Final report K5/1425

Daily Rainfall Mapping Over South Africa: Modelling

Geoff Pegram, Scott Sinclair and Stephen Wesson

University of KwaZulu-Natal

September 2005

Executive summary

The background to the study

At the beginning of 2003 the South African Weather Services (SAWS) website made a new contribution to water resources management, agriculture, flood forecasting and even locust swarm prediction. The site offered a detailed daily map of the previous day's rainfall over the whole of RSA and some neighbouring territories.

Called SIMAR (Spatial Interpolation and Mapping of Rainfall) the site has been visited many times since its inception. Probably most importantly, it has been criticized by its users, in particular the DWAF flood studies unit, where it has given rainfall estimates different from those recorded by independent organizations. This is not a negative thing; it is a positive one for two reasons: first the site is being used; second the criticism allows the designers/researchers involved the opportunity of feedback so they can improve the product. This last aspect creates an important mindset because it encourages accountability of those responsible; that is why this project was undertaken.

K5/1425: Daily Mapping of Rainfall over South Africa: Modelling, has a companion project K5/1426: Daily Mapping of Rainfall over South Africa: Infrastructure and Capacity Building which deals with the infrastructure of the precipitation measuring devices used by SAWS. This project (K5/1425) is also closely linked with another WRC project K5/1429: National Flood Nowcasting Initiative to which it feeds its products, especially the forecasting of rainfields.

During the course of the project, opportunities to form international linkages were realised. First SAWS spearheaded a contract with EUMETSAT for research in the utilization of METEOSAT Second Generation (MSG) over Southern Africa, with UKZN, ARC- Institute for Soil Climate and Water (ICSW) and CSIR- Satellite Application Centre (SAC) as partners. This provided us the facility to obtain MSG data directly from EUMETSAT. The figure below shows postgraduate students of Civil Engineering Programme of UKZN assisting CSIR technicians with the installation of the 2.5m satellite dish to receive MSG data.

Next, ESA's TIGER initiative allowed UKZN with SAWS ARC-ICSW, DWAF and NDMC as partners to enter into a contract to obtain data free from ESA for the purpose of soil moisture estimation by satellite. This initiative spawned the WRC consultancy with UKZN: K8/598: Soil Moisture Mapping from Satellites. International recognition of the work being done here via links with European Union researchers (in particular Professors Ezio Todini and Vincenzo Levizzani) resulted in two important invitations: the project leader was invited to the IPWG (closed) workshop in Monterey in October 2004 and then (with Dr Deon Terblanche and Ms Pearl Mngadi, both of SAWS) to the

GPM-GV (closed) workshop in Taipei in September 2005. The importance of these links is that SIMAR and its products and successors are seen as useful ground validation opportunities for NASA's, ESA's and EUMETSAT's programmes for the estimation of rainfall from space. Thus, not only is this WRC-sponsored research of great local importance, but it is also likely to have global impact.

Research Accomplishments of this Study

Six areas of improvement of SIMAR were earmarked for research in this study.

- Repair parts of radar rainfall images contaminated by reflectivity values which are not precipitation (e.g. ground clutter, anomalous propagation etc.)
- Extrapolate to ground level those precipitation estimates made aloft by radar and confirm that the estimates are valid. The figure below shows what the radar sees (left - in 1km layers) and after infilling (right); the exaggerated vertical scale (3km) versus the horizontal scale (200km) is for illustration purposes.

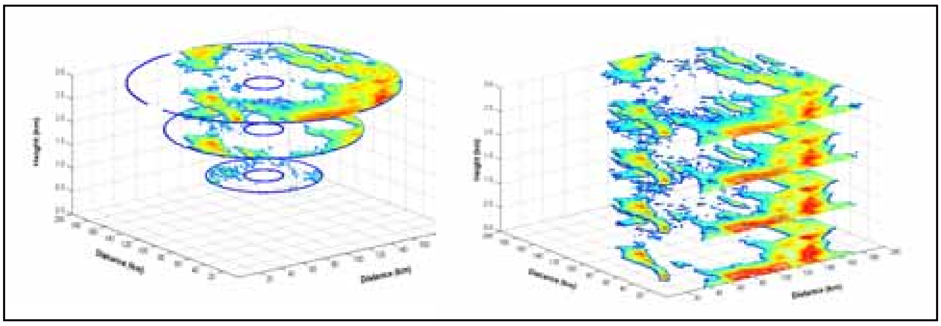

- Use instantaneous images of spatial rainfall captured by radar or satellites to compute meaningful accumulations of rainfall over periods from 5 minutes to 24 hours and longer; the upper and lower rows of images in the figure below are respectively operational and morphed accumulations of images; the methodology for the latter set was developed in this study.

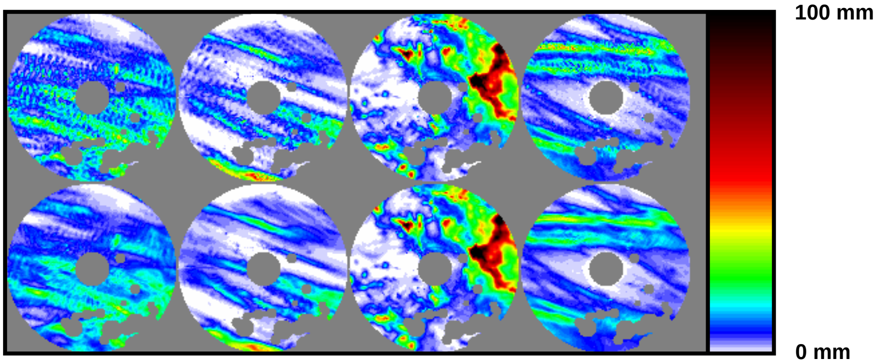

- Determine the optimum way of combining accumulated radar or satellite rainfall fields with rain gauge data; this led to the adoption of the conditional merging algorithm as the method of choice.

- Find a way of using the new MSG data to estimate rainfall, based on the methodology developed in SIMAR, which was appropriate for METEOSAT-7 data; this latter satellite is shortly to be decommissioned. The figure below (the day's rainfall on 22nd May 2005) shows the precision of the new measurement of rainfall using MSG.

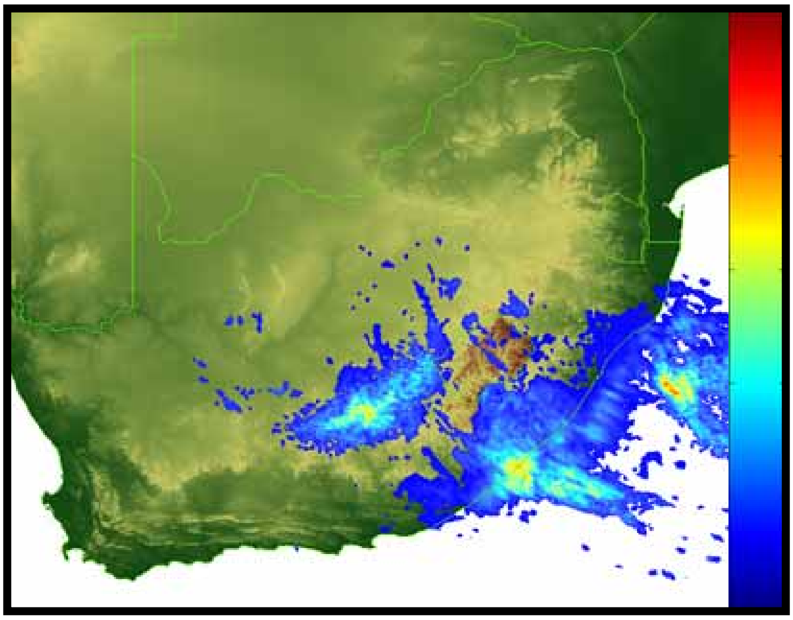

- Determine the optimal method for forecasting rainfall fields into the near future (up to 2 hours ahead). Candidates included:
    1. an adaptation of the String of Beads model (Clothier and Pegram, 2002),
    2. an algorithm S-PROG developed by Dr Alan Seed of the Australian Bureau of Meteorology for the 2000 Sydney Olympics and
    3. a novel method developed here using 2-dimensional Empirical Mode Decomposition in combination with Adaptive Time Series Forecasting.

These six research streams were all successful and are described in the main body of this report. The techniques have been combined in a new Improved Rainfall Merging Algorithm for producing optimally estimated rainfall fields called IRMA, which is also described in the report.

Knowledge Dissemination

A large proportion of the material appearing in this report has (i) already been published, (ii) been submitted for publication (iii) been presented at national or internationally conferences, symposia or workshops. A list of these publications and other contacts made by those involved in the projects follow.

Papers Published in Journals

Sinclair S. and Pegram G.G.S., (2005a), "Combining radar and rain gauge estimates using conditional merging", *Atmos. Sci. Let.*, **6**, 19-22.

Sinclair S. and Pegram G.G.S., (2005b), "Empirical Mode Decomposition in 2-D space and time: a tool for space-time rainfall analysis and Nowcasting", *Hydrol. Earth Syst. Sci.*, **9**, 127-137.

Wesson S.M. and Pegram G.G.S., (2004), Radar rainfall image repair techniques, *Hydrol. Earth Syst. Sci.*, **8(2)**, 220 – 234.

Papers Published on Websites

Pegram G.G.S., Deyzel I.T.H., Sinclair S., Visser P., Terblanche D. and Green G.C., (2004), Daily mapping of 24 hr rainfall at pixel scale over South Africa using satellite, radar and raingauge data, *2nd IPWG Workshop*, Naval Research Laboratory, Monterey, CA, USA, 25-28 October. http://www.isac.cnr.it/~ipwg/meetings/monterey/pdf/Pegram.pdf

Papers submitted for publication

Wesson, S.M. and Pegram, G.G.S., (2005), Improved Radar Rainfall Estimation at Ground Level, *Natural Hazards Earth System Sciences (NHESS)*, Submitted in August 2005.

Presentations at international symposia/conferences

Pegram G.G.S. and Sinclair S., (2004), National Flood Nowcasting System towards an integrated mitigation strategy in South Africa, *Proceedings of the 6th International Symposium on Hydrological Applications of Weather Radar*, Melbourne, Australia, 2-4 February.

Sinclair D.S., Ehret U., Bardossy A and Pegram G.G.S., (2003), Comparison of Conditional and Bayesian Methods of Merging Radar & Rain gauge Estimates of Rainfields, *Presentation at EGS - AGU - EUG Joint Assembly*, Nice, France, April.

Sinclair S. and Pegram G.G.S, (2004), Combining Radar and Rain Gauge Rainfall Estimates for Flood Forecasting in South Africa, *Proceedings of the 6th International Symposium on Hydrological Applications of Weather Radar*, Melbourne, Australia, 2-4 February.

Sinclair S. and Pegram G.G.S., (2005), Space-time rainfall analysis and nowcasting using Empirical Mode Decompostion in 2D, *EGU General Assembly*, Vienna, Austria, April.

Wesson S.M. and Pegram G.G.S., (2005), Repairing Radar Volume Scans, *EGU General Assembly*, Vienna, Austria, April.

Presentation at national symposia/conferences

Sinclair S. and Pegram G.G.S, (2003), Combining Traditional and remote sensing techniques as a tool for Hydrology, Agriculture and Water Resources Management, *Proceedings of the 11th South African National Hydrology Symposium*, Port Elizabeth, South Africa.

Wesson S. and Pegram G.G.S., (2003), Efficient cleansing of Radar Rainfall images, *Proceedings of the 11th South African National Hydrology Symposium*, Port Elizabeth, South Africa.

International visits and contacts

Scott Sinclair presented two of the three papers at EGU in Nice in 2003 and had the benefit of arguing his case with Professor Ezio Todini, face to face.

Professor Geoff Pegram attended EGS in Nice in 2003 and EGU in Vienna in 2005 where he made presentations and attended the editorial board meetings of HESS as Associate Editor.

He was Visiting Research Fellow in the Civil and Environmental Engineering Department, University of Melbourne, during 2003, 2004 and 2005 for 8 weeks, where he continued his research collaboration association in areas of precipitation modelling and water resources. The last two visits were spent exploring the benefits of Empirical Mode Decomposition and non-stationary Adaptive Time Series modelling, a point of departure for ***Scott Sinclair's*** work.

