as faithfully reproduced as the 128 km scale, but are still reasonable nevertheless. The autoregressive bead appears to out-perform the frequency-filtered bead for the finer spatial resolution, 15 minute accumulations, particularly for shorter lag times, while the frequency-filtered bead appears to out-perform the autoregressive bead for all spatial scales at the coarser temporal accumulations of 240 minutes and above. Closer investigations suggest that this is an artefact of masked re-sampling of the images. Re-sampling of the images within the mask can give rise to poor comparisons if significant wet portions of the field are masked out, leaving few non-zero data from which to compute the temporal correlation. Correlation for a sequence of zero values is not defined and is therefore ignored in the calculation of the auto-correlation function. The effects of masked re-sampling are exaggerated at the finer spatial and temporal scales.

For these individual tests, aggregations at the coarser temporal scales are more stable and therefore serve as a more reliable test for the bead process. In tests carried out thus far, the frequency-filtered bead consistently out-performs the autoregressive bead at all spatial scales, for all temporal scales of 240 minute or coarser. To investigate the performance of the model at higher temporal resolutions, a large number of simulations would be required, and their average behaviour compared to the observed event. This is a lengthy undertaking which really requires the bead calibration process to be automated. Qualitative judgement of the 5 minute animations of the simulated sequences, reveal a far more realistic looking evolution of the rainfall event is achieved by the frequency-filtered bead. Correlation coefficients for the first five lags of the underlying pixel scale sequences of these two beads are given in Table 5.1.

Table 5.1 - Correlation functions of autoregressive and frequency-filtered beads

Bead Type	ρ_0	ρ_1	ρ_2	ρ_3	ρ_4	ρ_5
Autoregressive	1.000	0.973	0.904	0.802	0.675	0.534
Frequency Filtered	1.000	0.995	0.994	0.992	0.989	0.986

Clearly the correlation function of the Frequency Filtered bead has a longer tail than that of the autoregressive bead. Although it may be possible to improve the performance of the autoregressive bead, the correlation structure achieved in the frequency-filtered bead by the use of $\beta_{time} = 4{,}0$ cannot be reproduced by a simple autoregressive model. Sequences with a spectral exponent β in excess of the dimension (time series are 1 dimensional) are not self-similar and do not have scaling moments (Menabde, 1998) and are therefore non-stationary. Pure autoregressive models are defined for stationary time series. To achieve the same characteristics, it may be necessary to replace the autoregressive component with an Autoregressive Moving average model (ARMA model), or even an autoregressive integrated moving average model (ARIMA model). This has not been investigated in the research to date and is left as a topic for future research. The β_{time} and $\phi_1 \ldots \phi_5$ parameters presented earlier in this chapter represent the best parameters that could be found in the available time, and as such they will be used for the remainder of this study.

An interesting aspect of the downscaling analysis is the effect of the mean field advection vector. Figure 5.3 considers the 15 minute spatially averaged temporal correlation functions of the same simulated event analysed in Figure 5.1, first with, and then without the advection component.

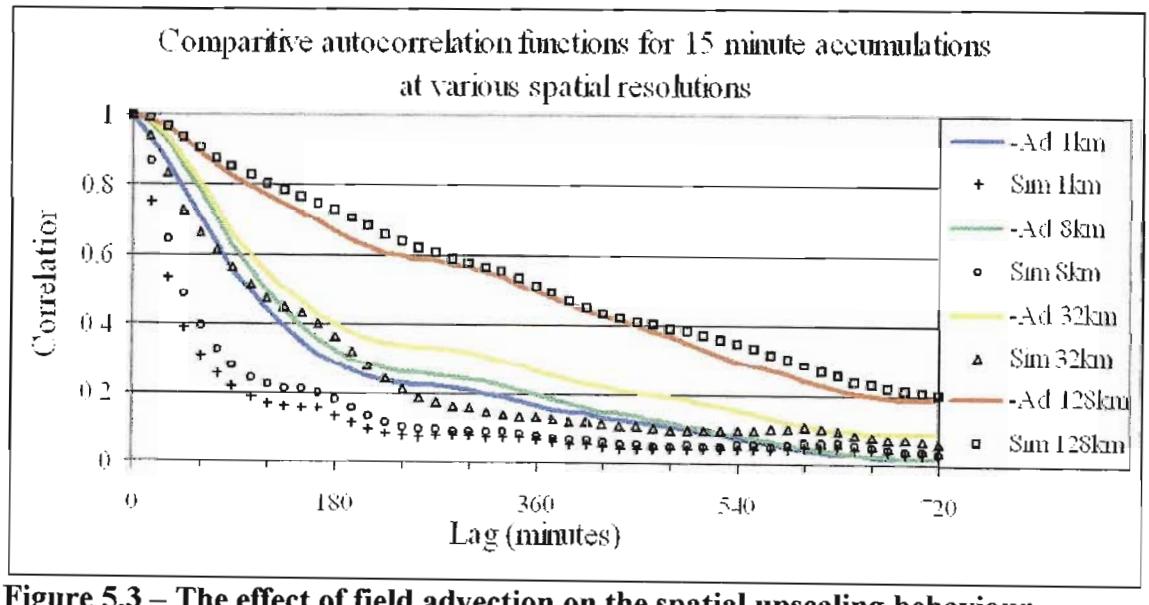

Figure 5.3 – The effect of field advection on the spatial upscaling behaviour

The two simulated sequences analysed in Figure 5.3 also have identical WAR_i and IMF_i image scale statistics, but one of the sequences (labelled "-Ad") is simulated with a zero advection vector. As expected, the 128 km scale temporal correlation is again

faithfully reproduced, however temporal correlation of the higher spatial scales are far higher for the zero-advection case. This emphasises the importance of faithful modelling of the advection component in the simulation.

The next aspect of space-time upscaling considered in Section 3.12, is the spatial correlation structure of the space-time aggregated images. This is characterised by the radially averaged, two-dimensional power spectrum. Figure 5.4 shows the power spectra of five space-time aggregated images sampled at corresponding points from the observed and the simulated aggregated sequences.

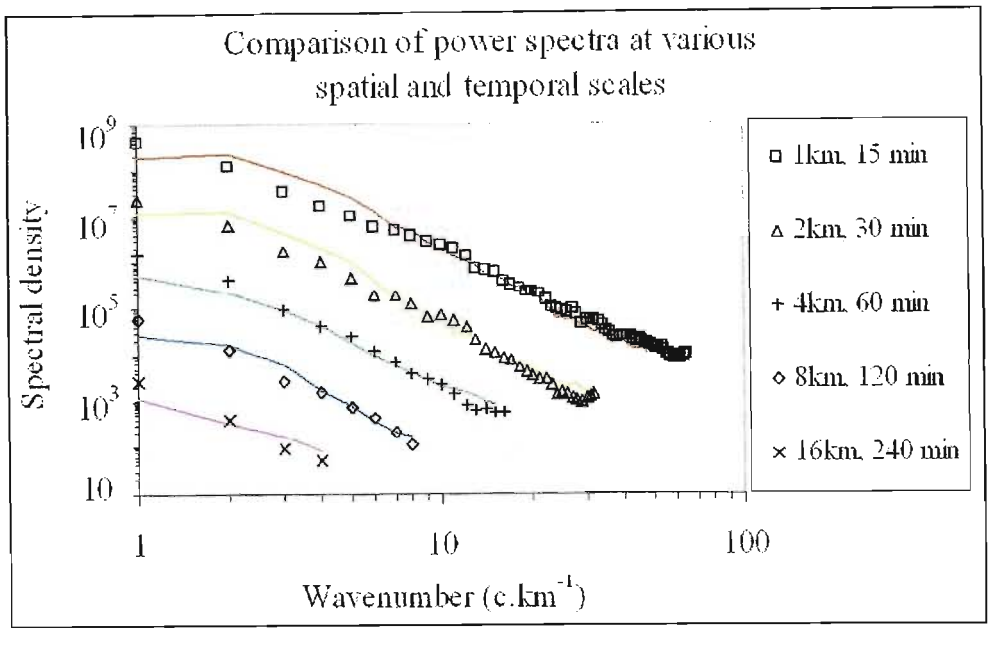

Figure 5.4 – Spatial power spectra for aggregated images (of various spatial and temporal scales) sampled at five corresponding points from the observed (solid lines) and simulated sequences (symbols)

The quality of fit is remarkably good and this is typical of the type of fit observed for most of the aggregated images in both the frequency-filtered and the autoregressive sequences, indicating good spatial upscaling properties. Of particular interest is the flattening off of the observed, aggregated power spectra at the lower wave number which corresponds to a 128 km wavelength. This is quite common in these aggregated images and particularly noticeable in the shorter duration aggregations. It is perhaps an indicator of the breakdown in long range, spatial dependence structure. That is, the limit of the spatial scaling assumption. It was this observation which prompted the technique of setting the edges of the Z(t) field to zero in the third stage of simulation, as shown in Section 4.5.3.

Figure 5.4 only shows the power spectra for one aggregated image sampled from each of the sequences of aggregated images. In order to investigate the performance of the model for all of the images in all of the sequences, power-law relationships were used to approximate their power spectra. Figure 5.5(a) compares the exponents of the power-law relationships fitted to the power spectra of the 1 km aggregated images of the observed sequence, to those of the autoregressive bead sequence, for the 15, 60 and 240 minute temporal scales, while Figure 5.5(b) compares the same to the frequency-filtered bead sequence.

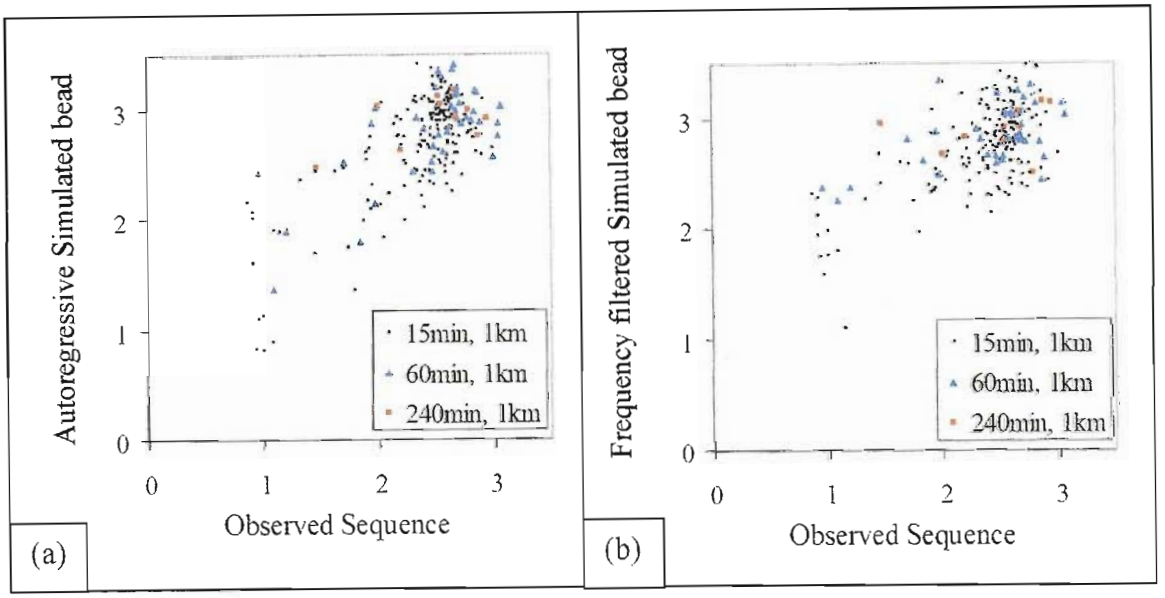

Figure 5.5 – Comparisons of spatial power spectral exponent for observed and simulated sequences, aggregated at 1 km – 15 min, 60 min and 240 min time scales.

The analyses of Figure 5.5 are encouraging since they show that the characteristics of the spatial structure are reproduced, without obvious bias, over a range of temporal aggregations at the 1km spatial scale for both types of bead. The power-law exponent of these aggregated sequences is not much different from that observed for the 1 km – 5 min images at approximately 2.5. Lower values are observed at a stage of the event which is relatively dry, so a large proportion of the aggregated image is zero. Another interesting point apparent in this analysis is the spatial smoothing effect as a result of the temporal aggregation. The average of the spectral exponents for the 1 km - 15 min sequence is 2.5, compared to 2.6 for the 1 km - 60 min sequence and 2.7 for the 1 km - 240 min sequence. Figure 5.6 considers the same analysis at the 8km spatial scale.

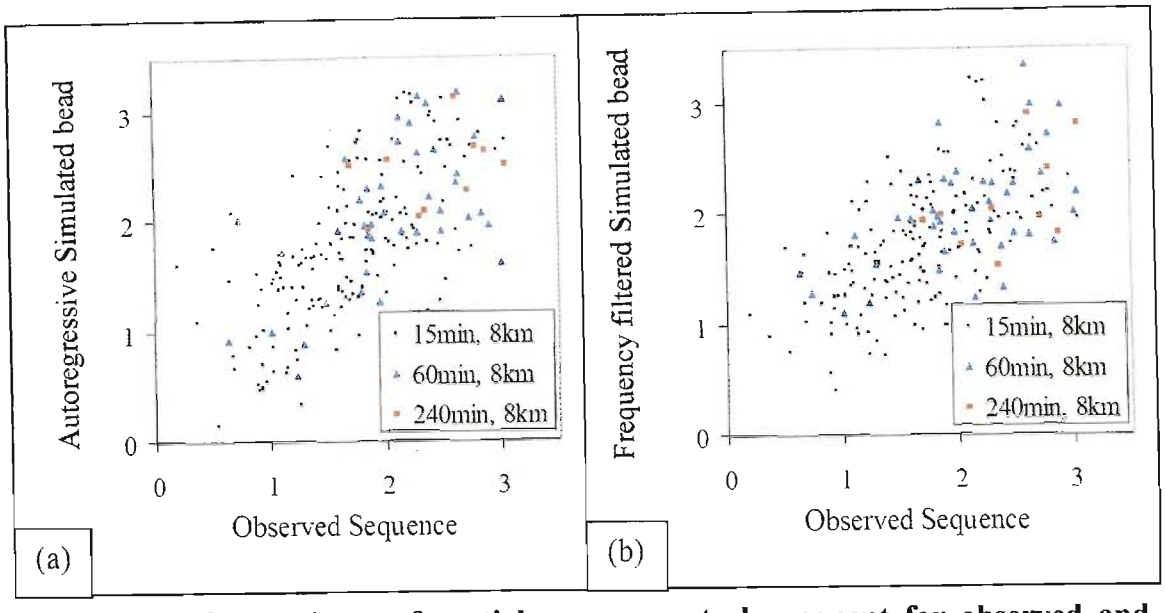

Figure 5.6 - Comparisons of spatial power spectral exponent for observed and simulated sequences, aggregated at 8 km - 15, 60 and 240 minute time scales.

At the 8 km spatial scale, the 128 km image is reduced to 16 x 16, 8 km blocks. Computation of the radially averaged, two dimensional power spectrum from such a small field of data is inaccurate and consequently the scatter of the spectral exponents is far greater at the 8 km scale than at the 1 km scale. Nevertheless, the analysis reveals a noisy, but unbiased representation of the spatial upscaling properties by both the frequency filtered and autoregressive models. The average of the spectral exponents for the 8 km - 15 min sequence is 1.7, compared to 2.2 for the 8 km - 60 min sequence and 2.4 for the 8km-240min sequence. Similar analyses for the 2 km, 4 km and 16 km spatial scales revealed similar trends to those presented in Figure 5.5 and Figure 5.6 for the 1 km and 8 km spatial scales.

In spite of manually fitting the PSI temporal parameters, $\phi_1 \ldots \phi_5$ and the β_{time} parameter, the upscaling analyses presented in this section reveal encouraging results. Automatic calibration of the bead process should result in better approximations for these parameters, further improving the upscaling performance of the *String of Beads* model at the pixel scale. It would appear that the bead process (non-zero space-time rainfall) is well modelled by both the frequency-filtered and the autoregressive processes, however the frequency-filtered methodology appears to out-perform the autoregressive methodology for coarser spatial aggregations.

5.3 THE 'STRING' CALIBRATION PROCESS

It was shown in Chapter 3 that the 54 pixel sample (simulating sampling by raingauges) taken from the radar image, when averaged, gave an imprecise, but unbiased estimate of the mean daily rainfall IMF_d, over the radar image. By a similar argument, it stands to reason that the 54 daily gauge data averaged, will also provide an imprecise but unbiased estimate of the mean daily rainfall on the image. In order to improve the precision of the estimate, longer samples of data should be used. Calibration of the *String* process exploits these observations.

The *String* process is generated through the use of the first and second stages of simulation, namely the event scale and the image scale processes as presented in Sections 4.6.1 and 4.6.2. Output from the second stage of simulation is a continuous time series of WAR_i and IMF_i and the incremental advection vectors in two directions.

Calibration of the *String* is an explicit process of estimating a normally distributed random variable to describe the number of typical rainfall events which arrive in any month (although on review, a Poisson distribution may have been more appropriate). Following the argument of Section 4.5.1, the first step is to determine the depth of rainfall ($\overline{R}_e$) for *each observed event* by aggregation of the IMF_i image scale statistics of the available radar data, according to Equation 4.11. From these data, the expected rainfall for a typical event, $E[IMF_e]$, can be computed. For the 405 observed radar events, $E[IMF_e]$ was estimated to be 3,58mm per event.

Next, the daily raingauge data are aggregated on a monthly basis, IMF_{m_j}, and the number of events which occurred in each month, Ne_{m_j}, is estimated from Equation 4.2.

$$Ne_{m_j} = \frac{IMF_{m_j}}{E[IMF_e]} \tag{4.2}$$

Estimates of the mean and standard deviation of the number of events in any month Me_m and Se_m, are then computed from Equations 4.3 and 4.4 repeated here:

$$Me_m = \frac{1}{Y}\sum_j Ne_{m_j} \tag{4.3}$$

and

$$Se_m = \sqrt{\frac{1}{Y}\sum_j \left(Ne_{m_j} - Me_m\right)^2} \tag{4.4}$$

where: $January \le m \le December$ and $1 \le j \le Y$

For the wet summer months, this Gaussian approach towards describing the distribution of Ne_m is very effective, however for the dry winter months, it should be used with care. As demonstrated in Chapter 3 (Figure 3.7), the normal distribution is not a very good descriptor of the distribution of rainfall totals in the dry months and not surprisingly, this is true also for the Ne_m. In order to achieve a sensible fit, the data require a small amount of pre-processing and an effective approach is simply to disregard any Ne_{m_j} which fall outside of the range

$$^1/_2 Me_m < Ne_{m_j} < {}^3/_2 Me_m \qquad (5.1)$$

The choice of these limits is somewhat arbitrary and comes from a subjective opinion of what appears to be reasonable. It is an aspect of the research which would benefit from some additional investigation. A Poisson distribution for Ne_{m_j} may be a more appropriate and elegant approach, although this has not yet been implemented or tested.

As an example, Table 5.2 presents the results of this initial calibration stage obtained through analysis of the 50 year daily raingauge dataset provided by the CCWR. Both the raw and the pre-processed analyses are shown in order to illustrate the effect of the pre-processing on the Me_m and Se_m parameters.

Table 5.2 -Number of rainfall events in any month

Month	Mean Monthly Rainfall $E[IMF_m]$	Expected number of rainfall events Me_m		Deviation in number of rainfall events Se_m	
		Raw	Pre-processed	Raw	Pre-processed
January	**104.01**	29.03	**28.53**	11.38	**7.34**
February	**83.21**	23.22	**21.20**	10.79	**4.54**
March	**76.58**	21.37	**20.91**	9.53	**4.38**
April	**50.37**	14.06	**13.01**	7.66	**2.12**
May	**21.73**	6.06	**3.13**	6.35	**0.55**
June	**8.89**	2.48	**0.89**	3.15	**0.14**
July	**9.20**	2.57	**0.82**	4.05	**0.08**
August	**15.15**	4.23	**1.81**	5.86	**0.28**
September	**29.68**	8.28	**5.37**	11.02	**1.05**
October	**69.71**	19.46	**17.13**	10.49	**3.55**
November	**86.79**	24.22	**23.50**	10.35	**5.51**
December	**98.91**	27.61	**27.68**	10.25	**6.64**
Total	654.23	182.59	163.98		

The pre-processed Me_m and Se_m presented in Table 5.2 become parameters of the SBM, which are used to generate a normally distributed random number of events $Ne_m \sim N(Me_m, Se_m)$ for any month m. For simulation purposes, Ne_m is rounded to the nearest integer value.

The most significant effect of the pre-processing is to remove the skewness in the distribution of Ne_m so that the remaining dataset can be sensibly described using a Gaussian distribution, however there are side effects. By discarding the rarer events, the expected number of events, $E[Ne_m]$ is reduced slightly and consequently the total number of events in a year is underestimated by the process as reflected in the total of the $E[Ne_m]$ of Table 5.2. To account for this, the randomly generated Ne_m is simply increased in the ratio of the Raw to the Pre-Processed totals of Table 5.2, $^{183}/_{164}$ in this case. Implementation of the Poisson distribution for Ne_{m_j} may eliminate the need for this crude adjustment.

5.3.1 Verification of the String process

Verification of the String is required in order to ensure that the long-term behaviour of the model is consistent with the observed climate. In particular, this is concerned with aggregations of image scale statistics over daily, monthly and yearly time periods, and analysis of these aggregations and comparison to the observed equivalent data recorded by the daily raingauge network. Since these statistics have been used in order to determine the parameters of the *String* process, this is a verification exercise to check that the model is able to reproduce the statistics on which it is calibrated. In essence, this is a measure of the coarse temporal scales of the 128km spatial scale.

In order to verify the *String* process, a 50 year *String* simulation was performed in order to compare the statistics to those of the 1948-1997 raingauge dataset. This is a one-dimensional simulation which can be executed very quickly as it does not include the computationally intensive *Bead* modelling process. The image scale statistics were then aggregated using Equation 4.11, in order to arrive at a 50 year time series of simulated daily rainfall data. These daily totals were then aggregated to arrive at monthly and annual totals.

Perhaps the most fundamental of rainfall statistics is the average annual rainfall for the study area. Figure 5.7 shows one of any number of time series of average annual rainfall for a simulated sequence which can be compared to Figure 3.3. Obviously each random scenario will yield completely different totals, but the statistics of the annual totals should remain reasonably constant.

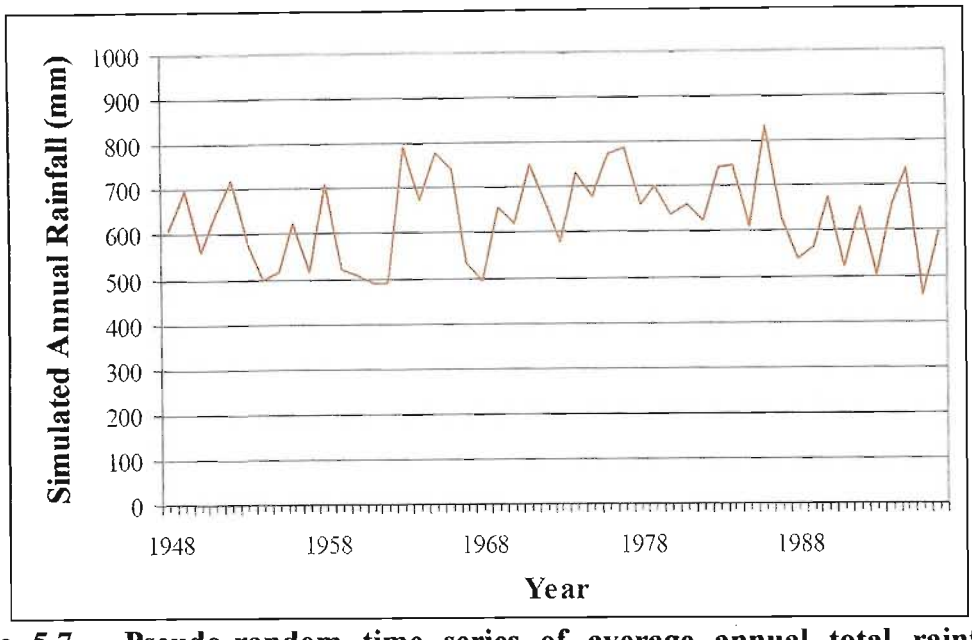

Figure 5.7 - Pseudo-random time series of average annual total rainfall as simulated by the string process

For this particular sequence, the mean annual rainfall was found to be 633mm with a standard deviation of 97mm and a skewness of 0,01 and kurtosis excess of 0,9. Analysis of the 50 year daily raingauge data in Chapter 3 revealed that the observed average annual rainfall for the region is 654mm per annum with a standard deviation of 108mm. The two series have very similar statistics. As for the observed series of Figure 3.3, there is no significant autocorrelation structure in the time series, and the low skewness and kurtosis excess suggest that the normal distribution is a reasonable descriptor of the distribution of annual totals. A small criticism of the time series of Figure 5.7 is that the number of years with less than 500 mm matches the observed sequence, however there is no occurrence of any year with more than 900 mm. This is probably due to cropping of the upper number of events, $Ne_{m_j} < {}^3/_2 Me_m$ which may be eliminated through the use of the Poisson distribution, however this investigation is deferred to later research.

The next temporal scale below annual is the seasonal variation in rainfall, in particular, the distribution of monthly totals. For the same 50 year simulation, the distributions of monthly totals are presented in Figure 5.8.

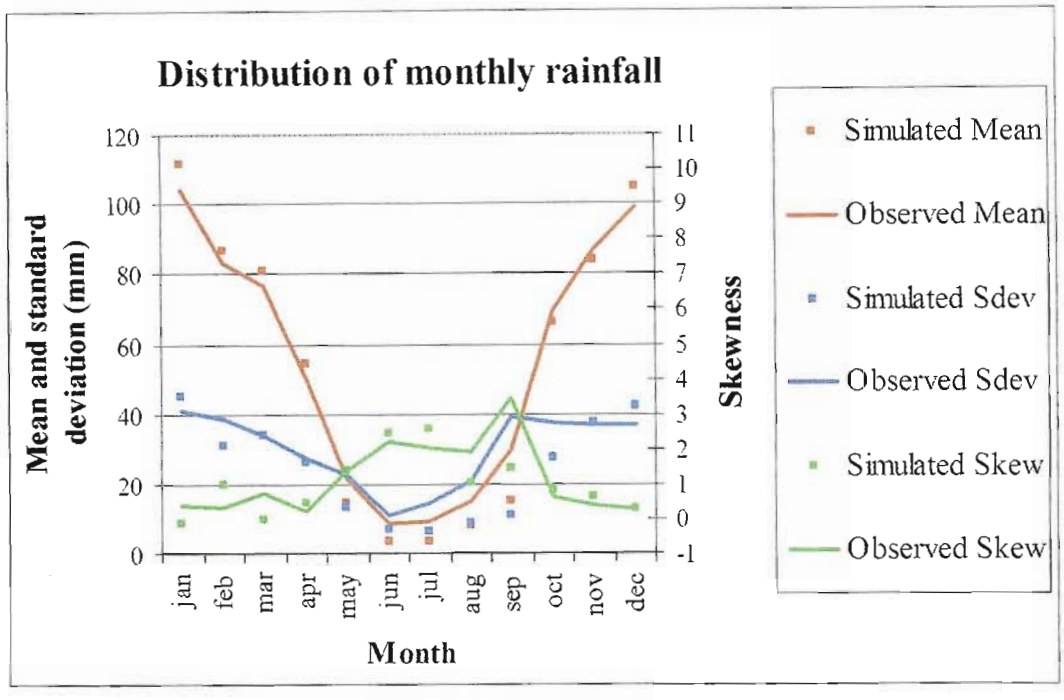

Figure 5.8 - Marginal distribution of monthly rainfall totals for simulated and observed sequences

These are encouraging results. The model has clearly reproduced the seasonal variation in rainfall. The mean and standard deviation serve as verification tests and they are sensibly recovered in the simulation process. Although these values were not input directly, the mean monthly totals were used to determine the mean and standard deviation of the number of events Ne_m for each month. The standard deviation of the months of August, September and October are noticeably lower than the observed values, suggesting that rainfall which occurs in these months might not be accurately represented by the 405 observed events. A possible explanation is that any rainfall experienced in these months is usually the result of large system rainfall – long duration, low intensity rather than the convective activity which is common in the summer months (there are more convective events than stratiform events in the 405 events). Another possibility is that the distribution of the number of events in each month (Ne_m) is not properly described by the cropped normal distribution as mentioned earlier.

Unlike the mean and standard deviation, the skewness was not used in the determination of the model parameters and is reproduced well, with an increased positive skewness in the winter and spring months as seen in the real data. This is useful validation of the normal approximation of Ne_m after filtering out extreme months. It shows that the

extreme rainfall observed in these months is in fact the result of extreme rainfall events rather than an extreme *number* of typical events.

Moving on to the daily statistics, Figure 5.9 shows the probability of exceedence of daily rainfall totals for each of the four seasons.

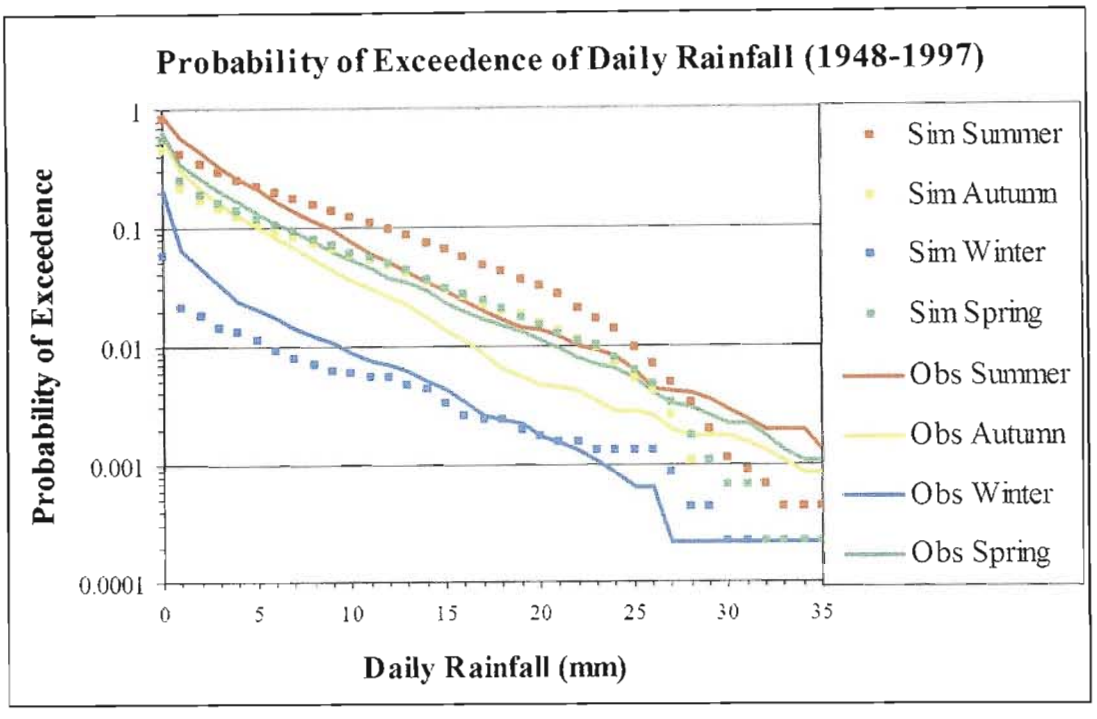

Figure 5.9 - Probability of exceedence of daily rainfall totals for observed and simulated 50 year sequences

Figure 5.9 is plotted to a logarithmic scale in order to critically examine the fit. In normal space, the simulated spring, autumn and winter probabilities are difficult to distinguish from the observed. Probability of exceedence in the summer months indicates that the model underestimates the number of light rainfall days (less than 5mm per day) and over estimates the number of heavy rainfall days. Also, in the winter months, the model underestimates the number of light rainfall days, although the probability of any rain in these months is very low so it does not have any significant effect on the model output. Although the exceedence probabilities are not reproduced exactly, they are reasonably similar. Furthermore, an encouraging aspect of the model is the trend in the statistics. They are the right order of magnitude, they decrease in a similar way and the relative positions of the seasonal grouping are consistent with the observed data.

The fact that these statistics are recovered in the *String* simulation stands testimony to combined validity of the lognormal approximation of the wet spell duration (Figure 3.50), the event renewal process (Figure 4.12), the modelled event scale parameters, M_{war*} , S_{war*}, M_{imf*} and S_{imf*} (Figure 4.13), the bivariate autoregressive WAR_i – IMF_i process (Figure 4.17) and the WAR_i and IMF_i target distributions (Figure 3.56 and Figure 3.57) .

5.4 MODEL FINE TUNING

In spite of the encouraging verification tests for both the *String* and the *Bead* process, the model output is not perfect and considerable benefit can be gained by fine tuning the parameters of the model.

