

**SPACE-TIME MODELLING OF
RAINFALL USING THE STRING
OF BEADS MODEL:
INTEGRATION OF RADAR AND
RAINGAUGE DATA**

AN Clothier • GGS Pegram

WRC Report No. 1010/1/02



Water Research Commission 

Disclaimer

This report emanates from a project financed by the Water Research Commission (WRC) and is approved for publication. Approval does not signify that the contents necessarily reflect the views and policies of the WRC or the members of the project steering committee, nor does mention of trade names or commercial products constitute endorsement or recommendation for use.

Vrywaring

Hierdie verslag spruit voort uit 'n navorsingsprojek wat deur die Waternavorsingskommissie (WNK) gefinansier is en goedgekeur is vir publikasie. Goedkeuring beteken nie noodwendig dat die inhoud die siening en beleid van die WNK of die lede van die projek-loodskomitee weerspieël nie, of dat melding van handelsname of -ware deur die WNK vir gebruik goedgekeur of aanbeveel word nie.

**Space-Time Modelling of Rainfall using the
String of Beads Model:
Integration of Radar and Raingauge Data**

Report to the Water Research Commission

Antony N Clothier and Geoffrey GS Pegram
Civil Engineering,
University of Natal,
Durban, South Africa

E:mail:
antony@clothier.za.net
peggram@eng.nu.ac.za

January 2002

WRC Report No 1010/1/02
ISBN No 1 86845 835 0

Space-Time Modelling of Rainfall using the String of Beads Model: Integration of Radar and Raingauge Data

Executive Summary

What was the rationale for the research?

The *String of Beads* model is a stochastic rainfall model based on the combined observations of a large network of daily raingauges and an S-band weather radar situated near Behtlehem, South Africa. The model draws on the ideas and observations of many other rainfall modelers (mentioned at appropriate places in the report) to form an efficient means of simulating rainfall over a wide range of spatial and temporal scales.

What were the aims of the research and were they achieved?

The aims in the original proposal were:

- Extend the model to where only rain-gauge data are available
- Validate the model against continuous raingauge data
- Extend the model's capacity to allow for topographic features and large meteorological systems
- Incorporate information from satellites
- Make the model available to the hydrological community as a simulation tool in a user-friendly package

Three of the five aims listed above were achieved. One of the exceptions was the intention to incorporate satellite data; this was transferred to another WRC contract, K5/1153, the Spatial Interpolation and Mapping of Rainfall (SIMAR) where progress is satisfyingly rapid. The other exception is the use of continuous raingauge data. As will be seen in the sequel, difficulty was encountered with the more recent radar data and its

correspondence with the historical raingauge data (the model's primary calibration source) so the continuous data were not relevant at this stage.

The most successfully achieved aim was the sharing of the strengths and simplicity of the String of Beads model (SBM) with the Hydrological, in particular, the international community. The list of technology transfer opportunities realised at the end of this Executive Summary attests to that. As a result of these initiatives, in July 2001, the rainfields on two foreign catchments (one in Australia the other in Switzerland) are being modelled with the SBM. In principle, the model can be adapted to varying topography and these ideas are under further development.

One of the unexpected successes of the model is that it can be used for short-term rainfall forecasting – more about this later in the summary – so although there is more to do in improving the calibration, there are already practical spinoffs which will bear fruit from this research contract.

What is this new version of the String of Beads Model?

Weather radar measures rainfall rates over a large area in fine spatial detail and very frequently. This translates into a circular area up to 200km across with 1km resolution every 5 minutes. In this study, the diameter of the circle is reduced to 128km for precision and manageability, even so, each image representing the estimated rain rate over the area is described by up to 12 868 numbers. A year's data can quickly fill a computer disk unless it is managed cleverly. What to do with all the information that these data represent? The answer is to summarize it using a model.

The String of Beads model was introduced in WRC contract K5/752 when the "Bead" (the model of the rainfall process while it is raining) was developed. In this study the model has been refined and extended to improve the "Bead" process and model the arrival of the beads of space-time rain on the "String" of time.

The new version of the model has the particularly useful property that it can be calibrated on historical daily raingauge data, even in the absence of a radar in the area,

although this would be an added advantage. The daily raingauges (if there is a reasonably good network covering more than 40 years) characterize the inter-annual and intra-annual variability of the rainfall. This means that the model can be used in places where radar data is short or even absent.

The model is designed to perform two important modelling tasks. The first is simulation of long sequences (of many years duration) to conduct “what if” studies of rainfall over the target area. The primary variables of interest are the mean rainfall rate over the area and the percentage of the area which is wetted by the rainfall (which are called the Image Mean Flux, or IMF, and Wetted Area Ratio, or WAR, in this report). The output from such simulations matches the refined gridded data of digital elevation models now being used in catchment rainfall/runoff studies. The second task to which the model can be adapted is to give short-term forecasts (nowcasts in the literature) of future rainfall for use in real-time flood forecasting in either rural or urban environments. Using the String of Beads model to forecast rainfall on a catchment up to one or two hours ahead has the benefit that managers (catchment, flood or disaster) have extra information to enable them to make anticipatory decisions - they can be proactive rather than reactive when a flood threatens.

Why is the String of Beads model good for rainfall modelling? What makes it a better model than others?

Space-time rainfall is a very complicated physical process and might seem to need a very complicated model to describe it. The advantage of the model described here is that it is relatively straightforward in that it is defined by a very few parameters and achieves remarkable realism in the synthetic sequences it generates. The result is that it is remarkably fast in execution, which is a necessary characteristic with all the detail it presents.

It was the first model of its sort to present spatial detail and realistic temporal development in one package. There are others which are based on point rainfall models which can capture the average behaviour over an area with time and those which can model rainfall events with similar detail, but none that did both before the String of

Beads model was devised. It is also adaptable to a variety of rainfall regimes - it has been fitted to South African, Swiss, German and English data with success and has become well known internationally as a serious competitor in the rainfield modelling environment.

What sorts of Simulations are envisaged?

The crucial characteristics that a space-time model of rainfall must capture are

- the storm arrivals,
- their duration and intensity while it is raining and also
- move realistically in the right directions and at the right velocities.

Each of these characteristics must be quantified and their mutual dependence and variability accounted for. The way the new model works is to examine the “String” of time and identify the alternating wet and dry periods on it. These lengths vary randomly (and we have discovered, independently of each other) but can be characterized by the lognormal probability distribution function whose two parameters vary periodically over the year. This is the basic process.

Once it is raining, the “String” supports a “Bead” of space-time rainfall. The “Beads” have the lengths of the wet periods on the “String” and their structure in space and time defines the behaviour of the rainfall event. These descriptors are its velocity and direction and whether it describes, on the one hand, patchy thunder-storm rainfall or, on the other hand, widespread rain of less variability but longer duration on average. The new model achieves all this and passes the stringent statistical and visual tests to check whether the result is good.

Can the “String of Beads” model be used for Forecasting?

Yes. The “String of Beads” model in its new form is particularly suitable for forecasting. This is because of its efficient structure based on time series ideas. In the previous formulation, the set of images comprising the “Bead” were all generated simultaneously using Fourier transforms in a box of space-time. Although mathematically elegant, this was not practical because it was difficult to generate wet

periods of the correct length and the method did not lend itself to forecasting applications.

The new model, by contrast, is based on time series ideas, in that images are constructed in series using linear combinations of the images in the recent past with added noise, both at image scale and at pixel scale. This is convenient for simulation because the wet period can be defined to be any length to suit the “String” process and the computation method is more efficient and quicker than previously. This is because the dependence structure is relatively short - only the previous five images are needed to construct the new one.

Forecasting is a natural extension of this formulation of the model. At any time that the data are available (which is possible with a radar dedicated to the catchment concerned) the previous five images are used to construct the possible future an image at a time until the forecast horizon has been achieved. It is found that forecasts are good to start with, deteriorating gradually so that they are still relatively useful an hour ahead, but useless after two hours. For urban applications in particular the methodology has particular promise for “buying time” to help managers to anticipate disaster.

What of Capacity Building?

This research was conducted in the Civil Engineering Programme of the University of Natal in Durban, in whose final year class of 46 in 2001, a mixed group of students can be found. The classification of the class is 5 Black (1 woman), 17 Indian (2 women) and 24 White (2 women). Two of the final year dissertations in 2001 involve aspects of rainfall measured and modelled by radar and all these students are taught about rainfall patterns and relationships based on the results of this research – this was not so 4 years ago. There are other spinoffs which feed other WRC research contracts, in particular SIMAR (K5/1153) and Flood Nowcasting in the Umgeni (K5/1217), recently commenced.

What of Technology Transfer?

The research outputs from this contract include:

International journal publications.

1. Pegram, GGS and AN Clothier, (2001). High Resolution Space-time Modelling of rainfall: The "String of Beads" Model. *Journal of Hydrology*,. (241)1-2, pp. 26-41
2. Terblanche, DE, GGS Pegram and MP Mittermaier, (2001). The Development of Weather Radar as a Research and Operational Tool for Hydrology in South Africa. *Journal of Hydrology*, (241)1-2, pp. 3-25
3. Pegram, GGS and AN Clothier, (2001). Downscaling Rainfields in Space and Time, Using the String of Beads Model in Causal Mode, *Hydrological and Earth Systems Sciences*, European Geophysical Society, vol 5 no 2.
4. Pegram, GGS, (2001). Kriging using the Fast Fourier Transform – The Large Data Mapping Problem. Submitted for publication in *Water Resources Research*.

Presentations at International Conferences/Symposia

1. Pegram GGS and Clothier AN (1999a). High Resolution Modelling of Rainfall: The String of Beads Model: New Results. European Geophysical Society 24th General Assembly, The Hague, Holland, April. Abstract in: Geophysical Research Abstracts, Vol.1, No.2, p.297
2. Pegram GGS and Clothier AN (1999b). Downscaling Rainfields in Space and Time with the "String of Beads" Model. *Invited* oral presentation at *AGU Fall meeting*, San Francisco, December. Abstract in EOS, Transactions, AGU Vol.80, No.46, p.419
3. Pegram GGS and Clothier AN (1999c). Nowcasting Rainfields Using the "String of Beads" Model, Poster presentation at *AGU Fall meeting*, San Francisco, December. Abstract in EOS, Transactions, AGU Vol.80, No.46, p.382
4. Pegram GGS and Clothier AN (1999d). The "String of Beads" rainfall model and its application in simulation and forecasting, *9th South African National*

Hydrology Symposium, University of the Western Cape, November. Paper in Proceedings, 11 pages.

5. Pegram GGS (1999). Continuous Space-time Modelling (and optimal conditional interpolation) of Rainfall Fields Using Non-fractal Methods. *Workshop on Scaling in Hydrology*, CRCCH and Bureau of Meteorology, Melbourne, Australia, June.
6. Pegram, GGS and Clothier, AN (2000a). The String of Beads – a Gaussian random field model of rainfall measured by radar. *Proceedings of International Conference on Applied Mechanics, SACAM 2000*. Adali S, EV Morozov and VE Verijenko (eds.), Durban, 11-13 January, pp 65-73.
7. Pegram, GGS and Clothier, AN, (2000b). Downscaling Rainfields in Space and Time Using the “String of Beads” Model. Oral presentation at *EGS Assembly*, Nice, France. April.
8. Pegram, GGS and Clothier, AN, (2000c). The String of Beads Model – A Tool for Rainfall Simulation and Short Term Forecasting. Oral presentation, *Fifth International Workshop on Rainfall in Urban Areas*, Pontresina, Switzerland, December.

Presentations at International Workshops/Study Courses

Professor Pegram facilitated the following presentations, by his PhD student Antony Clothier, at the universities of international colleagues in Europe during July 2000. The following presentation was made:

Application of the “String of Beads” model using European data sets. An interactive two-day workshop aimed at potential users of the rainfall model. A detailed description of the theory was presented and the model was calibrated on European data sets wherever available

Workshops were held at the following institutions:

1. Prof Ezio Todini, Dipartimento di Scienze della Terra e Geologico Ambientali, University of Bologna, Bologna, Italy.
2. Prof Andras Bardossy, Universität Stuttgart, Institut für Wasserbau, Stuttgart, Germany.

3. Prof Paulo Burlando, Inst. of Hydromechanics and Water Resources Management, ETH Hoenggerberg, Zurich, Switzerland.
4. Prof Enda O'Connell, Department of Civil Engineering, University of Newcastle upon Tyne, Newcastle upon Tyne, England.
5. Prof Ian Cluckie, Water Management Research Centre, University of Bristol, Bristol, England.
6. Prof Phillip O'Kane, Dept Civil and Environmental Engineering, National University of Ireland, UCC, Cork, Ireland.

Is the software available?

Yes, it is freeware and available on a Compact Disk at a nominal cost from Professor Geoff Pegram at the University of Natal for the next two years : peggram@nu.ac.za ; after that, enquiries should be addressed to the WRC.

WRC 752/1/99 Section Cross Reference

Some of the contents of this report have been repeated from WRC 752/1/99 for completeness and are cross referenced here.