In 2003, he was Invited to participate as rapporteur (and future full member of the Steering committee) at the European Union project: MUSIC / CARPE DIEM Joint Workshop with End Users, at Düsseldorf-Neuss, Germany: "CURRENT FLOOD FORECASTING PRACTICE IN EUROPE". This invitation was repeated for the final meeting in Helsinki in June 2004. These visits continued to maintain the good research collaboration relationship with Professor Todini, one of the foremost stochastic hydrologists in Europe.

In March 2004, he was Visiting Professor and Guest Lecturer: ENWAT: International Doctoral Program 'Environment Water' University of Stuttgart, hosted by Professor Andras Bardossy, where work on non-linear modelling was undertaken and published in Water Resources Research this year. This visit will be repeated in April 2006.

Then in July 2004 he was Visiting Professor at GRAHI (Group for Research Applications in Hydrology Institute), Barcelona Spain, hosted by Professor Daniel Sempere-Torres and working together with their Doctoral students on radar-hydrology, which gave great impetus to the work of ***Stephen Wesson*** in the UKZN group. This fruitful visit will be repeated in August 2006.

Later in October 2004, he was an invited participant at 2nd International Precipitation Working Group meeting (a closed workshop, sponsored by Coordination Group for Meteorological Satellites, WMO, NASA & EUMETSAT) in Monterey, CA. Here he presented a summary of the WRC research project K5/1153: SIMAR. These contacts have led to the invitation to the Global Precipitation Monitoring – Ground Validation workshop in September in Taipei, another closed workshop at which the current WRC research project will be showcased.

Capacity and Competency Building

Four post-graduate students have been involved with this project in the Civil Engineering programme of UKZN over the last two and a half years:

- ***Scott Sinclair*** has worked with the project leader since the beginning of 1999 when he started his Masters studies on catchment modelling research project K5/1005: A Linear Catchment Model for Real Time Flood Forecasting (Pegram and Sinclair, 2002). He went on to work (as an engineer employed by Umgeni Water) on the WRC project K5/1217: Umgeni Nowcasting Using Radar – An Integrated Pilot Study (Sinclair and Pegram, 2004b); he is now completing his PhD. He was responsible for the merging, accumulation and forecasting algorithms in this research project. He was also responsible for the installation of the hardware and software to receive and archive MSG data, then adapt the SIMAR algorithms to interpret the MSG data stream and convert it into rainfall estimates. In addition, he has worked in parallel on the WRC project: K5/1429

National Flood Nowcasting Initiative, has deputized for the project leader at several RUMSA/MTAP meetings and workshops and has filled in as the representative of the research group during the project leader's absence overseas. He has developed into a highly competent and knowledgeable scientist/engineer/manager who understands the intricacies of advanced computer coding, the workings of satellites and their data products, besides having mastered the applied mathematics and geostatistics necessary to develop the algorithms presented herein.

- ***Stephen Wesson*** is completing his MScEng degree which he started in 2003. He was responsible for developing the ground clutter removal algorithm and the Cascade Kriging technique designed to extrapolate information from radar measurements aloft in order to estimate rainfall at ground level. He has developed keen research skills using geostatistics and advanced computer programming and has mastered some difficult image processing algorithms. He has become scientifically mature, has shown encouraging sparks of originality and has made a sound contribution to this study.
- ***Mohamed Parak*** is completing his MScEng degree which he started in 2004. Until the beginning of 2005, he was working in parallel on a flood study for the department of Transport under the supervision of the project leader. Since then he has made a contribution to this study through his work on geospatial statistics and GIS modelling. His major contribution to the WRC projects is his adaption of the TOPKAPI model for distributed catchment modelling which has application in the partner project K5/1429: National Flood Nowcasting Initiative. He has shown a particular facility with the complexities of the mathematics of catchment modelling, its description and its local application.
- ***Ntokozo Nxumalo*** started his Masters studies at the beginning of 2005. He is working on a parallel WRC study K8/598: Soil Moisture Mapping from Satellites. His contribution to the rainfields study has been peripheral, however together with all the graduate students in the group he meets with the project leader/supervisor every Friday morning when the week's work is discussed. He is growing rapidly as a researcher with flair and is showing a marked facility with image analysis.

Recommendations

There was concern expressed at the final steering committee meeting that the access to the SIMAR web-site was limited to pictorial information, while numerical data of rainfalls on quaternary catchments was sought by some; indeed, gridded data at 1 km resolution would be preferred.

Two commitments rose from this concern: the first was to install the new software developed in this study; the second was to provide access to the data to those with a legitimate need, on approach to SAWS.

In more detail, to implement the first commitment, Stephen Wesson with the help of Scott Sinclair will adapt the research software to on-line code. Stephen will then spend as much time in Bethlehem as is necessary, before the end of 2005, to transfer the code and ensure with the help of Pieter Visser and Karel de Waal that it operates satisfactorily in on-line mode in real time.

Finally, it follows naturally from this study that there is a need for validation of the rainfall fields produced by the project: Daily Rainfall Mapping over South Africa: Modelling (DARAM). Towards this end, an invitation was sent to Professor Geoff Pegram by Drs Joe Turk and Peter Bauer, co-chairpersons of the International Precipitation Working Group (IPWG), to participate in the Pilot Evaluation of High Resolution Precipitation Products (PEHRPP), sponsored by the IPWG. From their letter of invitation: "PEHRPP is an effort to characterise as clearly as possible the errors in various high resolution precipitation products on many spatial and temporal scales over varying surfaces and climate regimes."

The work started in SIMAR and continued in this study will provide a platform to make a valuable contribution to the global effort initiated by IPWG, because it will force those of us responsible for the rainfall estimation to closely evaluate the errors in our own products. This can only be good for making future improvements to this important work.

Acknowledgements

The authors would like to express their thanks to the WRC for providing funding and in particular, the project steering committee for their valuable support, comments and suggestions throughout the life of this project. Specific thanks are also extended to SAWS, without their support in providing data and scientific comment many of the achievements presented here would not have been possible.

Table of contents

List of Figures

1 Introduction

The three WRC projects carried out under the umbrella programme called SIMAR (Kroese, 2004; Deyzel et al., 2004; Pegram, 2004) developed the infrastructure and computational procedures to provide a map of the previous 24 hour rainfall over the southern African region, on a daily basis, on the SAWS website. This product has been operating routinely since early 2003, providing detailed rainfall information over the subcontinent. In its conception, it was a brave and ambitious project, enabling those involved to learn much from its routine operation, particularly concerning its short-comings.

The present project K5/1425 aimed at improving the daily mapping of rainfall, in particular through the development of improved algorithms, designed to make the product more accurate and trustworthy. This project builds on SIMARs strengths and takes cognizance of the lessons learned so that it is more effective for the following reasons:

- The satellite Meteosat-8 (MET-8) was launched in August 2002 by the European Space Agency. When the data came on stream at UKZN in May 2005 (2 years later than envisaged), there was an evident improvement of the estimation of rainfall because the information sampling rate and the spatial precision are at least doubled when compared to Meteosat-7.
- There were recent publications in the Hydrometeorological literature whose adaption to the present project make the computations needed for the estimates faster, more accurate and objective.

There was considerable spin off and cross-pollination with other initiatives, especially the companion project K5/1429: a National Flood Nowcasting System. There is a natural progression from daily rainfall maps at 1km resolution over the whole country to smaller areas of concern at higher sampling frequency and even rainfield forecasting. Beneficiaries of these products will all be reliant on sound rainfall information for decision-making.

It is thus envisaged that catchment management authorities, local government disaster managers doing flood forecasting, agriculturists estimating drought and crop water demand, transport operators fearful of bad weather, eco-tourism operators checking the safety of their schemes to raft rivers and sports administrators dependent on the weather for their scheduling, will all benefit from the output from this project.

Turning to more technical detail, in the previous mapping initiative conducted by the project leader under the SIMAR umbrella (K5/1153), fields of rainfall were estimated from METEOSTAT-7 images using conditional Kriging, a method developed for the study. This technique was used to combine satellite-derived rainfields with available rain gauge data and use the resulting fields to adjust radar estimates of rainfall. To make Kriging over such large data sets manageable a technique was devised for Kriging large systems using the Fast Fourier Transform. Although innovative and successful, it was slow and there was a need to accelerate computation to make it more efficient in real time.

There have been recent advances in the mathematics of merging rainfields. Most exciting of these is the combination of Kriging and Kalman Filtering devised by Todini (2001) and adapted by Sinclair and Pegram (2004b) - this technique has been extended in this study. Seed and Pegram (2001) looked at techniques of 3D Kriging for infilling "poor" data in radar images with remarkable success; this start led to a thorough treatment of ground-clutter infilling and image repair techniques reported here. Ehret (2002) (in a PhD co-supervised and examined by the project leader) successfully applied the conditional Kriging merging method to compute rainfields for flood nowcasting in Germany. These inter-associated techniques found application and were successfully refined in this project.

In this report, an algorithm for merging satellite, radar and rain gauge daily totals is proposed, which draws on the combined efforts of the research team. The algorithm uses the decontaminated radar rainfall images extrapolated to ground level using a new technique called Cascade Kriging and the new Accumulation algorithm based on field advection measurements which gives smooth realistic daily totals of rainfall measured by radar and satellite. It combines these with the gauge data using the Conditional Merging technique as its basis and replaces the Explained Variance based algorithm currently used to produce the SIMAR product.

Finally, the stated aims of this research project were to:

1. Combine satellite, radar and rain gauge estimates of rainfall into a meaningful composite to provide an improved optimum rainfield to users via the Internet.
2. Improve the estimation methods initiated in SIMAR by using MSG (METEOSAT Second Generation) and novel methods of field merging, notably Kriging with the Kalman Filter and conditional Geostatistical merging.
3. Continue to seek numerical techniques to reduce the huge computational load and accelerate convergence in producing large maps.
4. Collect, combine, archive and disseminate the raw data and the product fields through the efficient use of the IT division of NDMC/DWAF:PSU (National Disaster Management Centre/Department of Water Affairs & Forestry: Public Safety Unit).

The first three have been achieved successfully; the fourth was not completed because of difficulties in establishing communication links, as well as defining responsibility for maintaining the information flow.

2 Quality control and infilling of radar data

Rain gauges are the traditional tool used for the recording of rainfall and are often regarded as the "true", or reference, rainfall estimates at ground level. Rain gauges have the advantage of being relatively cheap and easy to maintain and also of providing a direct estimate of the accumulated rainfall at a particular point. However there are various disadvantages associated with rain gauges. They tend to underestimate during heavy rainfall periods (Wilson and Brandes, 1979) and by providing a point estimate they can fail to capture the spatial variability of rainfall. Random errors are also present which can be demonstrated by the difference in measurement obtained between closely situated rain gauges (Ciach, 2002).

Weather radars overcome some of the disadvantages associated with rain gauges. Currently in South Africa the weather radar network of eleven C-band radars and one S-band radar provides images of the instantaneous reflectivity values at approximately five-minute intervals and at a one-kilometre horizontal resolution supplied in Cartesian co-ordinates on Constant Altitude Plan Position Indicators (CAPPIs) at 1km intervals above ground level. This type of data provides a detailed spatial representation of the rainfield in real time and over a large area. Figure 2.1 illustrates a typical instantaneous radar reflectivity image taken from the Bethlehem (South Africa) S-band weather radar where the white portions indicate no rainfall, grey portions indicate areas out of the weather radar's range (and also where there are no data available), and the black portions mark ground clutter locations.