One shortcoming of the research so far is that upscaling calibration of the bead modelling process discussed in Section 5.2.2, has been performed on a single rainfall event which occurred in February 1996. This is true also of the reversed bivariate autoregressive process used to describe the WAR_i-IMF_i time series. This was an extreme event which lasted 42 hours and displayed some unusual super-convective activity. At the time of the analysis, it was the only available significant event recorded by the Bethlehem Weather Radar. The super-convective activity resulted in unusually high intensity for such general rainfall and it is possible therefore that the bivariate correlation structure between the WAR_i and IMF_i time series, and the autocorrelation structure of the pixel scale intensity could have been distorted. Ideally these analyses should be performed for a large variety of observed events.

In the absence of sufficient radar data, coarser tactics can be employed to fine-tune the output of the model by direct manipulation of key parameters. These tactics range in complexity from simple linear scaling techniques to achieve the expected annual total through to complex approaches which manipulate several parameters of the model and attempt to take account of seasonal variation in the rainfall. For the purposes of this study only one of these has been tested, although several alterative options will be discussed briefly.

5.4.1 *Fine tuning techniques*

To begin, it is important to understand how key parameters of the model influence the output and how they may be adjusted. These key parameters are listed as:

- Ne_m can be adjusted to alter the frequency of rainfall events in any month. Scaling the expected number of events in any month will linearly alter the total amount of rainfall received in that month in this version of the SBM, since the bead process is assumed to be independent of season.
- D_w can be adjusted to alter the average duration of a rainfall event. Scaling of this parameter results in a non-linear reaction from the model, since this has an effect on all of the event scale parameters of the model as demonstrated in Figure 4.13. Nevertheless, it is a very effective means of altering the model output. Adjustments could be applied at a monthly scale to account for seasonal variation in the type of rainfall, or generally for season-independent changes.
- M_{imf^*} can be scaled to adjust the average intensity of the rainfall in an event. Increasing this parameter has the effect of making storms more intense with higher instantaneous rainfall rates. Adjustment of this parameter results in a non-linear reaction from the model since M_{imf^*} is defined in transformed space with reference to Section 3.10.2.
- M_{war^*} can be scaled to adjust the extent of the rainfall in an event. Increasing this parameter results in more general rainfall on the images. This parameter should be manipulated with care since the other event scale statistics are affected according to the relationships defined in Equations 4.5 through 4.8.
- The IMF_i time series can be adjusted in real space. With a fixed WAR_i, an increase of the IMF_i time series will linearly increase the aggregated output of the model and will give rise to more intense storm activity at all temporal scales. The converse is also true if less intense storm activity is required. Adjustment of this parameter is effectively altering the target distribution for IMF_i which is presented in Figure 3.57.
- The WAR_i time series can also be adjusted in real space. With a fixed IMF_i, this will not alter the aggregated output of the model such as daily, monthly and yearly totals, which makes the WAR_i a useful parameter for changing the *nature* of the simulated rainfall. Increasing the WAR_i results in more general, less intense rainfall (stratiform rainfall) and conversely reducing WAR_i results in

isolated, more intense rainfall (convective activity). Obviously, there is a limit to the amount that WAR_i can be used, since WAR_i is only defined on the range $0.01 < WAR_i \leq 1.00$. Adjustment of this parameter is effectively adjusting the target distribution for the WAR_i presented in Figure 3.56.

- The field advection. Increasing the speed of the storm results in a more spread and streaky cumulative rainfall pattern whereas reducing it results in isolated, intense cumulative rainfall patterns.

Adjustment of any of these parameters is simply a process of allowing the model to produce rainfall simulations which are consistent with the observed climate. Ideally, these should be *small* adjustments, however the more limited the available dataset, the more dramatic will be the fine tuning required to achieve the desired results.

5.4.2 Fine tuning for the Bethlehem climate

Several one-year, three-dimensional (128km x 128km, 5 minute resolution) simulations were generated to obtain five minute, daily, monthly and annual rainfall aggregations in space and time. These daily and monthly cumulative images were sampled at 54 positions which correspond to the 54 daily raingauges in the catchment, and their statistics were examined. Initial results of the simulation revealed some significant limitations in the modelling process, but by fine-tuning the model, it was possible to significantly improve the comparative statistics and the spatial distribution of rainfall on the daily cumulative images.

Of the techniques discussed in Section 5.4.1, it is a combination of the D_w, the Ne_m parameters and field advection which was selected for the fine-tuning of the model in this study. Adjustment of these parameters is no more or less valid than adjustment of any of the others, but they were selected because they do not interfere with the generation of the *Bead*. Adjustments to the WAR_i and IMF_i parameters or to the M_{war*} and M_{imf*} parameters can result in undesirable results. In particular they can give rise to unrealistic images at the 5 minute resolution such as extreme instantaneous rainfall leading to large "black spots". On the logarithmic colour palette defined in Figure 2.10, black pixels indicate rainfall in excess of 100mm/h. Although this intensity has been observed at isolated places on the radar images, it is unusual for large areas to experience storms of this intensity.

The comparisons were generally acceptable, but there were several instances of extreme daily rainfall which needed to be addressed. The absolute values of the extreme daily rainfall were not unrealistic, but they appeared too frequently in the sampled 54 pixel average. An example of an extreme daily rainfall event is given in Figure 5.10, in which the IMF_d is measured to be 45mm.

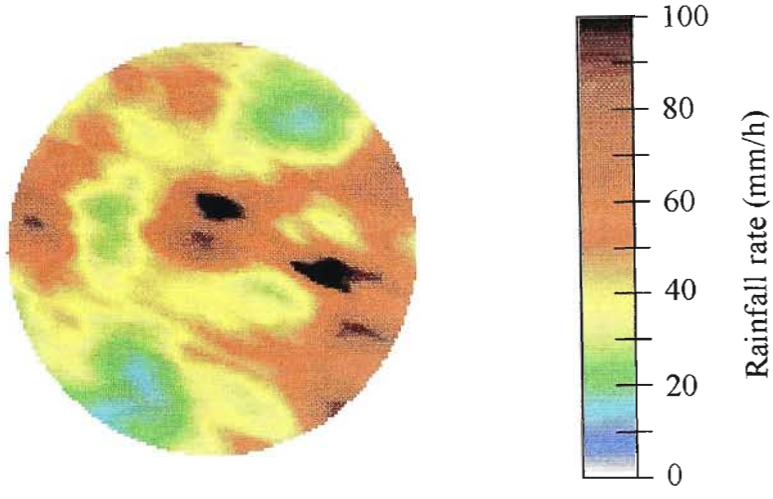

Figure 5.10 - An example of extremely intense (poor) daily rainfall aggregation

Although the IMF_i time series was not adjusted, one condition which has been imposed on the IMF_i time series is that the total rainfall averaged over the image in any event IMF_e should not exceed 30mm. If aggregated image scale statistics for the event exceed this threshold, the WAR_i–IMF_i bivariate process were re-sampled. This does not necessarily limit the average daily rainfall on the image to 30mm, since it is possible for more than one event to occur on a single day which would result in occasional daily rainfall totals in excess of this limit.

In addition to this limit, another condition was applied to the event scale statistics, this time to the event renewal process. Consecutive events with duration D_w in excess of 10 hours are not permitted to begin within 24 hours of each other. If this occurs in a month, the renewal process is re-sampled for that month. In the observed rainfall in the Bethlehem area, large systems pass over the region usually within a day. The next large system is unlikely to occur for several days afterwards. Convective activity generally dies out within the 10 hour period and is allowed to recur within 24 hours. This condition is specific to the Bethlehem climate and need not be applied for the calibration of the model on other datasets.

Elimination of the extreme events (IMF_e > 30mm) leaves a deficit in annual rainfall totals which must somehow be compensated for. To this end, the D_w and Ne_m parameters were adjusted (independent of season) to achieve the expected annual total. The idea was to increase the expected number of events in any month and also to increase the expected duration of a random event in any month. Although it is possible to achieve the same result by adjustment of either one of these parameters by a large amount, instead it was decided to increase both parameters by a smaller amount. Since annual simulations run over a period of approximately 6 hours on a modern PC, there is a limit to the number of tests which could be run. The most convincing result was obtained by increasing the expected number of events Ne_m by 7,5%, while increasing the D_w by 25%. After making these adjustments, another one-year simulation was run and the 54 pixel sample extracted and compared to an observed rainfall year with a similar annual rainfall total. As it turned out, the annual rainfall of the pseudo random simulation was measured to be 745mm – a wet year by Bethlehem standards. The most similar observed year amongst the 50 years of daily rainfall record was 1996 in which the region received 789mm. Figure 5.11 shows the time series of the 54 pixel sample from the simulated daily rainfall to that of the 1996 daily raingauge data

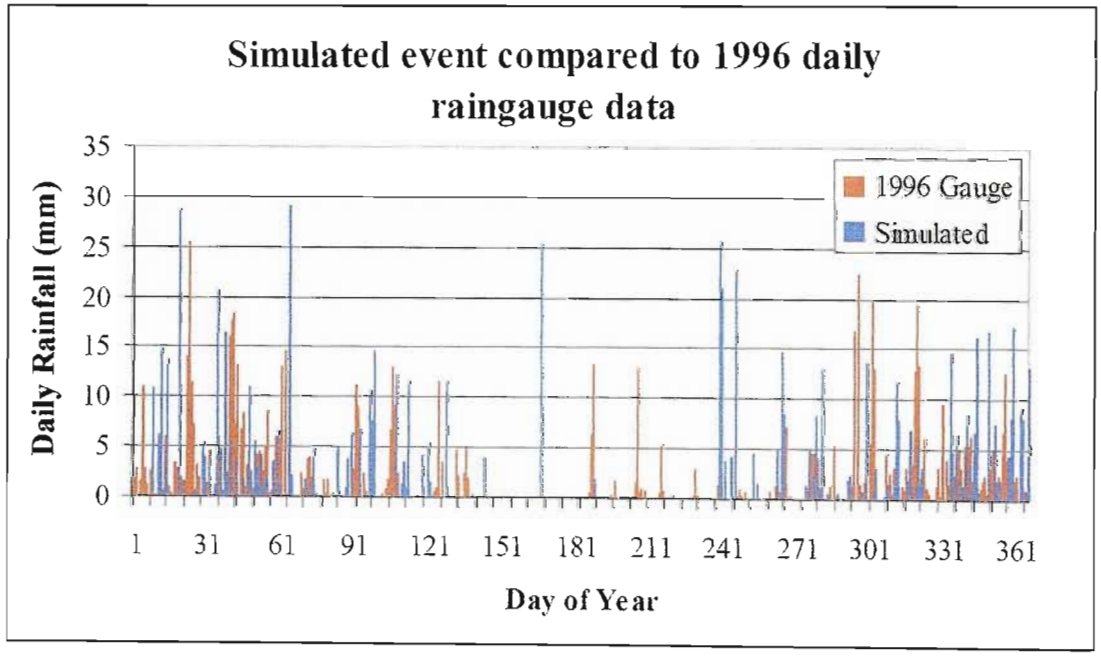

Figure 5.11 - Comparison of 54 gauge daily rainfall totals for observed and simulated years

The daily time series of Figure 5.11 serve to illustrate the distribution of rainfall during the year. There are very few wet daily rainfall records in the winter months, most of the rain occurs in the summer months. Dry spells between wet days are few and far between in the summer months becoming progressively more frequent in the autumn and spring.

A comparison of the marginal distribution of the daily rainfall is also encouraging, although it could possibly benefit from some additional fine tuning. Rainfall was recorded on 224 days in 1996 compared to 172 in the simulated data. In order to compare the marginal distribution, daily rainfall totals were ranked and are compared in Figure 5.12.

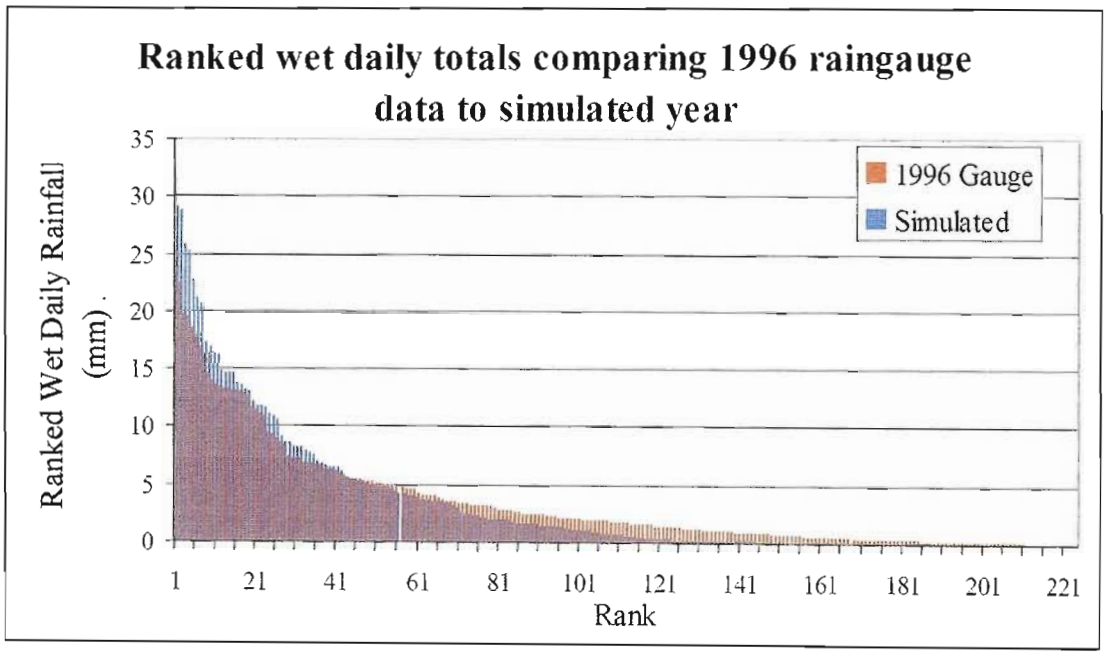

Figure 5.12 - Comparison of ranked daily rainfall for the 1996 daily raingauge data and a simulated year

The results of Figure 5.12 clearly indicate that the 1996 daily data comprises a larger number of wet days (210 as opposed to 144 for the simulated) and the rainfall depths of the 45 *most intense* days is higher for the simulated than the observed. The model over-estimates the number of heavy rain days and underestimates the number of light rain days, although these differences might be the result of random sampling. More simulations are required in order to confirm this.

To increase the number of general, light rainfall days, one possible fine-tuning approach would be to increase the expected number of events in any month (Ne_{m_j}), increase the WAR_i time series and reduce the IMF_i time series.

Another observation in the simulated data which required some attention was that some of the daily cumulative images appeared to lack large scale structure expected in these aggregations. Instead of large areas of similar rainfall, these comprised a large number of small intense areas – "blotchy images". An example of this type of image is given in Figure 5.13

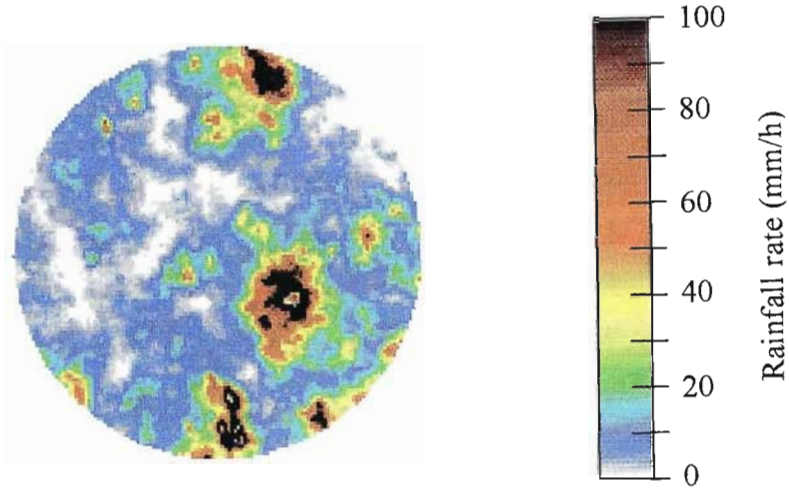

Figure 5.13 - An example of a daily cumulative image which lacks long range spatial structure

To eliminate these unrealistic spatial distributions of rainfall, the "blotchy" images, was a somewhat simpler task. For the field advection of events longer than 5 hours, a minimum speed of advection was set to 10km/h. Large events which move too slowly result in extreme rainfall over a small area. For convective activity this can occur, but stratiform activity is the result of large scale atmospheric process which are usually moving reasonably quickly.

It may be worth introducing a more intricate dependence structure between the duration of the storm and the speed at which it propagates. Nevertheless, this simple condition has proved very effective in the elimination of these unrealistic artefacts. Apart from realistic spatial distribution of rainfall, this also eliminates the large 'black spots' which are clearly evident in Figure 5.13.

An ideal verification test would be to perform a 50 year, three-dimensional simulation (128km x 128km, 5 minute resolution aggregated to daily images) and compare the

daily statistics to the 54 gauge dataset, however this simulation would take approximately 2 weeks on available computing equipment. For the purposes of this study, it has been demonstrated that the daily statistics can be fine-tuned to achieve the desired results.

The final verification of the *String* output which will be considered in this study is the marginal distribution of the image scale statistics for the instantaneous images. By making use of the 1999 radar dataset, it is possible to draw a comparison between the marginal distributions of the descriptive statistics of the instantaneous images. In particular, the marginal distributions of WAR_i, IMF_i. Figure 5.14 presents the marginal distribution of the observed *wet* WAR_i ($WAR_i > 1\%$) measured from the 1999 radar dataset, compared to those of a simulated equivalent year.

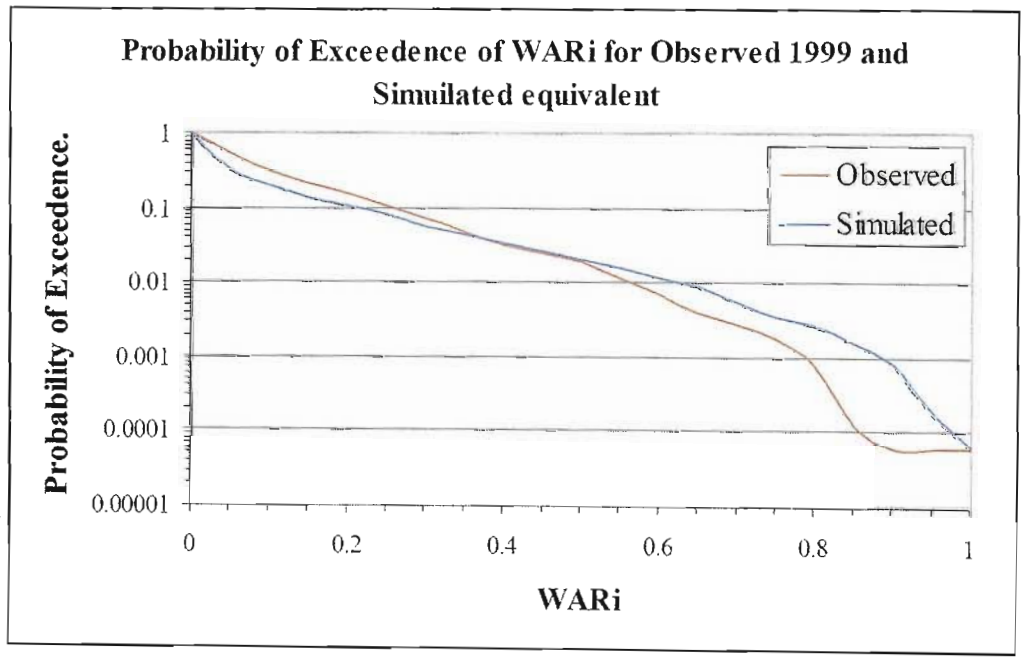

Figure 5.14 - Comparison of marginal distribution of WARi for the 1999 radar dataset to that of an equivalent simulated year

It is evident from Figure 5.14 that the marginal distribution of WARi can be sensibly recovered from a simulated annual record. From a critical standpoint, the number of isolated light rainfall is somewhat underestimated, although this could arguably be the result of random re-sampling – the simulated event comprised more general events and less isolated events. This could be altered by some fine tuning of the WAR_i time series,

however it would appear to be unnecessary in this instance. For the same annual sequence, the marginal distribution of the IMFi is presented in Figure 5.15.

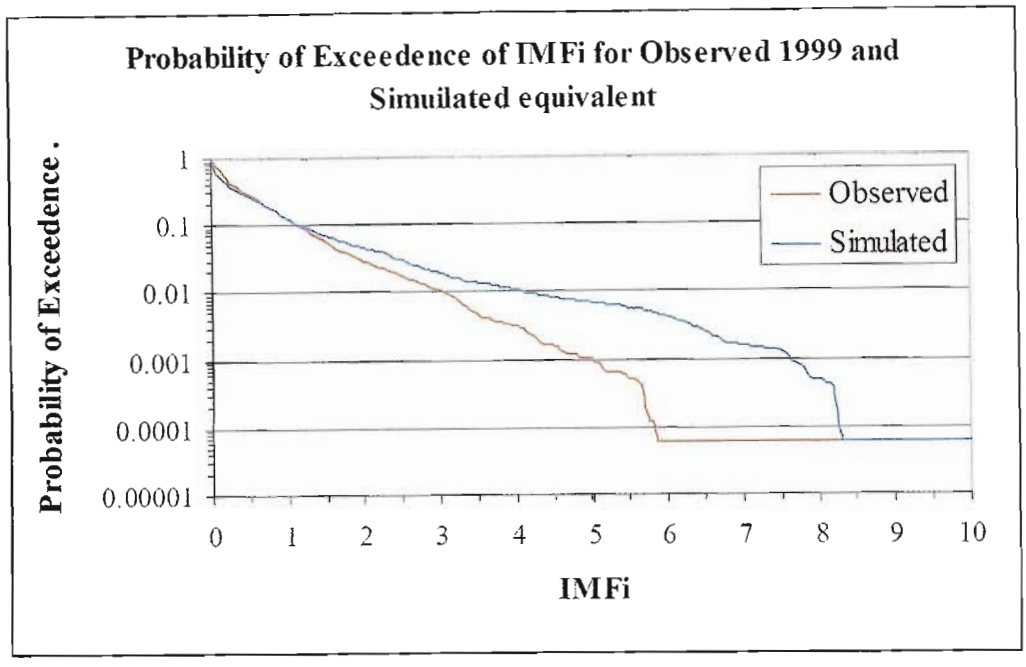

Figure 5.15 - Comparison of marginal distribution of IMFi for the 1999 radar dataset to that of an equivalent simulated year

On the natural scale, the probability of exceedence of IMF_i for the observed 1999 data and the simulated event are virtually indistinguishable. On the logarithmic scale however, it would appear that the simulated sequence slightly over estimates the IMF_i at the upper end of the range beyond the 1% exceedence probability. Again, this could be adjusted if necessary by manipulation of the IMF_i time series, but there appears to be no real need in this case.

The fine-tuning process has not been automated as it really requires user intervention and judgement to be applied. Although the initial results presented here are encouraging, the string calibration process is an aspect of the research which would benefit from further investigation.

5.5 SUMMARY

This chapter has demonstrated the processes whereby the *Bead* and *String* process can be calibrated, verified and fine-tuned. The results of the upscaling tests of the *Bead* modelling process are very encouraging over all of the observed space-time scales for

both the autoregressive and the frequency-filtered beads, verifying the *Bead* modelling process.

The *String* process, which combines the lognormal approximation of the wet spell duration (Figure 3.50), the event renewal process (Figure 4.12), the modelled event scale parameters, M_{war*}, S_{war*}, M_{imf*} and S_{imf*} (Figure 4.13), the bivariate autoregressive WAR_i – IMF_i process (Figure 4.17) and the WAR_i and IMF_i target distributions (Figure 3.56 and Figure 3.57), have been shown to adequately reproduce the daily, monthly and annual rainfall statistics when compared to the 50 year daily rainfall record for the region.

Finally, techniques of model fine tuning have been discussed and demonstrated in order to manipulate the model output so that it accurately reproduces the daily climate over a wide range of spatial and temporal scales.

Chapter 6

Model Output and Validation

6.1 INTRODUCTION

The purpose of this chapter is primarily to present examples of the model output and wherever possible, to devise tests which can validate the model. Where the verification of the model presented in Chapter 5 confirms that the intrinsic statistics which form the parameters of the model can be recovered by sampling the model output, the validation tests consider some extrinsic properties, not used in parameterisation of the modelling process. As a final test of the validity of the modelling process, the short-term forecasting abilities of the SBM are briefly demonstrated.

6.2 INSTANTANEOUS IMAGE SPATIAL STRUCTURE VALIDATION

Although the gradient of the radially averaged two dimensional power spectrum *(β)* is a sufficient descriptor of the spatial correlation structure of the rainfall intensities on an image, an important tool for validating the structure of the instantaneous images is the *Generalised Structure Function* introduced in Section 3.5.3. Unlike the power spectrum and the marginal distribution, the *Generalised Structure Function* is not used to determine the input parameters of the String of Beads model and it therefore serves as a useful validation test for the model output. Figure 6.1 presents an example of a real image and its simulated equivalent with the same statistics used to generated the simulated image. The re-sampled β_{space} is affected by the masking.

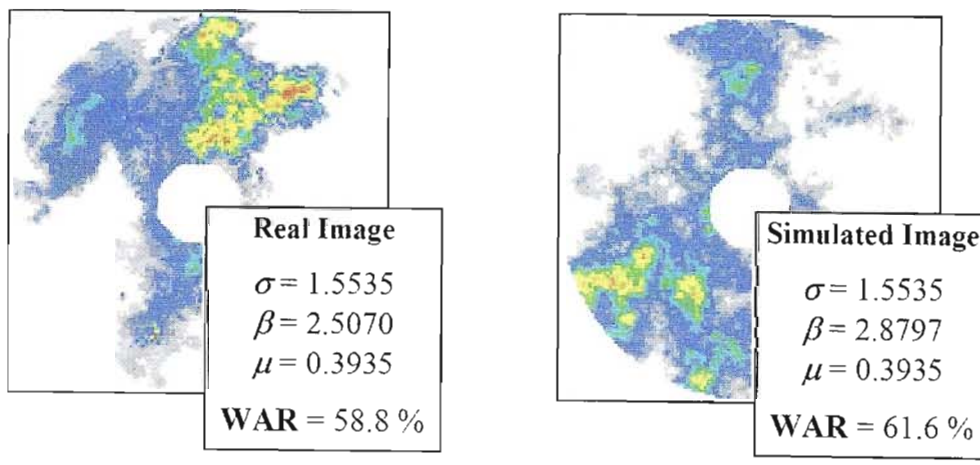

Figure 6.1 - Example of a real and equivalent simulated image

The $\zeta(q) \sim q$ functions of the real and simulated *double precision* images which were shown in Figure 6.1, are given in Figure 6.2. The two sampled $\zeta(q)$ functions compare reasonably well out to a range of $q = 5$. A difference appears at the lower end of the curve near the origin. In the simulated case the function plots nicely through the origin as expected, however in the real case the function appears to be discontinuous near the origin. This is an artefact of the radar data received from Bethlehem which is recorded to a precision of 0.1 dBZ. The truncation of the data in this way causes the introduction of some zero values and consequently the $\zeta(q)$ function to bend upwards at low values of q, with a discontinuity at the $\zeta(q)$ axis as explained in Section 3.5.3. This effect becomes more severe as the precision of the data becomes coarser.

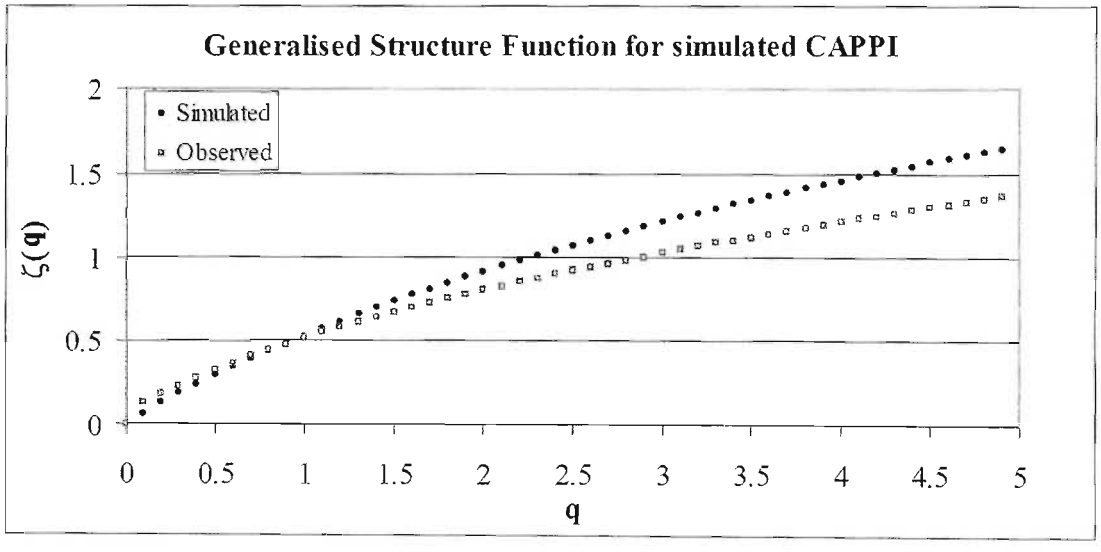

Figure 6.2 - GSF Exponent for Real and Simulated images of Figure 6.1

As discussed in Chapter 3, Menabde et al (1999) show that for a two dimensional non-stationary random field, the gradient β_{space}, of the radially averaged, two dimensional power spectrum (one of the three parameters of the String of Beads model) is related to $G_2(l)$ by Equation 3.20.

$$G_2(l) \propto l^{\beta - D} \tag{3.20}$$

Which implies that when $D = 2$ (as is the case for image data):

$$\beta = 2 + \zeta(2) \tag{3.21}$$

Figure 6.2 gives values of $\zeta(2)$ equal to 0.80 and 0.91 respectively, implying a β_{real} of 2.80 and a $\beta_{simulated}$ of 2.91. These compare quite well to the measured values of 2.51 and 2.88 presented in Figure 6.1.

The validity of comparing the spatial structure of two *integer precision* images using the Generalised Structure Function is questionable but is done here nevertheless for the five images of Figure 3.15 and their simulated equivalents. Plots of the gradients of the Generalised Structure Functions for these 10 images are presented in pairs of real and simulated images in Figure 6.3 through to Figure 6.7. Note the exaggerated step at the origin due to the integeristion of the data. Note also that the curves of the real and simulated images can lie either side of each other. This is due to the sampling of the Generalised Structure Function.