WRC 1010	Description	WRC 752
2.4.1	The MRL-5 S-Band radar near Bethlehem**	A.2
2.4.2	Methods of CAPPI extraction from volume scan data**	A.4
2.4.3	Quantitative limits of the radar data**	5
2.4.4	Sources of error in rainfall measurement using radar**	A.3
2.7.1	Structure of the windows bitmap**	A.6.1
2.7.2	Selecting the colour palette to display a radar rainfield**	A.6.2
3.4.1	Definition of the two parameter lognormal distribution**	C.2
3.4.2	Estimation of the parameters μ and σ^2 **	C.2.1
3.5.1	The Discrete Fourier Transform in one dimension**	C.3.2
3.5.2	The Fast Fourier Transform**	C.3.3
3.5.3	The Fast Fourier Transform in two or more dimensions**	C.3.4
3.5.4	Estimating the power spectrum and autocorrelation using the FFT**	C.3.5
4.2.1	Uniformly distributed random noise**	D.2.1
4.2.2	Percentage points of the normal distribution**	D.2.2
4.5.1	One-dimensional power-law filtering**	E.2
4.5.2	Two dimensional power-law filtering**	E.3

Table of Contents

Chapter 1 Introduction	1
-------------------------------------	----------

Chapter 2 Rainfall and Other Relevant Data	4
---	----------

2.1 INTRODUCTION.....	4
2.2 STUDY AREA.....	5
2.3 GAUGE DATA.....	6
2.3.1 Gauge data pre-processing.....	6
2.4 RADAR DATA	7
2.4.1 The MRL-5 S-Band radar near Bethlehem**	8
2.4.2 Methods of CAPPI extraction from volume scan data**	9
2.4.3 Quantitative limits of the radar data**	13
2.4.4 Sources of error in rainfall measurement using radar**	14
2.5 SPATIAL DATA STORAGE	17
2.5.1 Meteorological Data Volume (MDV) file format.....	17
2.5.2 Selective storage	18
2.5.3 Compression of radar data	19
2.5.4 Image File Nomenclature.....	19
2.6 A PAINFUL LESSON IN DATA QUALITY CONTROL.....	21
2.7 SPATIAL DATA VISUALISATION	22
2.7.1 Structure of the windows bitmap	23
2.7.2 Selecting the colour palette to display a radar rainfield**	24
2.8 UPPER LEVEL WIND DATA	26
2.9 SUMMARY	26

Chapter 3 Data Analysis – Statistics, Techniques and Results	27
---	-----------

3.1 INTRODUCTION.....	27
3.2 IMPORTANT PROPERTIES OF A SPACE -TIME RAINFALL MODEL	27

3.3	GAUGE ANALYSIS.....	30
3.3.1	A Question of Scale	31
3.3.2	Daily Gauge Network Statistics.....	34
3.4	MARGINAL DISTRIBUTION OF RAINFALL ON A RADAR IMAGE	42
3.4.1	Definition of the two parameter lognormal distribution**	42
3.4.2	Estimation of the parameters μ and σ^2 **	43
3.4.3	Relationship between WAR_i , IMF_i , μ and σ	49
3.5	SPATIAL STRUCTURE **	55
3.5.1	The Discrete Fourier Transform in one dimension**	55
3.5.2	The Fast Fourier Transform**	56
3.5.3	The Fast Fourier Transform in two or more dimensions**	57
3.5.4	Estimating the power spectrum and autocorrelation using the FFT**	58
3.6	FIELD ADVECTION	62
3.7	IMAGE SCALE ANALYSIS.....	64
3.7.1	Instantaneous image scale analysis.....	64
3.7.2	Cumulative image scale analysis	72
3.7.3	An empirical relationship between instantaneous and daily statistics.....	75
3.8	THE NORMALISING TRANSFORM.....	77
3.9	ALTERNATING WET AND DRY SPELLS	77
3.9.1	Marginal distribution of the wet and dry spell durations.....	78
3.9.2	Spell duration dependent structure	81
3.10	IMAGE SCALE TIME SERIES.....	83
3.10.1	Time series transformation	83
3.10.2	Event scale statistics	85
3.10.3	Correlation in time	91
3.11	PIXEL SCALE TIME SERIES	94
3.12	SPACE -TIME UPSCALING OF RAINFIELDS	97
3.13	SUMMARY	101

Chapter 4 The “String of Beads” Model Formulation..... 102

4.1	INTRODUCTION.....	102
4.2	PSEUDO-RANDOM NUMBER GENERATION**	102
4.2.1	Uniformly distributed random noise**	102

4.2.2	Percentage points of the normal distribution**	106
4.3	ALTERNATING RENEWAL PROCESS.....	108
4.4	AUTO-REGRESSIVE PROCESSES.....	109
4.4.1	Theory of autoregressive processes	109
4.5	POWER-LAW FILTERING OF GAUSSIAN NOISE**	111
4.5.1	One-dimensional power-law filtering**	112
4.5.2	Two dimensional power-law filtering**	118
4.6	SBM STRUCTURE.....	120
4.6.1	SBM First stage – event scale simulation	120
4.6.2	SBM Second stage – image scale simulation	128
4.6.3	SBM Third stage – pixel scale simulation	131
4.6.4	Summary of the SBM structure	135
4.7	COMPUTING CONSIDERATIONS.....	137
4.8	SUMMARY	138
 Chapter 5 Model Calibration and Validation.....		139
5.1	INTRODUCTION.....	139
5.2	BEAD CALIBRATION PROCESS	139
5.3	STRING CALIBRATION PROCESS	146
5.3.1	An unfortunate sadness	146
5.3.2	Run-length calibration	147
5.3.3	Integrated IMF _i calibration	155
5.4	FORECASTING USING THE STRING OF BEADS MODEL.....	157
5.5	SUMMARY	160
 Chapter 6 Conclusions		161
 References.....		162

Acknowledgements

The authors gratefully acknowledge the following organisations and people:

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,

Dr Deon Terblanche of the South African Weather Bureau,

Prof. Bruce Kelbe of the University of Zululand,

Prof. Dennis Hughes of Rhodes University,

Dr Mark Dent of the Computing Centre for Water Research,

Mr Stefan van Biljon of the Dept. of Water Affairs and Forestry,

Mr Steve Lynch of the University of Natal,

Dr Jeff Smithers of the University of Natal,

Mr Peter Visser of the South African Weather Bureau,

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,

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

Acknowledgements...

Alan Seed of the Australian Weather Service in Melbourne for fuelling our 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,

Professors Paolo Burlando, Ezio Todini, Andras Bardossy, Enda O’Connell, Phillip O’Kane, Ian Cluckie for their interest in the project, their input and their financial and moral support which provided the opportunity for valuable international exposure and information transfer between our research groups.

List of Figures

Figure 2.1 – The South African Radar Network.....	5
Figure 2.2 - Distribution of 54 daily raingauges over the study area	6
Figure 2.3 – The MRL-5 S-band radar near Bethlehem	9
Figure 2.4 - Projection technique used to extract CAPPIs from radar volume scans.....	10
Figure 2.5 – A CAPPI extracted using the projection technique. Discontinuities are observed in the reflectivity field (Mittermaier and Terblanche, 1997).....	10
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).....	12
Figure 2.7 - CAPPI generated from volume scan data using the DISPLACE averaging technique (Mittermaier and Terblanche, 1997).	13
Figure 2.8 - Masking unreliable data	14
Figure 2.9 - Structure of the Windows Bitmap file header.....	23
Figure 2.10 – Use of different colour palettes to display a radar rainfield	25
Figure 3.1 - An illustration of pixel sampling error measured from 300 days of daily radar rainfall accumulations.....	32
Figure 3.2 - The uncertainty of estimating the Daily Wet Area Ratio by using 54 of the 10900 pixels on the radar image	33
Figure 3.3 – Average annual total rainfall as sampled from the infilled 54 daily gauge network for the 50 year period 1948-1997	34
Figure 3.4 - Serial correlation of the time series of normalised, average annual rainfall total for 54 daily raingauges	35
Figure 3.5 - Power spectrum of the time series of normalised, average annual rainfall total for 54 daily raingauges, sampled over 50 years.	35
Figure 3.6 - Illustration of distribution of monthly average rainfall totals for 54 raingauges, sampled over 50 years	36
Figure 3.7 - Marginal distribution of monthly rainfall totals.....	37
Figure 3.8 - Probability of exceedance of Average Daily Rainfall, sampled by 54 daily raingauges over 50 years.....	38

Figure 3.9 - Probability of exceedence of the daily wet area ratio (3mm threshold), estimated from 54 daily gauges sampled over 50 years.	38
Figure 3.10 - Total number of dry days observed in each month in the 50 year period 1948 to 1997	39
Figure 3.11 - Summer wet day runs observed by the 54 raingauges over the 50 year period	40
Figure 3.12 – Autumn wet day runs observed by the 54 raingauges over the 50 year period	40
Figure 3.13 – Winter wet day runs observed by the 54 raingauges over the 50 year period	41
Figure 3.14 - Spring wet day runs observed by the 54 raingauges over the 50 year period	41
Figure 3.15 - Illustration of the types of rainfall experienced in Bethlehem.....	45
Figure 3.16 - Expected and observed histograms for image S1996-035237 depicted in Figure 3.15a	46
Figure 3.17 - Expected and observed histograms for image S1996-035358 depicted in Figure 3.15b	47
Figure 3.18 - Expected and observed histograms for image S1996-035671 depicted in Figure 3.15c	47
Figure 3.19 - Expected and observed histograms for image S1996-034491 depicted in Figure 3.15d	48
Figure 3.20 - Expected and observed histograms for image S1996-034802 depicted in Figure 3.15e	48
Figure 3.21 – WAR_i vs IMF_i for 24000 images illustrating lognormal and non- lognormal images.....	50
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.....	52
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.....	53
Figure 3.24 - Probability of exceedence of rainfall rate on 5 observed radar images	54
Figure 3.25 - An example of the power spectrum of a two dimensional field	60

Figure 3.26 - The Radially Averaged, Two Dimensional Power Spectrum. A one dimensional representation of the two dimensional power spectrum in Figure 3.25.....	60
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	62
Figure 3.28 - Algorithm used to measure mean field advection.....	63
Figure 3.29 - Time series of image scale statistics for 40 hour event.....	64
Figure 3.30 - Probability of exceedence of WAR_i as measured from wet images	66
Figure 3.31 - Probability of exceedence of IMF_i as measured from 24000 wet images	66
Figure 3.32 - Histogram to illustrate the probability density function of β_{space}	67
Figure 3.33 – β_{space} parameter as a function of WAR_i parameter.....	67
Figure 3.34 – β_{space} parameter as a function of IMF_i parameter.....	68
Figure 3.35 - Cumulative mean field advection vector for the 40 hour event of Figure 3.29	69
Figure 3.36 - Polar histogram of annual 500hPa wind for Bethlehem	70
Figure 3.37 - Annual cumulative frequency of 16 point 500Hpa wind direction.....	71
Figure 3.38 – Probability of exceedence of 500Hpa wind speed on the 16 point compass.....	72
Figure 3.39 - Annual radar rainfall accumulation for the 1999 calendar year.....	73
Figure 3.40 - 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)	74
Figure 3.41 - 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)	74
Figure 3.42 - Empirical relationship between $\overline{WAR_i}$ and WAR_d	75
Figure 3.43 - Probability of exceedence of WSD given that WSD > 60 minutes	78
Figure 3.44 - Lognormal approximation of the marginal distribution of the WSD.....	79
Figure 3.45 - Probability of exceedence of Dry Spell Duration	80
Figure 3.46 - Gamma and Lognormal approximations to the marginal Distribution of the Dry Spell Duration.....	80

Figure 3.47 - Wet Spell Duration dependence structure.....	81
Figure 3.48 - Dry Spell Duration dependence structure	82
Figure 3.49 - Dry Spell $\sim$ Wet Spell duration dependence structure.....	82
Figure 3.50 - The WAR_i normalising transform (24000 wet images).....	83
Figure 3.51 – The IMF_i normalising transform (24000 wet images)	84
Figure 3.52 - Transforming the WAR_i and IMF_i time series into WAR_i^* and IMF_i^* respectively	85
Figure 3.53 - Histogram of WAR_i^* for the event of Figure 3.52	86
Figure 3.54 – M_{war}^* as a function of the wet spell duration, D_W	87
Figure 3.55 - Histogram of $\hat{M}_{war}^*$ residuals and most likely Normal approximation ...	87
Figure 3.56 – S_{war}^* as a function of M_{war}^* for 140 observed events.....	88
Figure 3.57 – Histogram of $\hat{S}_{war}^*$ residuals and most likely Normal approximation	89
Figure 3.58 - The relationship between M_{war}^* and M_{imf}^*	89
Figure 3.59 – Histogram of $\hat{M}_{imf}^*$ residuals and most likely Normal approximation	90
Figure 3.60 - The relationship between S_{war}^* and S_{imf}^*	90
Figure 3.61 – Histogram of $\hat{S}_{imf}^*$ residuals and most likely Normal approximation.....	91
Figure 3.62 - Correlation of WAR_i^* and IMF_i^* time series (40 hour event of 1996)...	92
Figure 3.63 - Correlation envelopes for twenty-one analysed events.....	93
Figure 3.64 - Difficulty in estimating the correlation from normalised, integer data.....	96
Figure 3.65 - PSI auto-correlation function measured in Lagrangian space from normalised, integer precision data	97
Figure 3.66 - Illustration of spatial upscaling for a typical 1 km 5 minute image.....	97
Figure 3.67 - Mean and standard deviation of spatial spectra exponents for a range of space-time aggregations.....	99
Figure 3.68 – Spatially averaged temporal correlation functions for the 1, 8, 32 and 128 km spatial scales at 15, 60 and 240 minute temporal scales.....	100
Figure 4.1 - Cumulative frequency curve for the Standard Normal distribution.....	107
Figure 4.2 - Illustration of an alternating renewal process	108
Figure 4.3 - Plot of scalar coefficients for a typical power-law filter.....	112
Figure 4.4 - Sequence of standard, normally distributed random numbers.....	113