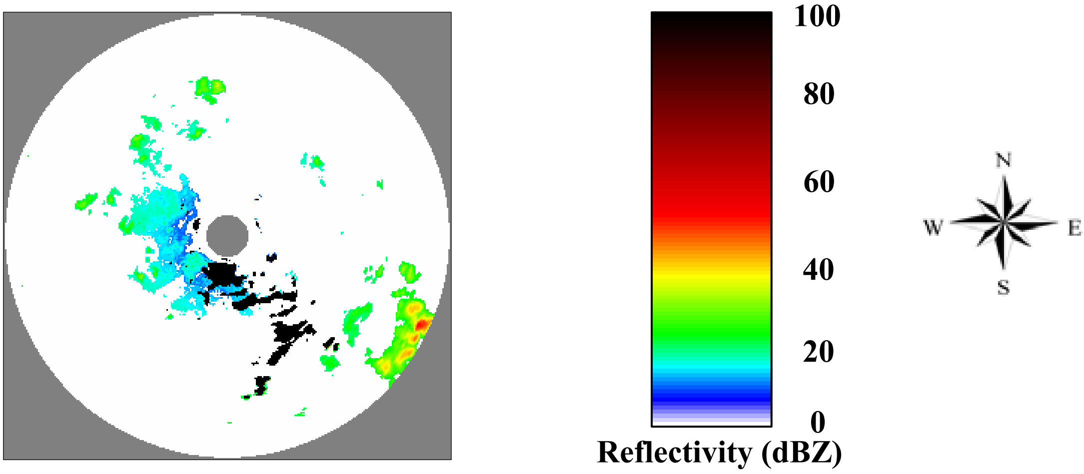

Figure 2.1 Typical radar reflectivity image taken from the Bethlehem (South Africa) weather radar on the 25 February 2003. The grey portions of the image indicate where no data are available and the black portions indicate ground clutter. The image is 300km square.

Rainfields estimated from weather radars experience various data quality problems such as ground clutter, anomalous propagation and beam blocking, to name a few (Terblanche et al, 2001). Another disadvantage of weather radars is that they provide an indirect measurement of precipitation intensity, so the returned power of measured reflectivity values has be converted to rainrate by an appropriate transformation, such as the Marshall Palmer formula (Marshall and Palmer, 1948) given by Eqn. (2.1), where Z is the reflectivity (dBZ) and R is the rainrate (mm/hr).

$$Z = 200 \cdot R^{1.6} \qquad (2.1)$$

The data in the CAPPIs are only available from one-kilometre above ground level and there are regions within the volume scan where the rainrate is unknown. In applications, such as disaster management, hydrology and agriculture, the rainfall estimates at ground level are of more interest and importance than the measurements aloft, which are unlikely to be an accurate indication of the rainfall at ground level (Jordan et al, 2000) since precipitation tends to be affected by a variety of atmospheric phenomena before reaching the earth's surface.

Attempts have been made to provide an improved rainfall estimate at ground level by taking into account atmospheric factors such as wind drift (Lack and Fox, 2004). A different approach, taken by Franco et al (2002), was to compute a Mean Apparent Vertical Profile of Reflectivity and by extending the profile shape obtain an estimate of the rainfall at ground level. Currently in South Africa the real-time radar rainfall estimates at ground level are simply taken as the maximum vertical reflectivity in a column. However this has been shown to provide an overestimation of the rainfall at ground level (Visser, 2003), even for convective precipitation.

Seed and Pegram (2001) proposed using 3D Kriging to estimate the rainfall at ground level and that paper is the source of some of the ideas in this report; which consists of the following sections. Firstly the details of the rainfall classification algorithm are outlined and the characteristics of the classified rainfall described. A bright band correction algorithm based on local gradients is then introduced and tested. The next part describes the empirical computation of the robust semivariogram from radar reflectivity data and outlines the use and validation of climatological semivariograms based on rainfall type. Kriging theory is then briefly discussed, where the methods designed to ensure computational efficiency and stability are outlined. 3D Universal and Ordinary Kriging follow in their application to Cascade Kriging, a novel idea for extrapolation. Finally a statistical comparison between the temporally accumulated radar estimates and the Block Kriged rain guage estimates is carried out over matching areas, for selected rainfall events, to determine the quality of the rainfall estimates at ground level.

The data from 1996 from the Liebenbergsvlei catchment near Bethlehem were selected as the test-set for this study. This was because there was a dense network of 45 tipping bucket rain gauges placed in a grid, evenly spaced at approximately 10km throughout the catchment at that time (since decommissioned). The Bethelehem S-band radar is also situated in close proximity to the catchment. Both of these factors are important in ensuring accurate estimates of rainfall accumulation values. The geographic position of the area and layout of the tipping bucket rain gauges are illustrated in Figure 2.2.

2.1 Rainfall Classification

The initial step, of the operational algorithm for infilling the missing or inaccurate data on a CAPPI, requires a simple classification and thresholding of the instantaneous radar volume scan data, pixel by pixel. A variety of rainfall classification algorithms has been suggested that involve subdividing the rainfall into three zones – convective, intermediate and stratiform rain. An example of a current rainfall classification algorithm employed by The Group of Applied Research on Hydrometeorology (GRAHI) (Sempere-Torres et al, 2000) also subdivides the rainfield into the same three separate groupings; the convective rain is identified by applying a threshold value of 40 dBZ; the stratiform rain is identified when a bright band is present; the rest of the rainfield classified as transitional. Another rainfall classification algorithm developed by Mittermaier (1999) was also designed along similar principles.

The above-mentioned algorithms are not suitable for the purpose of rainfall estimation at ground level over large areas. This is due to the fact that the data at the 2km CAPPI level (which with a base scan at 1.5º above horizontal has a range limited to 73km) plays a pivotal role in classifying the rainfall as stratiform or intermediate. This means that only a small proportion of the total volume scan data can be classified. Another drawback of the existing algorithms in real time applications is that they are computationally expensive, especially for large data sets.

The algorithm developed here, with assistance of meteorologists from METSYS (Meteorological Systems and Technology - a research division of the South African Weather Service), classifies the rainfall into only two groups: convective and stratiform and works as follows. The classification is rather like Steiner's (1995) algorithm and is done pixel by pixel throughout the volume scan data. The initial step of the algorithm works on a simple threshold where all reflectivity values less than or equal to 18dBZ are set to zero due to the fact that 18dBZ approximates to 0.50 mm/hr, a rate that can be considered negligible. The remainder of the reflectivity values above zero are then classified by the simple criterion: if the reflectivity value is 35dBZ or above, the rainfall is classified as convective; the reflectivity values in the range 18dBZ to 35dBZ are classified as stratiform rainfall.

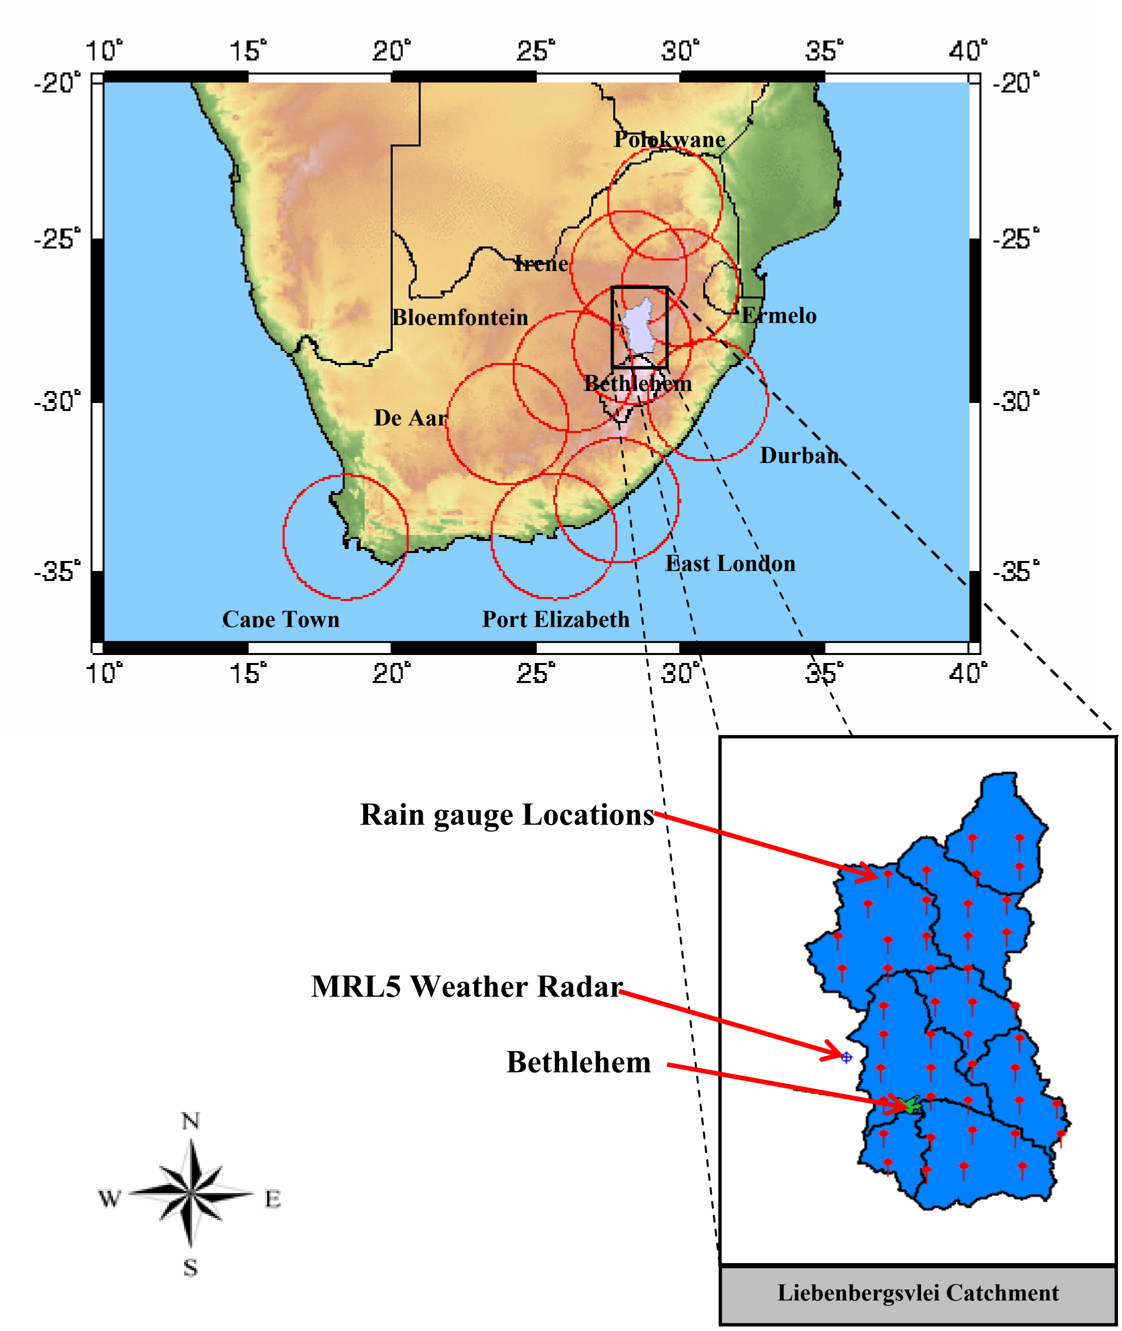

Figure 2.2 Map of South Africa illustrating the weather radar network coverage and the geographic location of the Liebenbergsvlei catchment. An enlarged image of the catchment shows the approximate positions of the rain gauges and the MRL5 weather radar.

A summary of the classification criteria is given in Table 2.1 and two examples of classified instantaneous images both 4km above ground level are illustrated in Figure 2.3. Two examples of vertical cross sections through the volume scan data are illustrated in Figure 2.4.

Table 2.1: Rainfall classification and threshold criteria to separate radar volume scan data into stratiform and convective rainfall domains.