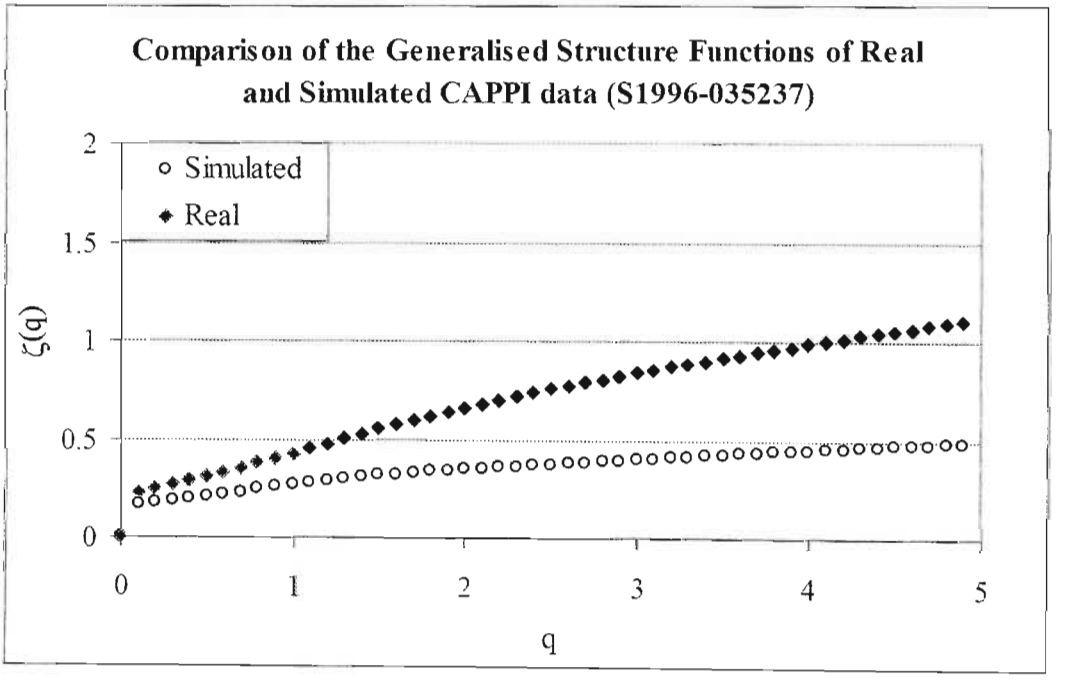

Figure 6.3 - Comparison of the GSF gradients of the real and simulated integer precision images of Figure 3.15(a)

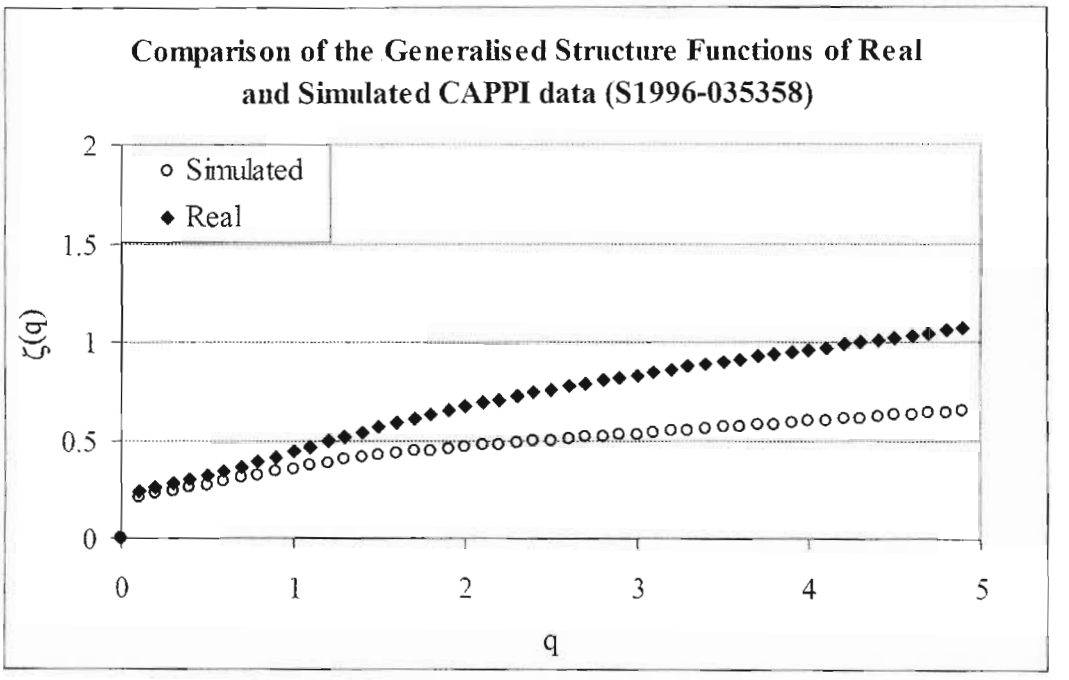

Figure 6.4 - Comparison of the GSF gradients of the real and simulated integer precision images of Figure 3.15(b)

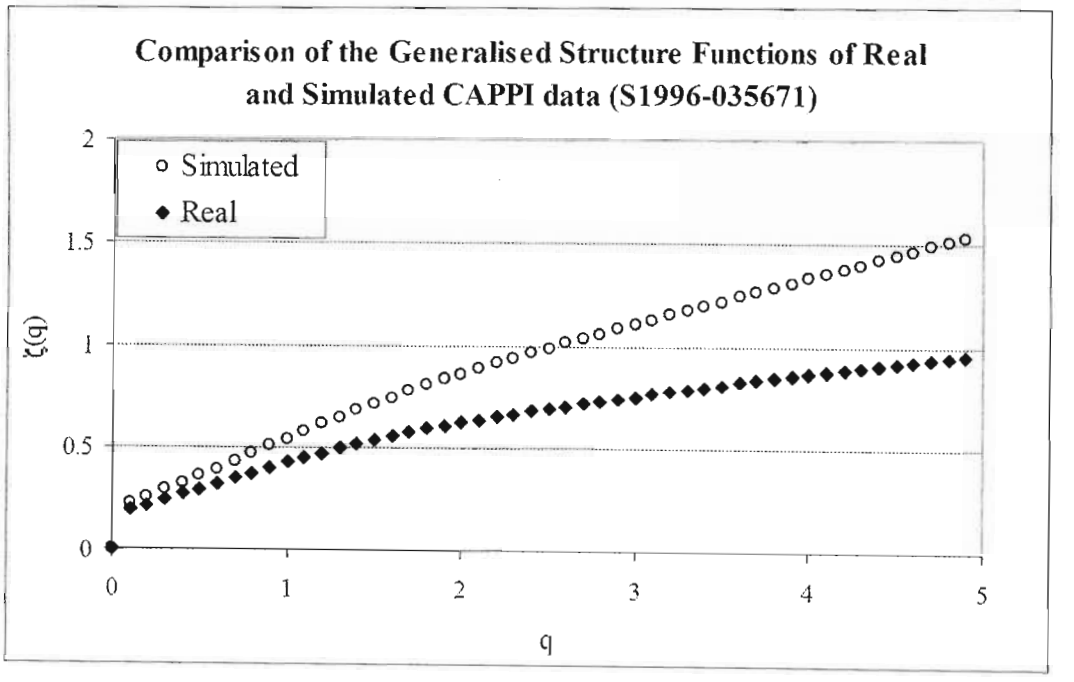

Figure 6.5 - Comparison of the GSF gradients of the real and simulated integer precision images of Figure 3.15(c)

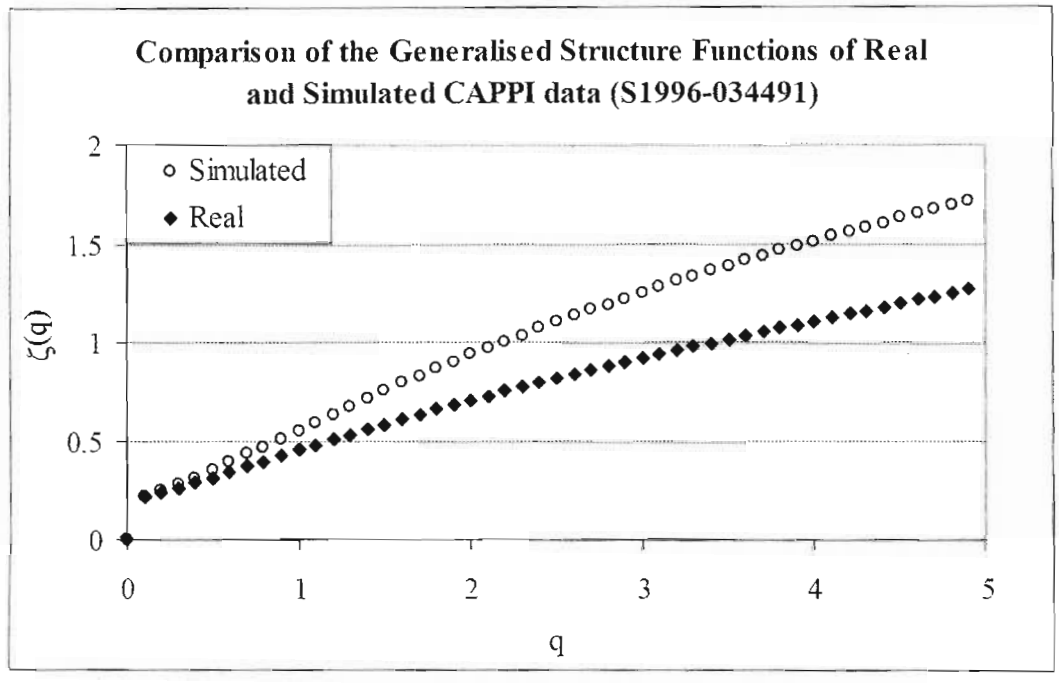

Figure 6.6 - Comparison of the GSF gradients of the real and simulated integer precision images of Figure 3.15(d)

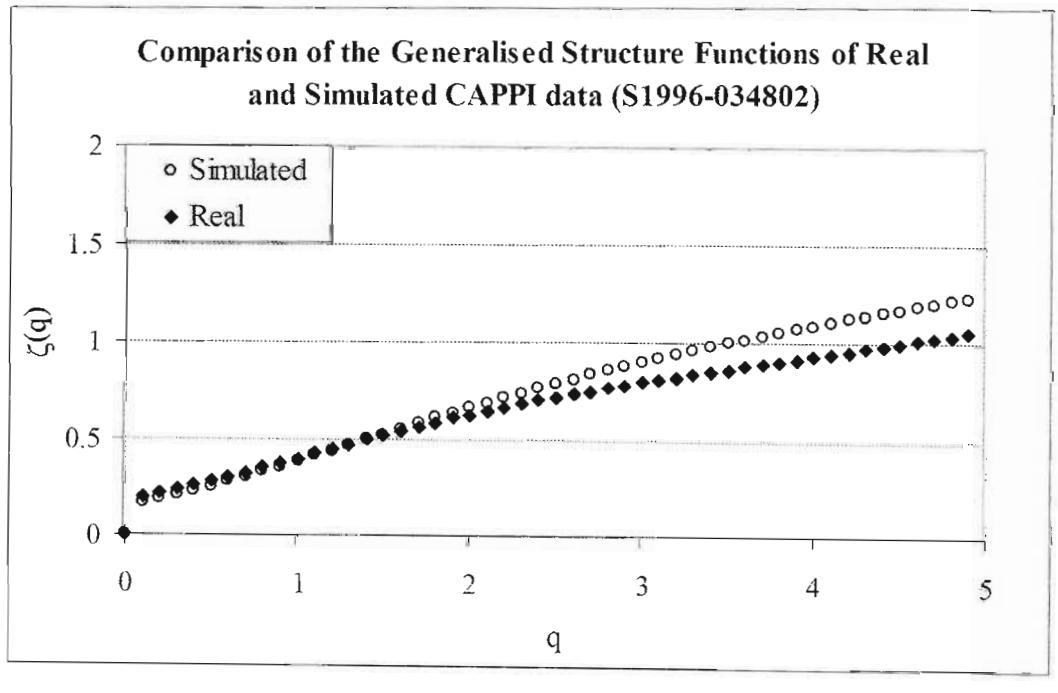

Figure 6.7 - Comparison of the GSF gradients of the real and simulated integer precision images of Figure 3.15(e)

The $\zeta(q)$ functions plotted in Figure 6.3 through to Figure 6.7 suggest that the Generalised Structure Function is better approximated for a higher WAR_i or for a lower number of zero values. The worst case is presented in Figure 6.3 for the image 035237 which has the lowest WAR_i and the best case is that of Figure 6.7 for image 034802 which has the highest. Apart from Figure 6.3, the $\zeta(q)$ functions are reasonably well approximated up to a q value of 1.0 but get progressively worse for the higher order moments.

When considering the *integer precision* radar data that has been analysed in this study, it is not possible to draw any firm conclusions regarding the spatial structure of the simulated (integer) images by the use of the Generalised Structure Function. With the limited amount of *double precision* radar data analysed thus far, the String of Beads model appears to perform well when simulating *double precision* images.

6.3 DESCRIPTIVE STATISTICS OF CUMULATIVE IMAGES

The descriptive statistics of the instantaneous images, WAR_i, IMF_i and β_{space}, are all intrinsic to the SBM and the verification tests of Chapter 5 have demonstrated that these are recoverable from the model output. Extrinsic to the modelling process however, are some descriptive statistics of the daily, monthly and yearly images, namely the WAR and the exponent of the radially-averaged two-dimensional power spectra. The IMF is intrinsic since it was used in the fine-tuning process discussed in Chapter 5. These two statistics therefore serve as a useful validation test for the model output. To this end, a one-year simulation was performed and the daily, monthly and yearly cumulative images were recorded.

6.3.1 Annual images

The annual cumulative images for the observed and simulated radar data are presented in Figure 6.8. The results of Figure 6.8 are very encouraging. Since the two modelled sequences were chosen to have similar annual totals to the observed, the IMF of the images will obviously be similar (620mm, 633mm, 650mm for the observed, *autoregressive* and *frequency-filtered* respectively), however it is the spatial distribution of rainfall which is important in this case. Qualitatively, this appears to be realistic and it is confirmed by comparison of the spatial power spectra of the images. The spectral

exponent (β_{space}) of the observed annual aggregation was found to be 2.53, while that of the *frequency-filtered sequence* is 2.72 and for the *autoregressive sequence* 2.68.

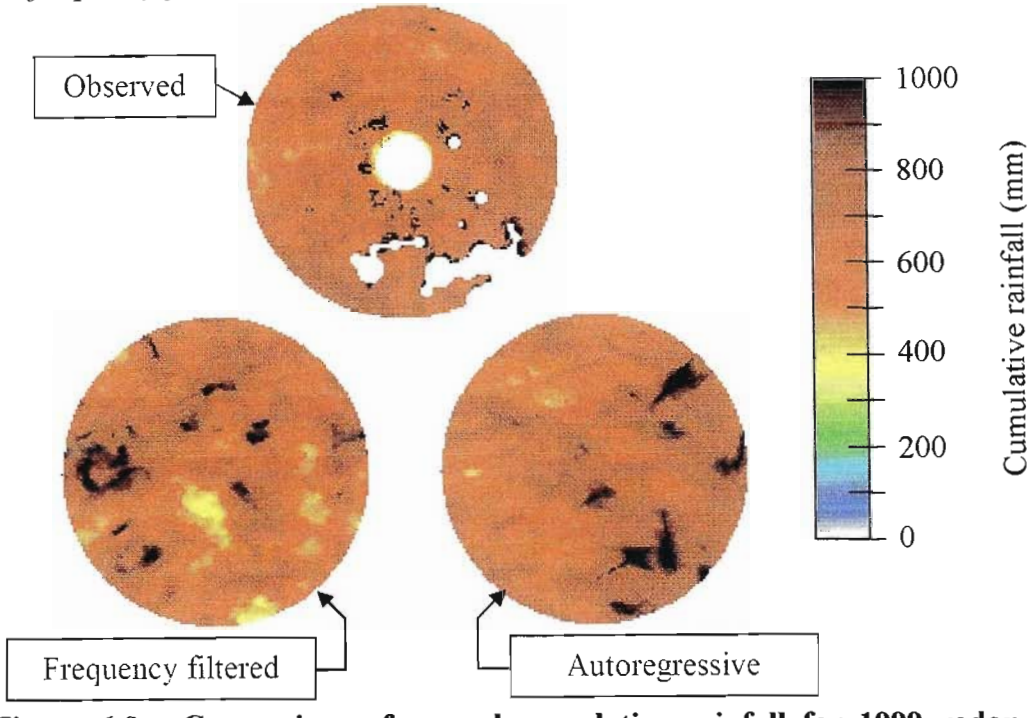

Figure 6.8 - Comparison of annual cumulative rainfall for 1999 radar data, a simulated sequence using the frequency filtered bead and a simulated sequence using the autoregressive bead

Since the simulated sequences are modelled using spatially homogeneous processes, the spatial variation in the image is the result of statistical sampling. Consequently, over very long simulations, these fields are expected to converge to a uniform value (within statistical uncertainties). By contrast, the observed rainfall field will tend to converge to a spatial structure which is consistent with the orographic and climatic influences on the rainfall field. Since the topography in the Bethlehem area is relatively flat, orographic effects on the rainfall field are assumed to be minimal apart from the areas near the Maluti Mountains which have been masked out for the purposes of this study. The study area has also been assumed to be small enough to discount spatial variation in the rainfall due to climatic variation. The only available estimate of the spatial variation of the rain field over a long period of time is obtained from the 54 gauge dataset. The standard deviation of the *average annual rainfall* recorded at the 54 gauges over the 50 year period is found to be 90mm, with a mean of 652mm. By interpolating a surface between these points (using multiquadrics or some other suitable method) it would be possible to warp the simulated field so that it converges to the shape of the observed raingauge data. This is left as a topic for future research.

6.3.2 *Monthly images*

Moving on to a finer temporal scale, Figure 6.9 , Figure 6.10 and Figure 6.11 show the monthly aggregations of the observed 1999 data, the frequency-filtered simulation and the autoregressive simulation.

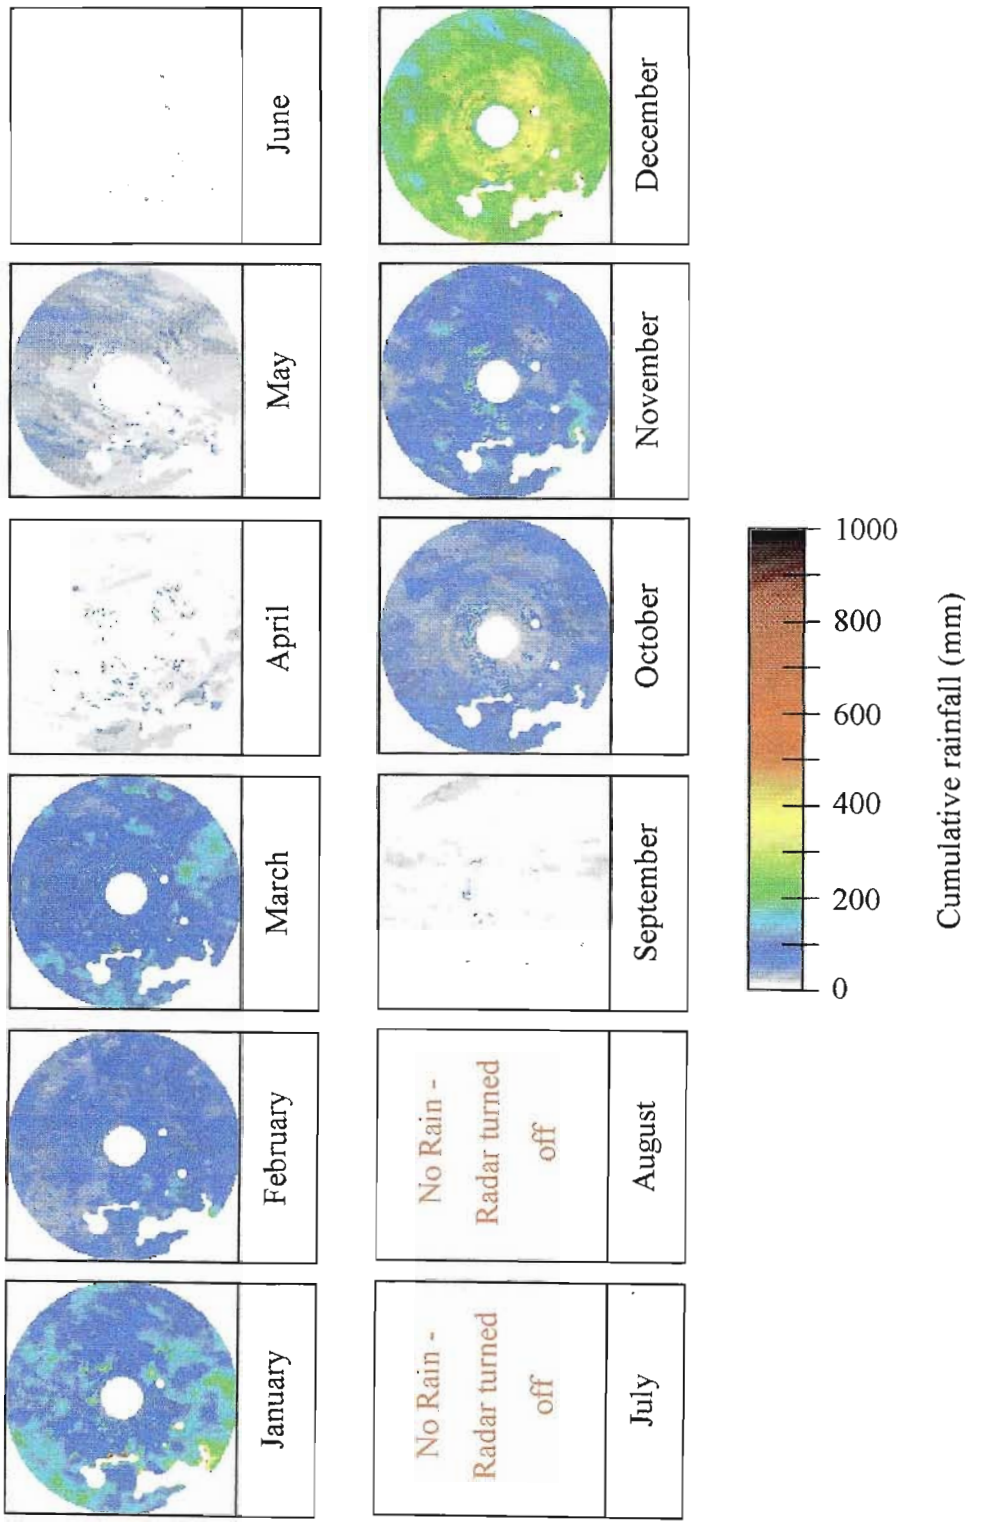

Figure 6.9 - Monthly aggregations of observed radar data for 1999

Image aligned for flip comparison to Figure 6.9 and Figure 6.11.

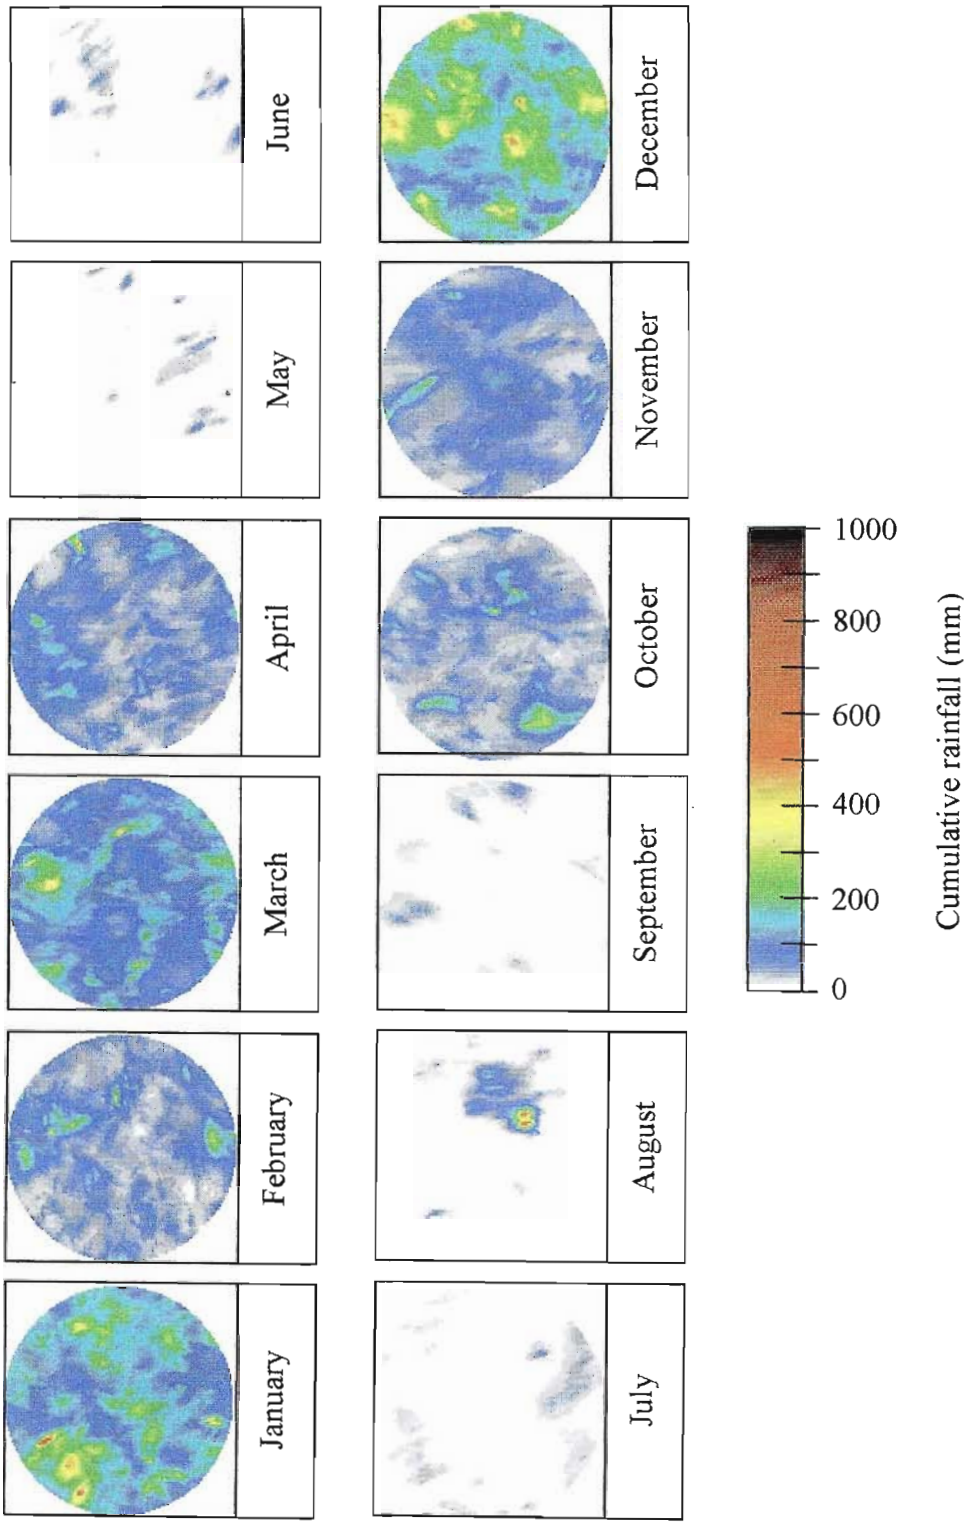

Figure 6.10 - Monthly aggregations for the simulated sequence using the frequency-filtered bead

Image aligned for flip comparison to Figure 6.9 and Figure 6.10.

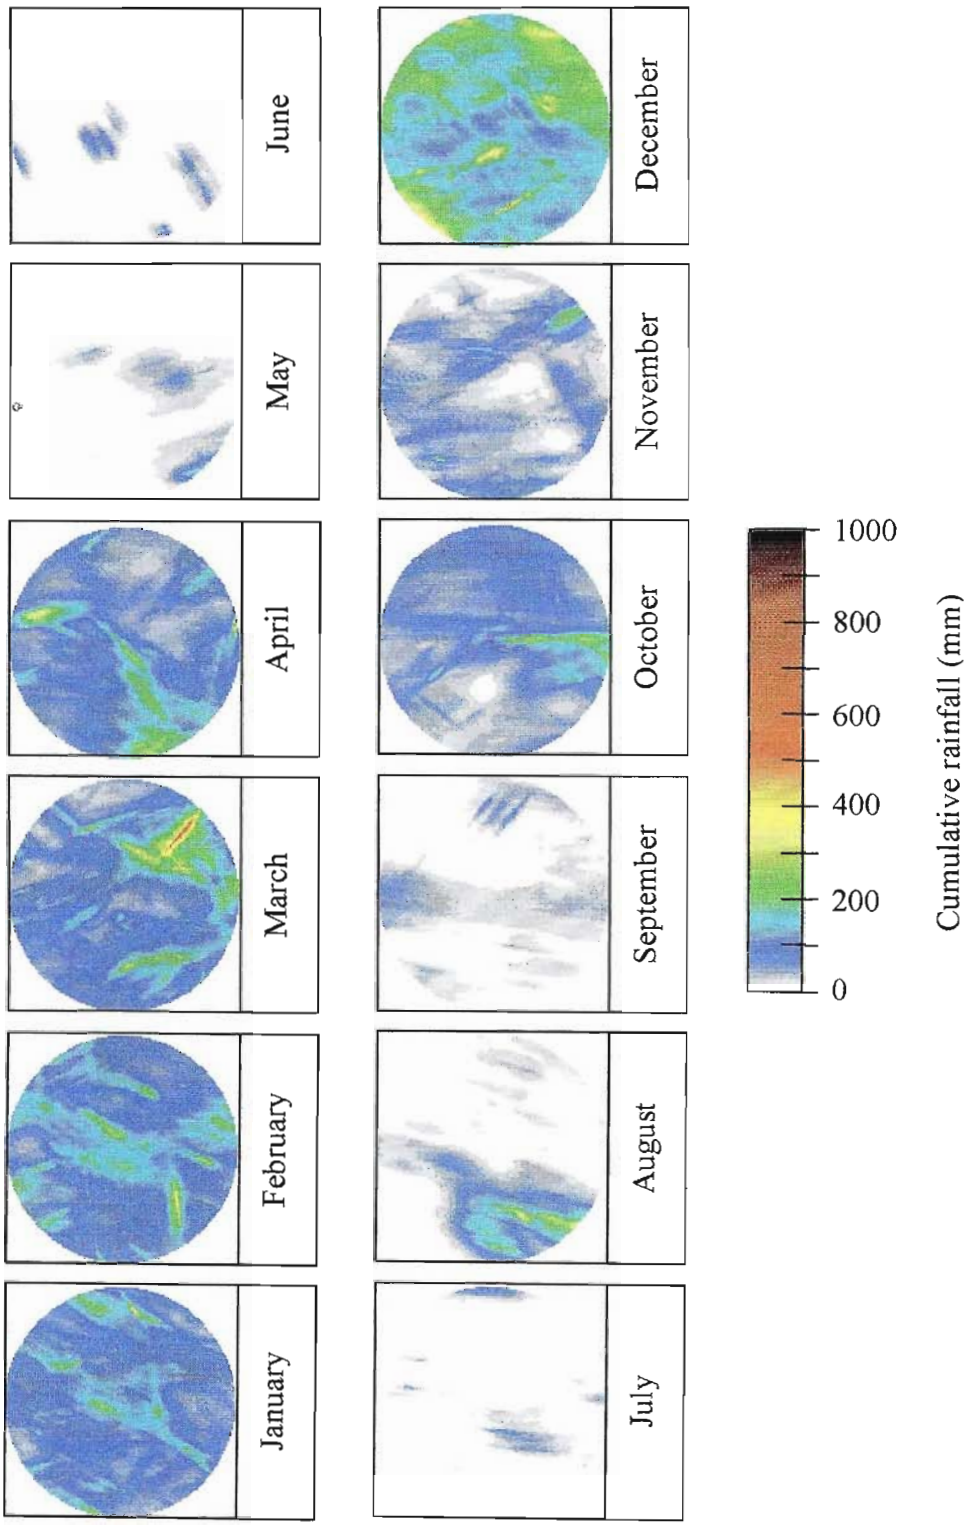

Figure 6.11 - Monthly aggregations for the simulated sequence using the autoregressive bead

Apart from the annual total rainfall, these three sequences are independent of each other. The distribution of rainfall throughout the year will be completely different and the three sequences of monthly aggregations are not expected to be any more similar than three independent years of annual rainfall. They are displayed in order to compare the qualitative properties of the three sequences.

Of the two simulated sequences, it is the frequency-filtered sequence which appears to produce the more realistic monthly and yearly aggregations. The autoregressive sequence results in streaky monthly aggregations which are not realistic. Quantitatively, the WAR_m, IMF_m and β_{space} of the three sequences are presented in Table 6.1.

Table 6.1 -Image scale statistics of the monthly aggregations shown in Figure 6.9, Figure 6.10 and Figure 6.11.

Month	WAR_m			IMF_m			β_{space}		
	Obs	AR	FF	Obs	AR	FF	Obs	AR	FF
January	1.00	1.00	1.00	109.50	86.02	147.59	2.37	2.74	2.87
February	1.00	1.00	1.00	50.82	96.37	52.86	2.49	2.74	2.70
March	1.00	1.00	1.00	84.32	94.77	107.78	2.46	2.95	2.81
April	0.12	1.00	1.00	2.47	65.11	56.05	1.14	2.66	2.80
May	0.95	0.10	0.06	19.99	1.92	1.07	2.17	2.41	2.18
June	0.01	0.05	0.11	0.18	1.67	1.93	0.98	2.59	2.58
July	n/a	0.06	0.20	n/a	1.57	2.77	n/a	2.35	2.39
August	n/a	0.36	0.21	n/a	13.30	11.74	n/a	2.71	3.28
September	0.12	0.41	0.15	1.48	8.17	2.63	1.65	2.57	2.56
October	1.00	1.00	1.00	43.64	47.90	38.65	2.15	2.52	2.90
November	1.00	0.91	1.00	72.45	35.12	42.32	2.30	2.70	2.70
December	1.00	1.00	1.00	235.69	181.54	184.85	2.55	2.71	2.76

At this coarse temporal scale, most of the image receives some rain and consequently the WAR statistic is mostly at 100%, the only exceptions being in the winter months. The IMF shows the monthly rainfall totals on each of the image. Finally the β_{space} is reasonably well reproduced in the simulated images. Ignoring the months with little or no rainfall, the β_{space} is over-estimated in the simulated sequences by around 10%. Both

forms of the model appear to perform well at the 1 km, one month space-time scale, however with very little to choose between the statistics, the frequency filtered approach produces more realistic looking images.

6.3.3 Daily images

Daily cumulative images for the year are too numerous to display in this document and are best viewed on computer, however Figure 6.12 compares the marginal distribution of the WAR_d for the 1999 radar data and the same two simulated sequences discussed above.