Figure 4.5 - Real and imaginary components of the product of the scalar filter of Figure 4.3 and the complex valued, Fourier transformed sequence of random noise.....	114
Figure 4.6 - Power-law filtered (Exponent $b = 1.9$) sequence of the standard, normally distributed random noise given in Figure 4.4	115
Figure 4.7 - Lognormally distributed, power-law filtered white noise.....	115
Figure 4.8 - Lognormally distributed, power-law filtered white noise, the greater β_{space} parameter produces a smoother sequence.....	116
Figure 4.9 - Power spectrum for sequence plotted in Figure 4.7	117
Figure 4.10 - Power spectrum for sequence plotted in Figure 4.8.....	117
Figure 4.11 - Comparison of two two-dimensional power-law filtered fields	118
Figure 4.12 - Averaged power spectrum for left hand image of Figure 4.11	119
Figure 4.13 - Averaged power spectrum for right hand image of Figure 4.11	119
Figure 4.14 - Algorithm used to generate event scale, wet spell duration and intensity statistics, showing cumulative bands of uncertainty.....	123
Figure 4.15 - Algorithm to simulate event cumulative advection vector	126
Figure 4.16 – Output of the <i>String of Beads</i> model at the end of the first stage of simulation.....	127
Figure 4.17 - Optimisation of the order of AR model	128
Figure 4.18 - Algorithm used to simulate image scale time series for an event.....	129
Figure 4.19 - Stack algorithm used to simulate pseudo-random images from image scale statistics.....	132
Figure 5.1 - Spatially averaged, temporal correlation functions for the 1,8,32 and 128 km spatial scales at 15, 60 and 240 minute temporal scales for simulated and observed events.....	141
Figure 5.2 – The effect of field advection on the spatial upscaling behaviour.....	142
Figure 5.3 - Power spectra for aggregated images (of various spatial and temporal scales) sampled at five corresponding points from the observed and simulated sequences	143
Figure 5.4 – Comparisons of power spectral exponent for observed and simulated sequences, aggregated at 1 km - 15, 60 and 240 minute time scales.....	144
Figure 5.5 - Comparisons of power spectral exponent for observed and simulated sequences, aggregated at 8 km - 15, 60 and 240 minute time scales.....	145

Figure 5.6 - 10 years of the 50 year daily time series of WAR_d as estimated from the 54 raingauge network.....	147
Figure 5.7 – Number of dry days per month observed during the 50 year calibration period	149
Figure 5.8 - Histogram of Summer Wet Run Lengths.....	150
Figure 5.9 - Histogram of autumn Wet Run Lengths	150
Figure 5.10 - Histogram of Winter Wet Run Lengths	150
Figure 5.11 - Histogram of Spring Wet Run Lengths.....	151
Figure 5.12 - Seasonally varying parameters of the alternating renewal process, obtained through minimization of the run length objective function	152
Figure 5.13 – Re-sampled wet and dry spell parameters.....	153
Figure 5.14 - Spatially averaged annual rainfall total for 11 year simulation	153
Figure 5.15 - Spatial distribution of rainfall on a simulated annual image	154
Figure 5.16 – Comparison of the mean monthly rainfall for the region as estimated from the 50 year raingauge network and the 11 year simulation	155
Figure 5.17 - Forecasting the WAR_t image scale statistic using the bivariate autoregressive process	158
Figure 5.18 - 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 5.17.....	159

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

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)

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
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
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

Chapter 1

Introduction

The *String of Beads* model is a stochastic rainfall model based on the combined observations of a large network of daily raingauges and an S-band weather radar situated near Behtlehem, South Africa. The model draws on the ideas and observations of many other rainfall modelers (mentioned at appropriate places in the report) to form an efficient means of simulating rainfall over a wide range of spatial and temporal scales. A review of recent papers submitted on this topic to American journals is given by Foufoula-Georgiou and Krajewski (1995).

This study is an extension of the work presented by the same researchers in the form of WRC report No 752/1/99:

High Resolution Space-Time Modelling of Rainfall: the *String of Beads* model.

Research at that stage was at a preliminary level and was focused on the analysis of a single rainfall event, observed by the South African Weather Bureau's MRL5 weather radar at 1 kilometer, five minute space-time resolution. Nevertheless it provides a good introduction to the topic of rainfall modelling and the concepts behind the *String of Beads* model.

On the strength of the results of the initial work, this follow up study was commissioned, with the primary objectives being:

1. to streamline the computational process
2. to investigate the forecasting abilities of the model
3. to extend the analysis to the improved and far bigger MDV radar dataset
4. to incorporate the long term temporal behavior into the model as observed by the tried and trusted raingauge network
5. to make the model portable to other regions by developing some form of calibration technique based on a combination of radar and raingauge data

All of these topics have been addressed during the course of this research and are presented in this document. Each will be briefly discussed here.

One of the limitations of the *String of Beads* model as presented in the preceding study, was that the simulation required a three dimensional Fourier transform for the entire space-time hypercube of the event. This was an extremely computer intensive process in terms of memory and processing power and therefore limited the size of the fields and duration of events which could be simulated. The fields were also Fourier wrapped (continuous over opposing edges) in space and in time, so the edges had to be sacrificed, a waste of computational effort.

Introduction of the autoregressive form of the *String of Beads* model eliminated the need to store in memory and manipulate the entire space-time hypercube for the event. Instead, only the most recent five images now need to be stored, drastically reducing the memory requirements of the model. Furthermore, the autoregressive form of the model affords a limited forecasting ability. Reasonable forecasts of up to an hour can be achieved and this can buy valuable time, particularly for urban flood forecasting where catchments are small and react quickly to intense rainfall.

Extension of the analysis over a larger dataset was successful and the definition of the model has been generalized in order to emulate properties which were observed in that dataset. Methods have been devised to incorporate the statistics of the daily raingauge network into the modelling process so that the model can be useful for long term simulations as well as short term. This makes the model useful for application in water resources planning.

The only disappointing aspect of the research has been the model calibration. The absence of a robust calibration technique is the result of the need to correct and reanalyze the entire MDV radar dataset which comprised over 98 500 images. A considerable amount of time and energy was lost when it was discovered that the MDV dataset misrepresented the rainfall rate by a factor of two. The first attempt at calibration has been unsuccessful, but is presented in this report nevertheless. The technique, stable on the corrupt data, proved unstable after correction of the dataset. A

more robust method of calibration has been suggested for the benefit of any future developments of the model.

This document begins by discussing the datasets used in this study and then moves on to the analysis of those datasets. Analysis is presented in an order which serves to explain the structure of the model, rather than chronological order. The formulation of the model is discussed in the fourth chapter and techniques of calibration, validation and forecasting form the penultimate chapter before concluding in the last.

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 daily raingauge data 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 is the radar data set which is maintained by the South African Weather Bureau (SAWB) 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. They have been used effectively in engineering design for the latter half of the twentieth century and it seems likely that they will continue to be used for many years to come. At this point, they are the only quantitative, long term rainfall measurements available.

Radar data lends itself to the communication of rainfall intensity over a large area through the use of images thereby exploiting the ability of the human mind to qualitatively assess the data much faster, and with far greater intuition, than can be done by 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 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. The radar cover is shown by blue circles up to a range of 100km and green circles up 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 km²) 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, a circular area of 12000km² 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 600km² which is maintained by the South African Weather Bureau (SAWB). The daily raingauge data for this region extends over approximately 80 years and the network of 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

Tipping bucket raingauge data are not used in the formulation of the model, but serve as a useful comparison for the short term performance of the model at a point. These gauges are located on a 10km square grid in the Liebenbergsvlei East of Bethlehem. The approximate position of Liebenbergsvlei is shown in Figure 2.2.

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 (1997a) 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 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 parts of it are reflected back to the antenna off 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 $1\mu\text{s}$ 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

It 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 MRL-5 has recently 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 CAPPIs considered in this study are extracted (by the SAWB) 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 kilometres for this study. The 2km level is chosen as a compromise between the lowest possible altitude with the least amount of ground clutter. CAPPIs 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 CAPPI of Figure 2.5.




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

A better method of extracting CAPPIs from volume scan data is that suggested by Mittermaier and Terblanche (1997). 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. This point corresponds to a point (r, ϕ, θ) in spherical co-ordinates and the relationships between the two frames of reference are defined by the four equations (2.1) through to (2.4).

$$r = \sqrt{(x^2 + y^2 + z^2)} \quad (2.1)$$

$$\theta = \arctan\left(\frac{x}{y}\right) \quad (2.2)$$

$$z_h = z - \left(\frac{x^2 + y^2}{aD}\right) \quad (2.3)$$

where a is the standard compensating tropospheric beam refraction factor of $4/3$ (Battan, 1973) and D is the approximate diameter of the Earth.

Finally,

$$\phi = \arcsin\left(\frac{z_h}{r}\right) \quad (2.4)$$

For each chosen Cartesian grid point, the eight closest radar grid points are identified (four on the elevation *cone* above the point and four on the *cone* below). This is illustrated in Figure 2.6. Mittermaier and Terblanche explain that since the radar grid points do not necessarily match the Cartesian grid points, six weighting factors have to be calculated, two for the slant range (r_1 and r_2), two for the azimuth (az_1 and az_2) and two for the elevation (el_1 and el_2). The reflectivity at the point (x, y, z) is then calculated by taking the weighted average of consecutive bins (r_i, r_{i+1}) in the ratio ($r_2 : r_1$), thus reducing eight points to four which surround (x, y, z) in a vertical plane. These four points are then averaged in azimuth in the ratio ($az_1 : az_2$) to give an average reflectivity directly above and directly below the point in Cartesian space. Finally, the reflectivity of the point above and the point below the point (x, y, z) are averaged in elevation ($el_1 : el_2$) to give the weighted average reflectivity at (x, y, z) .




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 CAPPI which does not exhibit the concentric rings observed in CAPPIs 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 CAPPI of Figure (2.5), is given in Figure (2.7).

Data received beyond a range of 67 km from the radar is not used in the generation of 2km CAPPIs when using this technique since the altitude of the base scan exceeds 2km at that range. Most of the data analysed in this study were CAPPIs 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 Fast Fourier Transform whose use will be 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 4 years 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 SAWB. For these reasons, CAPPIs in the first

part of the study were masked as shown in Figure 2.8 so that only data within the $\frac{3}{4}$ doughnut shaped sample area of 9128 km², are 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 a CAPPI. In a sequence of CAPPIs 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, 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).

Bright band

The so called *bright band* is a band of high reflectivity which 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 a CAPPI 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 weak that rainfall 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 Rayleigh 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. A simple analogy can be drawn to a straight stick when it is partly immersed in water. To the person holding the stick, it appears to be bent at the air-water interface. 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 a CAPPI 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.

The Z-R relationship

Much research has been done on the relationship between the observed reflectivity (**Z**) and the rainfall rate (**R**) recorded on the ground. The Z-R relationship $Z = 200R^{1.6}$ 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 CAPPI, 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 SAWB 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 SAWB 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 it 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 2.4.4, 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 reduces the annual storage requirements to $(128^2 \times 320 \times 365 =)$ 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 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.

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 in 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 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. 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 was:

- 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 PAINFUL LESSON IN DATA QUALITY CONTROL*

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 SAWB in the first half of 1999 and the first batch of Post-1997 MDV data (Oct to Dec 1998) was received from the SAWB in September 1999. Before the MDV data could be analysed, techniques for the management and analysis had to be developed since they were different from the Pre-1997 data not only in storage format, but also in precision and radar configuration. To add to this, the volume of data (some 27000 images) presented logistical problems with regards to data storage and extraction. The second batch of Post-1997 MDV data (an additional 71000 images) was received in March 2000.

Amongst the logistical maze of data storage and manipulation, a fundamental principle of scientific practice was forgotten – check the data. At this point in the research, the focus had shifted away from the details of individual rainfall events and onto statistics such as the marginal distribution of event duration, the seasonal variation thereof and the dependent structure of the wet and dry spells. The data were assumed to be correct and the results of the analysis appeared to be reasonable. Unfortunately this was not the case and without the benefit of any previous study on this dataset, there was nothing to indicate that there was a problem. Development of the long term, temporal aspects of the model went ahead on incorrect data.

It was late in June 2000 when the model finally took shape incorporating the “bead” modelling process developed on the Pre-1997 dataset and the “string” modelling process, developed on a combination of the erroneous Post-1997 data and the 50 year daily raingauge network data. The model appeared to calibrate sensibly not only on South African datasets, but also on Swiss, German and English datasets. Routines to validate the model output by checking extreme rainfall, as well as monthly and annual totals and spatial distribution were ready by September and testing began in earnest. Testing and re-testing revealed unexpected high annual rainfall totals and by the end of September it was established that the problem was not in the modelling process, but in the Post-1997 MDV dataset. The Post-1997 MDV radar data were misrepresented by a factor of 2. Since the String of Beads model is comprised of several non-linear processes, the repercussions of this discovery are severe and these will be discussed at appropriate stages in this document.

The data have since been corrected and re-analysed, but sadly, at the time of compilation of this document, there are still issues which must be resolved before the String of Beads model can be useful in an operational sense. Only the analyses of the corrected dataset are presented in this document.

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.

42 4D	36 44 00 00	00 00 00 00	36 04 00 00	28 00
00 00	80 00 00 00	80 00 00 00	01 00 08 00	00 00
00 00	00 40 00 00	00 00 00 00	00 00 00 00	00 01
00 00	00 01 00 00	FF FF FF 00	FF C8 C8 00	FF BA
BA 00	FF AB AB 00	FF 9D 9D 00	FF 8F 8F 00	FF 81

42 4D	Must be ASCII text "BM" (ASCII characters 66 and 77)
36 44 00 00	Size of the file (17462 bytes for 128x128 pixel image)
00 00 00 00	Reserved for future use
36 04 00 00	Byte offset to bitmap data (1078 bytes for 256 colour bmp)
28 00 00 00	Number of bytes in header (Currently 40 bytes)
80 00 00 00	Image width (128 pixels)
80 00 00 00	Image height (128 pixels)
01 00	Number of colour planes (must be 1)
08 00	Number of bits per pixel (8 bits = 2^8 = 256 colours)
00 00 00 00	Type of compression (0 for no compression, 1 = rle, 8 bit)
00 40 00 00	Size of the image (128x128 = 16384 bytes)
00 00 00 00	Horizontal resolution (pixels per meter)
00 00 00 00	Vertical resolution (pixels per meter)
00 01 00 00	Number of colour indices used by the bitmap (256)
00 01 00 00	Number of important indices for displaying bitmap (all 256)

Figure 2.9 - Structure of the Windows Bitmap file header

Each pair of 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 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 rainfall***

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 CAPPI 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 obtained in this regard.