CRITERIA	CLASSIFICATION
Pixel ≤ 18 dBZ	No Rainfall: Pixel set to 0 dBZ
18 dBZ < Pixel < 35 dBZ	Stratiform Rain
Pixel ≥ 35 dBZ	Convective Rain

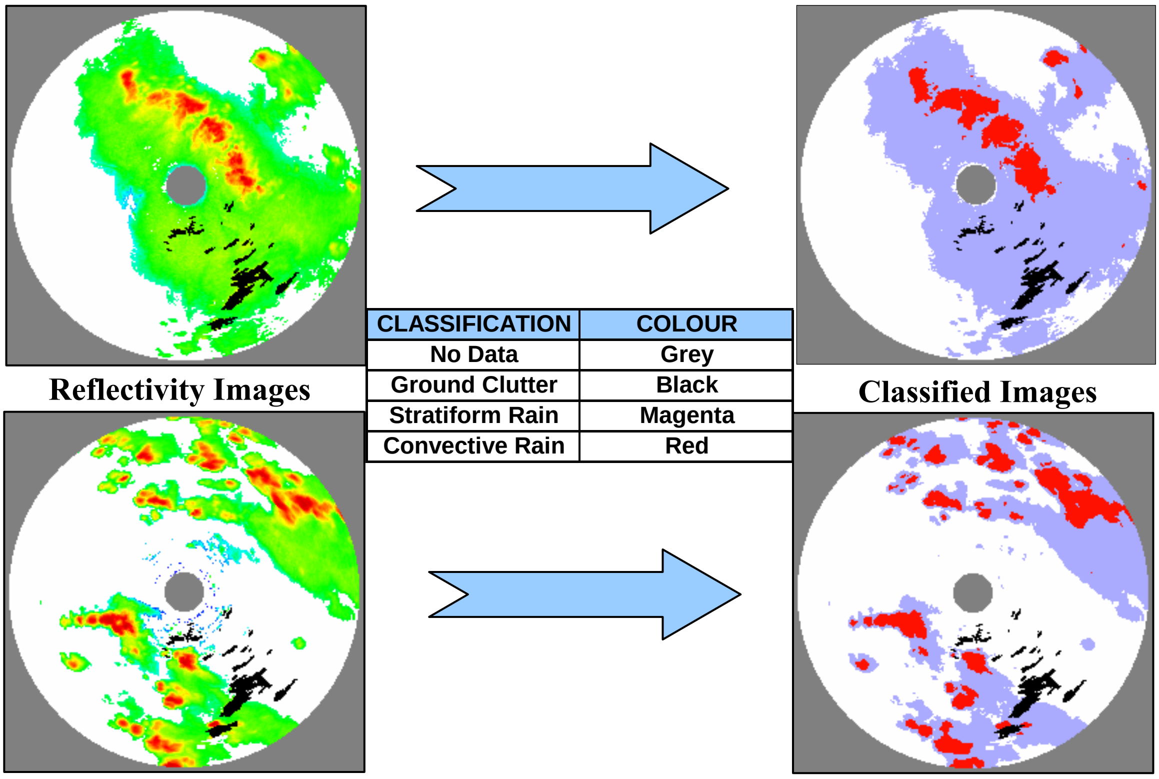

Figure 2.3 Radar reflectivity images from the Bethlehem weather radar (30 December 2001 and 8 January 2002) and the corresponding classified images. All images are 300km square.

The classification algorithms proposed by Steiner (1995) and Mittermaier (1999) work on the basis that the classification of convective rainfall on the 4km CAPPI level is applied to all levels below and that the 2km CAPPI level is used to classify stratiform rain. However the rainrates at a large number of pixels outside the range of the 4km and 2km CAPPI level need to be estimated and large portions of the image on the 4km and 2km CAPPI level may also be marked as ground clutter and therefore be contaminated. As indicated in Figure 2.4, where the convective rainfall approaches ground level (below 4km in altitude) the convective rain zones may increase substantially in the horizontal direction, which would then result in a scenario where rainfall of high intensity (e.g. 54 dBZ ≈ 84 mm/hr) is miss-classified as stratiform rainfall. For the above reasons a pixel by pixel classification approach was chosen throughout the CAPPI volume scan due to its greater suitability for the purpose of rainfall estimation at ground level.

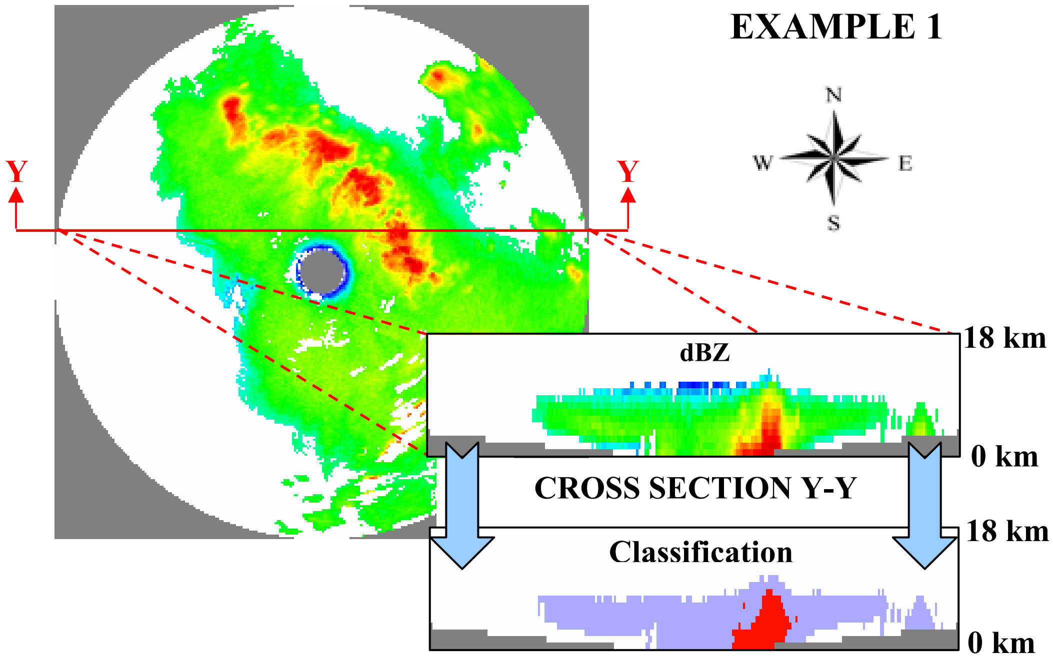

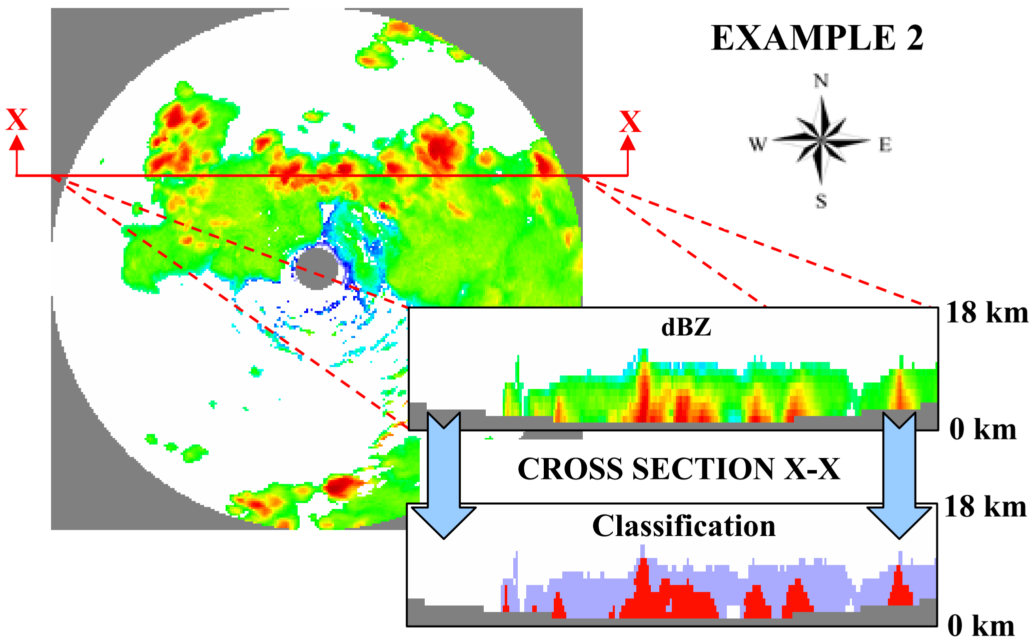

Figure 2.4 Cross sectional profile through two instantaneous reflectivity images. Indicated in the plan is the CAPPI level 4km above ground level and the position of the vertical cross section in the radar volume scan CAPPI. Example 1 is from the Bethlehem weather radar, 30 December 2001, and Example 2 from the Bethlehem weather radar, 24 January 2002. All images are 300km across. Note the effectiveness of the classification algorithm.

The main advantage of this algorithm is its simplicity, in that for a real time application it has little impact on the overall computation time; this is especially true for large data sets such as those returned from weather radar volume scans. Secondly the algorithm is not restricted to the 2km CAPPI level which is frequently the level of the melting layer in a South African context; thirdly, a far greater spatial range can be classified than that limited to 73km at the 2km level.

2.1.1 Characteristics of Classified Rainfall

The rainfall classification algorithm was tested on numerous images ranging over five different years (1995, 1996, 2000, 2001 and 2002) from two weather radars in South Africa: Bethlehem (mixed convective and stratiform rain, 1800m altitude, sub-tropical savannah) and Durban (mostly orographic warm rain, coastal, bush). The mean and standard deviation of the reflectivity values of selected portions (areas in size of approximately 600km^2 and greater) of the classified sections, which were respectively dominated by either convective or stratiform precipitation, were then computed and plotted as indicated in Figure 2.5. As shown, the statistical characteristics of the stratiform and convective rainfall are distinctly different. The stratiform rain tends to have a low and reasonably constant standard deviation of approximately 3dB and a mean value ranging from 18dBZ to 29dBZ. In the convective rainfall the standard deviation showed a far greater variation ranging from 2dB to 8dB, exhibiting a strong, steep linear trend where a higher mean value corresponds to higher standard deviation, to be expected from data sampled from a field which is approximately log-normally distributed (Pegram and Clothier, 2001).

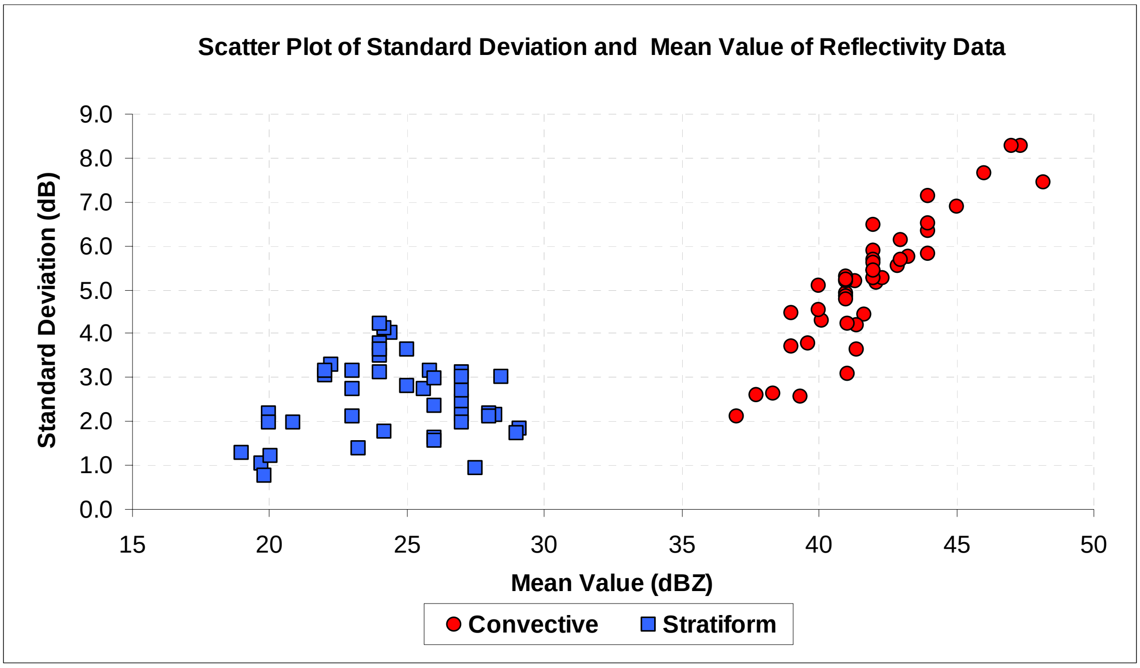

Figure 2.5 Scatter plot of the mean and standard deviation of rainfall zones classified as convective or stratiform rain indicating the different statistical properties of each rainfall type. The number of 1km pixels used in defining each point is approximately 600 or greater, from data over a five year period collected from the Bethlehem and Durban weather radars.

The vertical characteristics of the classified stratiform and convective rain were also examined by computing the mean vertical profile of reflectivity from selected images from the Durban and Bethlehem weather radars over the same five years as mentioned above. As shown in Figure 2.6 the stratiform rain generally has a low average vertical height, limited to 8km above ground level. The convective rainfall has considerable vertical extent ranging up to 14km above ground level and a bigger range in mean reflectivity values. Both the stratiform and convective profiles show a constant increase in rainfall intensity as ground level is approached, with the exception that in the stratiform profile, evidence of a bright band at the 2km level indicates that the stratiform rain has been classified correctly. The variability of the reflectivity in both are characterised by a standard deviation of 3.5dB or less at any level.