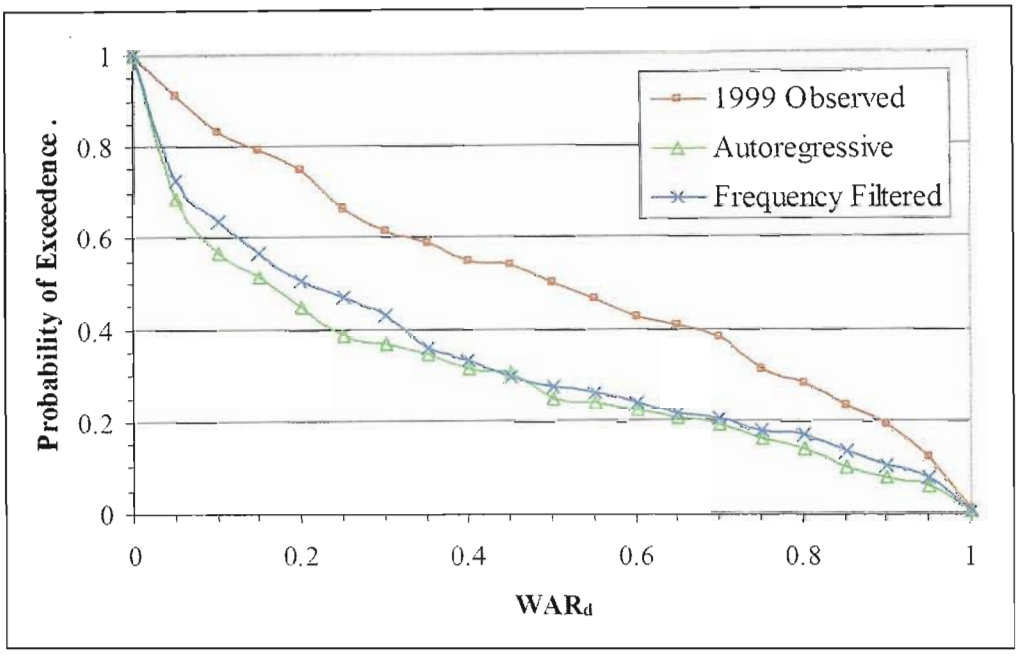

Figure 6.12 - Probability of exceedence of the WARd for the 1999 radar data and the frequency-filtered and autoregressive simulated sequences

The results of Figure 6.12 are slightly disappointing. When aggregated to daily images, the model appears to underestimate the number of general light-rain days. This phenomenon was also noticed in the comparison of the marginal distributions of the instantaneous images in Section 5.4.2. As stated previously, this could be partly as a result of random event sampling – the simulated images are simply made up of more intense storms – but this seems unlikely in the light of this result. Improvements could be made by fine-tuning of the WAR_i time series to encourage more general-light-rain images.

Figure 6.13 compares the marginal distribution of the IMF_d for the 1999 radar data and the two simulated annual sequences. The marginal distribution of the IMF_d as measured from the raingauge network was used in the fine-tuning process, but the radar data were not used in this fine-tuning process and therefore this is a reasonable validation test.

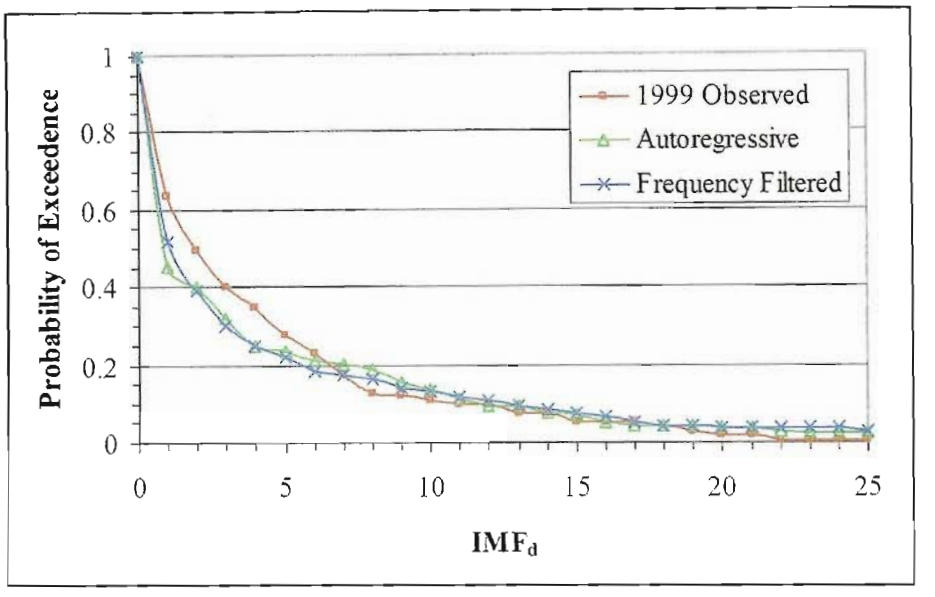

Figure 6.13 - Probability of exceedence of the IMFd for the 1999 radar data and the frequency-filtered and autoregressive simulated sequences

These are pleasing results as they show that the marginal distribution of the average daily rainfall for the region is convincingly captured by the model, although there is perhaps a small underestimation of the amount of rainfall that arrives in a day. Finally, the marginal distributions of the exponent of the radially averaged power spectra are compared in Figure 6.14.

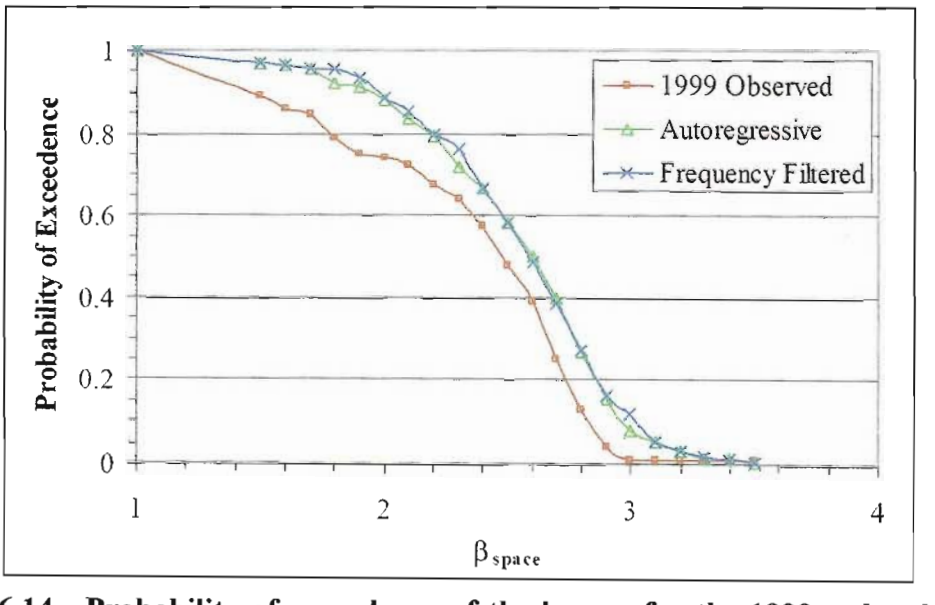

Figure 6.14 - Probability of exceedence of the bspace for the 1999 radar data and the frequency-filtered and autoregressive simulated sequences

Figure 6.14 indicates that the spectral exponent of the simulated daily images is slightly higher than that of the 1999 radar data, indicating that the spatial correlation at the 1km scale is exaggerated by the model. As for the monthly and yearly images, comparisons at coarser spatial scales are progressively more acceptable. Nevertheless, the results are reasonable.

6.4 *DISTRIBUTION OF WET AND DRY DAILY RUNS*

Recalling the definition of a wet day as a day with a WAR_d in excess of 10%, and making use of the empirical relationship between WAR_i and WAR_d presented in Figure 3.48, the daily runlengths were computed from the WAR_i image scale statistics of a 50 year string simulation in order to compare the daily runlengths to those of the daily raingauge data, presented in Figure 3.10 through Figure 3.14. The results of this analysis are presented in Figure 6.15 through Figure 6.19 with the simulated results shown as lines superimposed on the measured daily runs for the 1948 to 1997 raingauge data. The first of these figures shows the number of dry days in each month over the 50 year period.

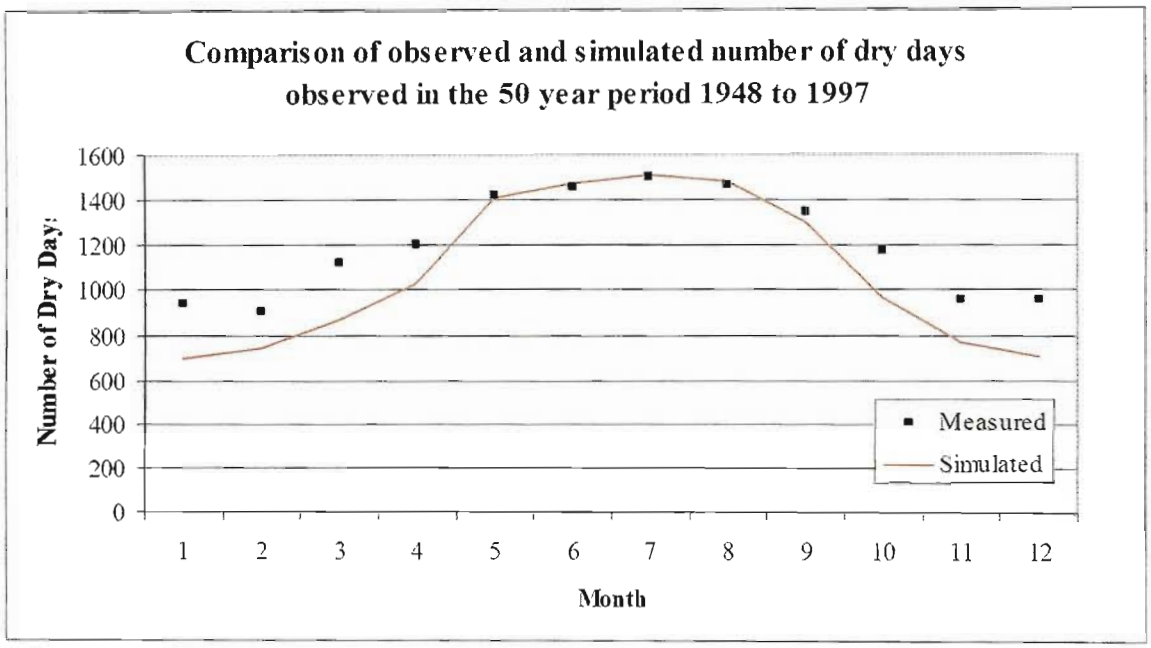

Figure 6.15 - Monthly comparison of the number of dry days for the simulated and observed 50 year period

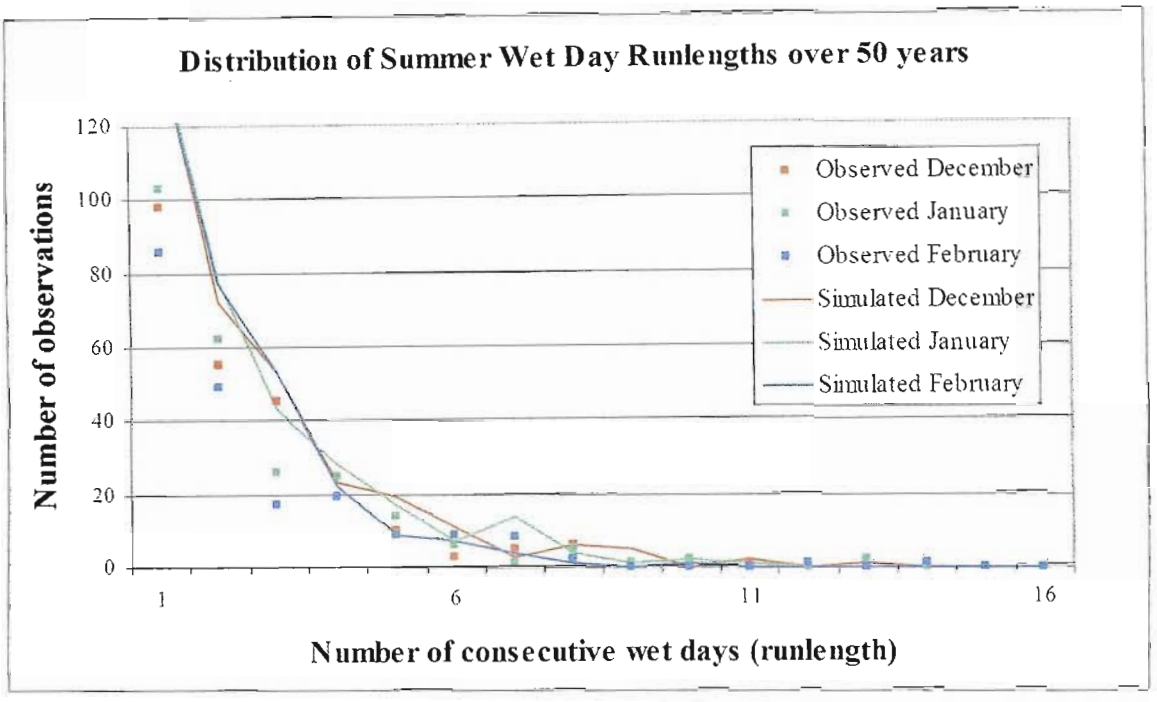

Figure 6.16 - Histogram of wet day runlengths for the summer months

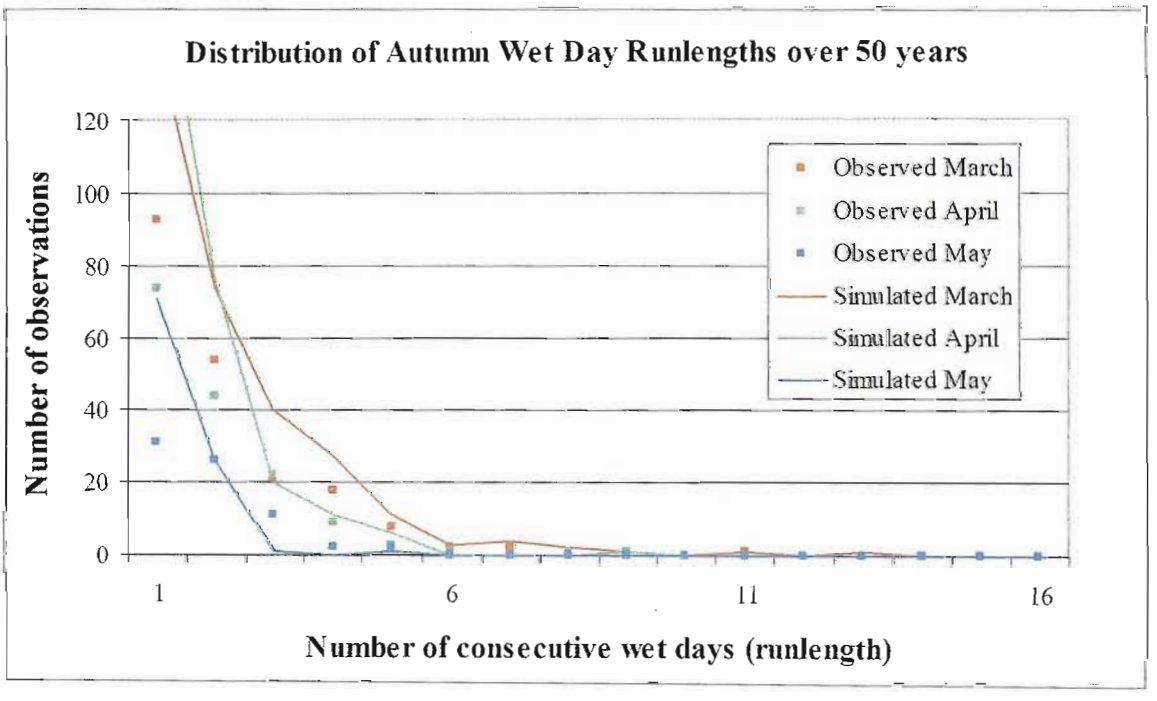

Figure 6.17 - Histogram of wet day runlengths for the autumn months

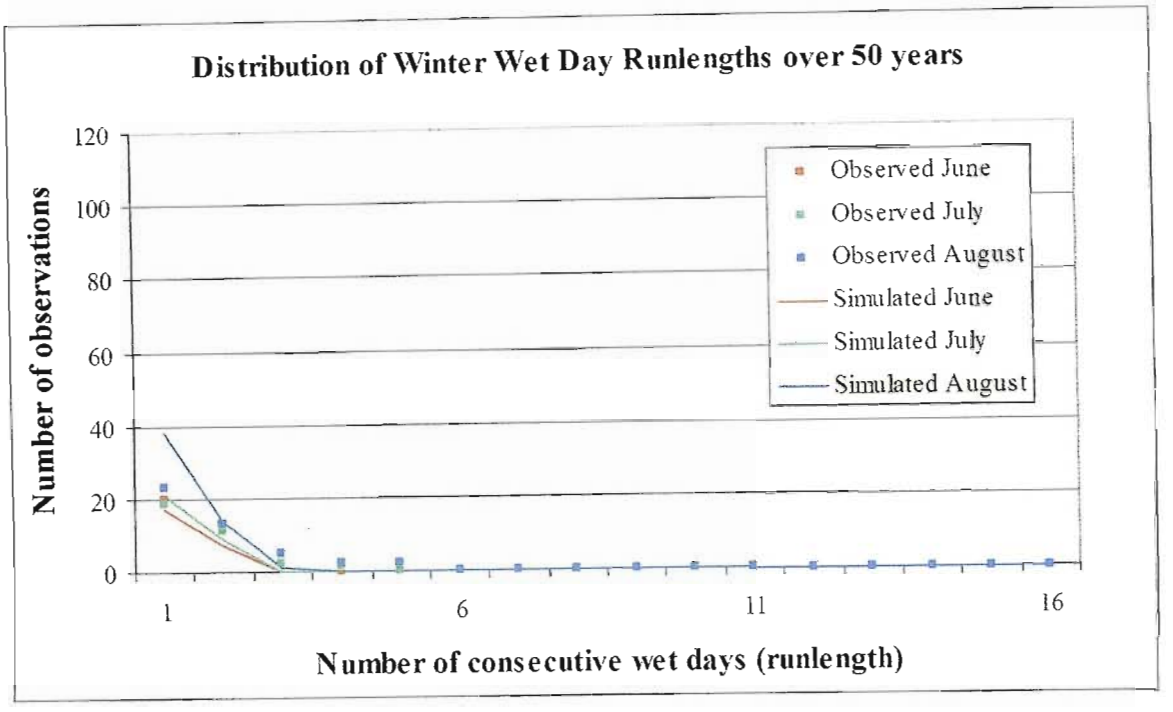

Figure 6.18 - Histogram of wet day runlengths for the winter months

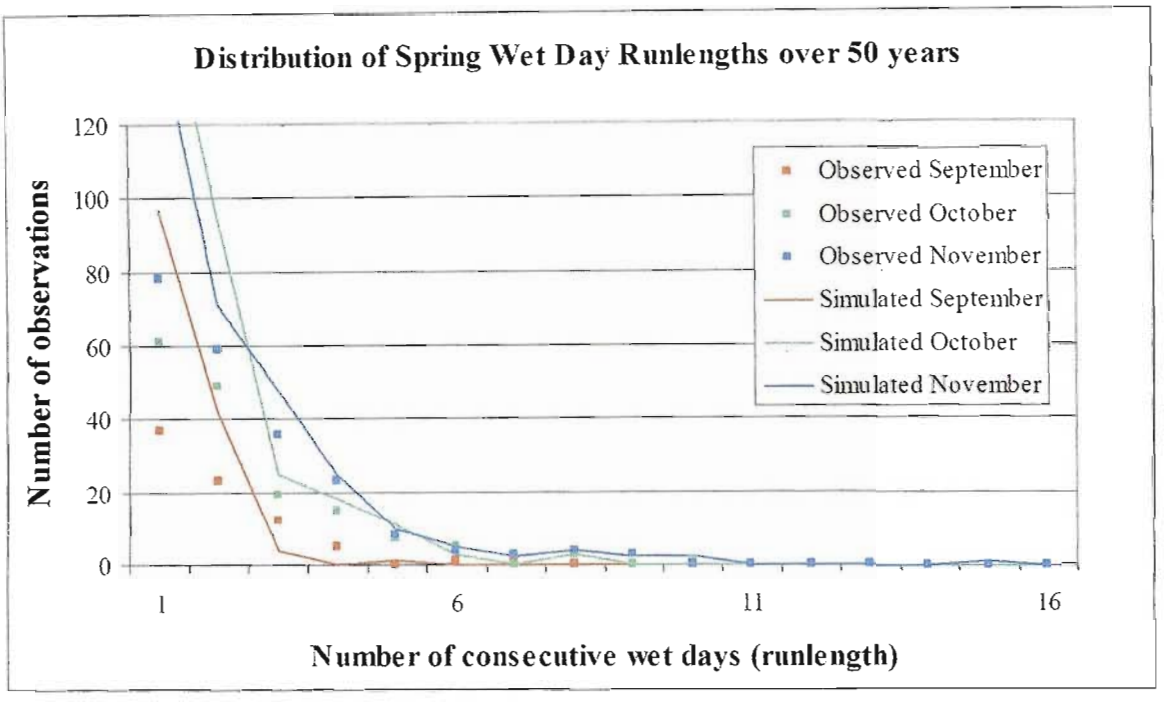

Figure 6.19 - Histogram of wet day runlengths for the spring months

Although the seasonal behaviour is reasonably well represented in these histograms, the model clearly underestimates the number of dry days (and overestimates the number of wet days) in the summer, spring and autumn months. Since the annual aggregations are accurately reproduced by the model, the overestimation of the number of wet days

seems to suggests that more rain should fall on individual days. This is consistent with observations made when comparing the IMF_d for the radar data to that of the model in Figure 6.13. Also, when these are considered in conjunction with the WAR_d analysis of Figure 6.12, the underestimation of the WAR_d also suggests more rain should fall on a day and it should be lighter, more general rain. Further model fine tuning to introduce more general, light rainfall may improve the model performance in all these respects. The number of dry days and winter runlengths are very well represented by the model

6.5 *Short-Term Forecasting Using the String of Beads model.*

The autoregressive processes used in the *String of Beads* model afford a limited forecasting ability which will be briefly discussed in this section. This was not the main focus of the research, but was investigated mostly as preliminary work for a follow on study for urban flood forecasting.

In its current form, to be used in Forecasting mode, the *String of Beads* model requires floating point radar rainfall images to be available in real time, at 1 kilometre 5 minute space-time resolution. Both the image scale and the pixel scale autoregressive processes are used, the second and third stages of the *String of Beads* model, to generate a range of possible future scenarios.

To begin, the WAR_i and IMF_i values of the observed data are first transformed into their Gaussian WAR_i^* and IMF_i^* using the target distribution functions shown in Figure 3.56 and Figure 3.57, according to the method illustrated in Figure 3.58. Using the WAR_i* and IMF_i* of the most recent five radar images, the bivariate autoregressive process for the WAR_i^* and IMF_i^*, illustrated in Figure 4.17, is then used to simulate a range of possible future WAR_i-IMF_i pairs. Figure 6.20 uses the image scale statistics of the 1996 event of Figure 3.35 to simulate three possible future scenarios for the WAR_i parameter, beginning at hour 10 in the event; a possible growth scenario, a possible decay scenario and the expected scenario. The M_{war*}, S_{war*}, M_{imf*} and S_{imf*} for the event are approximated using the WAR_i* and IMF_i* statistics from the start of the event, in this case 10 hours of historical image scale statistics.

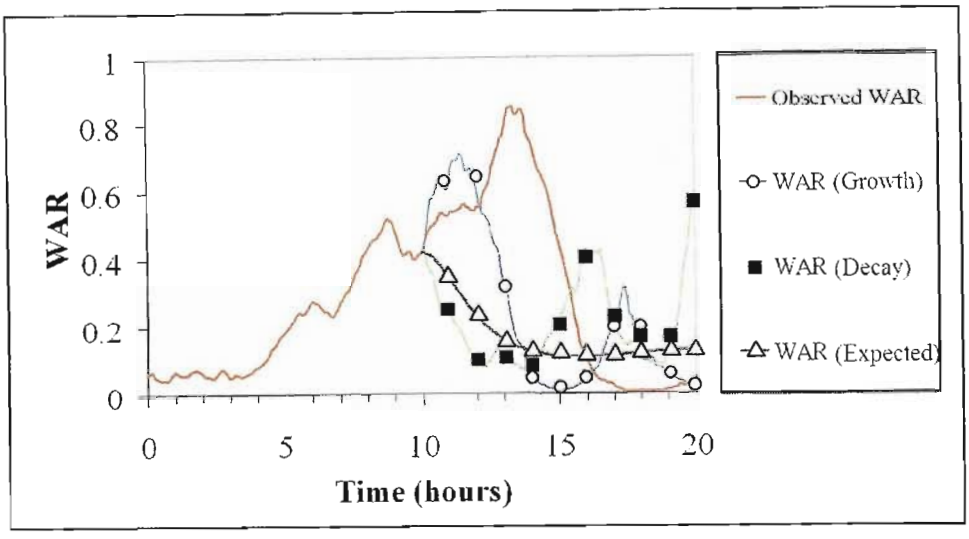

Figure 6.20 - Forecasting the WAR_i image scale statistic using the bivariate autoregressive process

The expected scenario is obtained by setting the shock term of the autoregressive process to zero, in which case the process will decay to its mean value, in this case M_{war*}. A similar picture is obtained for the IMF_i. Any number of these *what if* scenarios can be generated and the process is very fast. It is possible to generate hundreds of possible scenarios in the space of a minute, and then to select the critical scenarios to pass on to the pixel scale forecast routine (*String of Beads* model third stage).

Moving on to the pixel scale forecast, arguably the most important component is the mean field advection vector. This is measured in real time using the algorithm presented in Section 3.6. In its current form, the *String of Beads* model simply projects the most recent advection vector and allows the model to forecast with a constant vector. The PSI autoregressive process requires five historical fields whose marginal distributions must be transformed to standard normal variates. Assuming a two parameter lognormal distribution, $\Lambda(\mu, \sigma)$, on the radar rainfall images (a reasonable assumption for relatively wet images when forecasting is most critical), they can be transformed into the required Gaussian $N(0,1)$ space by taking logarithms of the PSIs, to acquire a $N(\mu, \sigma)$ field which can then be scaled and shifted to $N(0,1)$. These are then substituted into the $Z_{(t-5)}, Z_{(t-4)}, \ldots, Z_{(t)}$ arrays of the third stage of simulation shown in Figure 4.18. Possible future PSI fields are then obtained by simulation, using the forecast WAR_i-IMF_i image scale statistics obtained above. Figure 6.21 illustrates the results of this procedure this for a simulated sequence.

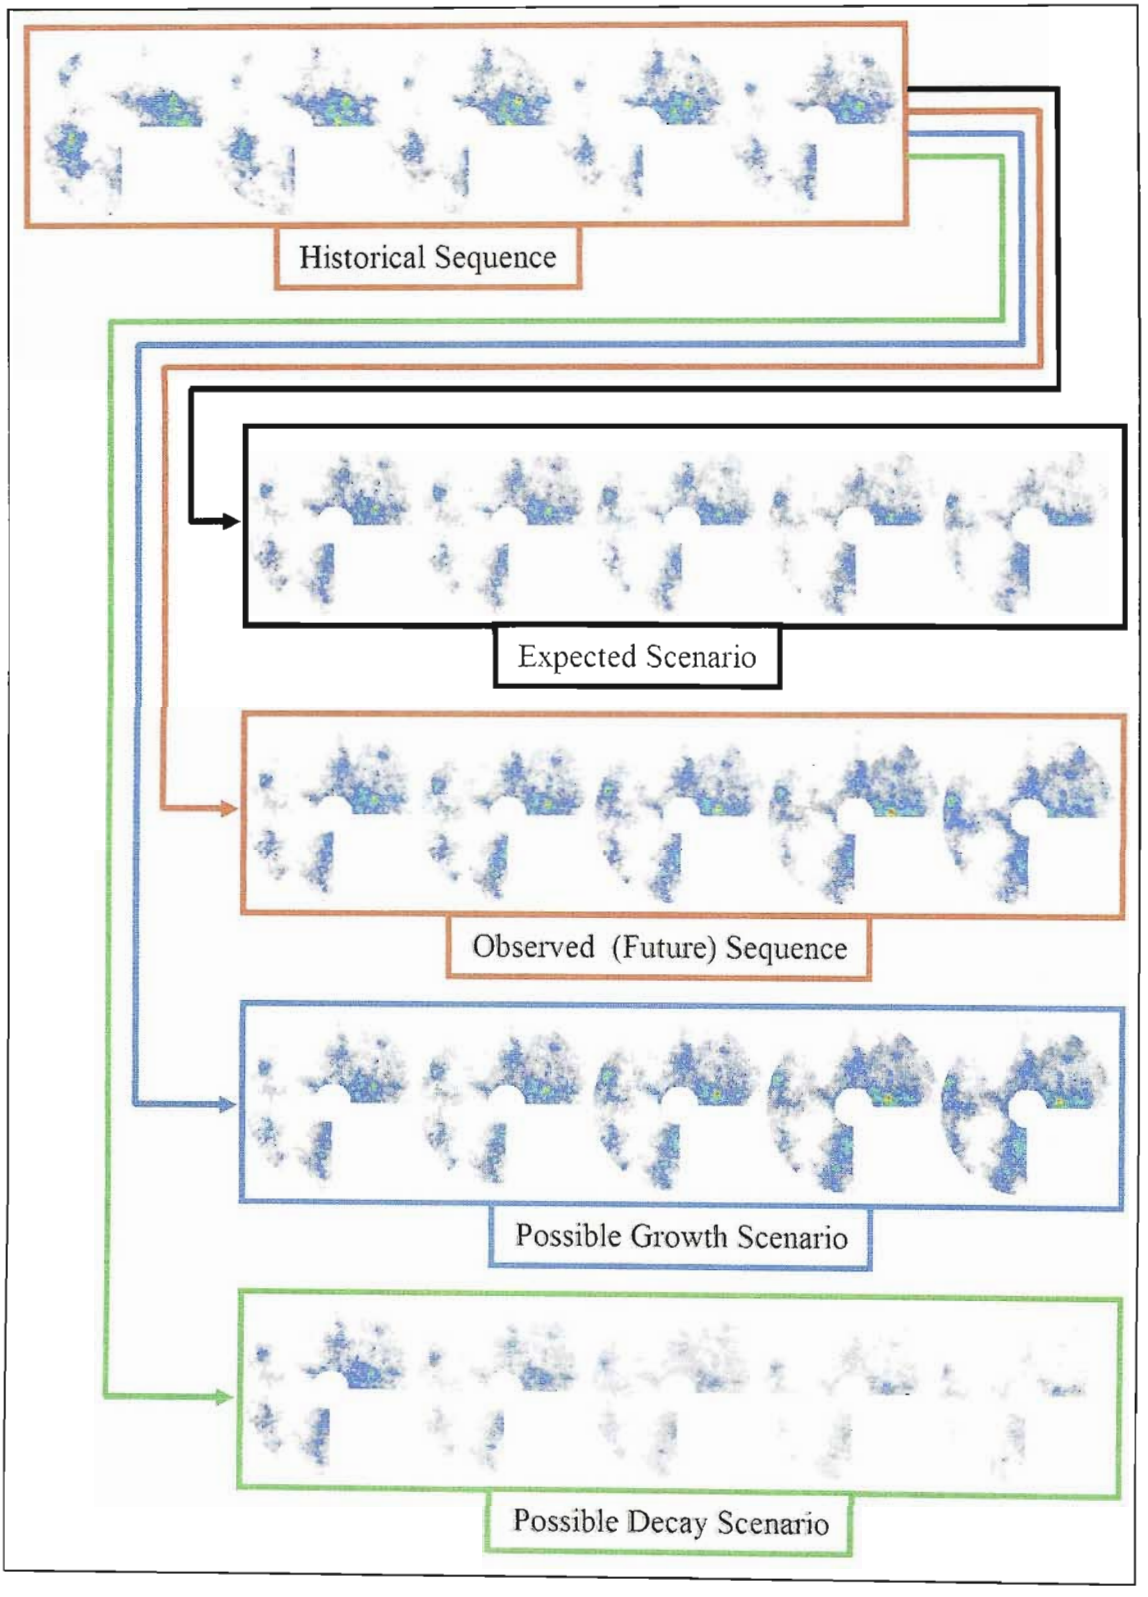

Figure 6.21 - Pixel scale forecasting (for a simulated sequence) comparing the expected, a possible growth and a possible decay scenario. Images shown at 10 minute intervals, colours correspond to Figure 6.20.

The sequences shown in Figure 6.21 were simulated using the WAR_i-IMF_i series of Figure 6.20 and the colours of the boxes around each sequence, corresponds to the line colour in Figure 6.20. The images are shown at 10 minute intervals. The historical sequence shows the information available at the time of forecasting and the observed sequence shows how the event eventually developed according to the WAR_i statistics of Figure 6.20. The expected scenario shows how the event decays if the shock term is omitted from the image scale autoregressive process and therefore the WAR_i* and IMF_i decay to their respective mean values of M_{war*} and M_{imf*}. The possible growth scenario and the possible decay scenario represent two of the *what if* future scenarios.

From these results, it appears that a reasonable forecast can be made of up to an hour, but it would be unrealistic to expect accurate forecasts for any longer period. In the example above, the observed sequence was ultimately more of a growth scenario than a decay scenario, but both were feasible outcomes based on the information available at the time of forecasting.