The advice offered by the meteorologists at the SAWB was to use the upper level wind data as an indicator of storm propagation. It was suggested that the 500 hecto-pascal wind data would be suitable for this purpose and this corresponds to a typical altitude of 5,8km. The 500 hecto-pascal 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 SAWB.

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 statistics 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 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 during the course of this study.

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 SAWB 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 an *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 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 as 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 of concentration of the runoff 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 statistics of 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 principal 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 unconvincing results.

In the same study, van Heerden and Steyn considered annual accumulations using a network of raingauges with a lot 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 variant 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. 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

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 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 of the points in Figure 3.1 will yield the relationship $y = 1.0296x - 0.1156$. 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.

An important model parameter defined in Section 3.2 is the *Daily Wet Area Ratio* (WAR_d), 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

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 Figures 3.1 and 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 corresponds -6 standard deviations, a probability of 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 Figures 3.4 and 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 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 statistical 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 Figures 3.6 and 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 exceedance for the *Daily Wet Area Ratio* (WAR_d).




Figure 3.9 - Probability of exceedance 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 than 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. Of particular interest in the context 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 too clearly shows the seasonal trend in the rainfall experienced in Bethlehem, July being the driest month and January the wettest. The seasons are shown in Figure 3.10 as labeled bands. A more detailed look at the daily wet-dry time series will reveal the runs of consecutive wet days and dry days. Figures 3.11 through 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 and it is interesting to note that the number of two day runs exceeds the number of one day runs in June and July. This is probably due to the fact that winter rain in Bethlehem almost always comes in the form of large system, stratiform rain which is likely to last for 12 to 24 hours as opposed to the convective rain which is prevalent in the warmer months and usually lasts for 2 to 5 hours.

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.

3.4.1 Definition of the two parameter lognormal distribution^{**}

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). A variate is considered to be *lognormally distributed* if the logarithms of the variate are described by the Normal Distribution. The formal definition and properties of the two parameter lognormal distribution are described by Aitchison and Brown (1957) and paraphrased here for completeness and clarity of exposition. 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 σ_y^2 . To simplify the notation, the y subscript will be dropped and the parameters μ_y and σ_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 $\Lambda(x | \mu, \sigma^2)$ and $N(y | \mu, \sigma^2)$ respectively so that

$$\Lambda(x | \mu, \sigma^2) = P\{X \leq x\} \quad (3.1)$$

$$\text{and} \quad N(y | \mu, \sigma^2) = P\{Y \leq y\} \quad (3.2)$$

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:

$$\mu_x = \exp(\mu + \frac{1}{2} \sigma^2) \quad (3.3)$$

$$\sigma_x = \exp(2\mu + \sigma^2) (\exp(\sigma^2) - 1) \quad (3.4)$$

and the median is at $x = \exp(\mu)$.

3.4.2 Estimation of the parameters μ and σ^2 **

Given a sample S_n of data consisting of n observations described by the real numbers $\{x_1, x_2, \dots, x_n\}$ sampled from $\Lambda(\mu, \sigma^2)$ a method is required for measuring the parameters μ and σ^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:

$$\frac{1}{\sigma^n \cdot (2\pi)^{n/2} \cdot \prod_{i=1}^n x_i} \exp \left\{ \frac{-1}{2\sigma^2} \sum_{i=1}^n (\log x_i - \mu)^2 \right\} \quad (3.5)$$

from which the maximum likelihood estimators m and s^2 are found to be:

$$m = \frac{1}{n} \sum_{i=1}^n \log x_i \quad (3.6)$$

and

$$s^2 = \frac{1}{n} \sum_{i=1}^n (\log x_i - m)^2 \quad (3.7)$$

Equations 3.6 and 3.7 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 given to its full precision). The data analysed in the latter part of this study were of this type.

The data analysed earlier in the study, however, were supplied by the SAWB at 8 bit precision (i.e. any integer value between 0 and 255 mm/h) and it was therefore necessary to adopt a slightly different approach when estimating the parameters of the continuous lognormal distribution of rainfall intensity on an image. The algorithm used 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 L shown in Equation 3.8 with respect to the parameters μ and σ for each image.

$$L = \sum_i n_i \log \left\{ N\left(\frac{y_i - \mu}{\sigma} \middle| (0,1)\right) - N\left(\frac{y_{i-1} - \mu}{\sigma} \middle| (0,1)\right) \right\} \quad (3.8)$$

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

$y_i = \log(\text{rainfall rate } i \text{ mm/h})$

$n_i = \text{number of pixels in the masked image recording a rainfall rate of } i \text{ mm/h}$

$\mu = \text{the mean of the logs of the rainfall rates in the image}$

$\sigma = \text{the standard deviation of the logs of the rainfall rates in the image.}$

The maximum value of Equation 3.8 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).

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 m and s values shown on each image are the most likely estimates of the μ and σ , obtained by optimising the log likelihood equation of Equation 3.8. 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 an increasing rainfall rate with range (and therefore altitude).

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 Figures 3.16 to 3.20. To enable the 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.

Figure 3.15a shows a *free convective system* with isolated convective cells which are less than 15km in diameter. 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 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 50km in diameter and the most intense rainfall experienced on the image is 75mm/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 Figures 3.16 to 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 at the tail end (higher rainfall intensities) it is important to note that these differences are of the order of 10 pixels out of a total of 9128.

3.4.3 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 3.3, it is related to the μ and σ parameters as defined in Equation 3.9.

$$IMF_i = \exp[\mu + \sigma^2 / 2] \quad (3.9)$$

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:

$$WAR_i = 100 \times \int_1^{\infty} f_X(x) dx \quad (3.10)$$

$$\text{where } f_X(x) = \frac{1}{\sqrt{2\pi}\sigma x} \exp[-\{(\ln x - \mu) / \sigma\}^2 / 2] \quad (3.11)$$

Using the relationships in Equations 3.9, 3.10 and 3.11, 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 Equations 3.9, 3.10 and 3.11.

With reference to Section 2.4.5, analysis of the uncorrected data set suggested that the two parameter lognormal was a good descriptor of the marginal distribution of the rainfall rate on most images. Analysis of the *corrected* data set however, revealed that the 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 order to give a better idea of the distribution of these parameters over the data set. As an aside, a linear regression of these points, provides a rule of thumb, the IMF_i has a numerical value which is approximately three 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. 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 3.12.

$$IMF_i = \frac{N_w \cdot 1,0}{N_w + N_d} \quad (3.12)$$

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:

$$WAR_i = \frac{N_w}{N_w + N_d} \quad (3.13)$$

so that Equations 3.12 and 3.13 suggest that the lowest value is given by:

$$IMF_i = WAR_i \cdot 1,0 \text{ mm/h}$$

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 $WAR_i - IMF_i$ 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 $N(0,1)$ 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 3.14.

$$\exp(Z_{WAR} \cdot \sigma + \mu) = 1,0 \quad (3.14)$$

This can be solved simultaneously with Equation 3.9 yielding Equations 3.15 and 3.16.

$$-Z_{WAR} \cdot \sigma = \mu \quad (3.15)$$

$$Z_{WAR} \pm \sqrt{Z_{WAR}^2 + 2 \cdot \ln(IMF_i)} = \sigma \quad (3.16)$$

Since σ must be real and positive,

$$Z_{WAR}^2 + 2 \cdot \ln(IMF_i) \geq 0 \quad (3.17)$$

resulting in the curve which separates the lognormal $WAR_i - IMF_i$ combinations from the non-lognormal ones in Figure 3.21. Furthermore, with reference to Equation 3.15, a real and positive σ also implies that the sign of Z_{WAR} is opposite to that of μ . Considering that $P(Z_i > Z_{WAR}) = WAR_i$, Z_{WAR} must be less than zero if WAR_i is greater than 50%, in which case μ must be greater than zero. With both μ and σ greater than zero, Equation 3.9 suggests 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 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 truncated 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 clutter mask which leaves 10 939 data for analysis as opposed to the 9128 of the $\frac{3}{4}$ 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 can be generated using the following algorithm:

1. Generate, a standard-normally distributed, pseudo random field, $Z_j \sim N(0,1)$, where $(j = 1, 2, \dots, n)$
2. Calculate the percentage point 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 will then be normally distributed, $N(-Z_{WAR}, 1)$.
4. Compute $IMF'' = \sum_{j=1}^n \exp(Z_j)$, for all $Z_j > 0$.
5. Compute $\varepsilon = |IMF'' - IMF_i|$
6. Scale the Z_i field 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 and 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 3.6 and 3.7, which were measured from the full precision image of Figure 3.22. This illustrates the limitation of the 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 3.18 for two binned datasets R_i and S_i , reveals that even the censored lognormal is an unlikely descriptor of the distribution.

$$\chi^2 = \sum_i \frac{(R_i - S_i)^2}{R_i + S_i} \quad (3.18)$$

The chi-square test appears to be very sensitive to differences in the tails of the distributions. 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 and 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 using the Fast Fourier Transform (FFT).

3.5.1 The Discrete Fourier Transform in one dimension**

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* $\mathbf{h}_k$ ($\equiv \mathbf{h}(\mathbf{t}_k)$), $k = 0, 1, 2, \dots, N-1$, where $\mathbf{t}_k = k\Delta$ 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 $\mathbf{H}(\mathbf{f})$ need only be estimated for the discrete values of $\mathbf{f}_n$ where

$$\mathbf{f}_n = \frac{n}{N\Delta}, \quad n = -\frac{N}{2}, \dots, \frac{N}{2} \quad (3.19)$$

The extreme values of n in Equation 3.19 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

$$H(f_n) = H_n = \sum_{k=0}^{N-1} h_k \cdot e^{2\pi \cdot i \cdot k \cdot n / N} \quad (3.20)$$

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 3.21 which will recover the *exactly* the h_k 's from the H_n 's.

$$h_k = \frac{1}{N} \sum_{n=0}^{N-1} H_n \cdot e^{-2\pi \cdot i \cdot k \cdot n / N} \quad (3.21)$$

3.5.2 The Fast Fourier Transform **

The Fast Fourier Transform 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 it therefore becomes important to consider the computational efficiency of the algorithm used to extract the Fourier Transform. 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. The computational effort for the Fast Fourier Transform algorithm is proportional to $N \cdot \log_2 N$ 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 $\mathbf{h}_{k(\text{even})}$ and odd indexed $\mathbf{h}_{k(\text{odd})}$, which are independently transformed and added together. Since $4N^2$ real valued multiplications are required to compute an N point Discrete Fourier Transform, $4(N/2)^2$ are required for each of the half sequences $\mathbf{h}_{k(\text{even})}$ and $\mathbf{h}_{k(\text{odd})}$, or a total of $2N^2$ for both. The two half sequences can themselves be split into two half sequences again reducing the multiplication count. Decimation can continue until $N = 1$. It is for this reason that the sequence to be transformed must be an integer power of 2 in length. A sequence of arbitrary length should be padded with a buffer of zeros to achieve this condition. This process is known as “zero padding”.

3.5.3 The Fast Fourier Transform in two or more dimensions**

Consider a complex valued function $\mathbf{h}(\mathbf{k}_1, \mathbf{k}_2)$ defined on a two dimensional space where $\mathbf{k}_1 = 0, 1, \dots, N_1 - 1$ and $\mathbf{k}_2 = 0, 1, \dots, N_2 - 1$, its complex valued two dimensional discrete Fourier Transform $\mathbf{H}(\mathbf{n}_1, \mathbf{n}_2)$ is given by Equation 3.22.

$$H(n_1, n_2) = \sum_{k_2=0}^{N_2-1} \sum_{k_1=0}^{N_1-1} \exp\left(\frac{2\pi.i.k_2.n_2}{N_2}\right) \cdot \exp\left(\frac{2\pi.i.k_1.n_1}{N_1}\right) h(k_1, k_2) \quad (3.22)$$

This can be rearranged in the form of Equation 3.23.

$$H(n_1, n_2) = \sum_{k_2=0}^{N_2-1} \exp\left(\frac{2\pi.i.k_2.n_2}{N_2}\right) \cdot \sum_{k_1=0}^{N_1-1} \exp\left(\frac{2\pi.i.k_1.n_1}{N_1}\right) h(k_1, k_2) \quad (3.23)$$

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 $\mathbf{h}(\mathbf{k}_1, \mathbf{k}_2)$ for the two dimensional case is given by:

$$\frac{1}{N_2.N_1} \sum_{k_2=0}^{N_2-1} \sum_{k_1=0}^{N_1-1} \exp\left(-\frac{2\pi.i.k_2.n_2}{N_2}\right) \cdot \exp\left(-\frac{2\pi.i.k_1.n_1}{N_1}\right) H(n_1, n_2) \quad (3.24)$$

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.