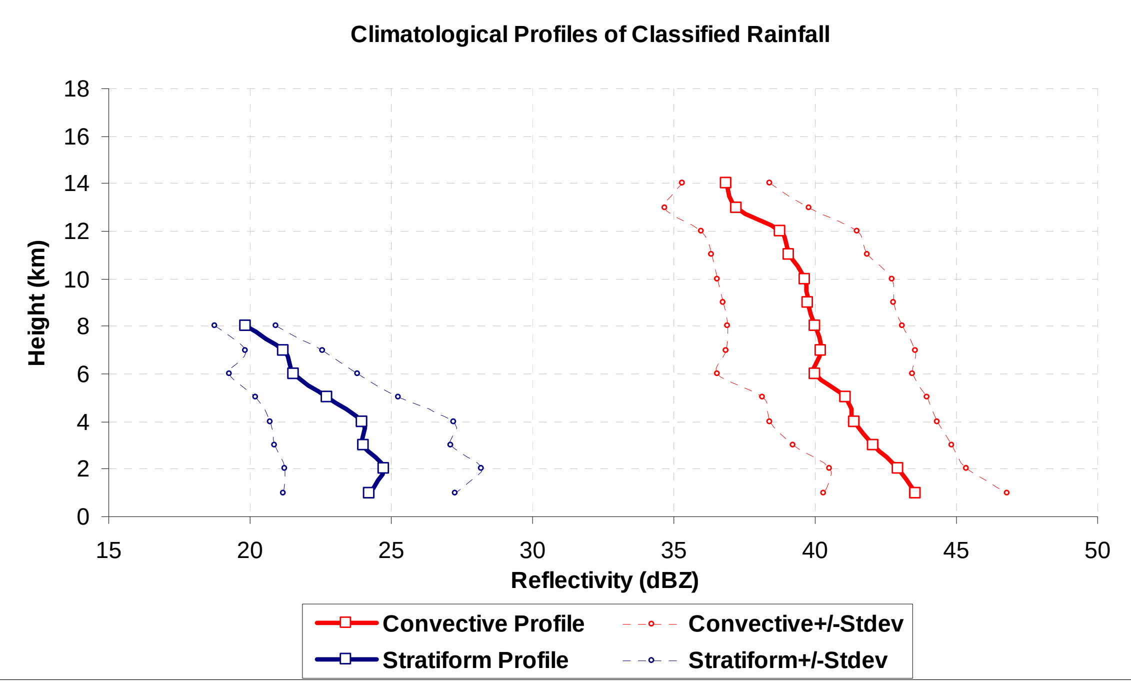

Figure 2.6 Average vertical profile of reflectivity for classified stratiform and convective rainfall. Profiles are computed from 20 different images ranging over a 5 year period from the Bethlehem and Durban weather radars.

2.2 Bright Band Correction

As snow and ice crystals drop through the 0^{o}C isotherm they begin to melt and are surrounded by a thin layer of water. The result is that, from the radar's point of view, the melting snow and ice crystals resemble large blobs of water and the reflectivity values for the melting level are enhanced greatly, resulting in an overestimation of rain (Sanchez-Diezma et al, 2000). This level is called the bright band. At Bethlehem, South Africa the bright band generally occurs 2km above ground level (Mittermaier, 1999) as can be seen by the stratiform climatological profile of reflectivity in Figure 2.6. This bright band effect results in a vertical profile of reflectivity that resembles Figure 2.7.

Predominantly two different approaches have been proposed, as described by Mittermaier (2003), in order to correct the problem of bright band: (1) A vertical profile is derived based on radar volume scan data and is used to correct the values and provide an estimate of the rainfall rate at ground level (e.g. Sanchez-Diezma et al, 2000); (2) A standard high resolution profile is used for correction with the height of the top of the bright band set by the height of the 0^{0}C wet bulb temperature derived from a model.

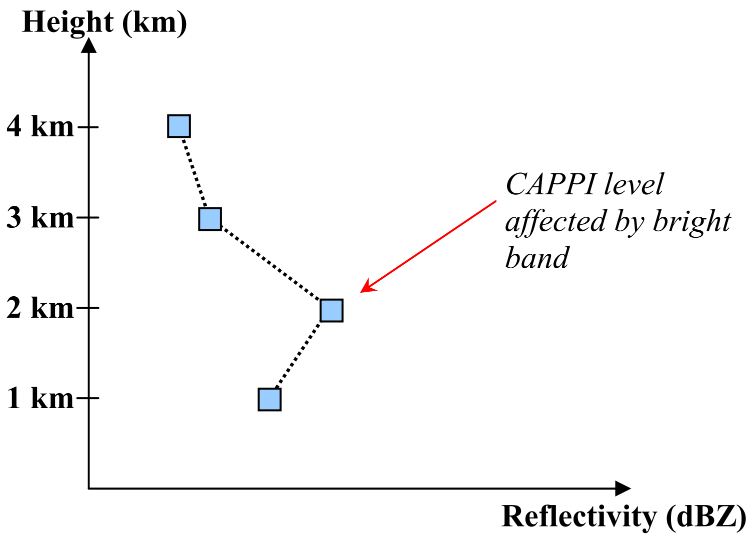

Figure 2.7 Idealised stratiform climatological profile of reflectivity with bright band affecting the reflectivity values situated on the 2km CAPPI level, indicated by the sharp increase in reflectivity at the 2km level.

2.2.1 Effect of Bright Band on Ground Level Estimates

The presence of a bright band has an adverse effect on the estimates at ground level when applying an extrapolation algorithm to volume scan radar reflectivity data. The reflectivity values at the 2km CAPPI level, which are on average higher than the values contained on the 1km CAPPI, create a trend of decreasing reflectivity values as ground level is approached. When extrapolating to ground level and pixels are selected from the 1km and 2km CAPPI levels, the trend of decreasing reflectivity is extrapolated to ground level resulting in lower estimates of stratiform rainfall at ground level than should be expected. There is therefore a need to adjust the reflectivity values at the 2km CAPPI level to provide an improved estimate of the rainfall at ground level.

2.2.2 Bright Band Correction Algorithm

The first step of the algorithm is to compute the wetted area ratio (WAR) of stratiform rainfall at the 3km, 2km and 1km CAPPI levels. If the WAR of stratiform rainfall is less than 10% at any of the CAPPI levels, then no correction takes place due to there being insufficient rainfall to determine if a bright band is present; if the WAR of stratiform rainfall is greater than 10% the 2km CAPPI level is then examined to detect evidence of a bright band.

The existence of the bright band is determined by comparing the temporally weighted mean values of reflectivity ($\tilde{x}_1$, $\tilde{x}_2$ and $\tilde{x}_3$) at the 1km, 2km and 3km CAPPI levels; the calculation is limited to the reflectivity values classified as stratiform rainfall on the image. If the computed temporally weighted mean values of the stratiform rainfall at each level are $\tilde{x}_1 > \tilde{x}_3$ and $\tilde{x}_2 \geq \tilde{x}_1$ then a bright band is considered to be present. Some memory of the mean value from previous time steps is retained by temporal weighting as indicated by Eqn. (2.2):

$$\tilde{x}_T = \lambda \cdot \bar{x}_T + (1-\lambda) \cdot \tilde{x}_{T-1} \qquad (\lambda = 0.50) \qquad (2.2)$$

where $\tilde{x}_T$ is the temporally weighted mean, λ represents the weighting value which is set to 0.50 (a value chosen to smooth the means, but not excessively), and $\bar{x}_T$ is the spatial mean value of reflectivity at a given level computed from the stratiform data at time T.

Throughout the rain event a second temporally weighted mean is also computed, as indicated by Eqn. (2.3), where the weighting value δ is set to 0.05; this has the result of producing a slowly adapting mean value w_T. If the WAR of the stratiform rainfall is below 10% then a break occurs and the $\tilde{x}_T$ value is now computed by Eqn. (2.4) where the $\tilde{x}_T$ is based on the final w_T value computed before the break which had the value w_{TB}:

$$w_T = \delta \cdot \bar{x}_T + (1-\delta) \cdot w_{T-1} \qquad (\delta = 0.05) \qquad (2.3)$$

$$\tilde{x}_T = \delta \cdot w_{TB} + (1-\delta) \cdot \tilde{x}_{T-1} \qquad (\delta = 0.05) \qquad (2.4)$$

an example of this is illustrated in Figure 2.8 where Eqn. (2.2), (2.3) and (2.4) were used to compute the mean values of reflectivity ($\tilde{x}_T$) by geometric weighting. The reflectivity data are from the Bethlehem weather radar, 25 January 1996. As indicated in Figure 2.8, in three periods, the WAR of stratiform rainfall is less than 10%; in these periods $\tilde{x}_T$ is computed from Eqn. (4), not (2), so that toward the end of the period, it tends to w_{TB}, which can be thought of as a locally adaptive historical mean value of reflectivity.

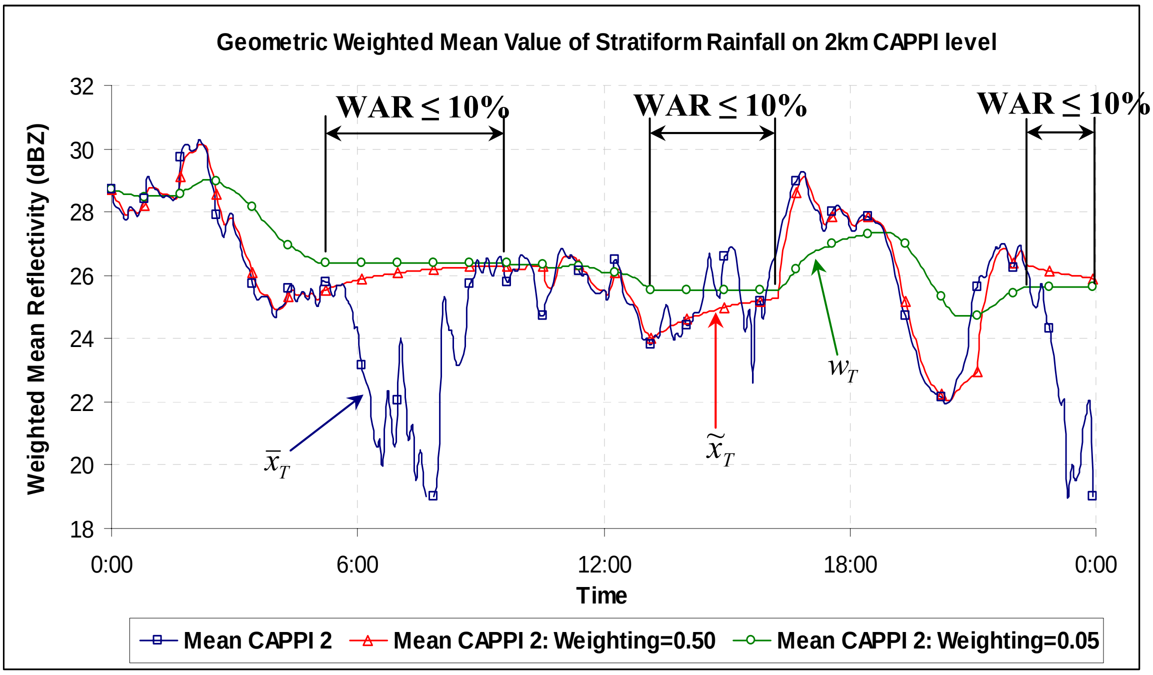

Figure 2.8 Weighted mean value of stratiform rain on 2km CAPPI level as computed by geometric weighting, Eqns. (2.2), (2.3) and (2.4). The data is taken from the Bethlehem weather radar (25 January 1996) over a 24 hour period.