6.6 *SUMMARY*

This chapter has demonstrated the ability of the SBM to produce realistic annual rainfall scenarios which scale sensibly from the 5-minute, 1km up to the 1year, 128km, space-time scales. Most of the validation tests presented in this chapter yield encouraging results, the only real exceptions being the marginal distributions of the daily WAR and the daily spatial spectra. Using the fine-tuning techniques of Chapter 5, it should be possible to improve the performance of the model by manipulation of the WAR_i time series. In addition to the simulation of long-term rainfall scenarios, the model has also been demonstrated to be an effective tool in short-term forecasting and the delivery in real time of "what if" scenarios.

The validation tests presented in this chapter are certainly not exhaustive. A considerable amount of work has been done on the space-time structure of rainfall fields by other authors and many of their methods would serve as useful validation tests for the output of the *String of Beads* Model. That said, the simple validation tests presented here have generally produced encouraging results and the concept and basic structure of the model appear to be sound and sensible.

Chapter 7

Conclusions and Recommendations

7.1 CONCLUSIONS

Four main sets of conclusions are presented here. The first is a summary of the model performance which discusses the results obtained in this study and the general usefulness of the *String of Beads* Model. The second is a brief discussion of some useful results which have been discovered in the course of this research. These are ideas or methodologies which were devised in order to address specific problems in the development of the model, many of which have much wider application than the specific purpose for which they were designed. The third is a discussion of the aspects of the model or the analysis which would benefit from further research. This highlights areas where the model fell short of expectations. The final section considers possible enhancements to the model. These are aspects of research which have not yet been investigated, but could be used to enhance the model performance to include locale specific properties of rainfields in the simulations.

7.1.1 Model performance

The *String of Beads* Model shows great promise as a stochastic space-time rainfall model. It incorporates information extracted from radar data as well as daily raingauge network data, in order to capture both the high resolution space-time structure and the seasonal trends in rainfall statistics for any region.

The model makes use of simple, tried and tested stochastic methods and assembles them in a pragmatic, easily understood combination. It is defined by a fair number of parameters, however these can be quickly and easily measured from the available data and adjusted in a parameters file during the calibration process. Many of these parameters are likely to be independent of location. Attempts were made at model calibration on German, Swiss and English datasets in the latter stages of the research, however at the time of analysis, the calibration process was still under development and

the calibration algorithm was unstable. This is an exercise which will be repeated in due course. With a reasonable size dataset such as one year of continuous radar data and fifty years of daily raingauge data over a network of gauges, model calibration can be achieved in a few hours. This is an iterative process which has not yet been automated. Estimates of the parameters will improve as an understanding of the model to various datasets grows.

In addition to the calibration process, the output can be fine tuned to achieve desired characteristics in the rainfall fields. The fine tuning process is an operational adjustment which caters for any shortcomings in the model calibration process. It can be used to influence the qualitative nature of the rainfall on the individual images, increasing or reducing convective activity, or increasing or reducing the extent of rainfall.

It is an extremely flexible model and has been designed in a modular way so that any module can be redesigned to meet specialist requirements of a user, without interfering with other parts of the model.

At the *image scale*, high resolution simulations, as fine as 1 kilometre 5 minute space-time resolution, can be efficiently simulated over a 128 kilometre field and output in a variety of useful formats which can be easily read and manipulated. The images are realistic in appearance and statistically consistent in marginal distribution and second order spatial structure, with observed weather radar images.

The image scale statistics *alone* can be useful as a one-dimensional long-term regional simulation tool, without the detailed spatial structure of the images.

At the *event scale*, when aggregated into coarser space and time scales, the simulated cumulative images preserve a reasonably realistic spatial and temporal correlation structure as demonstrated by comparison to cumulative images of real data. However, improvements in this respect could be made by additional research into the structure of the pixel scale time series.

At the *daily*, *monthly* and *yearly scales*, cumulative images maintain a reasonably realistic spatial and temporal structure and marginal distribution of rainfall, although the

statistics of these cumulative images would all benefit from additional research into the image scale time series and advection characteristics of the model.

Although the *String of Beads* Model has been designed to run on modest computing equipment, considerable benefit can be gained through the use of greater processing power and ideally, parallel processing. The architecture of the model lends itself to the use of massive parallel processing methods. A full space-time simulation for a single year (in South African conditions) will run over a period of 6 to 8 hours on a single modern personal computer. This can be reduced by 50% if two machines are used in parallel. In fact, the total amount of processing time for a simulation is inversely proportional to the number of available processors. In the course of this research, up to 10 computers were used in parallel over a network in order to speed up simulations for calibration purposes.

In addition to the simulation abilities, the model has been demonstrated to give useful short term forecasts of up to an hour at the 1 kilometre 5 minute space-time resolution – valuable time in the context of urban flood forecasting where catchments are small and react quickly to intense rainfall.

7.1.2 Some useful results with wider application

There are a few components of the model which required some unconventional thinking in order to achieve the desired result. With reference to the literature reviewed, these appear to be novel results some of which may have more general application. Specifically, these are:

- *Pseudo-random simulation of censored lognormal rainfall fields* – This is the technique of image simulation presented in Section 3.4.1. Given the WAR above a certain threshold and the and the IMF above that threshold, a censored lognormal distribution can be iteratively fitted to the field. This distribution has been demonstrated to be better able to describe the marginal distribution of rainfall rates over the field than the simple two parameter lognormal.
- *The influence of data precision on the generalised structure function* – This was discussed in Section 3.5.3. It clearly illustrates how data thresholding affects the estimation of the GSF.

- *Field advection measurement* – The two-dimensional bisection search routine on a square simplex demonstrated in Section 3.6 is extremely fast on fields which are unimodal since it only requires that absolute differences be computed over the field for a few shifts in the x and y directions, eliminating the need for continuous function minimisation algorithms.
- *Autoregressive processes with fixed start and finish values* – This was used in the simulation of the bivariate WAR_i - IMF_i process in Section 4.5.2. The problem was how to simulate a time series defined by a known autoregressive relationship while forcing the process to start and end at known values. Autoregressive processes are pseudo-random and in their conventional form there is no way of forcing the simulated time series to pass through a known point. Since this was a requirement of the WAR_i – IMF_i time series, an extended theory had to be developed.
- *Scattered event renewal process* – This is the seasonally varying, event renewal process used in Section 4.5.1 to define the binary wet-dry process. It is a continuous process which relies on a known number of independent identically distributed event durations arriving in a known time interval. It is particularly useful when the distribution of either one of the states (in this case the wet state) is better known than the other.
- *Conditional distribution cascade* – The techniques used in Section 3.10.2 and 4.5.1 (illustrated in Figure 4.13) to describe and model respectively the multivariable distribution of a complex process, specifically the image scale statistics.
- *Data manipulation and visualisation* – The techniques presented for the visualisation, management and manipulation of large sequences of spatial data discussed in Sections 2.5 and 2.7. The low-level manipulation of image data in this way is powerful and efficient and can be applied to almost any spatial dataset.
- *The String of Beads architecture* – Summarised in Chapters 4 and 5, which illustrate how the combination of these techniques (separate consideration of the binary wet and dry process from the detailed space-time event) has culminated in the calibration of a continuous space-time rainfield model onto historical daily raingauge data.

7.1.3 More research needed

In spite of the overall encouraging results of the *String of Beads* Model, there are components of the model and the analysis which require additional research.

Verified Radar Data

As stated in Section 2.6 aggregation of the available radar data revealed that they over-estimated the rainfall rate on an image by more than 100% due to miscalibration of the radar. In order to continue with the research, it was necessary to linearly scale the entire dataset in order to achieve agreement with the annual total as observed by a network of daily raingauges in the area. The validity of linearly scaling the entire dataset is questionable since the marginal distribution of the rainfall field is severely skewed. Ideally, the entire dataset should have been reprocessed as raw radar data and adjusted in reflectivity space, prior to being converted into rainfall rate. This was beyond the scope of this study so a simple adjustment was required in order to continue with research and the linear scaling was the suggestion made by the SAWS. There is a risk that the spatial and temporal correlation structure and marginal distribution of the rainfall fields will have been distorted in this process and for this reason the analysis and calibration should be repeated now that the data have been re-calibrated by the SAWS. The formulation of the *String of Beads* Model is believed to be robust enough to accommodate the re-computed parameters without compromising the principles on which the model has been built. Obviously, the uncertainty in the basic data introduces uncertainty in all of the verification and validation tests since there is no way of knowing if the radar data are actually representative of what fell, but the principles of the verification and validation tests are sound and can be applied to the corrected dataset. As stated in Chapter 2, the radar data are assumed to be correct for the purposes of the remaining discussion. In any event, the important touch-stone is that the behaviour of the raingauges can be recaptured – after all, that is our only ground "truth".

Bivariate autoregressive process for the WAR_i –IMF_i *time series.*

The parameters of the bivariate autoregressive process for the $WAR_i - IMF_i$ time series were computed for a single 42 hour event. It would be worth investigating the variability of these parameters for events of shorter duration. It is likely that the WAR_i and IMF_i image scale parameters will be more volatile in the short duration, convective events. Consequently the high autocorrelation and cross-correlation observed between

these two time series in the 42 hour event may be reduced for shorter duration events, giving rise to a wider range of $WAR_i - IMF_i$ combinations.

Daily, monthly and yearly aggregations

As observed in Chapters 5 and 6 aggregation of the model output into daily, monthly and yearly images reveals a slightly more variant field than the observed aggregations of real data. It would appear that the number of general, light rainfall days (High WAR_d, low IMF_d) is under-estimated by the model. This could be a consequence of the fact that the parameters of the bivariate $WAR_i - IMF_i$ autoregressive process were extracted from an extreme (42 hour) event in which these two statistics were both extremely high as a result of super-convective activity at the time. Revisiting these parameters as suggested above may introduce more of the high WAR_i low IMF_i time series required to improve the model performance in this respect. Alternatively, it may be necessary to revisit the relationship between the event scale statistics $MWAR_i$* and $MIMF_i$* or else to simply fine tune the model by linear scaling of the WAR_i time series as discussed in Section 5.4.2.

Pixel scale autoregression

The space-time upscaling analysis of Section 5.2.3 revealed shortfalls in the model performance for both the *autoregressive* and the *frequency-filtered* bead processes. This is certainly an aspect of the analysis which requires further attention. Attempts to measure the pixel scale correlation structure by normalising, removal of the field advection and pre-whitening of the images were unsuccessful. The autoregressive parameters and spectral exponent used in the *autoregressive* and *frequency-filtered* bead processes respectively were determined iteratively by comparison with the spatial and temporal correlation structure of untransformed simulated and observed sequences with identical image scale statistics. Due to the intricate nature of this comparison, the iterative search was done manually to save time. In order to improve the estimate of these parameters, a program might be written to automate the search for optimal parameters.

Observations of observed pixel scale time series reveal a long term dependence structure and non-stationary characteristics. It may be more effective to use an autoregressive moving average (ARMA) process to describe the pixel scale time series

in the autoregressive bead shown in Section 4.5.3. This would still afford some forecasting ability while providing the long term dependence structure achieved by using the *frequency-filtered* approach to bead modelling. As discussed in Section 3.11, it is difficult to extract the parameters of the pixel scale time series from real data, however the iterative fitting technique using the space-time upscaling characteristics of simulated beads could also be applied in order to determine the parameters of the ARMA model. It is likely that this approach would also result in faster simulation than the *frequency-filtered* methods.

7.1.4 Suggested enhancements for future development

The model formulation lends itself to some simple enhancements which would enable the *String of Beads* Model to be used in a wider variety of climates and locations and possibly over a wider range of scales.

Weather-type bead generation

As a first approximation, it has been assumed that rainfall events are the result of a single weather type which only varies in duration and intensity. In some climates this is clearly not true. A simple enhancement of the model would be to incorporate a weather-type selection which may be dependent on the season or the duration or both. Depending on the type of weather, parameters of the bead could be altered in such a way as to incorporate these weather types, possibly by adjusting the $WAR_i - IMF_i$ relationship or the storm advection. An approach used by Fowler et al. (2000) was to use a semi-Markov chain to switch between weather types and this approach could be easily incorporated into the *String of Beads* Model.

Autoregressive storm advection

Another simplistic approach to simulation in the *String of Beads* Model is the assumption that storms move in a straight line. This is not always true. A far better approach would be to make use of a bivariate autoregressive model of the north-south and east-west advective shifts. If necessary, the general direction of the storm could be controlled through the use of a two-dimensional version of the autoregressive process with fixed start and end values, as discussed in Section 4.5.2. Storms can also be made to rotate about a point in space (which needn't be on the image) to simulate cyclonic or anti-cyclonic weather systems. Improved field advection modelling is also likely to

improve the spatial characteristics of the cumulative images at the daily, monthly and yearly time scales.

Special effects

By biasing (in space) the random field in the bead modelling process (after field advection has been applied but prior to power-law filtering), it is possible to produce non-homogeneous fields in order to simulate effects such as the orographic influence on rainfall fields. The most effective template for biasing the field is an annual cumulative image.

Similarly, the random field can be biased in space prior to the application of field advection and power-law filtering, in order to generate frontal weather systems which appear as bands of rainfall.

Combined weather systems can be modelled by simultaneous modelling of independent beads and superposition of the independent fields generated from each bead.

Finally, it may be possible through modification of the spatial filter, to expand the *String of Beads* Model beyond the 128km scale. Consideration of the power law spectra suggests that the assumption of scaling begins to break down beyond this range, however a modified filter for longer wavenumbers may prove effective in modelling larger scale weather systems using radar and networks and satellite data.

In conclusion, the *String of Beads Model* shows great promise and its modular structure lends itself to a wide variety of enhancements.

Chapter 8

Appendix A – Early Modelling and Calibration Attempts

8.1 INTRODUCTION

The following sections serve to illustrate the development of the model. This is work which has now been superseded by more efficient or more realistic methods.

8.2 EARLY ATTEMPTS AT INDIVIDUAL IMAGE SIMULATION

The use of fractals to simulate natural phenomena has become extremely popular and many fractal models of rainfall have been developed in recent years. Mandelbrot (1986) tentatively defines a fractal as

A fractal is a shape made up of parts similar to the whole in some way

Arguably one of the most successful types of fractal rainfall model is the multiplicative cascade (Over and Gupta, 1994). Two fractal models were investigated in the early stages of this study and the first made use of a simple two-dimensional multiplicative cascade. The second of the fractal models that was investigated was a more complex version of the multiplicative cascade and can be visualised as having the structure of exploding fireworks. It will be referred to as the *fireworks cascade*. The other two methods discussed briefly in this section, the Direct method and the Turning Bands method, were considered as possibilities but ultimately rejected in favour of the more straight-forward power-law filtering approach, the first was rejected for reasons of efficiency and the second because it could not easily be extended to account for the temporal evolution of the rainfall event.

8.2.1 The multiplicative cascade - the α model

The multiplicative cascade falls into the category of the so-called *multifractal* models. Each element in the multiplicative cascade is a descendant of an element of higher iteration level, however the elements of any one level are not necessarily all the same size (or intensity). Figure 8.1 illustrates the construction of a simple one-dimensional multiplicative cascade.

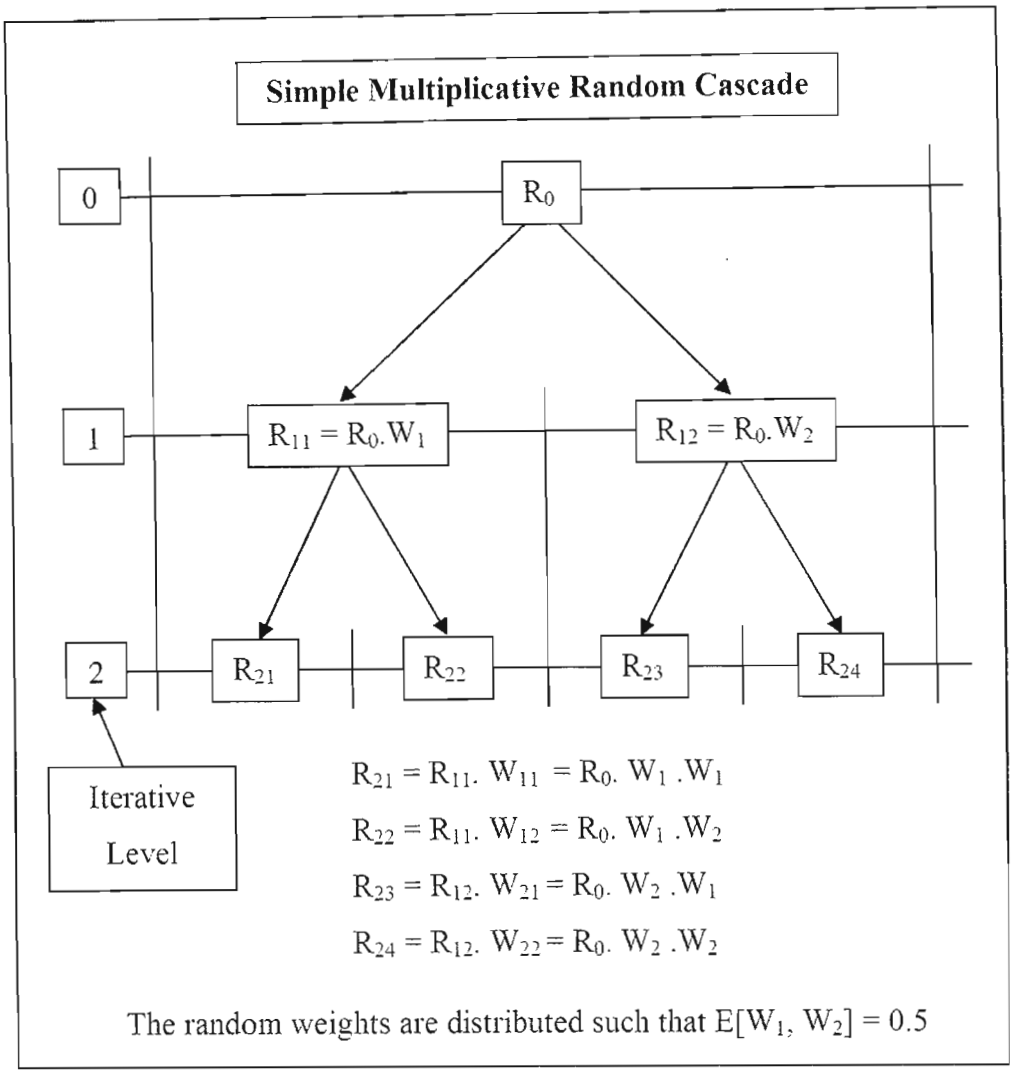

Figure 8.1 - Construction of the random multiplicative cascade

Beginning with an initial total rainfall R_0 to be distributed over the image, the first iteration level is computed by first generating two random weights W_1 and W_2 and then multiplying the total R_0 by each in turn to obtain the totals R_{11} and R_{12}. New values for W_1 and W_2 are then calculated to be multiplied by R_{11} to obtain the second level terms R_{21} and R_{22}. Again new values for W_1 and W_2 are calculated to be multiplied by R_{12} to obtain the second level terms R_{23} and R_{24}. This process can easily be extended to two (or more) dimensional space and can continue to any desired iterative level. An unfortunate artefact of rainfields which are generated using this technique, is that they tend to have a blocky appearance. The simulated image of Figure 8.2 was generated by Pegram and Seed (1992) using a two-dimensional multiplicative cascade with 6 computational iterations. Since 2^6 is 64, the size of this artificial image is 64 x 64 pixels.

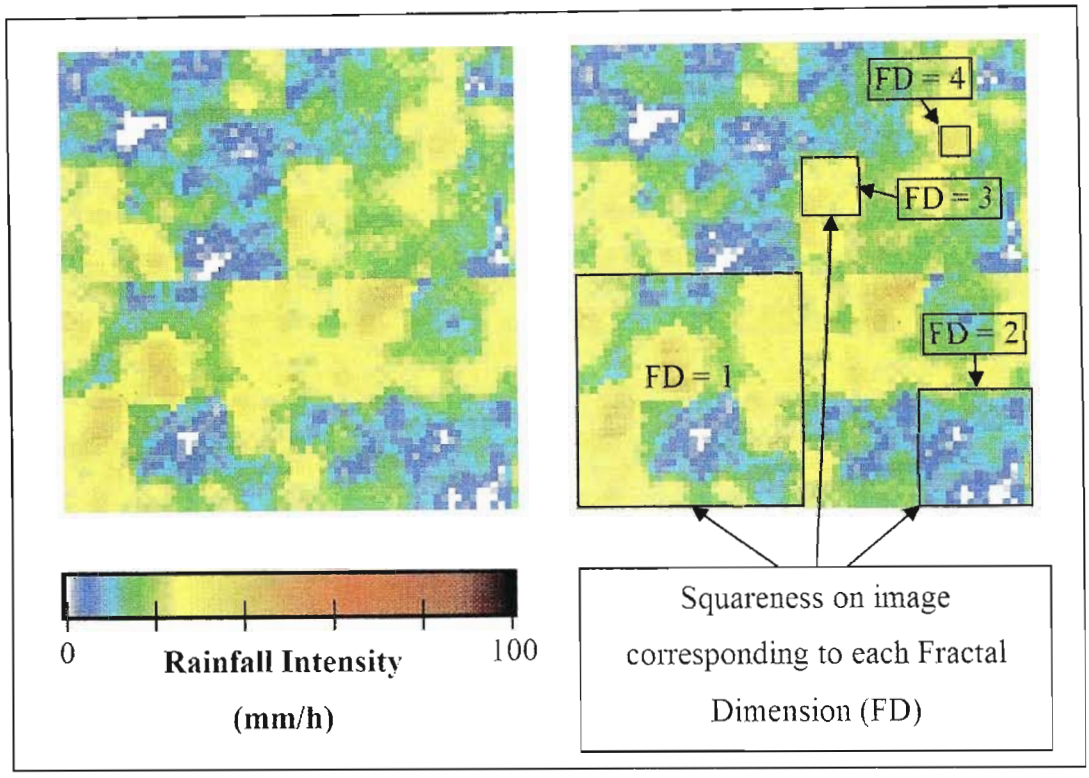

Figure 8.2 - An example of an image generated by Pegram and Seed (1992) using the multiplicative cascade. The image on the right highlights the square nature of the image on the left

Although the simple multiplicative cascade is able to reproduce the required marginal distribution and power spectrum, the artificial images are clearly unrealistic due to their blocky nature. Menabde et al. (1997) have developed a method which makes use of a discrete random cascade to generate a self-similar random field which they then power-law filter in Fourier space. Although this method produces images which do not have the obvious blocky nature of Figure 8.2, a simpler approach to image simulation, presented by Wilson et al. (1991) and discussed in detail in Section 4.4.2, was adopted for the *String of Beads model*.

8.2.2 The fireworks cascade

In the search for a fractal descriptor of rainfields which does not exhibit the square nature of the multiplicative cascade, the *fireworks cascade* was developed. The idea is similar to the multiplicative cascade only the fireworks cascade has no fixed shape or size. An illustration of the fireworks cascade is provided in Figure 8.3.

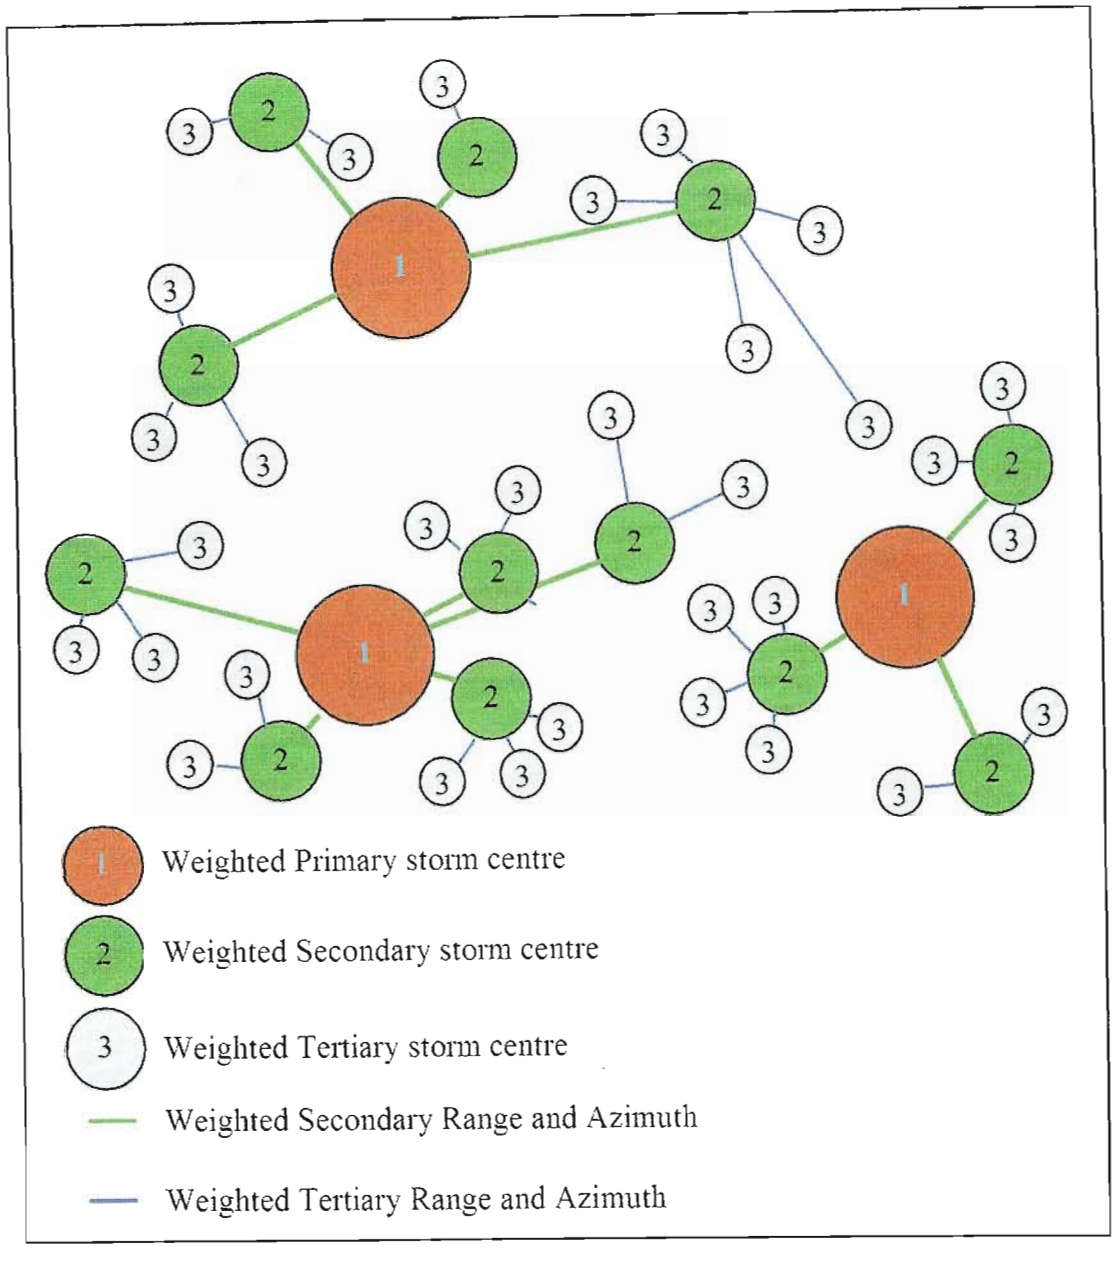

Figure 8.3 - Illustration of the fireworks cascade algorithm

Beginning with a total amount of rainfall to be distributed over the image (as in the multiplicative cascade) an exponentially distributed random number of *primary storm centres* is first generated. Depending on the type of weather system to be modelled, the mean of the exponential distribution of the number of storm centres will be adjusted.

These primary storm centres are then randomly distributed over the image. The total rainfall intensity is then distributed amongst them using random weights. The fewer the number of primary storm centres on the image, the more rainfall at each of them so that a image which resembles a convective weather system will have only a few, very intense primary storm centres. By contrast, a stratiform weather system will have a large

number of less intense primary storm centres so that the total rainfall is distributed more evenly over the area of the image. This completes the computation of the first iterative level.

Subsequent levels are computed slightly differently from the first. Four sets of exponentially distributed random weights are calculated for each *primary storm centre* in the first level. The first set is for the exponentially distributed random number of *secondary storm centres* or the number of points for the next computational level for each storm centre. The second and third sets of weights are a *random range* and *random azimuth* from each of the primary storm centres to each of their secondary storms. The random range is exponentially distributed with a mean which diminishes with increasing iterative level. The random azimuth is uniformly distributed. The final set is a random weight of rainfall which is used for the distribution of rainfall amounts to each secondary storm centre. The expected sum of the random weights of rainfall at each storm centre is unity so that the total amount of rainfall remains almost constant regardless of the number of iterations. If rainfall is assigned to a pixel which has already been assigned rainfall from another storm centre, the two rainfall rates are added for that pixel. Subsequent levels are computed in a similar fashion.

Two random fields generated using this algorithm are given in Figure 8.4. They are both computed using 5 computational iterations. The image on the left is generated with a mean number of 5 sub-storms giving a total of 5^5 (= 3125) distributed points, whilst the image on the right is generated with a mean number of 15 sub-storms giving a total of 15^5 (= 759375) distributed points. These fields are 200 x 200 pixels giving a total of 40000 points with intensities plotted to a logarithmic scale.

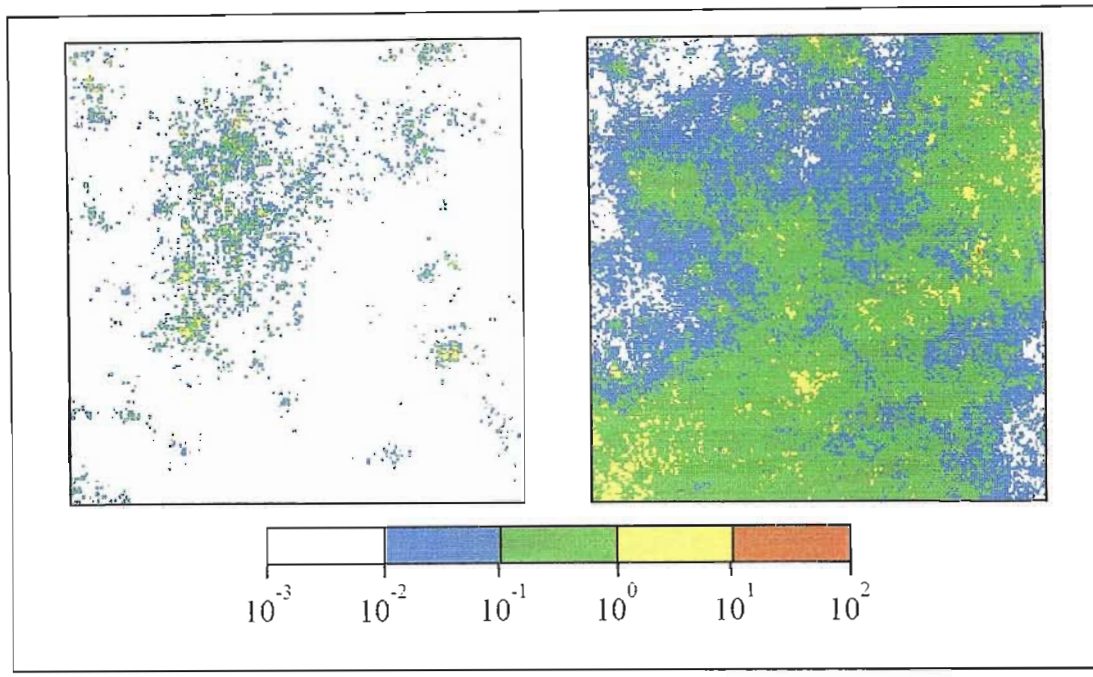

Figure 8.4 - Random fields generated using the fireworks algorithm. Left image has mean of 5 sub-storms, right image has a mean of 15 sub-storms.

These fields are clearly random in shape and have a clustering behaviour which resembles that of observed rainfields. The average size of the clusters can be adjusted by using a different mean when calculating the random range. The fields can be filtered to achieve a texture which more closely resembles real rainfields. Figure 8.5 gives an example of a random field before (left) and after (right) the application of a simple Laplacian filter. This field was also generated using a mean of 5 sub-storms and may be compared to the one in Figure 8.4.