3.5.4 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 $\mathbf{h}_k$ (where $k = 0, 1, \dots, N-1$), the one dimensional power spectrum is proportional to $|\mathbf{H}_n|^2$, where the $\mathbf{H}_n$ ($n = -N/2, \dots, N/2$) are defined by Equation 3.20 and represent the corresponding sequence in the frequency domain. That is to say that in order to estimate the power spectrum of a sequence of N complex valued points $\mathbf{h}_k$, the sequence is transformed into Fourier space via the FFT to get the transformed complex valued sequence $\mathbf{H}_n$ and the power spectral density $P(f_n)$ corresponding to each frequency (f_n) is obtained by squaring the absolute value of the $\mathbf{H}_n$ 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 3.25.

$$Corr(g, h) = \int_{-\infty}^{\infty} g(t + \tau).h(t).dt \quad (3.25)$$

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 3.26.

$$Corr(g, h) \Leftrightarrow G(f).H^*(f) \quad (3.26)$$

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:

$$\text{Corr}(g, g) \Leftrightarrow |G(f)|^2 \quad (3.27)$$

which is known 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 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. 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).




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.




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, indicating a power law function and thus the scale-invariant nature of the rain field 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 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 . This is one of the many motivations to move away from integer precision data.

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.

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 3.28.




Figure 3.28 - 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*.

Considering the literature reviewed so far, this appears to be a novel approach to measuring field advection and it is now published as part of a paper by Pegram and Clothier (2001b).

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 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 3.29 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.




Figure 3.29 - Time series of image scale statistics for 40 hour event

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 3.29.

- The μ and σ parameters are negatively correlated 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, 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.
- The 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 3.4.3.

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, Figures 3.30 and 3.31 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 3.30 - Probability of exceedence of WAR_i as measured from wet images




Figure 3.31 - 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. 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 3.32 presents a histogram to illustrate the probability density function of β_{space} .




Figure 3.32 - Histogram to illustrate the probability density function of β_{space}

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 3.29 suggested, the β_{space} parameter is almost completely independent of the WAR_i and IMF_i parameters. This is confirmed in Figures 3.33 and 3.34 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 3.33 – β_{space} parameter as a function of WAR_i parameter




Figure 3.34 – β_{space} parameter as a function of IMF_i parameter

The average β_{space} for all of the images is 2.45 and regression lines fitted through the data of Figures 3.33 and 3.34 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 3.35 illustrates the cumulative mean field advection vector for the 40 hour event whose image scale statistics are presented in Figure 3.29.




Figure 3.35 - Cumulative mean field advection vector for the 40 hour event of Figure 3.29

The cumulative mean field advection vector of Figure 3.35 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 Northerly direction (coming from the North East) and then swung to a North-Easterly direction where it remained until the thirtieth hour when it switched back to a Northerly direction. The storm speed appears to have been reasonably constant at approximately 20 km/h.

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 is the 500 hecto-pascal wind speed and direction (Steyn and Brintjes, 1990). Daily weather balloon data for Bethlehem (which include the 500hPa wind data) commencing in January 1980 were obtained from the South African Weather Bureau.

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 3.36.




Figure 3.36 - Polar histogram of annual 500hPa wind for Bethlehem

Clearly most of the 500hPa wind 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 3.37.




Figure 3.37 - Annual cumulative frequency of 16 point 500Hpa wind direction

Consistent with observations in Figure 3.36, most of the wind approaches between the west and west-south-west sectors, the point of maximum gradient on the curve of Figure 3.37. Perhaps more revealing than Figure 3.36, Figure 3.38 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. As an example with reference to Figure 3.38, 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 3.38). The complement of this is obviously that there is an 80% probability that the wind will be less than 7 km/h. In 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.




Figure 3.38 – Probability of exceedence of 500Hpa wind speed on the 16 point compass

These curves serve to illustrate the technique of analysis of the wind data. Although not included here, similar curves to those shown in Figures 3.37 and 3.38 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 Figures 3.36 through 3.38, it is interesting to note that the 1996 event of Figures 3.29 and 3.35 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 2.4.5. The *corrected* annual radar rainfall accumulation for the 1999 calendar year is shown in Figure 3.39.




Figure 3.39 - Annual radar rainfall accumulation for the 1999 calendar year

The average 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. 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.

The 98000 instantaneous images were accumulated to 307 daily images from which the WAR_d and IMF_d statistics were measured. Figures 3.40 and 3.41 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 Figures 3.8 and 3.9 earlier in this chapter.




Figure 3.40 - 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 3.41 - 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)

Figures 3.40 and 3.41 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

In order to incorporate both the long duration and short duration characteristics of rainfall into the "String of Beads" model, it was necessary to develop some efficient and sensible means of comparing the statistics of daily raingauge data to those of the instantaneous radar data. This becomes important in the calibration of the model which will be discussed in detail Chapter 5 - it eliminates the need to simulate full three dimensional (two-space and time) rainfall sequences over a long period and in doing so saves on time and computer resources.

Of particular interest in this context 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 3.42 illustrates this relationship.




Figure 3.42 - 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 3.42 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 calibration of the "String of Beads" model to be discussed in Chapter 5.

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. This 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 distribution, 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 3.28.

$$X_j^* = \Phi^{-1}\left(\frac{r_j}{n+1}\right) \quad (3.28)$$

where

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

r_j is the rank of X_j

$\Phi^{-1}(\cdot)$ 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 is 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 3.43.




Figure 3.43 - Probability of exceedence of WSD given that WSD > 60 minutes

Two series are plotted in Figure 3.43, 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 3.43 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 3.44 illustrates the observed histogram of the WSD and the probability density function of the most likely lognormal approximation to the distribution.




Figure 3.44 - 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 3.45. 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 senseless exercise which will not be attempted here.




Figure 3.45 - 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 3.46 illustrates the most likely approximations of these two distributions for the Bethlehem dataset.




Figure 3.46 - 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 $(k+1)^{\text{th}}$ 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 3.47 illustrates the (lack of) Wet Spell Duration dependent structure.




Figure 3.47 - Wet Spell Duration dependence structure

If it were present, dependence structure in the WSD would be revealed in Figure 3.47 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 appears 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 3.48 illustrates the (lack of) Dry Spell Duration dependence structure and Figure 3.49 illustrates the (lack of) Dry Spell ~ Wet Spell dependence structure.




Figure 3.48 - Dry Spell Duration dependence structure




Figure 3.49 - Dry Spell ~ Wet Spell duration dependence 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 $\sim$ 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 independently 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 3.50 illustrates the transformation which links the WAR_i to the normalised WAR_i^* .




Figure 3.50 - The WAR_i normalising transform (24000 *wet* images)

The fitted curve shown in Figure 3.50 is a scaled and shifted sigmoid function which has the form of Equation 3.29.

$$y = \left(\frac{1}{1 + \exp(b - a \cdot x)} \right) \cdot c + d \quad (3.29)$$

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 bound between c and $(c+d)$.

Similarly, the transformation which links the IMF_i to the normalised IMF_i^* is plotted in Figure 3.51 and another sigmoid function is fitted in this case.




Figure 3.51 – 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 3.30, 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 3.31, 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

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 3.52.




Figure 3.52 - Transforming the WAR_i and IMF_i time series into WAR_i^* and IMF_i^* respectively

In the example given in Figure 3.52, the WAR_i time series (upper right) of the 1996 event of Figure 3.29 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.

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 3.53, 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 illustrated in Figure 3.52. 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 3.53 - Histogram of WAR_i^* for the event of Figure 3.52

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 WSD of 2400 minutes – longer than the average, with reference to Figure 3.44 – 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 as a result of large systems and consequently higher WAR_i and conversely short duration events are usually as a result of isolated convective activity with lower WAR_i . To investigate this phenomenon further, Figure 3.54 considers the M_{war*} as a function of the WSD 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 3.54 – 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 3.55, along with the most likely Normal approximation to the distribution.




Figure 3.55 - 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 the WSD 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 3.56 considers this relationship.




Figure 3.56 – 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 3.57 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 3.57 – 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 3.58 and the corresponding distribution of residuals $\varepsilon_{imf*} = \hat{M}_{imf*} - M_{imf*}$, is plotted in histogram form in Figure 3.59 with its most likely Normal approximation, a mean of 0,01 and a standard deviation of 0,35.




Figure 3.58 - The relationship between M_{war*} and M_{imf*}




Figure 3.59 – Histogram of $\hat{M}_{imf^*}$ residuals and most likely Normal approximation

Finally, the relationship between S_{war^*} and S_{imf^*} is plotted in Figure 3.60. 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 3.61 with its most likely Normal approximation, a mean of 0,00 and a standard deviation of 0,11.




Figure 3.60 - The relationship between S_{war^*} and S_{imf^*}




Figure 3.61 – Histogram of $\hat{S}_{imf^*}$ residuals and most likely Normal approximation

The normal distribution is used effectively to approximate the residual histograms of Figures 3.55, 3.57, 3.59 and 3.61. Other distributions or even a non-parametric option could be considered, however the Normal distribution is a simple approach which is well understood and therefore a good starting point with this limited number of events.

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 3.25) 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):

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

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

$$WAR_i^* \sim IMF_i^* \text{ cross correlation, } Corr(WAR_{i(t)}^*, IMF_{i(t+\tau)}^*)$$

$$IMF_i^* \sim WAR_i^* \text{ cross correlation, } Corr(IMF_{i(t)}^*, WAR_{i(t+\tau)}^*)$$

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 $\tau = a \cdot \Delta t$ (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 3.29 are plotted in Figure 3.62.




Figure 3.62 - Correlation of WAR_i^* and IMF_i^* time series (40 hour event of 1996)

Figure 3.62 is plotted for positive values of τ , negative values can be inferred from these. An interesting feature of Figure 3.62 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 3.29 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 3.63.




Figure 3.63 - Correlation envelopes for twenty-one analysed events

Analyses of twenty-one events are presented in Figure 3.63 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 (EGS, 2000) 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 3.35) 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 $\frac{3}{4}$ 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 WAR_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 3.64 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 3.64, upper right). Applying the normalising transform of Section 3.8 yields an identical transform for the continuous and the integer time series (Figure 3.64, upper left). The problem arises when computing the correlation from the normalised, integer time series (Figure 3.64, 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 3.64, lower right.




Figure 3.64 - 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 3.65 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 3.65 - 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 3.29, 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 3.66 illustrates this process for a typical image.




Figure 3.66 - Illustration of spatial upscaling for a typical 1 km 5 minute image

The images in Figure 3.66 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 3.66. 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 3.66 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 3.5.4, however, with reference to Figure 3.66, 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 3.67.




Figure 3.67 - Mean and standard deviation of spatial spectra exponents for a range of space-time aggregations

The main points of interest noted in Figure 3.67 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.

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 3.68 shows the autocorrelation functions for the 1, 8, 32 and 128 km spatial scales at 15, 60 and 240 minute temporal scales.




Figure 3.68 – Spatially averaged temporal correlation functions for the 1, 8, 32 and 128 km spatial scales at 15, 60 and 240 minute temporal scales

The functions plotted in Figure 3.68 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 what is known as the ergodic theorem.

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 3.68, 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 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. One of the key ingredients in the model is a good *uniformly distributed pseudo-random number generator*. Such generators commonly available range from reasonable to awful. 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. Two algorithms are presented here; the first is *an efficient and portable pseudo-random number generator* (Wichmann and Hill, 1982); the second is to transform a sequence of uniformly distributed random numbers into a white noise sequence using, *the percentage points of the normal distribution* (Beasley and Springer, 1985). A quantitative measure of the quality of output for both algorithms is given.

4.2.1 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.

Strictly speaking, the sequences of uniformly distributed numbers used are *pseudo-random* and make use of an algorithm which calculates a number based on the previous number in the sequence. This ensures that any generated sequence of random numbers can be repeated when necessary by initialising the algorithm with the same seed. Obviously the random number generator should not refer to the time or date or any other volatile variable for its initial seed if the sequence is to be reproducible.

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 is outlined here. The *multiplicative congruential* sequence is a special case of the *linear congruential* sequence (X_n) defined by Equation 4.1.

$$X_{n+1} = (a.X_n + c) \bmod m, \quad n \geq 0. \quad (4.1)$$

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*; $m > X_0, m > a, m > c$

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

$$\begin{aligned} X_1 &= (7 \cdot 7 + 7) \bmod 10 \\ &= 56 \bmod 10 \\ &= 6 \end{aligned}$$

and the entire sequence X_n would be

$$X_n = 7, 6, 9, 0, 7, 6, 9, 0, 7, 6, 9, 0 \dots \quad (n \geq 0)$$

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 4.1 then reduces to Equation 4.2.

$$X_{n+1} = a \cdot X_n \bmod m, \quad n \geq 0. \quad (4.2)$$

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_θ 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 4.3.

$$x^k \bmod m = 1 \quad (4.3)$$

At least one solution (x) to Equation 4.3 exists for all prime numbers m .

As previously mentioned, Wichmann and Hill (1982) make use of three multiplicative congruential generators 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

	<i>modulus</i>	<i>multiplier</i>
(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 $R(0,1)$.

Wichmann and Hill (1982) 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. Due to the common factor of 2 in the periods of the generators, they have a combined period of one quarter of the product of the individual periods. That is a period exceeding 6.95×10^{12} . Considering that 6×10^6 random numbers are required for a single day's simulation at 128 x128 pixel, 5 minute resolution, about 3800 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 might possibly be observed in the images.

4.2.2 Percentage points of the normal distribution**

The *String of Beads model* makes use of *normally distributed* random noise in the generation of artificial images. 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 normal distribution for given mean μ and standard deviation σ is defined by Equation 4.4.

$$f(x; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right], \quad -\infty < x < \infty \quad (4.4)$$

The variable x is normally distributed, with mean μ and variance σ^2 and this is written as $x \sim N(\mu, \sigma^2)$. In the standard normal case $\mu = 0$ and $\sigma = 1$ and this equation reduces to (4.5).

$$f(x; 0, 1) = \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{x^2}{2}\right], \quad -\infty < x < \infty \quad (4.5)$$

Equation 4.5 describes the well known bell-shaped probability density function which can be represented in the cumulative form of Figure 4.1 which is described by Equation 4.6.

$$p = P(x < x_p) = \int_{-\infty}^{x_p} \frac{1}{\sqrt{2\pi}} \exp\left[-\frac{x^2}{2}\right] dx \quad (4.6)$$

where p is the probability that a randomly chosen number x is less than a chosen x_p .