To correct the bright band it is also necessary to compute the standard deviation (σ) of the stratiform rainfall on the 1km CAPPI level. Some memory of the σ value from the previous time step is retained by weighting as in Eqn. (2.2); to compute the weighted standard deviation, Eqn. (2.5) is used.

$$\tilde{\sigma}_T = \lambda \cdot \sigma_T + (1-\lambda) \cdot \tilde{\sigma}_{T-1} \qquad (\lambda = 0.50) \qquad (2.5)$$

where $\tilde{\sigma}_T$ is the weighted standard deviation, λ represents the weighting value which is set to 0.50 and σ_T is the standard deviation of reflectivity at time T. Once again if the WAR of the stratiform rainfall is less than 10% a break occurs and the $\tilde{\sigma}_T$ value is now computed by Eqn. (2.7) where the $\tilde{\sigma}_T$ is based on final s_T value computed by Eqn. (2.6) before the break which has the value s_{TB}:

$$s_T = \delta \cdot \sigma_T + (1-\delta) \cdot \sigma_{T-1} \qquad (\delta = 0.05) \qquad (2.6)$$

$$\tilde{\sigma}_T = \delta \cdot s_{TB} + (1-\delta) \cdot \sigma_{T-1} \qquad (\delta = 0.05) \qquad (2.7)$$

If a bright band is identified, all the values classified as stratiform rainfall are then evaluated on an individual pixel by pixel basis within a vertical column, to decide if the pixel at the 2km CAPPI is affected by the bright band. For each stratiform pixel location on the 2km CAPPI level the gradient ($\Delta_{critical} = \Delta dBZ/\Delta height$) is calculated between the pixel value (directly above it) on the 3km CAPPI and the value equal to $\tilde{x}_1 + \tilde{\sigma}_1$ on the 1km CAPPI level, as illustrated in Figure 2.9; the gradient is computed from Eqn. (2.8). The gradient ($\Delta_{observed}$) is also calculated between the pixel on the 3 km and 2km CAPPI level; as indicated by Eqn (2.9):

$$\Delta_{critical} = \frac{\Delta dBZ}{\Delta Height} = \frac{dBZ_3 - (\tilde{x}_1 + \tilde{\sigma}_1)}{2} \qquad (2.8)$$

$$\Delta_{observed} = \frac{\Delta dBZ}{\Delta Height} = \frac{dBZ_3 - dBZ_2}{1} \qquad (2.9)$$

If $|\Delta_{observed}| > |\Delta_{critical}|$ then the pixel on the 2km CAPPI level is classified as a bright band pixel and adjusted by altering the value of that pixel as shown in Eqn. (2.10):

$$New\, dBZ_2\ Pixel = \frac{dBZ_3 + \tilde{x}_1}{2} \qquad (2.10)$$

Figure 2.9 provides a schematic of a scenario where a pixel on the 2km CAPPI level is identified as being affected by bright band. The correction of the pixel is then also indicated as the new value is set to the midpoint of the pixel on the 3km CAPPI level and the mean value ($\tilde{x}_1$) of the reflectivity data on the 1km CAPPI level.

2.2.3 Testing the Bright Band Correction Procedure

To test the effectiveness of the proposed bright band correction procedure a period of rainfall from the Bethlehem weather radar was selected (17 December 1995) that consisted predominantly of stratiform rainfall and showed clear evidence of a bright band. Figure 2.10 illustrates the computed mean values $\tilde{x}_T$ at CAPPI levels 1, 2 and 3 throughout the day. The mean value of CAPPI 2 (green line with triangle markers) is consistently higher throughout the day than the mean values on the 1km and 3km CAPPI levels, indicating a definite bright band.

The bright band correction procedure described in the previous sub-section was applied to each 5-minute instantaneous image throughout the 24 hour period. The mean value of the adjusted reflectivity data on the 2km CAPPI level was then recomputed. Figure 2.10 illustrates the adjusted mean value for the 2km CAPPI level, which now lies at the approximate midpoint of the mean value for the 1km and 3km CAPPI level.

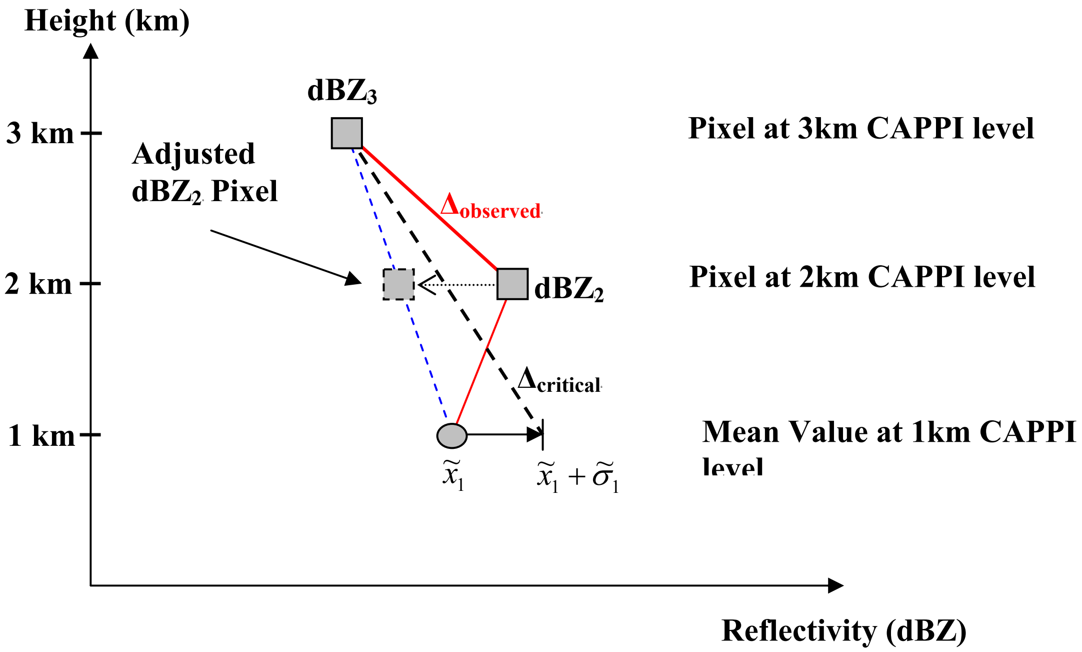

Figure 2.9 Adjustment of a stratiform pixel value on the 2km CAPPI level that is identified as being affected by bright band. The bright band pixel is adjusted so as to sit at the midpoint of the reflectivity value situated on the 3km CAPPI level and the mean value of the 1km CAPPI level.

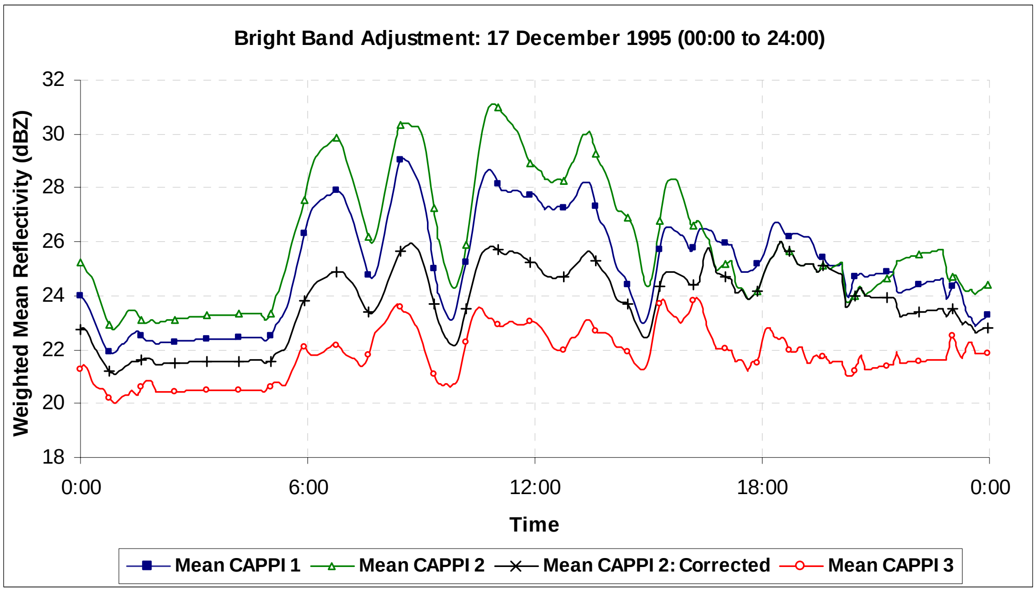

Figure 2.10 Bright band adjustment of the 2km CAPPI levels for the 17 December 1995 (00:00 to 24:00). The green line with triangle markers indicates the 2km CAPPI level affected by bright band and the black line with plus sign markers indicates the mean value of the corrected data on the 2km CAPPI level.

Figure 2.11 provides an example of an instantaneous reflectivity image that is affected by bright band. The northern half is only shown due to ground clutter contamination in the south of the image. The top image of Figure 2.11 is contaminated by bright band with the right hand side image indicating marked pixels that are affected by bright band. The bottom image of Figure 2.11 displays the corrected image.

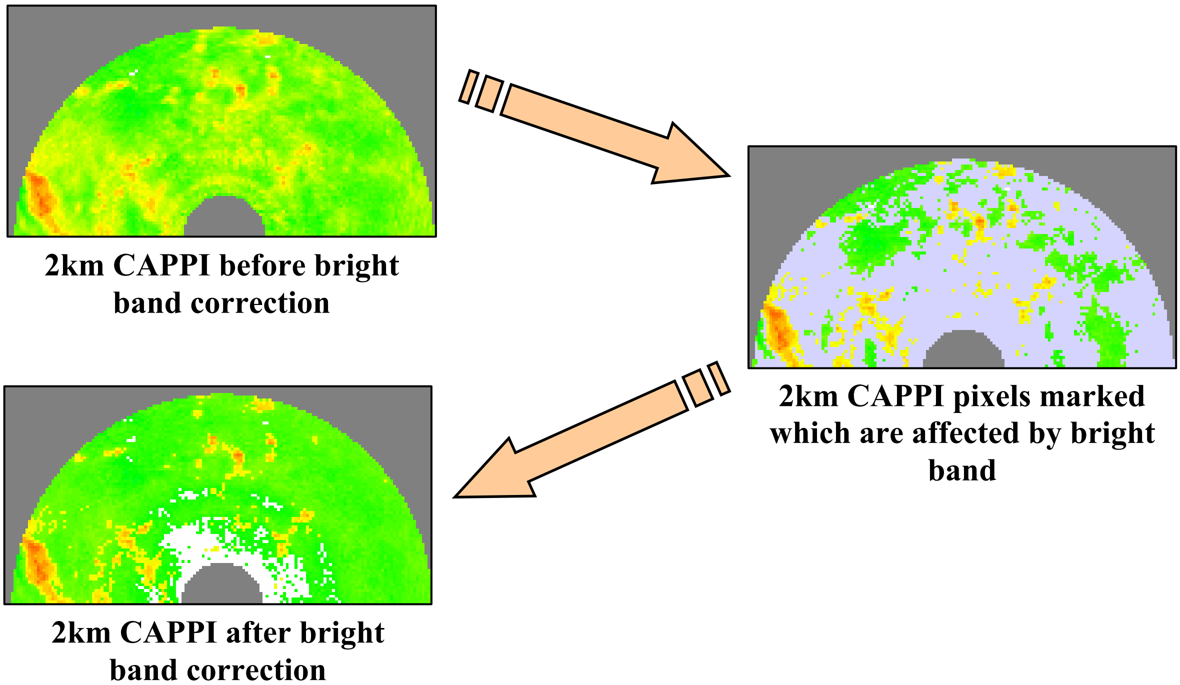

Figure 2.11 An example of bright band adjustment of an instantaneous image for the 2km CAPPI level for the 17 December 1995. The images have the dimensions of width = 150km and height = 75km.

To test the effectiveness of the bright band adjustment algorithm, a comparison was carried out between Block Kriged rain gauge estimates and the radar accumulation estimates before and after the bright band correction procedure. The comparison was carried out on a 24 hour rain event from the Bethlehem weather radar from the 17 December 1995 with a strong bright band evident, as indicated in Figure 2.11. The sum of square of errors (SSE) between the rain gauge and radar accumulations for a 24 hour period were computed and the r^2 values also estimated before and after the bright band correction. There was a notable improvement in the radar estimates as the radar accumulation values more closely matched the rain gauge estimates after the bright band was corrected. This is demonstrated by the lower SSE and higher r^2 value returned for radar estimates where the bright band correction was implemented as summarised and shown in Table 2.2.