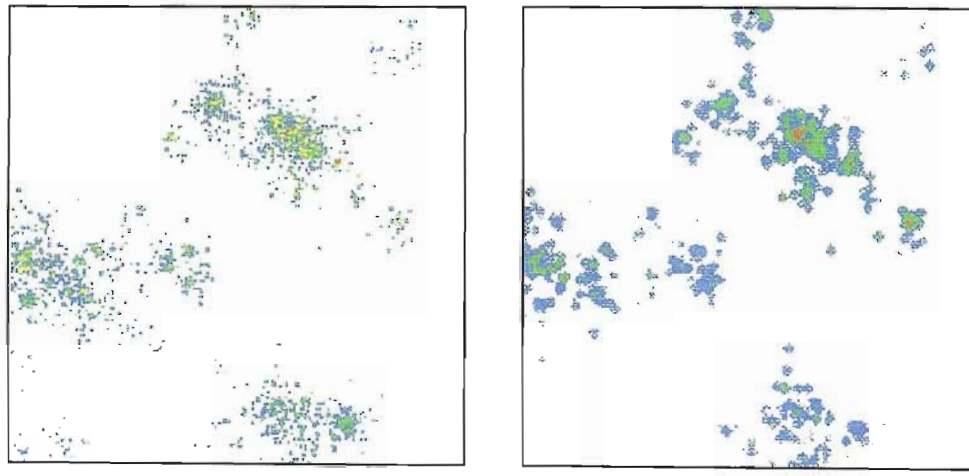

Figure 8.5 - Random field before (left) and after (right) Laplacian Filter

The simple Laplacian filter used in Figure 8.5 replaces the value at any pixel by a filtered pixel which is a linear combination of itself and the surrounding pixels. In this particular case, the linear weights of the filter were defined by the molecule:

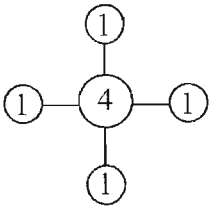

Clearly, this is not really effective in producing a realistic looking rainfield. Although other filters could be investigated, a power-law filter in Fourier space for example, this method of generating rainfields is extremely slow and was consequently abandoned in favour of the simpler, faster method proposed in Section 4.4.2.

8.2.3 The direct method

This method is likely to be the most commonly understood amongst time series analysts. Given an image, assuming that it has a lognormal distribution of rainfall intensities, the first step in the process would be to take logarithms of the rainfall intensities to obtain a normally distributed random field which could then be scaled an shifted to achieve a standard, normally distributed field which will contain the correlation structure of the image. From this field of N real numbers, it is possible to extract the correlation matrix $\boldsymbol{R}$ which will have the form

$$\boldsymbol{R} = \begin{bmatrix} 1 & \rho_{12} & \rho_{13} & \cdots & \rho_{1N} \\ \rho_{21} & 1 & \rho_{23} & \cdots & \rho_{2N} \\ \vdots & \vdots & \vdots & & \vdots \\ \rho_{N1} & \rho_{N2} & \rho_{N3} & \cdots & 1 \end{bmatrix}$$

We require a matrix $\boldsymbol{B}$ which, when multiplied by a noise matrix $\boldsymbol{a}$, will produce a field $\boldsymbol{X}$ which satisfies the relationship

$$\mathrm{Corr}[X X^{\mathrm{T}}] = R \qquad (8.1)$$

where

$$X = B.a \qquad (8.2)$$

and $\boldsymbol{a}$ is a NxN diagonal matrix with random, standard-normally distributed terms.

Substituting Equation 8.2 into 8.1 gives

$$R = \mathrm{Corr}[X X^{\mathrm{T}}] = \mathrm{Corr}[B.a.a^{\mathrm{T}}.B^{\mathrm{T}}] = B.B^{\mathrm{T}} \qquad (8.3)$$

Equation 8.3 can be solved for $\boldsymbol{B}$ using singular value decomposition and any number of random fields $\boldsymbol{X}$ can be generated with the required structure using Equation 8.2 and changing the random diagonal terms in the matrix $\boldsymbol{a}$. This process can be extended into three space but is computationally expensive since the number of calculations is proportional to N^2 in two space and N^3 in three space.

8.2.4 The Turning Bands method

This method of simulating fields was first presented by Matheron (1973) and is an extremely efficient method of generating random fields. In its two dimensional form, the process is to generate a number of lines which are concurrent through an arbitrary origin in the plane $\Re^2$. These lines are uniformly distributed in azimuth θ between 0 and 2π about the origin. Next, a one dimensional, correlated random sequence along each line - a *line process* - is generated. Each point in the simulated field is then computed as a linear combination of the values of the line processes. A schematic of this process, taken from Mantoglou and Wilson (1981) is provided in Figure 8.6.

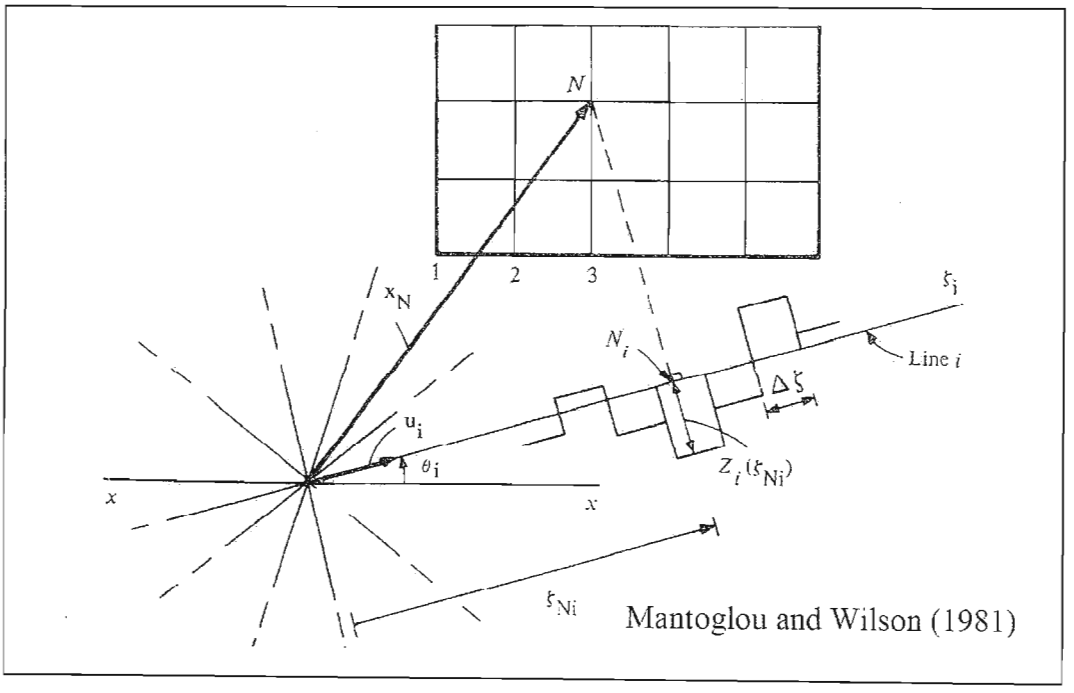

Figure 8.6 - Schematic of the Turning Bands method

Considering two points on the random field **R** with location vectors $\mathbf{x}_1$ and $\mathbf{x}_2$ respectively and $Z_R(\mathbf{x}_1)$ and $Z_R(\mathbf{x}_2)$ represent the intensity at those points, the covariance of the random field $C_R(\mathbf{x}_1, \mathbf{x}_2)$ is given by

$$C_R(v) = E[Z_R(\mathbf{x}_1).\, Z_R(\mathbf{x}_2)] \tag{8.4}$$

where v is the modulus of the vector $\overline{\mathbf{x}_1\, \mathbf{x}_2}$.

A one dimensional process, with points $Z_i(\zeta)$ and covariance $C_I(\zeta)$, is generated along each line ***i***. A point ***N*** in the random field with location vector $\mathbf{x}_N$, has a projected counterpoint on the line ***i*** denoted ζ_{Ni} and the value at that point is $Z_i(\zeta_{Ni})$. This can be written as $Z_i(\zeta_{Ni}) = Z_i(\mathbf{x}_N.\, \mathbf{u}_i)$, where $\mathbf{u}_i$ is the unit vector on the line ***i***. For each of the ***L*** lines, an independent value of the one dimensional process with covariance $C_I(\zeta)$ is generated and the point ***N*** is assigned the simulated value

$$Z_s(\mathbf{x}_N) = \frac{1}{\sqrt{L}} \sum_{i=1}^{L} Z_i\ (\mathbf{x}_N . \mathbf{u}_i) \tag{8.5}$$

According to Mantoglou and Wilson (1981), between 4 and 16 lines will provide sufficient accuracy in the two dimensional case.

Evidently, the crucial component in this method is the relationship between the covariance $C_I(\zeta)$ of the one dimensional line process and the covariance $C_S(v)$ of the random field to be simulated in two dimensions. For the two dimensional case, this relationship is given by

$$C_s(v) = \frac{1}{2\pi} \int_{\substack{unit \\ circle}} C_1(\mathbf{v} \,.\, \mathbf{u})\, du \tag{8.6}$$

which can be expressed as

$$C_s(v) = \frac{2}{\pi} \int_0^v \frac{C_1(\zeta)}{\sqrt{v^2 - \zeta^2}}\, d\zeta \tag{8.7}$$

$C_I(\zeta)$ cannot be expressed explicitly in terms of $C_S(v)$ and therefore Equation 8.7 must be solved numerically.

The process of simulating an image using the Turning Bands method would be:

- Take the logarithms of the image pixel intensities (lognormal distribution) and then normalise and standardise to achieve zero mean, unit variance.

  The covariance of the random field $C_R(v)$ could be computed from the correlation which in turn can be measured efficiently by employing the Fast Fourier Transform as discussed in Section 3.5.2.
- Substitute $C_R(v)$ for $C_S(v)$ and solve Equation 8.7 numerically for $C_I(\zeta)$.
- Select an arbitrary origin for the lines and generate a corresponding uniformly distributed azimuth θ for each of the 4 to 16 lines.
- Along each line, generate a random sequence of numbers with covariance function $C_I(\zeta)$
- Simulate a normally distributed random field with values Z_S using Equation 8.5
- Scale and shift the field to obtain the desired μ and σ
- Exponentiate the field.

Although this method is an efficient means of generating an image and can be easily extended into three space, it was rejected in favour of the power law filtering technique due to the fact that it is difficult to see how to build in the temporal evolution of the storm using turning bands.

8.3 EARLY ATTEMPTS AT MODELLING A RAINFALL 'BEAD'

The main differences between the earlier attempts at modelling the Bead process and those which are now part of the *String of Beads* Model are:

- algorithm efficiency which has been improved considerably by treating separately the temporal correlation structure and the spatial correlation structure
- forecasting abilities which are afforded by the use of the *autoregressive* bead simulation approach.

8.3.1 The three-dimensional space-time filter of Gaussian noise

The images presented in this document were generated using the same β_{space} exponent in the x and y directions. This lead to a field which had similar properties in both directions - the field was isotropic in this sense. Similar fields can be generated using a *warped*

power law filter which has a different β_{space} exponent in the two directions, β_x and β_y. This will result in a field which is more correlated in the direction of higher β_{space} and will appear to be banded in that direction. It was noted that instantaneous rainfields appear to be reasonably isotropic and homogeneous in space and therefore β_x and β_y are taken pragmatically as an average β_{space} when modelling images by power-law filtering normally distributed noise in the *String of Beads* model. The concept of a *warped* filter becomes a useful tool when the time dimension is introduced to the modelling process.

In keeping with the method of power-law filtering techniques discussed in Section 4.4, given the set of event scale parameters, the first step in the process was to generate a three dimensional field of normally distributed, real valued, random noise of size (2^n x 2^n x 2^t) pixels where n and t are integer values. The integer n determines the size of the rainfield and t determines the number of images that will be produced in the bead. The field is then power-law filtered with exponent β_{space} in space and β_{time} ($= c.\beta_{space}$ for some suitably chosen constant c) in time to give a real valued three-dimensional field of correlated, normally distributed noise. The three-dimensional field is then "sliced up" in time into its individual images which are then scaled and shifted to achieve the desired image scale statistics. These can then be assembled into an animated image after the work of Brenier (1990) and as such they simulate a simple rainfall event of 2^t images each 2^n x 2^n in size.

For general rainfall it is often not necessary to model the event to temporal resolution as high as 5 minutes, in which case it is possible to increase the time interval between images to 15, 30 or even 60 minute intervals and to reduce the correlation in time (β_{time} parameter) accordingly. This saves on computer memory requirements and hard disk storage space, two precious resources when working on a personal computer. A three-dimensional field of size 128 x 128 x 256 (1 day event at 5 min resolution) requires approximately 70MB of RAM to compute the double precision (8 byte) complex array.

A well known property of Fourier filtered images is their periodic (wrapped) nature. This was exploited in order to build in the advection of the field by cutting off the one edge of the image and pasting it onto its opposite edge as show in Figure 8.7. In this

way the storm was made to drift with any velocity provided that the distance moved by the storm between time steps is small compared to the size of the image.

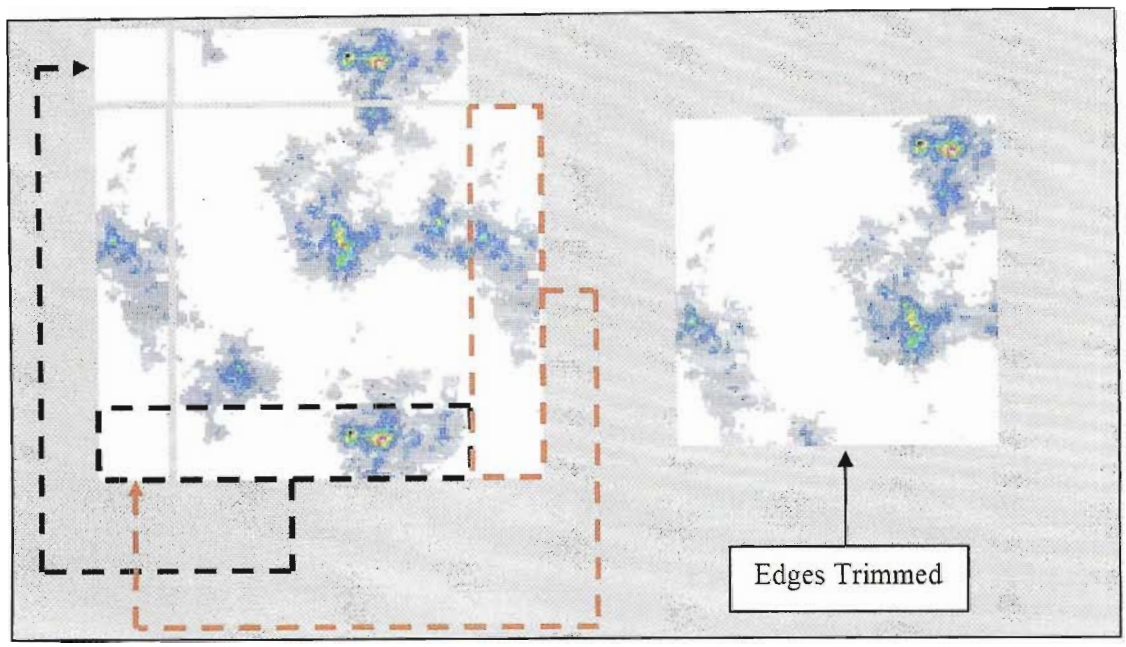

Figure 8.7 - The disused method accounting for storm advection. The trimmed image shows that Fourier periodicity is not obvious if the edges are discarded

The final stage of this modelling process was to discard the edges so that the wrapped nature of the images was no longer obvious. The approach suggested by Bell (1987) was to use only a quarter of the artificial image for simulation, thereby minimising the correlation between the opposite edges. However, since the filter decays exponentially from a point and a typical correlation length for rainfall is 16km, it is very difficult to detect any relationship between the opposite edges of a simulated (128 x 128 km) image after discarding a border as small as $^{1}/_{16}^{th}$ of its width reducing it by 16 pixels in each direction. This is illustrated in the right hand image of Figure 8.7. Since this three-dimensional Fourier-filtered event will also wrap in time, it was also necessary to discard the first or last few images in the generated sequence, in order to ensure that the first and last images of the event were not too similar.

Although this technique produced excellent space-time rainfall events, the algorithm was extremely demanding on memory and processing power and consequently very slow; it was also wasteful because the length of the *bead* had to be simulated as an

integer power of two and then trimmed in time to match the generated wet spell duration (D_w), often discarding large numbers of images.

8.4 *EARLY ATTEMPTS AT MODELLING THE 'STRING'*

Before the development of the renewal process with exponential waiting times described in Chapter 4, only one other model was considered as a serious contender to describe the continuous *String* process and this was the Alternating Renewal process, although several variations of this model were tested.

8.4.1 *Theory of the alternating renewal process*

Consider two sequences of mutually independent and identically distributed random variables $\{X_1, X_2, \ldots\}$ and $\{Y_1, Y_2, \ldots\}$ with probability density functions $f(x)$ and $f(y)$ respectively. If random variables are sampled alternately from the two sequences starting at time 0 with an interval X_1, ending at time X_1 in an *X-to-Y* event, followed by an interval Y_1, ending at time $(X_1 + Y_1)$ in a *Y-to-X* event ...and so on, this is known as an Alternating Renewal Process and is illustrated in Figure 8.8.

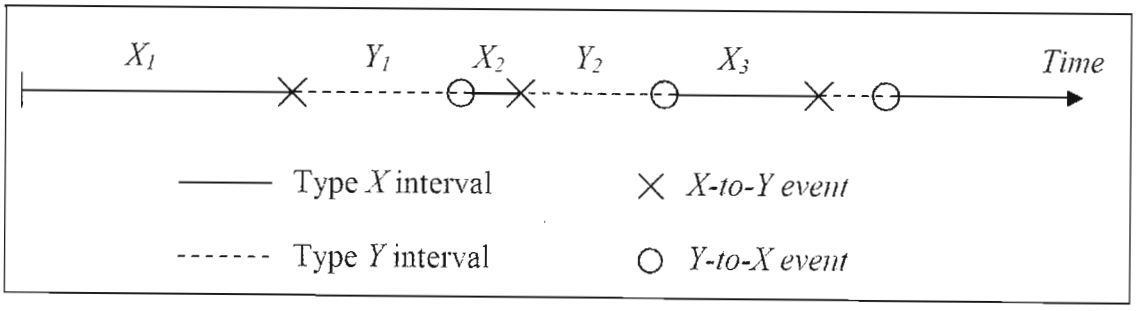

Figure 8.8 - Illustration of an alternating renewal process

The $(2r\text{-}1)^{th}$ interval is of type X and begins after time $(X_1 + \ldots + X_{(r-1)} + Y_1 + \ldots + Y_{(r-1)})$ and the $2r^{th}$ interval is of type Y and begins at time $(X_1 + \ldots + X_r + Y_1 + \ldots + Y_{(r-1)})$, where $r = 1,2,\ldots$.

8.4.2 *Conventional fitting of the alternating renewal process*

If the univariate wet-dry process is truly an alternating renewal process, the wet and dry spell durations can be alternately sampled from two mutually independent processes.

Following the ideas of Harberlandt (1998), the sequence of wet spell duration (D_w) and their complementary dry spell durations (D_d) is an alternating renewal process.

Harberlandt (1998) found that the lognormal distribution was well suited to describing the distribution of wet spell durations in Germany, while the gamma distribution was the best descriptor of the distribution of dry spell durations. The analyses of Section 3.9.1 revealed that these are also reasonable descriptors of the wet and dry spell durations in South Africa. Figure 3.52 also reveals that there is very little difference between the gamma and lognormal approximations of these distributions. For the sake of simplicity, it was assumed that both are lognormally distributed where:

$$D_w \sim \Lambda(\mu_w, \sigma_w)$$

fitted to the wet spell duration and

$$D_d \sim \Lambda(\mu_d, \sigma_d)$$

fitted to the dry spell duration where:

D_w is the duration of the wet spells

D_d is the duration of the dry spells

$\Lambda(\mu, \sigma)$ is the lognormal distribution with parameters μ and σ

μ_w represents the mean of the logs of D_w

σ_w represents the standard deviation of the logs of D_w

μ_d represents the mean of the logs of D_d

σ_d represents the standard deviation of the logs of D_d

With the strong seasonal variation experienced in South Africa, the basic alternating renewal model is over simplistic. It is necessary to account for the fact that, at the very least, the dry spell duration is dependent on the season.

One possible approach would be to allow the parameters of the lognormal distributions, μ_w, σ_w, μ_d and σ_d to vary on a seasonal basis, following the ideas of Woolhiser and Pegram (1979). With enough high resolution rainfall data, from either radar data or an extensive network of tipping bucket raingauges (or similar device), it would theoretically be possible to measure D_w and D_d on a monthly basis and therefore to

extract approximations for μ_w, σ_w, μ_d and σ_d. In practice however, this does not work, at least not for South African data. This became apparent in several failed attempts at calibration of the model which are discussed in the next section.

8.5 *EARLY ATTEMPTS AT 'STRING' CALIBRATION*

The following sections describe failed attempts at calibration of the *String* process. They are included to illustrate the evolution of the model and to demonstrate what does not work and why.

8.5.1 *Run-length String calibration*

The idea behind the earlier run-length calibration technique came from work by Pegram and Seed (1998) who used the statistics of runs of wet and dry days to successfully define and validate a three-state Markov-chain model which classified days as Dry days, Scattered Rain days or General Rain days for a network of raingauge data.

Instead of three states, the run-length calibration algorithm for the *String of Beads* model used only two – wet and dry. A wet day is classified as a day in which WAR_d exceeds 10% where WAR_d is defined as the proportion of the image which receives rainfall in excess of 3,0 mm during the day. Using this threshold, a 50 year time series was extracted from the infilled daily raingauge data received from the Computing Centre for Water Research at the University of Natal, Pietermaritzburg, (since disbanded). Figure 8.9 shows 10 years of the 50 year daily time series of WAR_d as estimated from the 54 raingauge network.

The assumption made in the extraction of this time series was that WAR_d is adequately approximated by the percentage of raingauges which received more than 3mm of rainfall during the day. Based on the analysis presented in Figure 3.2, comparing the WAR_d as estimated from an entire radar image to the WAR_d estimated using only 54 pixels at the positions of the raingauges, this seemed to be a reasonable assumption.

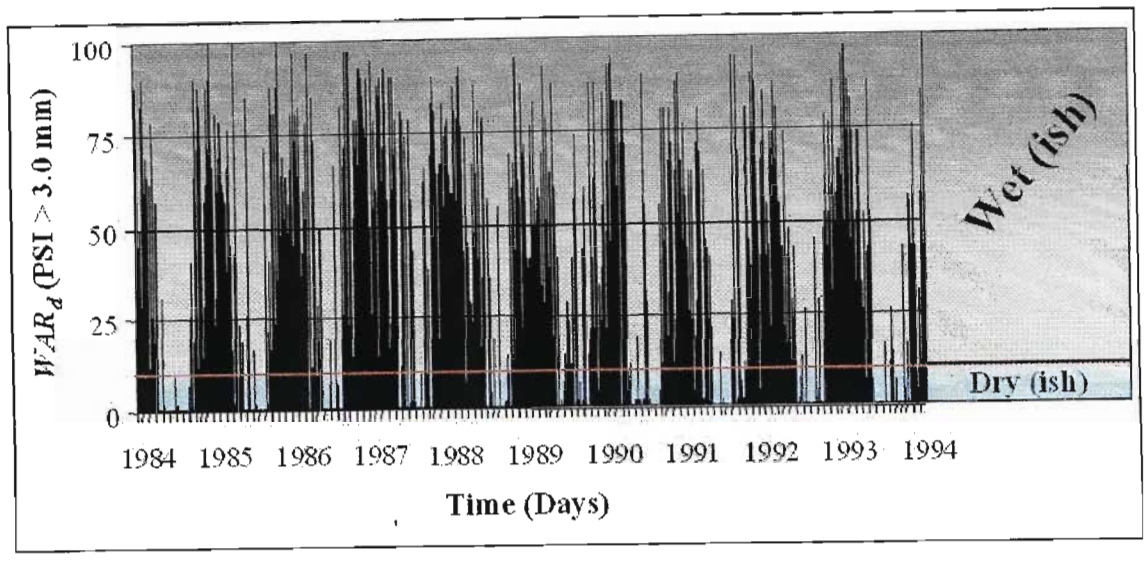

Figure 8.9 - 10 years of the 50 year daily time series of WAR_d as estimated from the 54 raingauge network

Using the wet-dry threshold of 10% exceeding 3mm, as indicated in Figure 8.9, the WAR_d time series can be converted into a 50 year binary time series of wet and dry days. Consecutive wet days are called *Wet Runs* and the number of consecutive wet days in a *Wet Run* is the *Wet Run Length*. Twelve histograms, showing the *observed* frequency of wet runs of a given run length originating in each of the twelve months, were extracted from the binary time series and these are presented in Figure 3.11 through Figure 3.14. These histograms formed one half of the objective function used to calibrate the *String* process.

The second half of the objective function was computed from the WAR_i and IMF_i time series, output at the end of the second stage of simulation. Analysis presented in Section 3.7.3 showed an empirical relationship between WAR_d and $\overline{WAR_i}$, the daily averaged WAR_i. It suggested that the wet or dry daily classification, based on the 10% WAR_d threshold, would be ascertained with a 97% certainty by considering the $\overline{WAR_i}$ for that day. A $\overline{WAR_i}$ below 0.007 suggests a WAR_d below 10% and therefore a *dry day*, and conversely, a $\overline{WAR_i}$ above 0.007 suggests a WAR_d above 10% and therefore a *wet day*. Using this argument, the simulated $WAR_i - IMF_i$ time series output by the second stage of simulation, can be converted to a binary time series of wet and dry days by averaging the WAR_i on a daily basis. As for the raingauge data, twelve histograms can be

extracted from the simulated binary time series, showing the frequency of *simulated* wet runs of a given run length originating in each of the twelve months.

The objective function (ψ) for the calibration process given by Equation 8.8.

$$\psi = \sum_{k=1}^{12} \sum_{j=0}^{25} \left(n_s^{jk} - n_o^{jk} \right)^2 \qquad (8.8)$$

where:

n_s^{jk} represents the number of *wet runs* of length j, which start in month k, measured from the binary time series of the simulation

n_o^{jk} represents the number of *wet runs* of length j, which start in month k, measured from the binary time series of the raingauge data

and $j = 0$ is a dry day.

Recalling the formulation of the first stage of the *String of Beads* model as discussed in Section 4.5.1, the parameters of the model which are calibrated during this iterative *String* calibration process, are the Fourier coefficients which define the seasonal variation of the μ_w, σ_w, μ_d and σ_d. These define the lognormal distributions of the wet and dry spell durations and hence the alternating renewal process. Good starting values for these parameters was thought to be the parameters of the lognormal distributions fitted to all of the observed events, plotted in Figure 3.50 and Figure 3.52. That is, starting values of:

$$\mu_w = 5{,}4 \qquad \sigma_w = 1{,}3 \qquad \mu_d = 5{,}5 \qquad \sigma_d = 1{,}7$$

as measured for South African, summer conditions.

Minimization of the objective function in Equation 8.8 gave reasonable results which are plotted in Figure 8.10 through Figure 8.14.

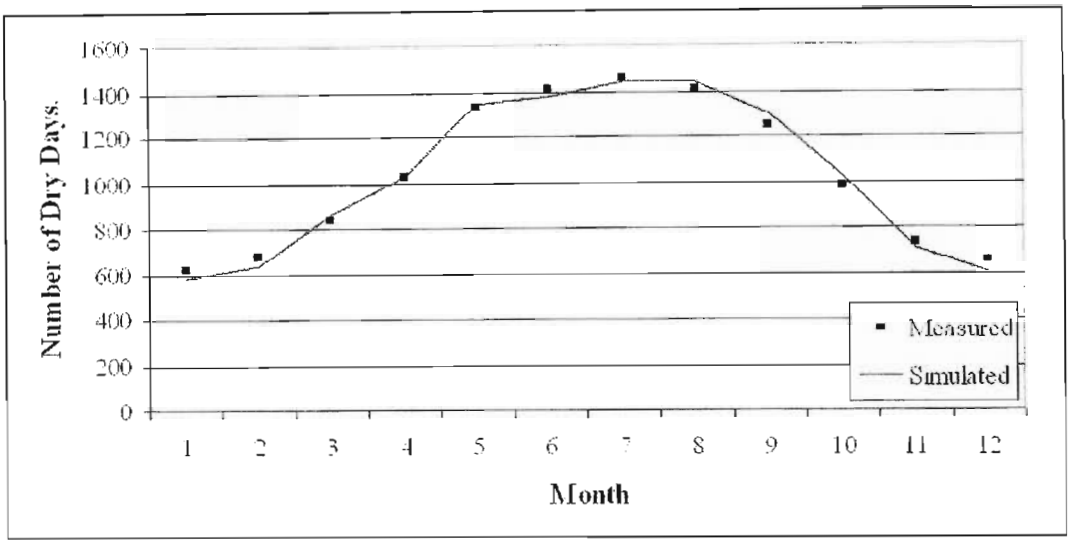

Figure 8.10 – Number of dry days per month observed during the 50 year calibration period

Figure 8.10 shows the comparison between the number of dry days observed in each of the twelve months, over the 50 year calibration period. The complement of the total number of dry days observed in each month is the total number of wet days and clearly, this is reproduced very well by the model. Figure 8.11 through Figure 8.14 show histograms of the number of wet runs of a given length, which originate in each of the twelve months during the 50 year calibration period.

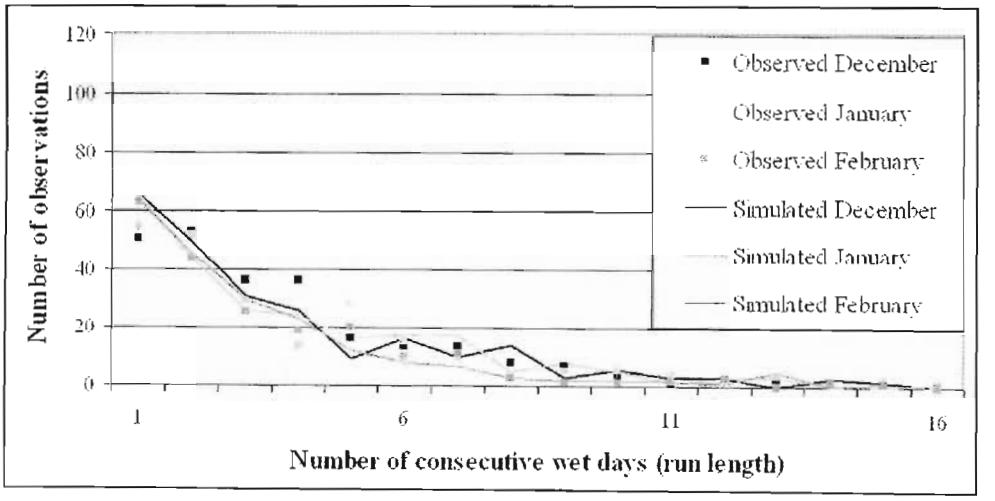

Figure 8.11 - Histogram of Summer Wet Run Lengths

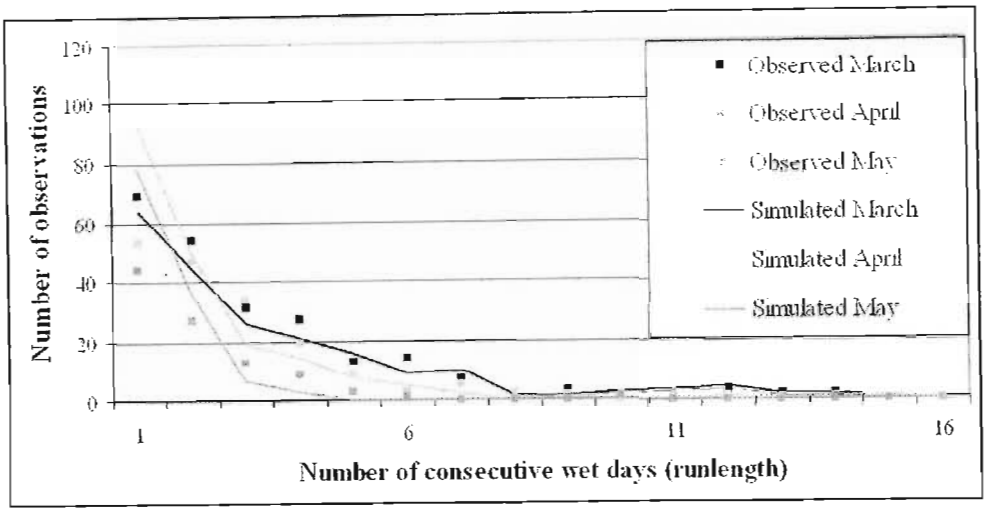

Figure 8.12 - Histogram of autumn Wet Run Lengths

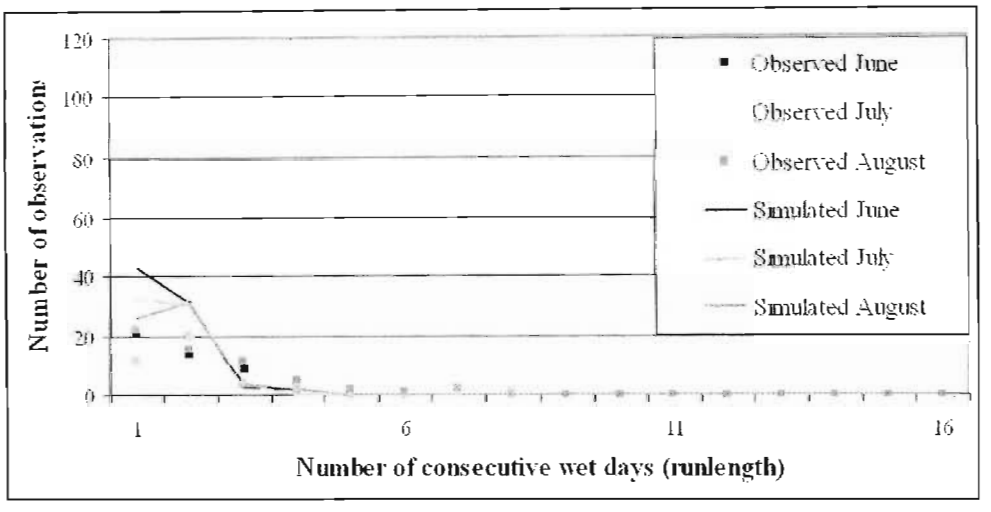

Figure 8.13 - Histogram of Winter Wet Run Lengths

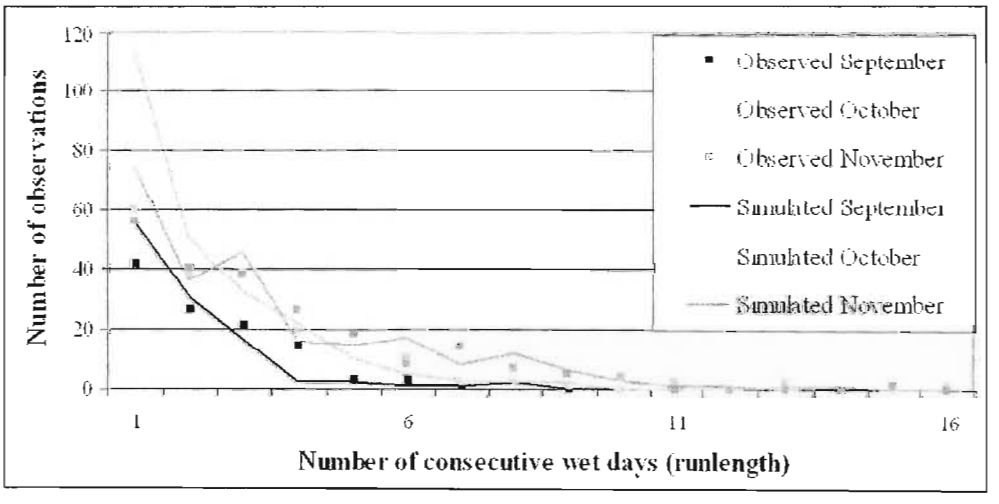

Figure 8.14 - Histogram of Spring Wet Run Lengths

Figure 8.11 to Figure 8.14 are plotted to the same scale to enable easy comparison between the seasons. The model appeared to over-estimate the number of one-day runs in most of the months, but then to compensate by slightly under-estimating the number of longer wet runs.