Figure 4.1 - 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 4.6 and solving for $X = x_p$. Equation 4.6 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). Paraphrasing from that text, the routine replaces the p of Equation 4.6 by a dummy variable $q = p - 0.5$ and then compares $|q|$ with 0.42. If $|q| \leq 0.42$, x_p is obtained through the rational approximation of Equation 4.7.

$$x_p = q \cdot A(q^2)/B(q^2) \quad (4.7)$$

where A and B are polynomials of degrees 3 and 4 respectively. If $|q| > 0.42$, another dummy variable $r = \sqrt{\ln(0.5 - |q|)}$ is introduced and x_p is calculated from a different rational approximation given by Equation 4.8.

$$x_p = \pm C(r)/D(r) \quad (4.8)$$

where C and D are polynomials of degrees 3 and 2 respectively, and the sign is taken to be that of q . 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

$$|p' - p| < \epsilon$$

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 20×2^{-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.

4.3 ALTERNATING RENEWAL PROCESS

Consider two sequences of mutually independent and identically distributed random variables $\{X_1, X_2, \dots\}$ and $\{Y_1, Y_2, \dots\}$ 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 4.2.




Figure 4.2 - Illustration of an alternating renewal process

The $(2r-1)^{\text{th}}$ interval is of type X and begins after time $(X_1 + \dots + X_{(r-1)} + Y_1 + \dots + Y_{(r-1)})$ and the $2r^{\text{th}}$ interval is of type Y and begins at time $(X_1 + \dots + X_r + Y_1 + \dots + Y_{(r-1)})$, where $r = 1, 2, \dots$.

4.4 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).

4.4.1 Theory of autoregressive processes

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 will be outlined here, following the text of Wei (1990). The multivariate theory presented here can be easily reduced to the univariate case.

Consider the zero mean, unit variance case, $\dot{Z}_t = \frac{Z_t - \mu_z}{\sigma_z}$, the p^{th} order, q dimensional autoregressive model, $AR(p, q)$, is defined by:

$$\dot{Z}_t = \Phi_1 \dot{Z}_{t-1} + \Phi_2 \dot{Z}_{t-2} + \dots + \Phi_p \dot{Z}_{t-p} + A_t \quad (4.9)$$

where $\dot{Z}$, Φ_p and a_t are $[q \times 1]$, $[q \times q]$ and $[q \times 1]$ matrices respectively. For example, for the $AR(1,2)$ case:

$$\dot{Z}_t = \Phi_1 \dot{Z}_{t-1} + A_t \quad (4.10a)$$

or, close up:

$$\begin{bmatrix} z_1(t) \\ z_2(t) \end{bmatrix} = \begin{bmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{bmatrix} \begin{bmatrix} z_1(t-1) \\ z_2(t-1) \end{bmatrix} + \begin{bmatrix} a_1(t) \\ a_2(t) \end{bmatrix} \quad (4.10b)$$

For the general case, the Φ_p matrices can be computed by solving the generalized Yule-Walker equations:

$$\begin{bmatrix} \Gamma_{(0)} & \Gamma_{(1)} & \Gamma_{(2)} & \cdots & \Gamma_{(p-1)} \\ \Gamma_{(1)}^T & \Gamma_{(0)} & \Gamma_{(1)} & \cdots & \Gamma_{(p-2)} \\ \Gamma_{(2)}^T & \Gamma_{(1)}^T & \Gamma_{(0)} & \cdots & \Gamma_{(p-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{(p-1)}^T & \Gamma_{(p-2)}^T & \Gamma_{(p-3)}^T & \cdots & \Gamma_{(0)} \end{bmatrix} \begin{bmatrix} \Phi_1 \\ \Phi_2 \\ \Phi_3 \\ \vdots \\ \Phi_p \end{bmatrix} = \begin{bmatrix} \Gamma_{(1)}^T \\ \Gamma_{(2)}^T \\ \Gamma_{(3)}^T \\ \vdots \\ \Gamma_{(p)}^T \end{bmatrix} \quad (4.11)$$

where $\Gamma_{(k)}$ represents the covariance matrix of $\dot{Z}$ for lag k and will be $[q \times q]$.

$$\Gamma_{(k)} = \begin{bmatrix} \gamma_{11}(k) & \gamma_{12}(k) & \cdots & \gamma_{1q}(k) \\ \gamma_{21}(k) & \gamma_{22}(k) & \cdots & \gamma_{2q}(k) \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{q1}(k) & \gamma_{q2}(k) & \cdots & \gamma_{qq}(k) \end{bmatrix} = \text{Corr}(\dot{Z}_{t-k}, \dot{Z}_t) \quad (4.12)$$

and the $\gamma_{xy}(k)$ are computed as $\text{Corr}(z_x(t), z_y(t-k))$. Solving Equation 4.11 for Φ_p terms:

$$\begin{bmatrix} \Phi_1 \\ \Phi_2 \\ \Phi_3 \\ \vdots \\ \Phi_p \end{bmatrix} = \begin{bmatrix} \Gamma_{(0)} & \Gamma_{(1)} & \Gamma_{(2)} & \cdots & \Gamma_{(p-1)} \\ \Gamma_{(1)}^T & \Gamma_{(0)} & \Gamma_{(1)} & \cdots & \Gamma_{(p-2)} \\ \Gamma_{(2)}^T & \Gamma_{(1)}^T & \Gamma_{(0)} & \cdots & \Gamma_{(p-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{(p-1)}^T & \Gamma_{(p-2)}^T & \Gamma_{(p-3)}^T & \cdots & \Gamma_{(0)} \end{bmatrix}^{-1} \begin{bmatrix} \Gamma_{(1)}^T \\ \Gamma_{(2)}^T \\ \Gamma_{(3)}^T \\ \vdots \\ \Gamma_{(p)}^T \end{bmatrix} \quad (4.13)$$

The shock term of Equation 1, A_t , can be expressed as:

$$A_t = \Sigma \bar{e} = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1q} \\ \sigma_{21} & \sigma_{12} & \cdots & \sigma_{2q} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{q1} & \sigma_{q2} & \cdots & \sigma_{qq} \end{bmatrix} \begin{bmatrix} e_1(t) \\ e_2(t) \\ \vdots \\ e_q(t) \end{bmatrix} \quad (4.14)$$

where $e_1(t) \dots e_m(t)$ are independent, identically distributed variables $N(0,1)$ and the Σ matrix is found by Equation 4.15:

$$\Sigma = \Gamma_{(0)} - \begin{bmatrix} \Phi_1 & \Phi_2 & \Phi_3 & \cdots & \Phi_p \end{bmatrix} \begin{bmatrix} \Gamma_{(0)} & \Gamma_{(1)} & \Gamma_{(2)} & \cdots & \Gamma_{(p-1)} \\ \Gamma_{(1)}^T & \Gamma_{(0)} & \Gamma_{(1)} & \cdots & \Gamma_{(p-2)} \\ \Gamma_{(2)}^T & \Gamma_{(1)}^T & \Gamma_{(0)} & \cdots & \Gamma_{(p-3)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \Gamma_{(p-1)}^T & \Gamma_{(p-2)}^T & \Gamma_{(p-3)}^T & \cdots & \Gamma_{(0)} \end{bmatrix} \begin{bmatrix} \Phi_1^T \\ \Phi_2^T \\ \Phi_3^T \\ \vdots \\ \Phi_p^T \end{bmatrix} \quad (4.15)$$

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:

$$AIC_c = N \log\left(\frac{S}{N}\right) + \frac{N(N+p)}{N-p-2} \quad (4.16)$$

where:

$$S = \sum_{i=1}^N \varepsilon_i^2 \quad (4.17)$$

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, $\dot{X}_i$ (where $i = 1, 2, \dots, N$) using Equation 4.13. A sequence of residuals (ε_i) are then extracted by rearranging Equation 4.9 so that

$$E_i = \begin{bmatrix} \varepsilon_1(i) \\ \varepsilon_2(i) \\ \vdots \\ \varepsilon_q(i) \end{bmatrix} = \dot{X}_i - \{\Phi_1 \dot{X}_{i-1} + \Phi_2 \dot{X}_{i-2} + \dots + \Phi_p \dot{X}_{i-p}\} \quad (4.18)$$

If the AR model is a good descriptor of the process, the ε_i should, ideally, be normally distributed.

4.5 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.

4.5.1 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* noise process is uncorrelated and therefore has a β_{space} of zero. By power-law filtering a field of random data it is possible to increase the correlation of the data from zero to any desired level β_{space} . This is done by transforming the sequence into Fourier space via the Fast Fourier Transform, multiplying (pointwise) the transformed sequence 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 $\beta_{space}/2$) from the first frequency to the Nyquist critical frequency and is then reflected about that point. The *scalar* coefficients of a typical filter are plotted in Figure 4.3.




Figure 4.3 - Plot of scalar 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 4.4.




Figure 4.4 - 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 4.5. 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 4.5 - Real and imaginary components of the product of the scalar filter of Figure 4.3 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 4.5 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 4.6.




Figure 4.6 - Power-law filtered (Exponent $b = 1.9$) sequence of the standard, normally distributed random noise given in Figure 4.4

The sequence of numbers plotted in Figure 4.6 has the same marginal distribution as that of Figure 4.4, however the spatial correlation has been introduced by power-law filtering the sequence and thereby increasing the β_{space} parameter from 0 in Figure 4.4, to 1.9 in Figure 4.6. 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 4.7, with marginal distribution $\Lambda(\mu, \sigma)$ and a spatial correlation defined by β_{space} .




Figure 4.7 - 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 4.8 is a plot of a sequence which was generated using the set of white noise plotted in Figure 4.4, 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 4.7. The two sequences have been scaled and shifted so that they have identical marginal distributions.




Figure 4.8 - Lognormally distributed, power-law filtered white noise, the greater β_{space} parameter produces a smoother sequence

Figures 4.7 and 4.8 are plotted to the same scale and illustrate how the higher exponent power-law filter smoothes out the extreme points in the sequence. The *wrapped 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 power spectra of these two sequences are plotted in Figures 4.9 and 4.10 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 4.9 - Power spectrum for sequence plotted in Figure 4.7




Figure 4.10 - Power spectrum for sequence plotted in Figure 4.8

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.

4.5.2 Two dimensional power-law filtering**

Following the format adopted for the one dimensional case in Section 4.5.1, 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 4.11, 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 4.11 - 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 wrapping 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 Figures 4.12 and 4.13 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 imperfect




Figure 4.12 - Averaged power spectrum for left hand image of Figure 4.11




Figure 4.13 - Averaged power spectrum for right hand image of Figure 4.11

The two spectra shown in Figures 4.12 and 4.13 are plotted up to a wave number 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 4.5.1, 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 $\mathcal{A}(\mu, \sigma)$ and a correlation structure defined by a single parameter β_{space} .

4.6 SBM STRUCTURE

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.

4.6.1 SBM First stage – event scale simulation

Two independent sets of 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 revealed that there is no significant dependence structure in the durations of consecutive wet and dry spells. The wet and dry spell durations can thus be alternately sampled from two mutually independent processes. Following the ideas of Harberlandt (1998), the principal component in the simulation of the wet spell duration (D_w) and its complementary dry spell duration (D_d) is an alternating renewal process.

The theory for this type of model is presented in Section 4.3, however 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 the wet and dry spell durations are dependent on the season. Harberlandt (1998) found that the lognormal distribution was well suited to describing the distribution of dry 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 3.46 also reveals that there is very little difference between the gamma and lognormal approximations of these distributions. For the sake of simplicity, it will be assumed that:

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

and

$$D_d \sim \Lambda(\mu_d, \sigma_d)$$

where

D_w is the duration of the wet spells

D_d is the duration of the dry 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

μ_d represents the mean of the logs of D_d

σ_d represents the standard deviation of the logs of D_d

To account for the seasonal variation of the wet and dry spells in the *String of Beads* model, the parameters of the lognormal distribution, μ_w , σ_w , μ_d and σ_d are allowed 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 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 . Unfortunately, for this study, there were insufficient data for this purpose. The approximations of Figures 3.44 and 3.46 were for an entire year's data and provide a good starting point, but the seasonal variations for μ_w , σ_w , μ_d and σ_d have to somehow be inferred from the only long term dataset available – the daily raingauge dataset. This is an issue of model calibration which will be discussed in Chapter 5. Let it suffice at this point, to say that μ_w , σ_w , μ_d and σ_d are inferred, on a monthly basis, from the daily raingauge dataset and are then allowed to vary continuously over the year by fitting a 2 harmonic Fourier series to the monthly approximations of each variable.

For Bethlehem conditions, rain falls on most days in the summer months and it is virtually dry during the winter months so the μ_d parameter is expected to be relatively high in the winter months and low in the summer. Being a second order moment, behaviour of the σ_d parameter is slightly more difficult to predict, but analysis of the daily raingauge data presented in Section 3.3.2 suggest that the spring months are particularly prone to large variations in rainfall totals, so a relatively high σ_d is likely to be observed in these months.

Rainfall generating mechanisms in Bethlehem vary from predominantly widespread stratiform rain in the winter, to a mixture of isolated convective activity and widespread stratiform rain in the summer. Consequently, a relatively high μ_w would be expected in the winter - although it rarely rains in winter, when it does, it is usually as a result of a large weather system resulting in a high D_w . The opposite scenario would be expected for the σ_w , with a relatively high σ_w in the summer due to the mixture of stratiform and convective rainfall resulting in a wide range of D_w and a low value in the winter.

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 4.14 illustrates the algorithm for simulating these statistics. The technique of independently modelling the event arrival, duration and intensity is similar in many ways to the Neymann-Scott Shot noise model proposed by Cowpertwait (1994).