Table 2.2: Sum of square of errors (SSE) and r^2 value between radar and rain gauge estimates before and after bright band correction.

	Before Bright Band Correction	**After Bright Band Correction**
SSE	16 601	13 161
r^2	0.34	0.43

The proposed bright band correction procedure provides an effective means of adjusting the 2km CAPPI level reflectivity values when affected by bright band so as to provide an improved estimate of the rainfall at ground level. The bright band correction algorithm is computationally inexpensive and therefore highly suitable for real time applications.

2.3 Semivariogram Estimation

One of the advantages of Kriging compared to other interpolation/extrapolation techniques is that the basis function is determined by the actual data set and takes into account the spatial structure of the data. The basis function in Kriging is determined by the computation and fitting of an appropriate model to the computed semivariogram. The variogram is defined as the expectation of the square of the differences of the field variables separated by a specified distance, as given by Eqn. (2.11):

$$2\cdot\gamma(x,h)=E\left\{\left[Z(x)-Z(x+h)\right]^2\right\} \tag{2.11}$$

where the semivariogram is defined as γ(x, h) (Journel and Huigbregts, 1978: 11). The semivariogram provides a way of measuring the spatial dependence that exists amongst the variables in a stationary random field. The Kriging computation can be carried out using either a covariance or a semivariogram function. However in this application the semivariogram is used in preference due to its more robust properties as outlined by Cressie (1993: 70-73).

2.3.1 Empirical Computation of the Semivariogram

The semivariogram is most commonly computed by the Classical Variogram estimator as was proposed by Matheron (1962) and is given by Eqn. (2.12):

$$2\cdot\hat{\gamma}(h)=\frac{1}{|N(h)|}\sum_{N(h)}\left[Z(s_i)-Z(s_j)\right]^2 \tag{2.12}$$

where $\hat{\gamma}(h)$ is the sample semivariogram at lag *h* (the specified lag distance), $Z(s_i)$ and $Z(s_j)$ are the values of the variables at the specified locations, s_i and s_j which are *h* apart and $N(h)$ is the number of pairs separated by lag *h*.

Unfortunately the classical estimator of Eqn. (2.12) is badly affected by non-typical observations which can be attributed to the $(\cdot)^2$ term in the summand (Cressie, 1993: 40). The effect of non-typical data points can have a dramatic effect on the semivariogram since they are used numerous times in the calculation at different lag intervals. This can result in peaks or shifting of the entire semivariogram upwards (Sabyasachi et al, 1997). It was also noted that the distribution of points at each lag distance (*h*) computed from radar rain fields was highly skewed and did not approximate a normal distribution.

It was thus decided that the Robust Variogram, proposed by Cressie and Hawkins (1980) was the more appropriate model to use and is given by Eqn. (2.13):

$$2\cdot\bar{\gamma}(h)=\left\{\frac{1}{|N(h)|}\sum_{N(h)}\left|Z(s_i)-Z(s_j)\right|^{1/2}\right\}^4 \Bigg/ \left(0.457+\frac{0.494}{|N(h)|}\right) \tag{2.13}$$

By computing the sum of the absolute difference of the pairs of data points, $Z(s_i)$ and $Z(s_j)$ in the square root domain and then raising the result to the power of four dramatically reduces the effect of uncharacteristic observations. It was found that the skewness of the points at each lag (*h*) was greatly reduced and also that the mean and the median value of the sets of points for each lag interval more closely coincided, both indicating that the distribution of points at each lag (*h*) more closely approximates a normalized distribution.

An example of the difference between the robust and classical estimator is given in Figure 2.12. These two semivariograms were computed from a segment (30km square in size) of an instantaneous reflectivity image from the Bethlehem weather radar (24 January 2002) in the horizontal direction 4km above ground level, previously classified as stratiform rainfall. As can be seen, there are significant differences between the two estimates, especially (and importantly) near the origin. Computing the empirical semivariogram using the classical estimator would result in an underestimation of the "actual" correlation length and overestimation of the "actual" alpha parameter described in Eqn. (2.14) below.

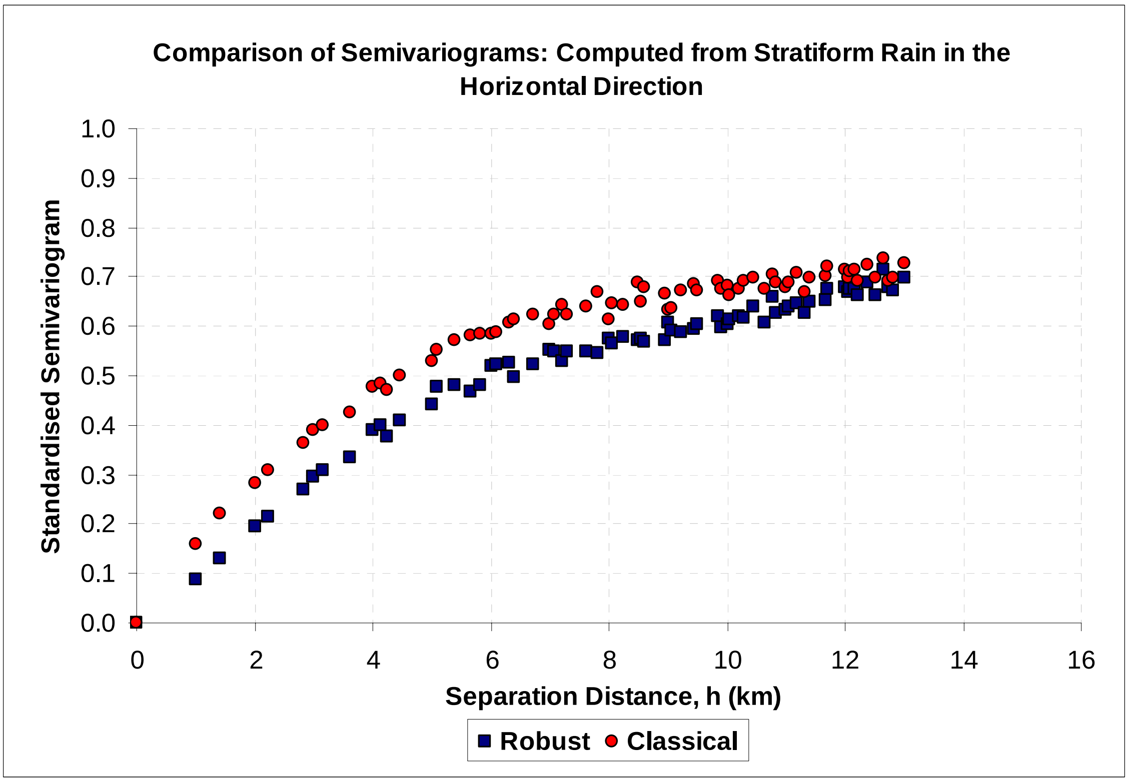

Figure 2.12 Comparison of robust (Eqn. (2.13)) and classical (Eqn. (2.12)) semivariograms computed from a segment (900km^2) of stratiform rainfall in the horizontal direction. Differences between the two semivariogram estimates are evident, especially at close range (h).

2.3.2 Semivariogram Model Fitting

The two-parameter isotropic exponential model was chosen as the model to fit to the empirical semivariogram values, as defined by Eqn. (2.14):

$$g(h) = 1 - \exp[-(h/L)^{\alpha}] \qquad (2.14)$$

where h is the Euclidian distance between data points, L the correlation length and α the shape parameter which lies in the range $0 < \alpha \leq 2$. As recommended by Journel and Huijbrechts (1978: 194), the model was fitted using up to only half the maximum possible lag and only to the bins at each lag distance that contained thirty or more points. The α and L parameter values were found by minimising an objective function U defined in Eqn. (2.15) by a weighted least squares approach:

$$U = \sum_{j=1} N(h_j) \cdot \left[\frac{\hat{\gamma}(h_j)}{\gamma(h_j, b)} - 1 \right]^2 \tag{2.15}$$

where $N(h_j)$ is the number of points in bin j, $\hat{\gamma}(h_j)$ are the empirical semivariogram values and $\gamma(h_j, b)$ are the model values (Cressie, 1993: 97).

2.3.3 Example of Semivariogram Computation

An example of computing a semivariogram is illustrated in Figure 2.13. A 30km square section of radar reflectivity data classified as stratiform rainfall was selected from the Durban weather radar (11 December 2000). Initially the selected reflectivity values were standardized by use of Eqn. (2.16):

$$Z = \frac{x - \mu}{\sigma} \tag{2.16}$$

where Z is the standardised variable, μ the sample mean value of the selected data, x the selected variable and σ the sample standard deviation of the selected data. Standardising the field has the effect of producing a scaled and dimensionless field to ease the fitting of the climatological semivariograms. If the field were not to be scaled by σ, this factor would in any case cancel from both sides of the Kriging matrix equations, Eqn. (2.21). As can be seen in Figure 2.13 the standardized semivariogram increases smoothly with lag h until the sill is reached, whose value is equal to 1.0, the variance of the selected data. The range is also indicated, which is the distance beyond which the variables are considered to be uncorrelated. The correlation length in this case is L = 8.31km, the shape parameter is α = 1.38 and the nugget is zero.

2.3.4 Parameter Fitting to Rainfall Types

Computing the empirical semivariogram values and then applying a model fitting routine to solve for the α and L parameters in Eqn. (2.14) is a computationally burdensome and time-consuming task, especially if this is done each time an unknown reflectivity value is estimated. This problem is exacerbated in the context of real time calculation, when the images are collected every 5 minutes. An alternative is therefore needed that is computational efficient yet does not compromise the accuracy of the estimates. Reducing the computational burden of computing the semivariogram by the use of a climatological semivariogram was employed for rain gauge interpolation by Lebel et al. (1987) in order to improve computational efficiency. This idea was explored in the context of radar rainfields.

A set of instantaneous rainfall images was selected from the Bethlehem and Durban weather radars ranging over a period of several years (1995, 2000, 2001 and 2002). The images consisted of a variety of rainfall types ranging from images containing solely stratiform rain to images containing a combination of both stratiform and convective rainfall.

The rainfall classification algorithm of Table 2.1 was applied to the CAPPI level at 4km above ground level. From the images various rectangular portions were selected that contained either predominantly stratiform or convective rainfall. These regions were approximately 1200km^2 (40km by 30km) in area or larger. From the selected regions the Robust semivariogram was computed and the two parameter exponential model (Eqn. (2.14)) fitted.