The attempts at run length calibration process proved to be unstable. A wide range of starting values were tried and the calibration was found to be sensitive to the order of optimisation of the Fourier coefficients. After more than 30 attempts at calibration using the run length objective function of Equation 8.8 (each taking 8 to 10 hours), it was concluded that this objective function should be abandoned. Of the 30 attempts, the best one yields seasonally varying μ_w, σ_w, μ_d and σ_d parameters as shown in Figure 8.15.

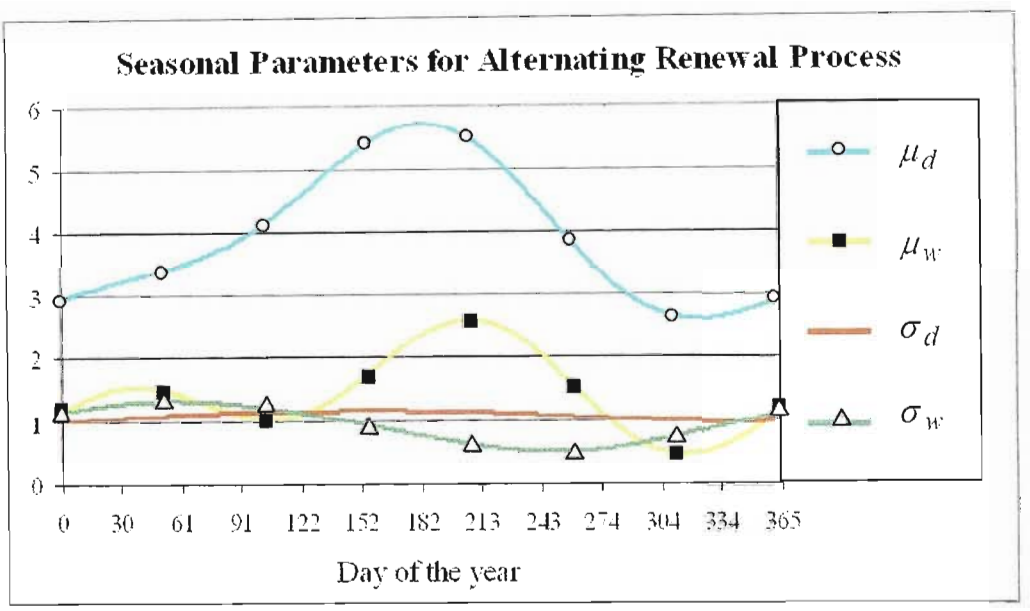

Figure 8.15 - Seasonally varying parameters of the alternating renewal process, obtained through minimization of the run length objective function

At first glance, these parameters appear to be consistent with the type of weather experienced in Bethlehem. The μ_d parameter (the mean of the logs of the durations of the dry spells) is seen to be far greater in the dry winter months than in the wet summer months, as would be expected. The μ_w parameter (the mean of the logs of the durations of the wet spells) increases in the winter months, consistent with the observation that wet spells in the winter months tend to be as a result of large systems. Re-sampling of the wet and dry spell durations after simulation reveals D_w and D_d statistics which are consistent with the calibrated parameters as shown in Figure 8.16. This is a verification

of the model integrity and confirms that the calibrated model is performing as it is supposed to.

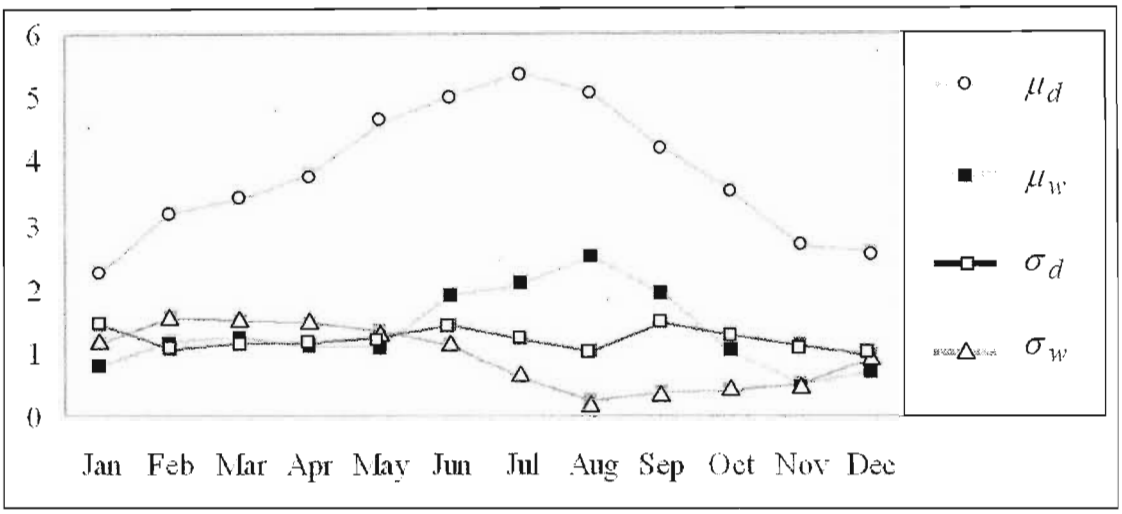

Figure 8.16 – Re-sampled wet and dry spell parameters

Sadly, quantitative validation of the model revealed more desperate shortcomings of the run length calibration process. An eleven year simulation was run at 5 minute resolution using these best available parameters. The idea was to test the output of the model in terms of the distribution (probability of exceedence of a given rainfall rate) of pixel scale intensity, the distribution of monthly rainfall, the distribution of annual rainfall. These were to be compared to the analysis of the radar and raingauge network rainfall data presented in Chapter 3.

Integration of the IMF_i according to Equation 4.11 to obtain the average annual and monthly totals over the study area, (a validation exercise since these were not used as input for the model), showed an under-estimation of the total annual rainfall. Figure 8.17 shows the spatially averaged annual rainfall total for the study area.

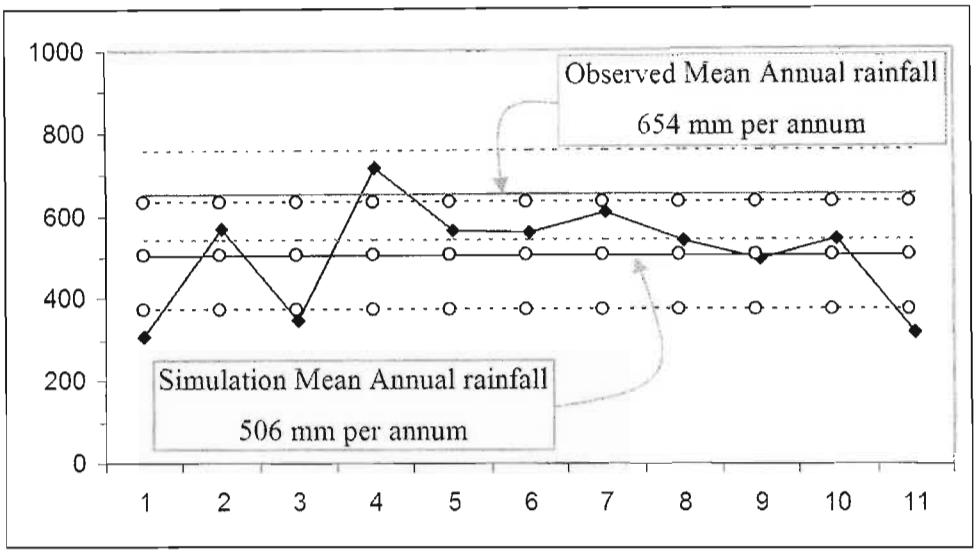

Figure 8.17 - Spatially averaged annual rainfall total for 11 year simulation

The eleven-year simulation would suggest a mean annual rainfall for the region of 506 mm compared to the observed value (as measured by raingauge network) of 654 mm. The two means are marked on Figure 8.17, as well as their corresponding standard deviations (dotted lines) of 129 mm^2 and 108 mm^2 respectively. These compare poorly. Figure 8.18 shows the spatial distribution of rainfall on the image, to be compared to that of Figure 3.45.

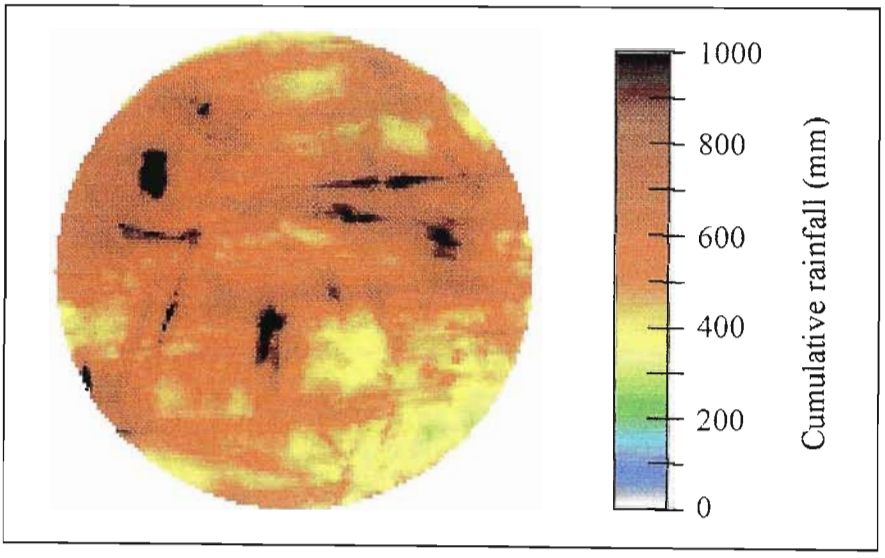

Figure 8.18 - Spatial distribution of rainfall on a simulated annual image

The spatial distribution does have a realistic structure, although the spatial variation is considerably higher in the simulated image of Figure 8.18 when compared to Figure 3.45. Nowhere on Figure 3.45 does the annual rainfall drop below 400 mm as observed in the South East corner of Figure 8.18. Streaks across the image are as a result of very

high intensity simulated rainfall cells, an aspect of the *Bead* simulation process which was later resolved.

Moving on to the monthly statistics, this was the most disappointing aspect of the simulation at that time, and ultimately the reason for abandoning the run length calibration process. Figure 8.19 shows the mean monthly rainfall for the region, averaged over the 11 year simulation.

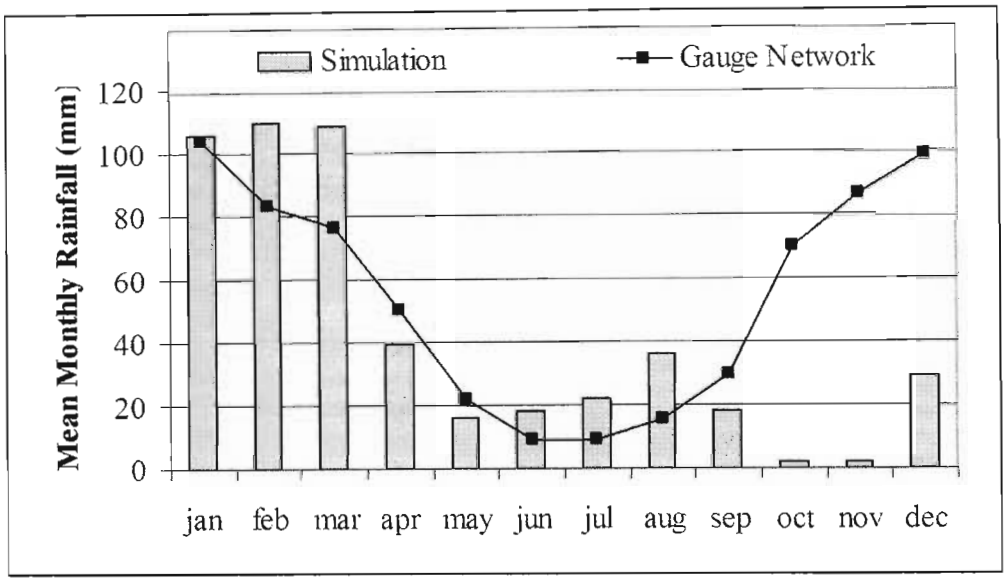

Figure 8.19 – Comparison of the mean monthly rainfall for the region as estimated from the 50 year raingauge network and the 11 year simulation

Simulation between January and September appears to be reasonable, but between September and December it is very poor. With reference to Figure 8.16, the parameters of the wet spell duration, μ_w and σ_w, are very low at that time of the year and the result is a series of extremely short duration (and consequently low IMF) events, resulting in very low rainfall totals. The calibration process using run-lengths was abandoned.

8.5.2 Integrated IMF_i calibration

The next attempted method of calibrating the μ_w, σ_w, μ_d and σ_d parameters which was tested made use of the fact that monthly rainfall totals can be approximated by integrating the IMF_i parameter according to Equation 4.11. The calibration technique was to change the run length objective function of Equation 8.8, to a distribution based objective function as shown in Equation 8.9.

$$\psi = \sum_{k=Jan}^{Dec} \sum_{j=1}^{N} \left(x_s^{jk} - x_o^{jk} \right)^2 \qquad (8.9)$$

where:

k is the month

N is the number of years in the calibration period

j is the rank of the monthly rainfall total (from 1 to N)

x_s^{jk} is the j^{th} largest *Simulated* rainfall total recorded in the k^{th} month

x_o^{jk} is the j^{th} largest *Observed* rainfall total recorded in the k^{th} month

This objective function is distribution independent - an important point when dealing with highly variable marginal distributions such as the monthly rainfall totals. The technique for calibration was then:

1. For the N available years of daily raingauge data, compute the observed monthly rainfall totals for the region by spatially averaging the monthly rainfall totals, measured by a network of daily raingauges which are evenly spread over the region
2. Choose suitable starting values for the μ_w, σ_w, μ_d and σ_d parameters. This is done indirectly by choosing the corresponding Fourier coefficients which define the continuous seasonal variation of these parameters
3. Simulate the WAR_i-IMF_i image scale statistics for a period of N years (50 years in this study) using the *String of Beads* model
4. Compute the total monthly rainfall for the region by integrating the simulated IMF_i for each month in the simulation using Equation 4.11
5. For each calendar month, determine the rank j (j = 1,2,...N) of the rainfall totals recorded for that month, from the simulated and observed sequences
6. Compute the objective function (ψ)
7. Iterate from step 3, to determine the minimum of the ψ objective function by adjusting the Fourier coefficients which determine the seasonally varying μ_w, σ_w, μ_d and σ_d parameters.

This technique eliminated the need for the rather arbitrary classification of wet and dry days and it also eliminated the need for the empirical relationship between WAR_d and $\overline{WAR_i}$. This too proved to be a fruitless exercise.

8.5.3 *Inferred continuous Markov process*

The final abortive attempt at fitting a pure alternating renewal process was fuelled by some encouraging work by Pegram et al (2003), who made use of a result by Small and Morgan (1986) which linked the continuous time Alternating Renewal model to the transition probabilities of a discrete Markov chain. Small and Morgan (1986) consider Markovian transitions of daily raingauge rainfall data between a wet and dry state, and derive closed-form solutions for the continuous Alternating Renewal process in which D_w and D_d are exponentially distributed random variables. Small and Morgan show that in this case:

$$E[D_d] = \frac{-T}{\ln(p_{dd})} \tag{8.10}$$

and

$$E[D_w] = \frac{T\, p_{dw}\, (p_{dw} - p_{ww})}{p_{wd} \ln(p_{dd})} \tag{8.11}$$

where:

T	represents the period of the discrete Markov chain (24 hours in this case)
p_{ij}	represents the transition probability, given that the current state is i, that the following state will be j
d	represents the dry state
w	represents the wet state

Small and Morgan also derived a general solution which is independent of the distribution of D_w and D_d, and although not in closed form as presented for the Markov case in Equations 8.10 and 8.11, it is in a form which is amenable to numerical solution.

Pegram et al. (2003) make use of this result to infer wet and dry spell durations from continuous pluviograph rainfall data recorded over an area 300km in diameter centred on Melbourne, Australia. They investigated the use of exponential, Weibull, Gamma and lognormal distributions as descriptors of the wet spells and dry spells for this purpose and showed that there was very little difference between the exponential and Weibull fits although the lognormal was a poor model for either the dry or wet spells. They showed that wet and dry spell durations were independent of eachother and thus that an Alternating Renewal model is suitable.

Following this lead, in search of a theoretical link between the daily raingauge data and the continuous wet-dry process, transition probabilities for the 50 year daily raingauge dataset at Bethlehem were measured based on the definition of a wet day being one in which WAR_d exceeds 10% where WAR_d is defined as the proportion of the image which receives rainfall in excess of 3,0 mm during the day. These probabilities are presented in Table 8.1 with the corresponding $E[D_w]$ and $E[D_d]$ computed from Equations 8.10 and 8.11.

Table 8.1 -Transition probabilities for the daily wet dry process from Bethlehem raingauge data

	p_{dd}	p_{dw}	p_{wd}	p_{ww}	$E[D_d]$ (hours)	$E[D_w]$ (hours)
January	0.633	0.367	0.228	0.772	**67.54**	**31.29**
February	0.701	0.299	0.275	0.725	**73.55**	**25.59**
March	0.722	0.278	0.321	0.679	**136.15**	**27.05**
April	0.838	0.162	0.376	0.624	**335.72**	**24.51**
May	0.931	0.069	0.452	0.548	**786.38**	**21.17**
June	0.970	0.030	0.512	0.488	**815.06**	**23.92**
July	0.971	0.029	0.483	0.517	**546.83**	**28.00**
August	0.957	0.043	0.437	0.563	**269.30**	**28.90**
September	0.915	0.085	0.405	0.595	**115.31**	**38.57**
October	0.812	0.188	0.292	0.708	**69.89**	**32.75**
November	0.709	0.291	0.272	0.728	**52.49**	**28.40**
December	0.633	0.367	0.256	0.744	**67.54**	**31.29**
				Average	**297.11**	**28.20**

Trends in the monthly $E[D_d]$ and $E[D_w]$ would appear reasonable with a relatively constant $E[D_w]$ throughout the year and a seasonal variation in $E[D_d]$ which is consistent with the observed seasonal behaviour showing significantly longer dry periods in the dry winter months than in the summer months, however the actual values are unrealistic. Wet spells in excess of 20 hours are extremely rare in this region and this number is certainly not the expected value which is estimated at 8,6 hours for the 1998/1999 radar data based on a $WAR_i > 0.01$. Further, dry spells in the summer

months of the order of 70 hours would be unusual as rainfall is experienced somewhere on the catchment almost every day. In the winter months, although rainfall is unusual, a dry spell of the order of 700 hours (29 days) is at the upper end of the observed range rather than the mean.

By adjusting the criteria which define the wet/dry day such as the WAR_d and the rainfall threshold, solutions were obtained which produced realistic distributions of D_w, but these solutions yielded expected values of D_d which were even more unrealistic than those presented in Table 8.1. $E[D_d]$ for these solutions during the dry season were in excess of 140 days with some up to 490 days. A range of thresholds were investigated with no success and the quest was ultimately abandoned.

The problem of fitting a continuous time renewal process to the Bethlehem rainfall data lay in the fact that during the winter months, dry spells in excess of 30 days are common. Consequently the expected value of the fitted distribution was high and the seasonally varying alternating renewal process frequently produced dry spells of hundreds of days – eliminating the possibility of rainfall in the remaining months of the year. The shortfalls of the seasonally varying, alternating-renewal, *String* modelling process, meant that it was ultimately abandoned in favour of calibrating the Scattered event Alternating Renewal process presented in Section 4.5.1.

Chapter 9

Appendix B – Relevant Theory

9.1 INTRODUCTION

This appendix contains the formal theory which has been used in the analysis and definition and formulation of the SBM. It is well documented and included here for ease of reference from the main document. During the course of this work, it has been necessary to fully understand these techniques and their application.

9.2 THE TWO-PARAMETER LOGNORMAL DISTRIBUTION

The two parameter lognormal distribution plays a crucial role in the SBM, primarily in the description of the marginal distribution of rainfall flux on a radar image.

9.2.1 Definition of the two parameter lognormal distribution

The formal definition and properties of the two parameter lognormal distribution are described by Aitchison and Brown (1957) and paraphrased here. Using their notation, consider a positive variate X ($0<x<\infty$) such that $Y = \log X$ is normally distributed with mean μ_y and variance $\sigma_y{}^2$. To simplify the notation, the y subscript will be dropped and the parameters μ_y and $\sigma_y{}^2$ will be written as μ and σ^2 respectively. X is then said to be lognormally distributed, or X is a Λ-variate written as X is $\Lambda(\mu, \sigma^2)$ and correspondingly Y is $N(\mu, \sigma^2)$. The distribution of X is therefore specified by the two parameters μ and σ^2. Clearly X must be positive and non-zero as the transformation $Y = \log X$ is not defined for zero or negative X. The distribution functions of X and Y are then given by $\Lambda(x \mid \mu, \sigma^2)$ and $N(y \mid \mu, \sigma^2)$ respectively so that

$$\Lambda(x \mid \mu, \sigma^2) = P\{X \leq x\} \tag{9.1}$$

and

$$N(y \mid \mu, \sigma^2) = P\{Y \leq y\} \tag{9.2}$$

where $P\{X \leq x\}$ and $P\{Y \leq y\}$ represent the probabilities that $X \leq x$ and $Y \leq y$.

The two parameters of the lognormal distribution μ and σ^2 are related to the mean μ_x and standard deviation σ_x of the untransformed X-variate by the equations:

$$\mu_x = \exp(\mu + {}^1/_2\sigma^2) \tag{9.3}$$

$$\sigma_x = \exp(2\mu + \sigma^2)(\exp(\sigma^2) - 1) \tag{9.4}$$

and the median is at $x = \exp(\mu)$.

9.2.2 *Estimation of the parameters μ and σ^2 for continuous data*

Given a sample S_n of data consisting of n observations described by the real numbers $\{x_1, x_2, \ldots x_n\}$ sampled from $\Lambda(\mu, \sigma^2)$ a method is required for measuring the parameters m and s^2, the most likely estimators of the parameters μ and σ^2. This is achieved by finding the maximum value of the Likelihood function of the sample given as:

$$\frac{1}{\sigma^n.(2\pi)^{n/2}.\prod_{i=1}^{n} x_i}\exp\left\{\frac{-1}{2\sigma^2}\sum_{i=1}^{n}(\log x_i - \mu)^2\right\} \tag{9.5}$$

from which the maximum likelihood estimators $\boldsymbol{m}$ and $\boldsymbol{s^2}$ are found to be:

$$m = \frac{1}{n}\sum_{i=1}^{n}\log x_i \tag{9.6}$$

and

$$s^2 = \frac{1}{n}\sum_{i=1}^{n}(\log x_i - m)^2 \tag{9.7}$$

Equations 9.6 and 9.7 can be used explicitly to find the maximum likelihood estimation of the parameters for continuous data sets.

9.2.3 *Estimation of the parameters μ and σ^2 for integerized data*

The maximum likelihood estimators for the parameters of the lognormal given for the continuous data are not suitable for integer precision data since the lack of precision introduces a significant bias. Instead, the algorithm used for integer precision data was that described by Aitchison and Brown (1957) for fitting the lognormal distribution to grouped data. The multi-dimensional downhill simplex method of Nelder and Mead (1965) as given by Press et al. (1992) was used to maximise the log likelihood function $\boldsymbol{L}$ shown in Equation 9.8 with respect to the parameters μ and σ for each image.

$$L = \sum_i n_i \log\left\{N\left(\frac{y_i - \mu}{\sigma}\middle|(0,1)\right) - N\left(\frac{y_{i-1} - \mu}{\sigma}\middle|(0,1)\right)\right\} \qquad (9.8)$$

where $i = 1, 2, \ldots, 255$

y_i = log (rainfall rate i mm/h)

n_i = number of pixels in the masked image recording a rainfall rate of i mm/h

μ = the mean of the logs of the rainfall rates in the image

σ = the standard deviation of the logs of the rainfall rates in the image.

The maximum value of Equation 9.8 can be found by changing its sign and using a function minimisation routine such as the *Downhill Simplex* method devised by Nelder and Mead (1965), the algorithm for which is given by Press et al. (1992).

9.3 THE FOURIER TRANSFORM

The Fourier transform in its various forms is an essential component of the SBM used for the measurement of the correlation structure (in the form of power spectra) of multi-dimensional fields of data. Its main advantage over traditional time series analysis is the efficiency with which it can be computed.

9.3.1 The continuous Fourier Transform in one dimension

The Fourier Transform from a function of a complex variable t in the *time domain* $h(t)$ to its corresponding function in the *frequency domain* $H(f)$ is defined by Equation 9.9.

$$H(f) = \int_{-\infty}^{+\infty} h(t).e^{2\pi.i.f.t}\,dt \qquad (9.9)$$

Where $e^{2\pi.i.f.t}$ is a complex multiplier which can be represented as

$$e^{i\theta} = \cos\theta + i\sin\theta \qquad \text{where } \theta = 2\pi ft$$

The inverse transform, from the *frequency domain* into the *time domain* is given by Equation 9.10.

$$h(t) = \int_{-\infty}^{+\infty} H(f).e^{-2\pi.i.f.t}\,df \qquad (9.10)$$

9.3.2 *The Discrete Fourier Transform in one dimension*

The rainfall data measured by radar are obtained in discrete space-time intervals and it is therefore necessary to consider the Fourier transform in its discrete form. Following Press et al. (1992), consider the case of a sample set containing N consecutive values in the *time domain* h_k ($\equiv h(t_k)$, k = 0, 1, 2, ..., N-1 , where $t_k = k\Delta$ and Δ is the sampling interval). Assuming that N is even as will be the case in radar image data, with N number of inputs, no more than N independent numbers of output will be produced so that the Fourier Transform $H(f)$ need only be estimated for the discrete values of f_n where

$$f_n = \frac{n}{N\Delta}, \qquad n = -\frac{N}{2},\ldots,\frac{N}{2} \tag{9.11}$$

The extreme values of n in Equation 9.11 correspond to the lower and upper limits of the Nyquist critical frequency (f_c) range. The Fourier Transform for frequencies beyond these limits cannot be calculated due to the fact that critical sampling of a sine wave is two sample points per cycle. The discrete form of the Fourier Transform is given by

$$H(f_n) = H_n = \sum_{k=0}^{N-1} h_k . e^{2\pi . i . k . n / N} \tag{9.12}$$

Where H_n is the *Discrete Fourier Transform* of the N points h_k. The Discrete Fourier Transform maps N complex numbers (h_k 's) in the time domain onto N complex numbers (H_n 's) in the frequency domain independently of the time interval between the sampled points. The inverse of the Discrete Fourier Transform is calculated using Equation 9.13 which will recover the *exactly* the h_k 's from the H_n 's .

$$h_k = \frac{1}{N} \sum_{n=0}^{N-1} H_n . e^{-2\pi . i . k . n / N} \tag{9.13}$$

9.3.3 *The Fast Fourier Transform*

The Fast Fourier Transform is a means of computing the Discrete Fourier Transform using far fewer calculations thereby saving on computing resources. It has been shown that for a sample of N points, the computational effort required for the Fourier Transform is proportional to N^2 calculations. By contrast, the computational effort for

the Fast Fourier Transform algorithm is proportional to $Nlog_2 N$ calculations. Savings made by using this method mean the difference between months and minutes of computation for the large fields of data used in the SBM. The derivation of the *Decimation in Time Fast Fourier Transform algorithm* is covered in many texts (Kuc (1988), for example) and will not be repeated here, but it takes advantage of certain symmetries in the transform by splitting the sequence h_k into two half sequences, even indexed $h_{k(even)}$ and odd indexed $h_{k(odd)}$, which are independently transformed and added together.

9.3.4 The Discrete Transform in two or more dimensions

For the two dimensional case, consider a complex valued function $h(k_1, k_2)$ defined on a two dimensional space where $k_1 = 0,1,\ldots, N_1 - 1$ and $k_2 = 0,1,\ldots, N_2 - 1$, its complex valued two dimensional discrete Fourier Transform $H(n_1, n_2)$ is given by Equation 9.14.

$$H(n_1,n_2) = \sum_{k_2=0}^{N_2-1} \sum_{k_1=0}^{N_1-1} \exp\left(\frac{2\pi.i.k_2.n_2}{N_2}\right).\exp\left(\frac{2\pi.i.k_1.n_1}{N_1}\right).h(k_1,k_2) \qquad (9.14)$$

This can be rearranged in the form of Equation 9.15.

$$H(n_1,n_2) = \sum_{k_2=0}^{N_2-1} \exp\left(\frac{2\pi.i.k_2.n_2}{N_2}\right).\sum_{k_1=0}^{N_1-1} \exp\left(\frac{2\pi.i.k_1.n_1}{N_1}\right).h(k_1,k_2) \qquad (9.15)$$

In this way the two dimensional transform can be computed by taking sequential one dimensional Fast Fourier Transforms on each index of the original function (in this case 1 and 2). The inverse transform $h(k_1, k_2)$ for the two dimensional case is given by:

$$\frac{1}{N_2.N_1} \sum_{k_2=0}^{N_2-1} \sum_{k_1=0}^{N_1-1} \exp\left(-\frac{2\pi.i.k_2.n_2}{N_2}\right).\exp\left(-\frac{2\pi.i.k_1.n_1}{N_1}\right).H(n_1,n_2) \qquad (9.16)$$

The theory presented here for the two dimensional case can be shown to extend to multi-dimensional space. Efficient programming of this algorithm is not a trivial exercise. The routine given by Press et al. (1992) was used for this study.

9.4 RANDOM NUMBER GENERATION

The purpose of this section is to revisit this well worn topic and indicate the care taken to ensure that the simulations of the String of Beads model do not suffer from a poor generator.

9.4.1 *Uniformly distributed random noise*

The randomness of the sequence of pseudo-random numbers produced by the random number generator is of great importance in the simulation process. Algorithms which are built into the compiler software cannot always be relied upon to produce sufficiently random numbers and for this reason one devised by Wichmann and Hill (1982), was adopted to generate *uniformly distributed* random noise.

Several factors were considered in the selection of the pseudo-random number generator. The algorithm needed to be *portable* in the sense that it should be machine independent. Routines which make use of low-level commands, such as bit shifts and other binary arithmetic, rely on a format of data storage which is specific to the computer operating system, the programming language and the brand of compiler. If the routine is to be portable these low-level commands should be avoided. The *randomness* of the original uniformly distributed sequence is of great significance in the model. Simulation of the rainfall for a single day at 128x128 pixel, five minute resolution requires 5 million random numbers. The consequences of using a poor random number generator would be that patterns would soon become apparent in the simulated rainfields. Another important factor considered in the selection of the algorithm was the speed or efficiency. An algorithm which involves too many arithmetic operations would result in a very slow program when generating long sequences of numbers.