Figure 4.14 - Algorithm used to generate event scale, wet spell duration and intensity statistics, showing cumulative bands of uncertainty

The arrival date and time of the wet spell is announced by the alternating renewal process in the form of *the end of a dry spell*. Depending on the day of event arrival, suitable parameters μ_w and σ_w are computed from the respective Fourier series which describes their seasonal trend. Together, these define the cumulative density function (CDF) for the lognormal distribution of D_w beginning on that particular day of the year. Next, a uniformly distributed pseudo-random number on the range (0,1) is generated. This corresponds, via the CDF, to a wet spell duration, D_w , in hours, from which it is possible to compute an appropriate M_{war}^* from the relationship given in Equation 4.19.

$$M_{war}^* = 0,54.Ln(D_w) - 0,98 + 0,57.a \quad (4.19)$$

This relationship was devised from the analysis of D_w which was presented in Figures 3.54 and 3.55. Figure 3.54 presents the regression relationship between M_{war*} and D_w while Figure 3.55 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 4.19 is a shock term which is distributed as $N(0,1)$.

Dependent on the M_{war*} the terms M_{imf*} , S_{war*} and indirectly S_{imf*} are simulated via the Equations 4.20 , 4.21 and 4.22 respectively.

$$M_{imf*} = 0,89.M_{war*} + 0,04 + 0,34.a \quad (4.20)$$

$$S_{war*} = 0,16.M_{war*} + 0,58 + 0,29.a \quad (4.21)$$

$$S_{imf*} = 1,00.S_{war*} + 0,01 + 0,11.a \quad (4.22)$$

Again, the a terms in these equations are independent shock terms distributed as $N(0,1)$. These three relationships were extracted from the analyses presented in Figures 3.56 through 3.67.

The band of uncertainty for a single standard deviation is represented in Figure 4.14 as dotted lines on either side of these relationships. The arrows illustrate the progress of the algorithm for an arbitrary, 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.

Event advection

With only one complete year of 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 was suggested by researchers at the SAWB that the 500HPa wind speed and direction is usually a good indicator of storm speed and direction. 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 is

perhaps a naive assumption which may require refinement as more radar data are made available for analysis.

Acknowledging these limitations, simulation of the *event advection vector* is achieved via a very simple non-parametric approach. Analyses of the 500HPa wind direction and speed were presented in Figures 3.37 and 3.38 respectively. The distribution of 500HPa wind direction is shown in Figure 3.37 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 3.37 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 4.15 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} .

To simulate the first component, 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 4.15, the event starts in January. Next, a uniformly distributed random number is generated on the interval (0,1) using the multiplicative congruential generator discussed in Section 4.2.1. 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 4.15, 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 4.15, the second random number was 0,45, which corresponds to an advection *speed* for the Western sector, in January, of 12 knots.




Figure 4.15 - Algorithm to simulate event cumulative advection vector

In its current configuration, the "String of Beads" model uses twelve (monthly) CDFs for the event advection direction, each binned into sixteen compass sectors, 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 4.16.




Figure 4.16 – Output of the *String of Beads* model at the end of the first stage of simulation

The seasonally varying 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 4.14 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 4.15.

4.6.2 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 the β_{space} , the WAR_i and the IMF_i . The 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 Figures 3.33 and 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 was outlined in Section 4.4 and pertinent analysis in this regard was presented in Section 3.10.3.

The ε_i were extracted using Equation 4.18 for a range of p and the optimum order of the autoregressive model was found to be a 5th order process. This is illustrated by Figure 4.17 in which the AIC_c is plotted for increasing order, p .




Figure 4.17 - 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 4.18.




Figure 4.18 - Algorithm used to simulate image scale time series for an event

The parameters of the bivariate autoregressive process are computed using the Yule-Walker relationship given in Equation 4.11, 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 3.63. 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. 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 4.18.

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 4.23.

$$w_3 = \frac{1}{\sqrt{\left(1 - \frac{2t}{D_w}\right)^2 + \left(\frac{2t}{D_w}\right)^2}} \quad (4.23)$$

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

Symbolically, the linear combination of the time series is represented in Equation 4.24.

$$WAR_i^*(t) = (w_1(t).FWAR_i^*(t) + w_2(t).RWAR_i^*(t)).w_3(t) \quad (4.24)$$

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 so far, 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 3.50 for the WAR_i and Figure 3.51 for the IMF_i . This transformation process is illustrated (in reverse) in Figure 3.52.

Second stage simulation output

At 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, $\bar{R}$ (mm), for an event of duration D_w , simulated at a temporal resolution of δt minutes, can be approximated by Equation 4.25.

$$\bar{R} = \frac{\delta t}{60} \sum_{k=1}^n IMF_i(k) \quad (k = 1, 2, \dots, n) \quad (4.25)$$

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

$$n = \frac{D_w}{\delta t} \quad (4.26)$$

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 essential result in the context of the *String of Beads* model calibration, which will be discussed in Chapter 5.

4.6.3 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 4.19 and requires careful explanation to avoid confusion. There are two main underlying process in this stage, a univariate autoregressive process and a two-dimensional power-law filtering process. The first generates temporally correlated, pseudo-random fields of pixels and the second imposes a realistic spatial correlation structure on the fields. Theory pertinent to this stage of simulation was presented in Sections 4.4 and 4.5 and relevant analysis was presented in Sections 3.4.3, 3.7.1 and 3.11.




Figure 4.19 - Stack algorithm used to simulate pseudo-random images from image scale statistics

The univariate autoregressive model 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, $Z_{(t-5)}, Z_{(t-4)}, \dots, Z_{(t)}$, each 256 x 256 pixels in size. To start the process off at the beginning of an event, the arrays $Z_{(t-5)}, \dots, Z_{(t-1)}$ are filled with normally distributed pseudo-random noise. The $Z_{(t)}$ array is then computed at each of its 65536 pixels, z_t , independently using Equation 4.27.

$$z_t = \phi_1 \cdot z_{t-1} + \phi_2 \cdot z_{t-2} + \dots + \phi_5 \cdot z_{t-5} + a_t \quad (4.27)$$

where

a_t is a shock term

$\phi_1 \dots \phi_5$ are autoregressive parameters

The autoregressive parameters 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.

To homogenise the temporal structure of the $Z_{(t-5)}, \dots, Z_{(t-1)}$ arrays, the autoregressive process is allowed to run for several iterations before image simulation begins. Thereafter, at each time step, t , the $Z_{(t)}$ array is spatially shifted according to the mean field advection vector output by the first stage of simulation. In its current state, this is a constant direction, B_{adv} , and speed, C_{adv} , however a more sophisticated approach could be adopted if required. Array edges are refilled with noise when the field is shifted.

The shifted $Z_{(t)}$ array (comprised of $z_t \sim N(0,1)$) 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. This is to avoid introducing long range spatial dependent structure at scales in excess of the 128km wave number. Qualitative analysis of images simulated without zeroing the edges, indicate that it is perhaps naïve to suggest that the rainfields scale for wave numbers much larger than 128km. 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 4.5. 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 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 3.4.3, a threshold value, T_r^* , for which the ratio

$$\frac{number(r_t^* \geq T_r^*)}{number(r_t^*)} = WAR_i \quad (4.28)$$

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

$$\frac{number(r_t^* \geq 0)}{number(r_t^*)} = WAR_i \quad (4.29)$$

Scaling of the $R^*_{(t)}$ field has to be done iteratively by first computing IMF'' using Equation 4.30,

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

where λ is a scalar term to be determined iteratively. The objective function for the iterative solution, ε , is then given by Equation 4.31.

$$\varepsilon = | IMF'' - IMF_i |$$

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 form of the third stage of simulation discussed 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.

4.6.4 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 4.1. 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. 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. 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 and requires the full set of image scale statistics from an observed event. Calibration modes will be discussed in Chapter 5.

Table 4.1 - Summary of the structure of the *String of Beads* model

Simulation Stage	First Stage	Second Stage	Third Stage
User input	String of Beads Simulation: - Simulation start date - Simulation stop date - Random seeds - Output options String calibration: (1, 2 and 3) + 5. 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 Alternating 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 - 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_{var}^{**}, S_{var}^{**}, M_{imp}^{**} and S_{imp}^{**}	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 4.6.2.

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.

4.7 COMPUTING CONSIDERATIONS

Until recently, 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 computing requirements would be a 300MHz 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, on this minimum specification machine, requires less than 1 seconds. The second stage would be complete in under 10 seconds and the third stage 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 5 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.

4.8 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* and *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 Validation

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, validation of the model output, after calibration, will be presented.

5.2 BEAD CALIBRATION PROCESS

The bead calibration is the calibration of the third stage of the *String of Beads* model, and involves determining the parameters $\phi_1 \dots \phi_5$ of the univariate autoregressive process defined in Equation 4.27 of Section 4.6.3. These define 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. It may be worth attempting the PSI correlation analysis outlined in Section 3.11 on a more precise dataset. In the absence of a more precise dataset, the parameters $\phi_1 \dots \phi_5$ were *inferred* from the Pre-1997 data through a process of iterative fitting.

The technique involves the spatial upscaling analysis of a rainfall event, presented in Section 3.12, and the third stage of simulation, presented in Section 4.6.3. After grasping the concepts presented in these two sections, the bead calibration process is simple:

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. 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.
4. 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).
5. Using the same routines as in step 2, quantify the space-time upscaling properties of the simulated event (Section 3.12).
6. 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 3.68.
7. Iterate from step 3, adjusting the model PSI parameters $\phi_1 \dots \phi_5$ to obtain their 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, univariate process which describes the temporal correlation structure of the (spatially uncorrelated) PSIs in the normalised, Lagrangian reference frame as described in Section 4.6.3. The 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, due to time constraints, has not been fully automated at the time of writing this document. Instead, the first three of the $\phi_1 \dots \phi_5$ parameters were manually manipulated to obtain a reasonable solution and are given as:

$$\phi_1 = 1,6 \quad \phi_2 = -0.4 \quad \phi_3 = -0.2 \quad a = 0,1$$

The ϕ_4 and ϕ_5 terms are set to zero. Using these *best-guess* parameters for simulation, the comparative downscaling analysis for the 1996 observed sequence, whose image

scale statistics are presented in Figure 3.29, and the *best-guess* simulated sequence, are presented in the analyses which follow beginning with Figure 5.1.




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

Bypassing the automation of the Bead calibration process was an operational decision in an attempt to get the *String of Beads* model working. The intention was to revisit the topic once the String modelling problem had been addressed and the model had been tied together, however this has not been done due to time constraints on the research.

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 simulated equivalents. The two series are identical in their WAR_i , IMF_i and advection vector image scale statistics. 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. Higher spatial resolution aggregations are not as faithfully reproduced as the 128 km scale, but they are reproduced well never-the-less, validating the ability of the model to downscale the image scale parameters.

An interesting point in this regard is the effect of the mean field advection vector. Figure 5.2 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 5.2 – The effect of field advection on the spatial upscaling behaviour

The two simulated sequences analysed in Figure 5.2 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 the advection component in the simulation. Differences between local storm advection and the mean field advection may also have a significant effect on the temporal correlation structure, although this has not been considered in this study. Automation of the bead calibration process would facilitate a more extensive investigation into these matters.

The other 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 5.3 shows the power spectra of five space-time aggregated images sampled at corresponding points from the observed and the simulated aggregated sequences.




Figure 5.3 - Power spectra for aggregated images (of various spatial and temporal scales) sampled at five corresponding points from the observed and simulated sequences

The quality of fit is remarkably good and this is typical of the type of fit observed for most of the aggregated images in the two 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 4.6.3.

Figure 5.3 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 5.4 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 simulated sequence, for the 15, 60 and 240 minute temporal scales.




Figure 5.4 – Comparisons of power spectral exponent for observed and simulated sequences, aggregated at 1 km - 15, 60 and 240 minute time scales.

The analyses of Figure 5.4 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. The power-law exponent of these aggregated sequences is not much different from that observed for the 1km-5min 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 1km-15min sequence is 2,5, compared to 2,6 for the 1km-60min sequence and 2,7 for the 1km-240min sequence. Figure 5.5 considers the same analysis at the 8km spatial scale.




Figure 5.5 - Comparisons of power spectral exponent for observed and simulated sequences, aggregated at 8 km - 15, 60 and 240 minute time scales.

At the 8km spatial scale, the 128 km image is reduced to 16 x 16, 8km 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 8km scale than at the 1km scale. Nevertheless, the analysis reveals a noisy, but unbiased representation of the spatial upscaling properties by the model. The average of the spectral exponents for the 8km-15min sequence is 1,7, compared to 2,2 for the 8km-60min sequence and 2,4 for the 8km-240min sequence. Similar analyses for the 2,4 and 16 km spatial scales revealed similar trends to those presented in Figures 5.4 and 5.5 for the 1 and 8km spatial scales.

In spite of manually fitting the PSI autoregressive parameters, $\phi_1 \dots \phi_5$, the upscaling analyses presented in this section reveal encouraging results. Automatic calibration of the $\phi_1 \dots \phi_5$ parameters should result in better approximations for these parameters, further improving the upscaling performance of the *String of Beads* model at the pixel scale. From these encouraging results, we can conclude that the bead process (non-zero space-time rainfall) is well modelled by the proposed structure.

5.3 *STRING CALIBRATION PROCESS*

There are two components to the String calibration process. The first is to measure the event scale statistics as shown in Sections 3.9 and 3.10 and to extract the Cumulative Distribution Functions (CDFs) for the 500Hpa wind speed and direction as shown in Section 3.7.1. Once these parameters of the model have been measured from the data and the *Bead* calibration process is complete, the iterative calibration of the *string* process can begin. It is this iterative *String* calibration process which will be described here. Two algorithms will be presented. The first which describes the *String* calibration process in its current form, which unfortunately does not work, and the second which has not yet been tried, but which should work.