The method chosen was that suggested by Wichmann and Hill (1982) who make use of a combination of three *multiplicative congruential* generators. The algorithm AS 183 is published in Applied Statistics Algorithms (Griffiths and Hill, 1995). Uniformly distributed random number generation is discussed at length by Knuth (1969) and his description of the *linear congruential* generator, and the Wichmann and Hill algorithm are outlined here.

The *multiplicative congruential* sequence is a special case of the *linear congruential* sequence (X_n) defined by Equation 9.17.

$$X_{n+1} = (a.X_n + c) \bmod m, \qquad n \geq 0. \tag{9.17}$$

where

X_0	is the starting value or *seed*;	$X_0 \geq 0$
a	is the *multiplier*;	$a \geq 0$
c	is the *increment*;	$c \geq 0$
m	is the *modulus*;	$m > X_0, \quad m > a, \quad m > c$

and the "mod" function is the remainder when a number is divided by a divisor.

It can be shown that all such sequences will eventually repeat themselves for a large enough n known as the *period*. A larger period implies a *more random* sequence of numbers. The choice of these four parameters is critical to the period of the sequence and an example given by Knuth (1969) of a *poor* choice of parameters is when $X_0 = a = c = 7$, $m = 10$. In this case the second value of the sequence (X_1) would be calculated as

$$\begin{aligned} X_1 &= (7.7 + 7) \bmod 10 \\ &= 56 \bmod 10 \\ &= 6 \end{aligned}$$

and the entire sequence X_n would be

$$X_n = 7, 6, 9, 0, 7, 6, 9, 0, 7, 6, 9, 0 \ldots \qquad (n \geq 0)$$

This sequence has a period of length 4 and is therefore not useful as a random number generator.

For the special case in which $c = 0$, the number generation is slightly faster with a small sacrifice in the length of the period. This is often referred to as the *multiplicative congruential* method and is the only case that will be considered in this document. Equation 9.17 then reduces to Equation 9.18.

$$X_{n+1} = a.X_n \bmod m, \qquad n \geq 0. \tag{9.18}$$

The modulus m should be chosen as a large number since the period can never have a value larger than m. Furthermore, if m is a prime number the maximum period possible is $(m - 1)$ and this can be achieved provided that two conditions are satisfied. These are

(1) X_0 is relatively prime to m

(2) a is a *primitive root* of m

If m is chosen as a prime number then X_0 will automatically be prime relative to m. The number x is said to be a *primitive root* of m if there exists an integer value k which satisfies the relationship given in Equation 9.19.

$$x^k \bmod m = 1 \qquad (9.19)$$

At least one solution (x) to Equation 9.19 exists for all prime numbers m.

Wichmann and Hill (1982) make use of three multiplicative congruential generators in combination each of which uses a prime number for its modulus and a corresponding primitive root for its multiplier. The modulus and its corresponding multiplier for each of the generators are listed as

	modulus	*multiplier*
(1)	30269	171
(2)	30307	172
(3)	30323	170

Although the sum of n independent rectangular distributions will tend to a normality as n tends to infinity, the *fractional part* of a sum of rectangular distributions will remain rectangular for all n. This property is exploited by Wichmann and Hill (1982) in order to compensate for imperfections in the randomness of the individual generators. Each generator is seeded with an initial value between 1 and 30000, X_{10} , X_{20} , X_{30} respectively. The results of the three generators X_{1n} , X_{2n} , X_{3n} are then added and the fractional part of the sum is taken as the uniformly distributed pseudo-random number $R(0,1)$. They claim that since each generator has a period of 1 less than its prime modulus, the periods are all of even length and therefore have a common factor of 2 but no other common factors exist for the values chosen. For this reason the three generators are not completely independent but the results have been well tested and perform very well.

Although the multiplicative congruential generator is cyclic in nature, the Wichmann and Hill algorithm has been well tested and they claim a period exceeding 6.95×10^{12}. Considering that ($256 \times 256 \times 12 \times 24 =$) 1.9×10^7 random numbers are required for a single day's simulation at 128 x128 pixel, 5 minute resolution, about 1000 years of simulated rainfall data of a *continuous* wet period could be generated before the pseudo random sequence starts to repeat itself and a cyclic behaviour would be observed in the images.

9.4.2 *Percentage points of the normal distribution*

Having established an effective means of generating uniformly distributed random noise, the next step is to transform it into normally distributed noise by means of an inverse normal transformation. The probability density function for the standard normal distribution (zero mean, unit variance) can be represented in the cumulative form of Figure 9.1 which is described by Equation 9.20.

$$p = P(x < x_p) = \int_{-\infty}^{x_p} \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{x^2}{2}\right] dx \qquad (9.20)$$

where p is the probability that a randomly chosen number x is less than a chosen x_p.

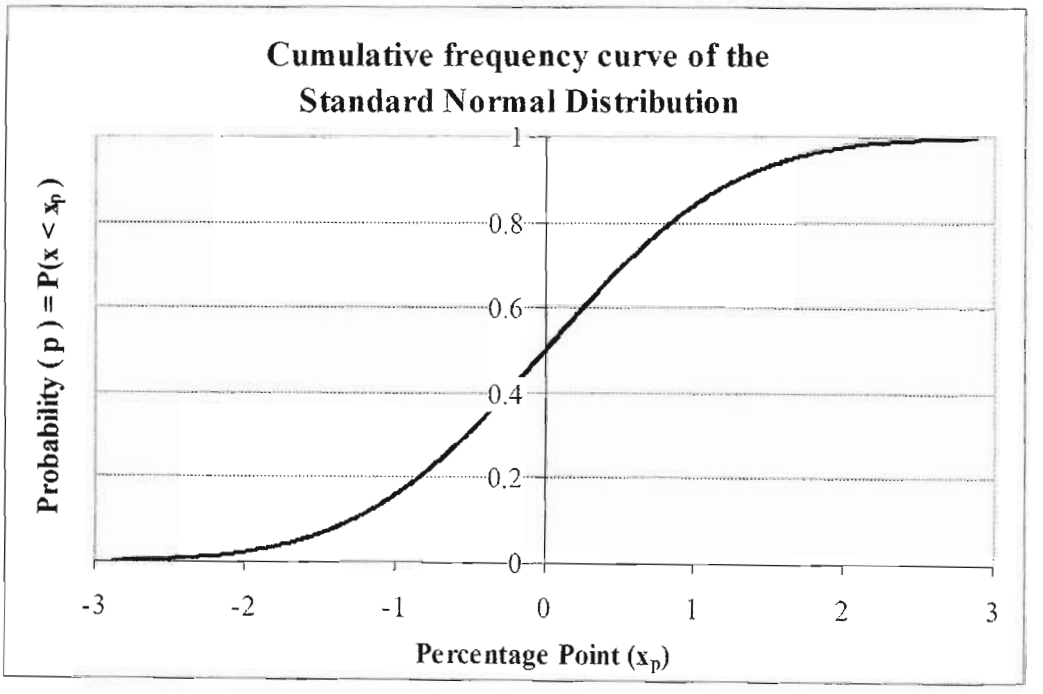

Figure 9.1 - Cumulative frequency curve for the Standard Normal distribution

A uniformly distributed random variate $Y \sim U(0,1)$ can be converted to a standard normally distributed random variate $X \sim N(0,1)$ by setting $Y = p$ in Equation 9.20 and solving for $X = x_p$.

Equation 9.20 has no known solution and must therefore be approximated. A good approximation is given by Beasley and Springer (1977) and published as AS111 in Applied Statistics Algorithms (Griffiths and Hill, 1995). They claim that in the absence of rounding error, the x_p calculated for a given p will correspond to a true value p' which will satisfy the relationship

$$|p' - p| < 2^{-31}$$

In the presence of rounding error, they claim an accuracy which will satisfy the relationship

$$|p' - p| < \varepsilon \qquad (9.21)$$

Where ε is of the order of 20 times the value of the smallest digit in the mantissa. The routine was re-written in the C programming language using 64 bit double precision with an 11 bit exponent and a 52 bit mantissa, the remaining bit used for sign. In this case the value of ε would be 20 x 2^{-52} which is much less than 2^{-31} so rounding error can be neglected in this case and the algorithm is used to its full potential.

9.5 *MULTIVARIATE AUTOREGRESSIVE PROCESSES*

Both the univariate and the multivariate autoregressive processes are used in the SBM. The multivariate theory presented here can be easily reduced to the univariate case. Following the text of Wei (1990), consider the zero mean, unit variance case, $\dot{Z}_t = \dfrac{Z_t - \mu_z}{\sigma_z}$, the p^{th} order, q dimensional autoregressive model, AR(p,q), is defined by:

$$\dot{Z}_t = \Phi_1 \dot{Z}_{t-1} + \Phi_2 \dot{Z}_{t-2} + \cdots + \Phi_p \dot{Z}_{t-p} + A_t \qquad (9.22)$$

where $\dot{Z}$, Φ_p and a_t are [q x 1], [q x q] and [q x 1] matrices respectively. For example, for the AR(1,2) case:

$$\dot{Z}_t = \Phi_1 \dot{Z}_{t-1} + A_t \qquad (9.23)$$

or, close up:

$$\begin{bmatrix} z_1(t) \\ z_2(t) \end{bmatrix} = \begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix} \begin{bmatrix} z_1(t-1) \\ z_2(t-1) \end{bmatrix} + \begin{bmatrix} a_1(t) \\ a_2(t) \end{bmatrix} \qquad (9.24)$$

For the general case, the Φ_p matrices can be computed by solving the generalized Yule-Walker equations:

$$\begin{bmatrix} \Gamma_{(0)} & \Gamma_{(1)} & \Gamma_{(2)} & \cdots & \Gamma_{(p-1)} \\ \Gamma_{(1)}^T & \Gamma_{(0)} & \Gamma_{(1)} & \cdots & \Gamma_{(p-2)} \\ \Gamma_{(2)}^T & \Gamma_{(1)}^T & \Gamma_{(0)} & \cdots & \Gamma_{(p-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{(p-1)}^T & \Gamma_{(p-2)}^T & \Gamma_{(p-3)}^T & \cdots & \Gamma_{(0)} \end{bmatrix} \begin{bmatrix} \Phi_1 \\ \Phi_2 \\ \Phi_3 \\ \vdots \\ \Phi_p \end{bmatrix} = \begin{bmatrix} \Gamma_{(1)}^T \\ \Gamma_{(2)}^T \\ \Gamma_{(3)}^T \\ \vdots \\ \Gamma_{(p)}^T \end{bmatrix} \qquad (9.25)$$

where $\Gamma_{(k)}$ represents the covariance matrix of $\dot{Z}$ for lag k and will be [q x q].

$$\Gamma_{(k)} = \begin{bmatrix} \gamma_{11}(k) & \gamma_{12}(k) & \cdots & \gamma_{1q}(k) \\ \gamma_{21}(k) & \gamma_{22}(k) & \cdots & \gamma_{2q}(k) \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{q1}(k) & \gamma_{q2}(k) & \cdots & \gamma_{qq}(k) \end{bmatrix} = Corr(\dot{Z}_{t-k}, \dot{Z}_t) \qquad (9.26)$$

and the $\gamma_{xy}(k)$ are computed as $Corr(z_x(t), z_y(t-k))$. Solving Equation 9.25 for Φ_p terms:

$$\begin{bmatrix} \Phi_1 \\ \Phi_2 \\ \Phi_3 \\ \vdots \\ \Phi_p \end{bmatrix} = \begin{bmatrix} \Gamma_{(0)} & \Gamma_{(1)} & \Gamma_{(2)} & \cdots & \Gamma_{(p-1)} \\ \Gamma_{(1)}^T & \Gamma_{(0)} & \Gamma_{(1)} & \cdots & \Gamma_{(p-2)} \\ \Gamma_{(2)}^T & \Gamma_{(1)}^T & \Gamma_{(0)} & \cdots & \Gamma_{(p-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{(p-1)}^T & \Gamma_{(p-2)}^T & \Gamma_{(p-3)}^T & \cdots & \Gamma_{(0)} \end{bmatrix}^{-1} \begin{bmatrix} \Gamma_{(1)}^T \\ \Gamma_{(2)}^T \\ \Gamma_{(3)}^T \\ \vdots \\ \Gamma_{(p)}^T \end{bmatrix} \qquad (9.27)$$

The shock term of Equation 9.23, A_t, can be expressed as:

$$A_t = \Sigma.\bar{e} = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1q} \\ \sigma_{21} & \sigma_{12} & \cdots & \sigma_{2q} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{q1} & \sigma_{q2} & \cdots & \sigma_{qq} \end{bmatrix} \begin{bmatrix} e_1(t) \\ e_2(t) \\ \vdots \\ e_q(t) \end{bmatrix} \qquad (9.28)$$

where $e_1(t) \ldots e_m(t)$ are independent, identically distributed variables N(0,1) and the Σ matrix is found by Equation 9.29:

$$\Sigma = \Gamma_{(0)} - \begin{bmatrix} \Phi_1 & \Phi_2 & \Phi_3 & \cdots & \Phi_p \end{bmatrix} \begin{bmatrix} \Gamma_{(0)} & \Gamma_{(1)} & \Gamma_{(2)} & \cdots & \Gamma_{(p-1)} \\ \Gamma_{(1)}^T & \Gamma_{(0)} & \Gamma_{(1)} & \cdots & \Gamma_{(p-2)} \\ \Gamma_{(2)}^T & \Gamma_{(1)}^T & \Gamma_{(0)} & \cdots & \Gamma_{(p-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{(p-1)}^T & \Gamma_{(p-2)}^T & \Gamma_{(p-3)}^T & \cdots & \Gamma_{(0)} \end{bmatrix} \begin{bmatrix} \Phi_1^T \\ \Phi_2^T \\ \Phi_3^T \\ \vdots \\ \Phi_p^T \end{bmatrix} \qquad (9.29)$$

The optimal order of the autoregressive model (p) can be chosen by minimising the Corrected Akaike Information Criterion (Hurvich and Tsai, 1989) which is defined by the relationship:

$$AIC_c = N \log\left(\frac{S}{N}\right) + \frac{N(N+p)}{N-p-2} \qquad (9.30)$$

where:

$$S = \sum_{i=1}^{N} \varepsilon_i^{\,2} \qquad (9.31)$$

which is the sum of the squares of the residual values (ε_i) which represent the difference between the fitted and the observed sequences.

Optimisation is achieved by computing the Φ_p terms of an AP(p) model from an observed sequence of N values, $\dot{X}_i$ (were i = 1,2,..., N) using Equation 9.27. A sequence of residuals (ε_i) are then extracted by rearranging Equation 9.22 so that

$$E_i = \begin{bmatrix} \varepsilon_1(i) \\ \varepsilon_2(i) \\ \vdots \\ \varepsilon_q(i) \end{bmatrix} = \dot{X}_i - \left\{\Phi_1 \dot{X}_{i-1} + \Phi_2 \dot{X}_{i-2} + \cdots + \Phi_p \dot{X}_{i-p}\right\} \quad (4.18)$$

If the AR model is a good descriptor of the process, the ε_i should, ideally, be normally distributed.

9.6 *VARIANCE CORRECTION FOR THE FORWARD/REVERSE COMBINED AUROREGRESSIVE PROCESS*

The average of two independent time series results in a distorted variance when considering the combined time series. This is important when considering the forward/reverse combined autoregressive process used in the simulation of the string process.

Consider two autoregressive sequences X_t and Y_t, both of which are N(0,1).

$$X_t = \phi_1.X_{t-1} + \ldots + a_t.\sigma_\varepsilon \qquad \text{(forward sequence)}$$

$$Y_t = \phi_1.Y_{t-1} + \ldots + b_t.\sigma_\varepsilon \qquad \text{(reverse sequence)}$$

Where a_t and b_t are the are the shock terms which are N(0,1) and are multiplied by σ_ε, the required standard deviation of the noise term obtained explicitly from the ϕ_i poles and $\sigma_x = \sigma_y = 1$.

Consider the linear combination of these two sequences Zt such that

$$Z_t = \gamma.X_t + (1\text{-}\gamma)\, Y_t \qquad 0 \le \gamma \le 1$$

Dropping the t subscript, Z has a variance $E[Z^2].\sigma_x = E[Z^2]$

$$\begin{aligned} E[Z^2] &= E[(\gamma.X + (1\text{-}\gamma)\, Y)^2] \\ &= E[(\gamma^2.X^2 + 2\gamma.(1\text{-}\gamma).XY + (1\text{-}\gamma)^2.Y^2] \\ &= \gamma^2.E[X^2] + (1\text{-}\gamma)^2.\, E[Y^2] \qquad \text{(since } X \text{ and } Y \text{ are independent)} \end{aligned}$$

Consequently

$$E[Z^2] = \gamma^2 + (1-\gamma)^2 \quad (\text{since } E[X^2] = E[X^2] = 1)$$
$$= 1 - 2\gamma + 2\gamma^2$$

Therefore the combined sequence Z should be scaled by $\sqrt{1 - 2\gamma + 2\gamma^2}$ in order to achieve the desired variance in the sequence and the combined sequence is given by:

$$Z = (\gamma X + (1-\gamma)Y)w_3$$

where

$$w_3 = \frac{1}{\sqrt{1 - 2\gamma + 2\gamma^2}}$$

Chapter 10

References

Aitchison, J. and J.A.C. Brown. (1957). *The lognormal distribution with special reference to its uses in economics*, University of Cambridge, Department of Applied Economics, Monograph: 5, Cambridge University Press.

Anselmet F., Y. Gagne, E.J. Hopfinger, and R. Antonia. (1984). High-Order Velocity Structure Functions in Turbulent Shear Flows. *J. Fluid Mech.*, Vol. 140, p63.

Austin, G. L. (1998). History of Radar and Radar Meteorology, pp 3-15 in *Radar Hydrology for Real Time Flood Forecasting* - R J Griffith, I D Cluckie, G L Austin and D Han editors - European Commission Publications (EUR 19888) ISBN 92-894-1640-8, 2001.

Bardossy, A. and E.J. Plate.(1992). Space-Time Model for Daily Rainfall using Atmospheric Circulation Patterns. *Water Resources Research*. Vol. 28. pp 1247-1259.

Basson M.S., R.B. Allen, G.G.S. Pegram and J.A. van Rooyen.(1994). *Probabilistic Management of Water Resource and Hydropower Systems*. Water Resources Publications, Highlands Ranch, Colorado, USA.

Battan, L.J. (1973). *Radar Observation of the Atmosphere.* University of Chicago Press, Chicago, USA.

Beasley, J.D. and S.G. Springer. (1985). Algorithm AS111 - The Percentage Points of the Normal Distribution, *Applied Statistics Algorithms,* P Griffiths and I D Hill (eds). Royal Statistical Society.

Bell, T.L. (1987). A Space-Time Stochastic Model of Rainfall for Satellite Remote Sensing Studies, *J. Geophys. Res.*, Vol. 92, No. D8, p 9631.

Benzi, R., L. Biferale, A. Crisanti, G. Paladin, M. Vergassola and A. Vulpiani. (1993). A Random Process for the Construction of a Multiaffine Field, *Physica* D, Vol. 65, p353.

Box, G.E.P. and G.M. Jenkins. (1970). *Time Series Analysis: Forecasting and Control.* Holden Day, San Francisco.

Bras, R.L. and I. Rodriguez-Iturbe. (1985). *Random Functions and Hydrology.* Addison-Wesley.

Brenier, P. (1990). *Simulations Dynamique Multifractale Des Nuages.* MS Thesis. Ecole Normale de Science et Technologie, Paris, France.

Cooley, J.W. and J.W. Tukey. (1965). An algorithm for the machine calculation of complex Fourier series. *Math. Comp*. 19.

Court, A.P. (1979). Rainfall Characteristics of Classification Systems used by the BEWMEX project. *Bethlehem Weather Modification Experiment.* Progress report no 14. Weather Bureau, Pretoria, South Africa.

Cowpertwait, P.S.P. (1994). A Generalised Point Process Model for Rainfall. *Proc. Royal Soc. London,* A, Vol. 447, pp 23-37.

Cowpertwait, P.S.P. (1995). A Generalised Point Process Model for Rainfall. *Proc. Royal Soc. London,* A, Vol. 450, pp 163-175.

Cowpertwait, P.S.P, P.E. O'Connell, A.V. Metcalfe and J.A. Mawdsley . (1996a). Stochastic Point-process modelling of rainfall. I. Single Site Fitting and Validation. *Journal of Hydrology*. Vol 175. pp17-46.

Cowpertwait, P.S.P., P.E. O'Connell, A.V. Metcalfe and J.A. Mawdsley . (1996b). Stochastic Point-process modelling of rainfall. II. Regionalisation and disaggregation. *Journal of Hydrology*. Vol 175. pp47-65.

Cowpertwait, P.S.P. and P.E. O'Connell. (1997). A Regionalised Neyman-Scott Model of Rainfall with Convective and Stratiform Cells. *Hydrology and Earth System Sciences*. Vol. 1. pp 71-80.

Cox, D.R. and V. Isham. (1988). A simple spatial-temporal model of rainfall. *Proc. Roy. Soc. London.*, Vol A415, pp317-328

Crane, R.K. (1990). Space-Time Structure of Rain Rate Fields. *J. Geophys. Res.*, Vol. 95, No. D3, pp 2011-2020.

Cressie, N.A.C. (1991). *Statistics for spatial data*. Wiley. Chichester, UK

Dempster A.P., N.M. Laird and D.B. Rubin. (1977). Maximum likelihood estimation from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*. Series B, Vol. 39, p1-38.

Dixon, M. and G. Wiener. (1993). TITAN: Thunderstorm Identification, Tracking, Analysis and Nowcasting - a Radar Based Methodology. *J. Atmos. Ocean. Tech.*, Vol. 10, No.6, pp785-797.

Foufoula-Georgiou, E., and W. Krajewski. (1995). Recent Advances in Rainfall Modelling, Estimation and Forecasting. *Reviews of Geophysics*, Supplement, pp1125-1137, July 1995.

Fowler, H.J., C.G. Kilsby and P.E. O'Connell. (2000). A stochastic rainfall model for the assessment of regional water resource systems under changed climatic conditions. *Hydrology and Earth System Sciences*, Vol. 4, No. 2, pp263-282.

Gabriel, K.R. and J. Neumann. (1992). A Markov Chain Model for Daily rainfall occurrences at Tel-Aviv. *Q. J. R Meteorol. Soc.* Vol 88. pp 85-90.

Griffiths, P. and I.D. Hill. (1995). (eds). *Applied Statistics Algorithms*. Royal Statistical Society.

Gupta, V.K. and E.C. Waymire. (1987). On Taylor's Hypothesis and Dissipation in Rainfall. *J. Geophys. Res.*, Vol. 92 , No. D8, pp 9657-9660.

Gupta, V.K. and E.C. Waymire. (1990). A Statistical Analysis of Mesoscale Rainfall as a Random Cascade. *J. App. Meteorology*, Vol. 32, pp251-267.

Haberlandt, U. (1998). Stochastic processes using regionalized model parameters. *J. Hydrologic Eng.* 3 (3). p160-168.

Han, D., I.D. Cluckie, R.J. Griffith and G.L. Austin. (2000). Using Weather Radars to Measure Rainfall in Urban Catchments. *J. Urban Technology*, Vol 7, No. 1, pp85-102.

Hurvich, C.M. and C.L. Tsai. (1989). Regression and Time Series Model Selection in Small Samples, Biometrika, v76, n2, p297-307

Jenkins, G.M. and D.G. Watts. (1967). *Spectral analysis and its applications*. Chap 5, Holden-day, Oakland, Calif.

Katz, R.W. and M.B. Parlange. (1995). Generalisations of chain-dependent processes: applications to hourly precipitation. Water Resources Research. Vol 31. pp1331-1341.

Knuth, D.E. (1969). *The Art of Computer Programming. Vol. 2 / Seminumerical Algorithms*. pp 9-34, Addison-Wesley.

Kuc, R. (1988). *Introduction to Digital Signal Processing*. McGraw Hill, New York.

Le Cam, L. (1961). A stochastic description of precipitation. In *Proceedings of the fourth Berkley symposium on Mathematics, statistics and probability*. Vol. 3. University of California, Berkley CA. pp 165-186.

Makhuva T., G.G.S. Pegram, R. Sparks and W. Zucchini. (1997a). Patching rainfall data using regression methods. 1. Best subset selection, EM and pseudo-EM methods: theory. *Journal of Hydrology*, 198, p289-307

Makhuva T., G.G.S. Pegram, R. Sparks and W. Zucchini. (1997b). Patching rainfall data using regression methods. 2. Comparisons of accuracy, bias and efficiency. *Journal of Hydrology*, 198, p308-318

Mandelbrot, B.B. (1983). *The fractal Geometry of Nature*. Updated and augmented ed. Freeman, San Fransisco.

Marshall, J.S. and W.M. Palmer. (1948). The Distribution of Raindrop Size. *J. Meteor.* Vol. 5, p165.

Marsan, D., D. Schertzer and S. Lovejoy. (1996). Causal space-time multifractal processes: Predictability and forecasting rain fields. *J. Geophys. Res.*, 101, pp26333 - 26346.

Matheron, G. (1973). The Intrinsic Random Functions and Their Applications. *Adv. Appl. Prob.* 5: pp439-468.

Menabde, M. (1998). Bounded Lognormal Cascades as Quasi-Multiaffine Random Processes. *Nonlinear Processes in Geophysics*, Vol 5, pp 63 – 68.

Menabde, M., D. Harris, A.W. Seed and G. Austin. (1997). Self-similar Random Fields and Rainfall Simulation. *J. Geophys. Res.* Vol. 102. D12. p13509.

Menabde, M., A.W. Seed, D. Harris and G. Austin. (1999). Multiaffine Random Field Model of Rainfall. *Water Resour. Res.* Vol. 35, No. 2, pp509-514.

Mittermaier, M.P. (1999). *Investigating the Characteristics of the Radar Reflectivity Profil*e. MScEng dissertation, Civil Engineering, University of Natal, Durban

Mittermaier, M.P. and D.E. Terblanche. (1997). Converting Weather Radar Data into Cartesian Space: A New Approach Using DISPLACE Averaging. *Water SA*, Vol. 23, No. 1, pp46-50.

Nelder, J.A. and R. Mead. (1965). A Simplex Method for Function Minimisation. *Computer Journal*, Vol. 7, pp. 308-313.

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

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

Over, T.M. and V.K. Gupta. (1994). Statistical Analysis of Mesoscale Rainfall: Dependence of a Random Cascade Generator on Large Scale Forcing. *J. Appl. Meteorology*. Vol. 33. pp1526-1542.

Press, W., S.A. Teukolsky, W.T. Vetterling and B.P. Flannery. (1992). *Numerical Recipes in C - The art of scientific computing*. 2nd ed. Cambridge University Press.

Pegram, G.G.S. (1980). An Auto-regressive Model For Multi-lag Markov Chains. *J. Applied Probability.* Vol. 17. pp350-362.

Pegram, G.G.S. (1997). Patching rainfall data using regression methods. 3. Grouping, patching and outlier detection. *Journal of Hydrology*, 198, p319-334

Pegram, G.G.S. and A.N. Clothier. (1998). Space-Time Modelling of Rainfall in Fine Intervals: The "String of Beads" model. *Proceedings of 4th Int. Symp. on Hydrological Applications of Weather Radar*, San Diego, CA, 5th to 9th April.

Pegram, G.G.S. and A.N. Clothier. (1999). *High resolution Space-Time Modelling of Rainfall: The "String of Beads" model*. WRC Report No. 752/1/99. Report to the Water Research Commission, Pretoria, South Africa..

Pegram, G.G.S. and A.N. Clothier. (2001a). High resolution Space-Time Modelling of Rainfall: The "String of Beads" model. *J. Hydrology*. V 241. p26-41. (Journal of Hydrology reprint included in final pages of this document).

Pegram, G.G.S. and A.N. Clothier. (2001b). Downscaling Rainfields in Space and time, using the String of Beads Model in Causal Mode. *Hydrology and Earth System Sciences,* Vol. 5, No. 2, pp175-186. (HESS reprint included in final pages of this document).

Pegram, G.G.S. and M.P. Mittermaier. (1998). Estimating Rainfall at Low Altitude at Distance, From Radar CAPPI Data. *Proceedings of 4th Int. Symp. on Hydrological Applications of Weather Radar,* San Diego, CA, 5th to 9th April.

Pegram, G.G.S. and A.W. Seed. (1998). *The Feasibility of Stochastically Modelling the Spatial and Temporal Distribution of Rainfields.* WRC Report No. 550/1/98. Report to the Water Research Commission, Pretoria, South Africa.

Pegram, G.G.S., A.W. Seed and A.N. Clothier. (1997). Temporal and Spatial Modelling of Rainfields Using Fractals. *Proceedings of the 8th South African National Hydrology Symposium* (17th - 19th November 1997), Pretoria, South Africa.

Pegram, G.G.S., A.W. Seed, R. Srikanthan, and Chiew. (2003). Modelling Areal Rainfall with Pluviograph Data.

Rodrigues-Iturbe, I., D.R. Cox and V. Isham. (1987). Some Models for Rainfall based on Stochastic Point Processes. *Proc. Roy. Soc. London.* A410, pp269-288.

Schertzer, D. and S. Lovejoy. (1987). Physical Modelling and Analysis of Rain and Clouds by Anisotropic Scaling Multiplicative Processes. *J. Geophys. Res.*, Vol. 92, No D8, pp 9693-9714.

Seed, A.W. (1992). *Generation of a Spatially Distributed Daily Rainfall Database for Various Weather Modification Scenarios.* WRC 373/1/92. Report to Water Research Comm., Pretoria, South Africa.

Seed, A.W., R. Srikanthan and M. Menabde. (2001). A space and time model for design storm rainfall. *J. Geophys. Res.*, pp 9693– 9714.

Small, M.J. and D.J. Morgan, (1986). The relationship between a continuous-time renewal model and a discrete Markov chain model of precipitation occurrence, *Water Resources Research*. Vol. 22, no 10, pp 1422-1430.

Smithers, J.C. and R.E. Schulze. (2000). *Development and evaluation of techniques for estimating short duration design rainfall in South Africa*. WRC Report No. 681/1/00. Report to the Water Research Commission, Pretoria, South Africa.

Stern, R.D. and R. Coe. (1984). A model fitting analysisof daily rainfall data (with discussion). *Journal of the Royal Statistical Society*. Series A, 147, pp1-34.

Steyn, P.C.L. and R.T. Bruintjes. (1990) Convective cloud characteristics for the Bethlehem area. *Water SA*, 16, p115 - 118

Taylor, G.I. (1938). The Spectrum of Turbulence. *Proc. R. Soc. London, Ser.* A, Vol. 164, pp476-490.

Terblanche, D.E. (1996). A simple Digital Signal Processing Method to Simulate Linear Quadratic Responses from a Radar's Logarithmic Receiver. *J. Atmos. and Oceanic Tech*, Vol. 13, No. 2, pp533-538

Terblanche, D.E, G.G.S. Pegram and M.P. Mittermaier. (2001). The development of weather radar as a research and operational tool for hydrology in South Africa, *J. Hydrology*.,Vol. 241, pp3-25

Todini, E. (2001). A Bayesian technique for conditioning rdar precipitation estimates to raingauge measurements. *Hydrology and Earth System Sciences,* Vol. 5, No. 2, pp187-199.

Uijlenhoet, R. (2001). Raindrop Size Distributions and the Z-R Relationship. *Hydrology and Earth System Sciences,* Vol. 5, No. 4, pp615-627.

Van Heerden, J. and P.C.L. Steyn. (1999). Weather radar measurements of rainfall for hydrological and other purposes. WRC 693/1/99. Report to Water Research Comm., Pretoria, South Africa.

VanMarke, E. (1983). *Random Fields: Analysis and Synthesis*. MIT Press. Cambridge MA.

Wei, W.W.S. (1990). *Time series analysis, univariate and multivariate methods*. Addison Wesley.

Wichmann, B.A. and I.D. Hill. (1982). A Pseudo-Random Number Generator. NPL Report DITC 6/82. Algorithm AS183 in *Applied Statistics Algorithms,* P Griffiths and I D Hill (eds). Royal Statistical Society.

Wilson, J., D. Schertzer and S. Lovejoy. (1991). Continuous Multiplicative Cascade Models of Rain and Clouds in *Non-Linear Variablility in Geophysics : Scaling and Fractals.* D Schertzer and S Lovejoy (eds). Kluwer Academic.

Wheater. H.S., V.S. Isham, D.R. Cox, R.E. Chandler, A. Kakou, P.J. Northrop, L. Oh, C. Onof and I. Rodriguez-Iturbe. (2000). Spattial-Temporal Rainfall Fields: Modelling and Statistical Aspects. *Hydrology and Earth System Sciences,* Vol. 4, No. 4, pp581-601.

Woolhiser, D.A. and G.G.S. Pegram. (1979). Maximum Likelihood of Fourier Coefficients to Describe Seasonal Variations of Parameters in Stochastic Daily Precipitation Models. *J. Applied Meteor,* Vol. 18, pp 34-42.

Zucchini, W.S. and P.T. Adamson. (1984). *Assessing the risk of deficiencies in streamflow*. Report No. WRC 91/2/84. Report to the Water Research Commission, Pretoria, South Africa.