5.3.1 *An unfortunate sadness*

With reference to Section 2.6, the *String* of the model was developed using the Post-1997 radar dataset which, it was discovered late in the research, misrepresented the PSI by a factor of 2. Initial results, using the *incorrect* radar dataset, were very encouraging. The least-squares objective function, used to compare the radar dataset to the raingauge dataset, optimised beautifully and qualitative analysis of the optimised *String* process seemed to reveal the type of seasonal variation which would be expected in the Bethlehem region, based on the 50 year raingauge record. Testing of the model on Swiss, German and English datasets revealed the first signs of trouble, when the model could not find a solution for the English dataset. After development of model validation routines, testing revealed the error in the radar dataset.

Once the error was discovered, correction of the data was a relatively simple process, however the repercussions on the calibration of the *String of Beads* model were severe. For a start, all of the analyses of the radar dataset had to be repeated in order to obtain the corrected event scale statistics now presented in Chapter 3. Being a highly non-linear process, the result of all the adjustments, was a poorly defined objective function. Attempts to alter starting values and constrain the parameters of the model unfortunately failed and in its current form the *String* component of the *String of Beads* model does not calibrate. In order to make it work, the objective function must be redefined and this requires extensive re-programming. Time constraints on this project have prevented completion of this aspect of research.

5.3.2 Run-length calibration

The idea behind run-length calibration 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 a 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. Figure 5.6 shows 10 years of the 50 year daily time series of WAR_d as estimated from the 54 raingauge network.




Figure 5.6 - 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.

Using the wet-dry threshold of 10% exceeding 3mm, as indicated in Figure 5.6, 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 Figures 3.11 to 3.14. These histograms form one half of the objective function used to calibrate the *String* process.

The second half of the objective function is 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 suggests that the wet or dry daily classification, based on the 10% WAR_d threshold, can 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 is given by Equation 5.1.

$$\psi = \sum_{k=1}^{12} \sum_{j=0}^{25} \left(n_s^{jk} - n_o^{jk} \right)^2 \quad (5.1)$$

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 4.6.1, 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 would be the parameters of the lognormal distributions fitted to all of the observed events, plotted in Figures 3.44 and 3.46. That is, starting values of:

$$\mu_w = 5,4 \quad \sigma_w = 1,3 \quad \mu_d = 5,5 \quad \sigma_d = 1,7$$

as measured for South African, summer conditions.

Minimization of the objective function in Equation 5.1 gave reasonable results which are plotted in Figures 5.7 through 5.11.




Figure 5.7 – Number of dry days per month observed during the 50 year calibration period

Figure 5.7 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. Figures 5.8 through 5.11 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 5.8 - Histogram of Summer Wet Run Lengths




Figure 5.9 - Histogram of autumn Wet Run Lengths




Figure 5.10 - Histogram of Winter Wet Run Lengths




Figure 5.11 - Histogram of Spring Wet Run Lengths

Figures 5.8 to 5.11 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 run length calibration process is unstable. Where the model converged quickly on the incorrect dataset, once the data had been corrected, calibration required considerable nursing to ensure reasonable results. 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 5.1 (each taking 8 to 10 hours), it must be 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 5.12.




Figure 5.12 - 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. Similar qualitative analysis made sense when the model was calibrated on the Swiss, German and English datasets. 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 5.8. This is a verification of the model integrity and confirms that the calibrated model is performing as it is supposed to.




Figure 5.13 – 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 of PSI (probability of exceedence of a given rainfall rate), the distribution of monthly rainfall, the distribution of annual rainfall. These would be compared to the analysis of the radar and raingauge network as presented in Chapter 3.

Integration of the IMF_i according to Equation 4.25 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, show an under-estimation of the total annual rainfall. Figure 5.14 shows the spatially averaged annual rainfall total for the study area.




Figure 5.14 - 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 5.14, as well as their corresponding standard deviations (dotted lines) of 129 mm² and 108 mm² respectively. These show a poor fit, although it is of the right order of magnitude.

Figure 5.15 shows the spatial distribution of rainfall on the image, to be compared to that of Figure 3.39.




Figure 5.15 - 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 this simulated image when compared to Figure 3.39. Nowhere on Figure 3.39 does the annual rainfall drop below 400 mm as observed in the South East corner of Figure 5.15. Streaks across the image are as a result of very high intensity cells, an aspect of the Bead simulation process which may require further investigation.

Moving on to the monthly statistics, this is the most disappointing aspect of the simulation and another good reason to abandon the run length calibration process. Figure 5.16 shows the mean monthly rainfall for the region, averaged over the 11 year simulation.




Figure 5.16 – 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 5.12, 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 encouraging aspect of the results shown here is that the model is behaving as it was designed, since the parameters of the wet and dry spell durations can be used to adjust the mean monthly rainfall. If a better method of calibrating these parameters can be developed, the model should hopefully perform well.

5.3.3 Integrated IMF_i calibration

With hindsight, there is a far simpler method of calibrating the μ_w , σ_w , μ_d and σ_d parameters which has not been tested at the time of writing this report. The monthly rainfall totals can be approximated by integrating the IMF_i parameter according to Equation 4.25. The proposed calibration technique is to change the run length objective function of Equation 5.1, to a distribution based objective function as shown in Equation 5.2.

$$\psi = \sum_{k=Jan}^{Dec} \sum_{j=1}^N (x_s^{jk} - x_o^{jk})^2 \quad (5.2)$$

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. Figure 3.7 showed that although the rainfall totals of the wetter months can be approximated by a Gaussian distribution, the rainfall totals of the drier months are far from Gaussian. The technique for calibration would then be:

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 4.25
5. For each calendar month, determine the rank j ($j = 1, 2, \dots, N$) 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 eliminates the need for the rather arbitrary classification of wet and dry days and it also eliminates the need for the empirical relationship between WAR_d and $\overline{WAR_i}$. It is expected to be considerably more robust and will ensure that the monthly and annual rainfall totals are reproduced after calibration. If necessary, it could be combined with the run-length objective function of Equation 5.1. Time constraints on this study have prevented this technique from being implemented as a result of the need to correct and re-analyse the radar dataset.

5.4 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 are first transformed into their Gaussian WAR_i^* and IMF_i^* using the target distribution functions shown in Figures 3.50 and 3.51, according to the method illustrated in Figure 3.52. 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 4.18, is then used to simulate a range of possible future WAR_i - IMF_i pairs. Figure 5.17 uses the image scale statistics of the 1996 event of Figure 3.29 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 5.17 - 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 into $N(0,1)$. 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 $Z_{(t-5)}, Z_{(t-4)}, \dots, Z_{(t)}$ arrays of the third stage of simulation shown in Figure 4.19. Possible future PSI fields are then obtained by simulation, using the forecast WAR_i - IMF_i image scale statistics obtained above. Figure 5.18 illustrates the results of this procedure this for a simulated sequence.




Figure 5.18 - 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 5.17.

The sequences shown in Figure 5.18 were simulated using the WAR_i - IMF_i series of Figure 5.17 and the colours of the boxes around each sequence, corresponds to the line colour in Figure 5.17. The images are shown at a 10 minute interval. 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 5.17. 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.

5.5 SUMMARY

This chapter has demonstrated the processes whereby the Bead process and the String process can be iteratively calibrated. The results of the upscaling tests of the Bead modelling process are very encouraging over all of the observed space-time scales. Development of the calibration technique for the String process was unfortunately hindered by a corrupt dataset and requires additional work, however an improved technique has been suggested for the benefit of future research into this matter. This technique should be faster, simpler and more robust than the method currently in place. Finally, the *String of Beads* model can be used as a short term forecasting tool and to simulate a range of *what if* scenarios, based on available radar information.

Chapter 6

Conclusions

The *String of Beads* model continues to show great promise as a high resolution space-time rainfall model. The advances in this study have made it useful both as a simulation tool and as a short term forecasting tool, which can be run on a modern personal computer.

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. The model can incorporate the information extracted from a daily raingauge network, in order to simulate the seasonal trends in rainfall statistics for any region. 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.

It was intended that the model be self calibrating, requiring only readily available datasets in text format to adjust its parameters, but unfortunately the first design of the calibration module failed. Time constraints on the research, attributable to a corrupt radar dataset, have prevented the implementation of a more robust technique of calibration, although a method has been proposed and is described in the latter stages of this report for the benefit of any future developments of the *String of Beads* model.

Finally, 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

References

- Aitchison, J. and JAC 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.
- Basson MS., RB Allen, GGS Pegram and JA van Rooyen.**(1994). *Probabilistic Management of Water Resource and Hydropower Systems*. Water Resources Publications, Highlands Ranch, Colorado, USA.
- Battan, LJ.** (1973). *Radar Observation of the Atmosphere*. University of Chicago Press, Chicago, USA.
- Beasley, JD. and SG 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, TL.** (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, GEP. and GM Jenkins.** (1970). *Time Series Analysis: Forecasting and Control*. Holden Day, San Francisco.
- Brenier, P.** (1990). *Simulations Dynamique Multifractale Des Nuages*. MS Thesis. Ecole Normale de Science et Technologie, Paris, France.
- Court, AP.** (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, PSP.** (1994). A Generalised Point Process Model for Rainfall. *Proc. Royal Soc. London, A*, Vol. 447, pp 23-37.
- Crane, RK.** (1990). Space-Time Structure of Rain Rate Fields. *J. Geophys. Res.*, Vol. 95, No. D3, pp 2011-2020.
- Dempster AP., NM Laird and DB Rubin.** (1977). Maximum likelihood estimation from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B*, Vol. 39, p1-38.
- Foufoula-Georgiou, E., and W Krajewski.** (1995). Recent Advances in Rainfall Modelling, Estimation and Forecasting. *Reviews of Geophysics, Supplement*, pp1125-1137, July 1995.
- Griffiths, P. and ID Hill.** (1995). (eds). *Applied Statistics Algorithms*. Royal Statistical Society.
- Gupta, VK and EC Waymire.** (1987). On Taylor's Hypothesis and Dissipation in Rainfall. *J. Geophys. Res.*, Vol. 92 , No. D8, p 9657-9660.
- Gupta, VK and EC Waymire.** (1990). A Statistical Analysis of Mesoscale Rainfall as a Random Cascade. *J. App. Meteorology*, Vol. 32, p251-267.
- Haberlandt, U.** (1998). Stochastic processes using regionalized model parameters. *J. Hydrologic Eng.* 3 (3). p160-168.
- Hurvich, CM and CL Tsai.** (1989). Regression and Time Series Model Selection in Small Samples, *Biometrika*, v76, n2, p297-307
- Knuth, DE.** (1969). *The Art of Computer Programming. Vol. 2 - Seminumerical Algorithms*. p 9-34, Addison-Wesley.
- Kuc, R.** (1988). *Introduction to Digital Signal Processing*. McGraw Hill, New York.
- Makhuva T, GGS 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, GGS 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
- Marshall, JS. and WM Palmer.** (1948). The Distribution of Raindrop Size. *J. Meteor.* Vol. 5, p165.
- Menabde, M., A Seed, D Harris and G Austin.** (1999). Multiaffine Random Field Model of Rainfall. *Water Resour. Res.* Vol. 35, No. 2, pp509-514.
- Mittermaier. MP.** (1999). *Investigating the Characteristics of the Radar Reflectivity Profile*. Unpublished MScEng dissertation, Civil Engineering, University of Natal, Durban
- Mittermaier, MP. and DE Terblanche.** (1997). Converting Weather Radar Data into Cartesian Space: A New Approach Using DISPLACE Averaging. *Water SA*, Vol. 23, No. 1, pp46-50.
- Nelder, JA. and R Mead.** (1965). A Simplex Method for Function Minimisation. *Computer Journal*, Vol. 7, pp. 308-313.
- Press, W., SA Teukolsky, WT Vetterling and BP Flannery.** (1992). *Numerical Recipes in C - The art of scientific computing*. 2nd ed. Cambridge University Press.
- Pegram GGS.** (1997). Patching rainfall data using regression methods. 3. Grouping, patching and outlier detection. *Journal of Hydrology*, 198, p319-334
- Pegram, GGS. and AN 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, GGS. and AN Clothier.** (2001a). High resolution Space-Time Modelling of Rainfall: The "String of Beads" model. *J. Hydrology*. V 241. p26-41.
- Pegram, GGS. and AN Clothier.** (2001b). Downscaling Rainfields in Space and time, using the String of Beads Model in Causal Mode. *HESS*. (in press).

- Pegram, GGS. and AW 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.
- 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, p9693-9714.
- Smithers, JC and RE 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.
- Seed, AW.** (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.
- Steyn, PCL and RT Brintjes.** (1990) Convective cloud characteristics for the Bethlehem area. *Water SA*, 16, p115 - 118
- Taylor, GI.** (1938). The Spectrum of Turbulence. *Proc. R. Soc. London, Ser. A*, Vol. 164, pp476-490.
- Uijlenhoet, R.** (1998). Raindrop Size Distributions and the Z-R Relationship. *Proceedings of Advanced Study Course on Radar Hydrology for Real Time Flood Forecasting*, University of Bristol, England, 24th June to 3rd July.
- Van Heerden, J. and PCL Steyn.** (1999). Weather radar measurements of rainfall for hydrological and other purposes. WRC 693/1/99. Report to Water Research Comm., Pretoria, South Africa.
- Wei, WWS.** (1990). *Time series analysis, univariate and multivariate methods*. Addison Wesley.
- Wichmann, BA. and ID 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 Variability in Geophysics : Scaling and Fractals*. D Schertzer and S Lovejoy (eds). Kluwer Academic.
- Woolhiser, DA. and Pegram GGS.** (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.