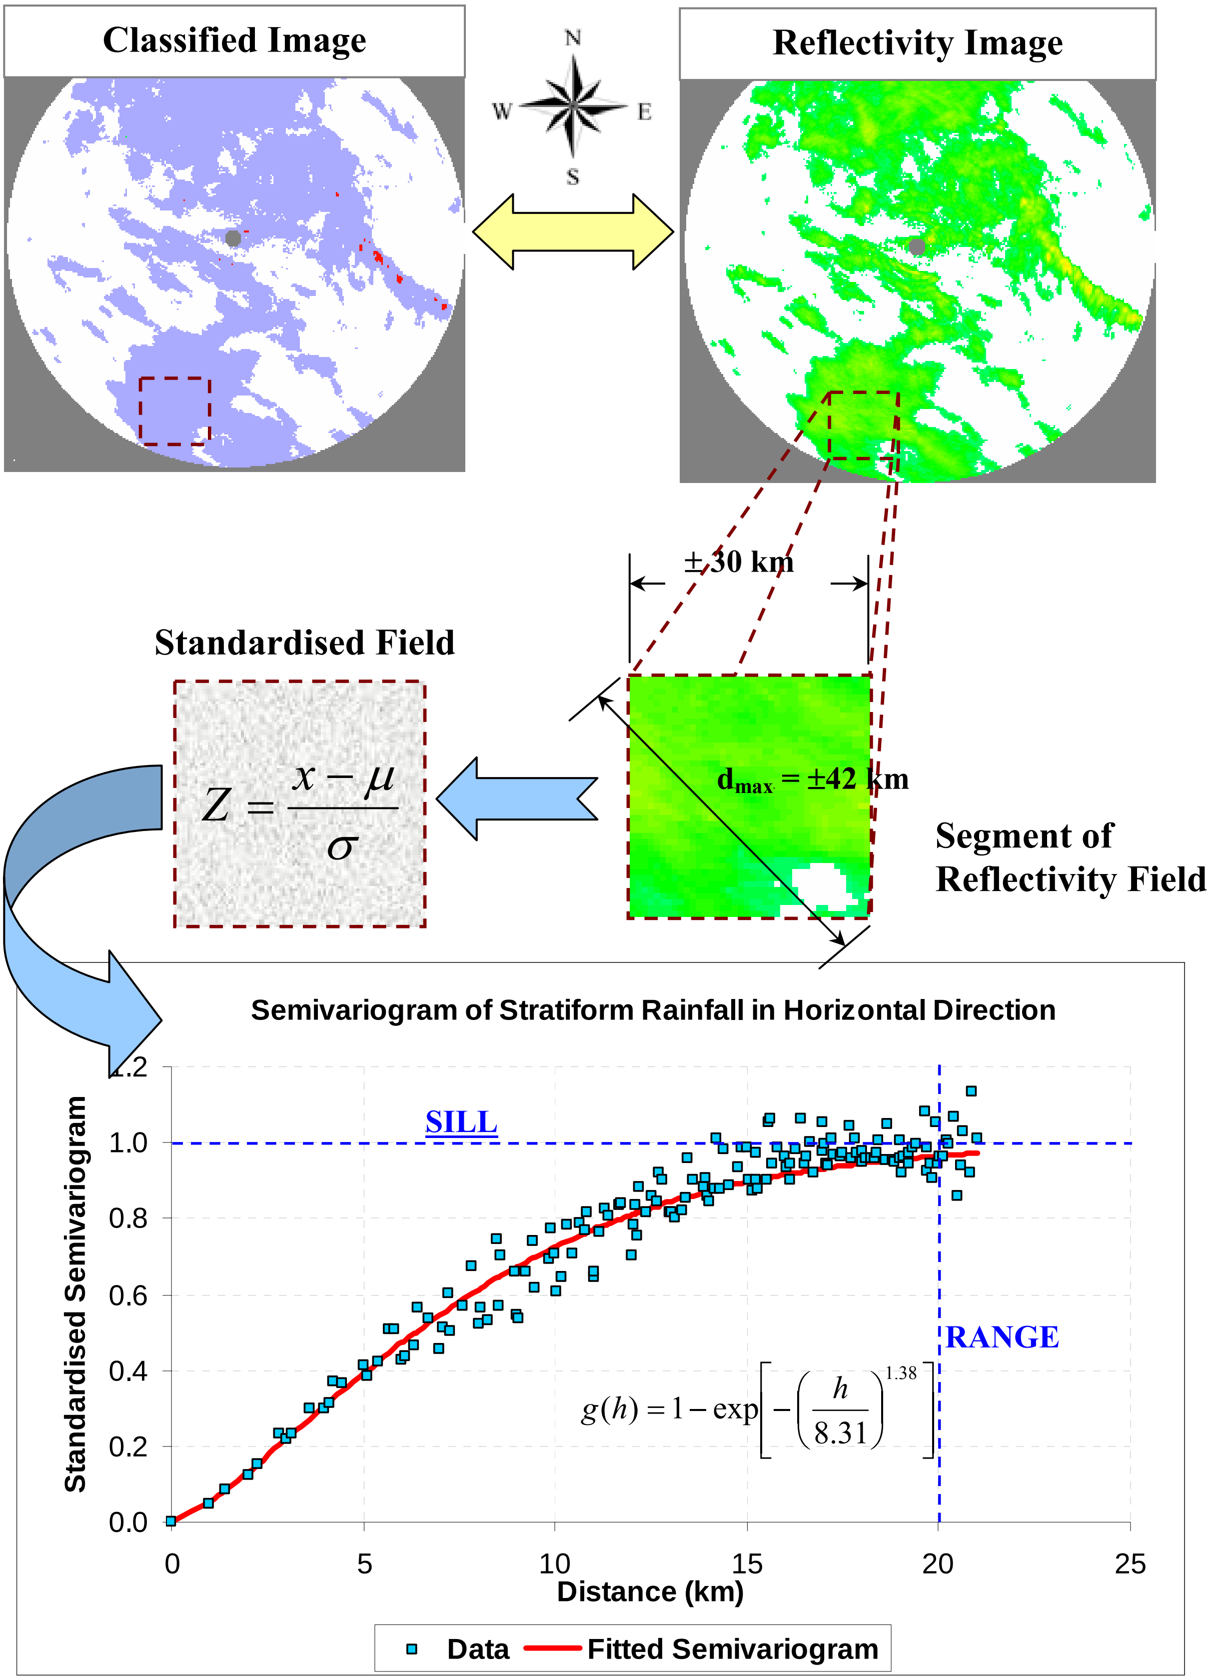

Figure 2.13 Typical semivariogram computed from a selected portion of stratiform rain with the sill and approximate range indicated. The maximum distance (d_{max}) is 42km, *h* is limited to 21km as recommended by Journel and Huijbrechts (1978: 194).

This was done separately in both the horizontal and vertical direction. The two parameters, α and L estimated from each field, were then plotted against one another in a scatter plot. This is shown in Figure 2.14 for the horizontal direction and Figure 2.15 for the vertical direction.

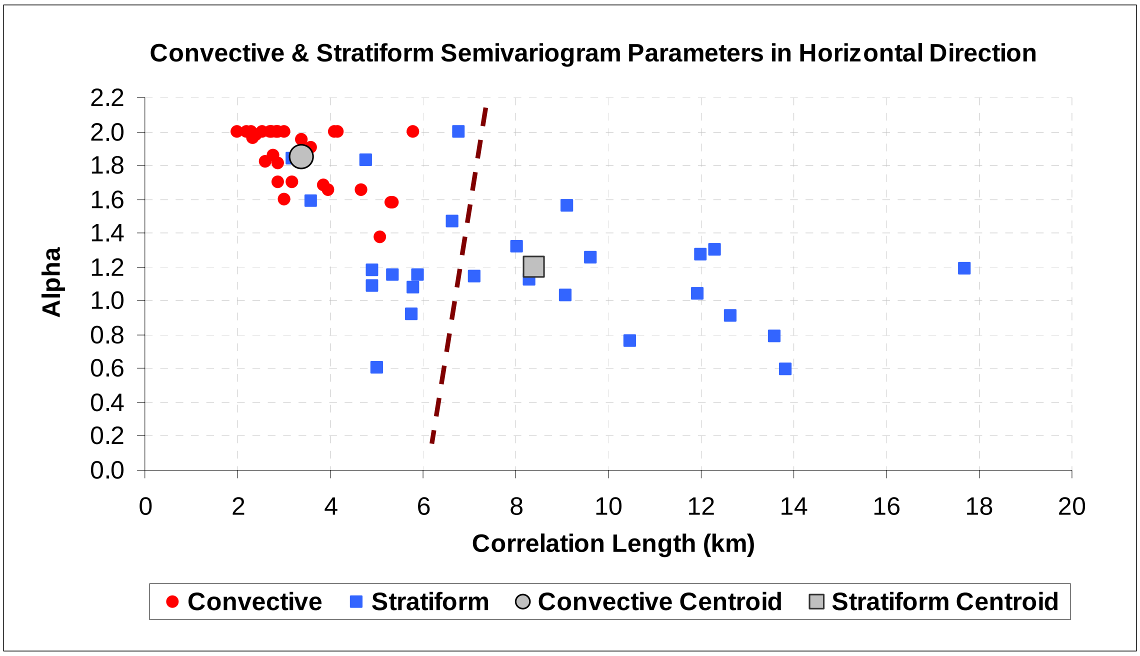

Figure 2.14 Scatter plot of alpha (α) and correlation length (L) parameters for stratiform and convective rainfall in a horizontal direction. The dashed line indicates the separation boundary of the two clusters as identified by a Fuzzy C-Means cluster algorithm.

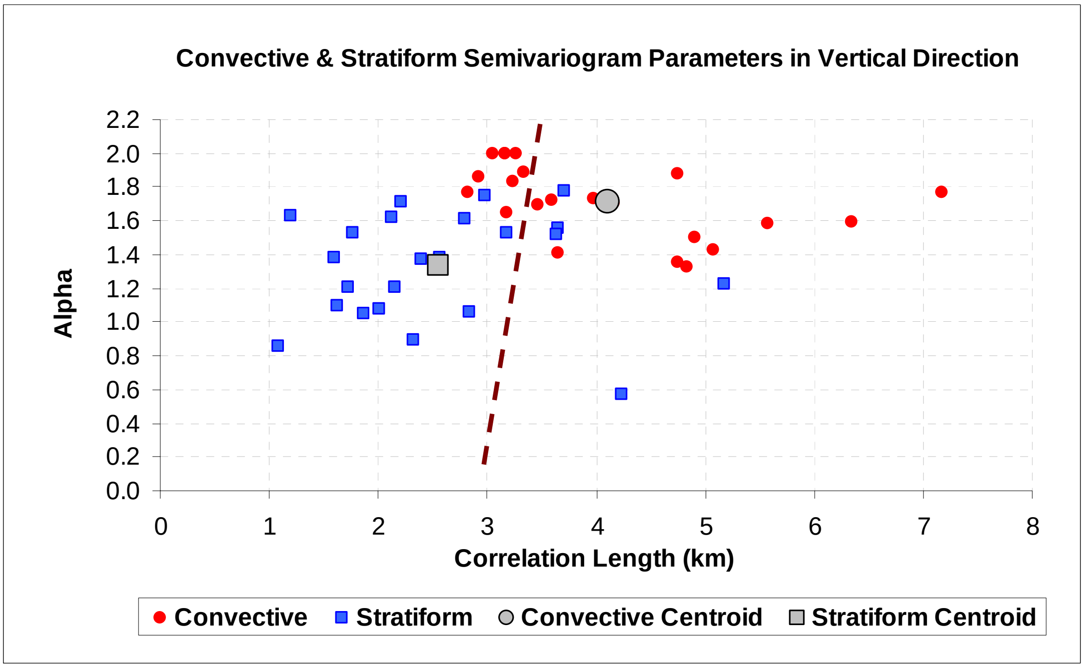

Figure 2.15 Scatter plot of alpha (α) and correlation length (L) parameters for stratiform and convective rainfall in a vertical direction. The dashed line indicates the separation boundary of the two clusters as identified by a Fuzzy C-Means cluster algorithm.

The scatter plots in Figure 2.14 and 2.15 illustrate that the semivariogram parameters tend to cluster in a particular region according to the type of rainfall, indicating that a fixed set of parameters can comfortably be used depending on the rainfall type.

2.3.5 Cluster Analysis

To test if a natural clustering of the variables occurs, and also to ascertain the degree of clustering in the vertical and horizontal directions, two cluster analysis techniques were utilised. Initially a Fuzzy C-Means clustering algorithm (Gordon, 1981: 58-60) was run on the scatter plot of the data, where the specification set in the algorithm was that the data should be divided into two groups. In the horizontal direction the Fuzzy C-means cluster algorithm split the groups roughly along the 6.5-kilometre correlation length. In this instance 24% of the points were misclassified and associated with the incorrect rain type. When the clustering algorithm was run on data in the vertical direction it resulted in a division running approximately along the 3.5-kilometre correlation length, in this instance 30% of the points were misclassified. The divisions for the clusters in the vertical and horizontal direction determined by the Fuzzy C-means cluster algorithm are indicated on Figures 2.14 and 2.15 by the dashed line.

The cophenetic correlation coefficient (c) (Gordon, 1981: 125) was also used to determine the degree of clustering of each group. If the value of the cophenetic correlation coefficient is close to one it indicates a high degree of clustering, while a value close to zero indicates a low degree of clustering (Gordon, 1981: 125). As indicated in Table 2.3 the cophenetic correlation coefficient shows a high degree of clustering as all the values are close to one. Also indicated by the values in Table 2.3 is that the convective clusters show a higher degree of clustering than the stratiform clusters in both the horizontal and vertical directions.

Table 2.3: Cophenetic correlation coefficient (c) (Gordon, 1981: 125) for convective and stratiform semivariogram parameter clusters in the horizontal and vertical direction.

	HORIZONTAL		**VERTICAL**	
	Convective	**Stratiform**	**Convective**	**Stratiform**
c	0.842	0.724	0.849	0.752

Because the semivariogram parameters naturally cluster around a centroid depending on the rainfall type, it seems to be a safe assumption to use fixed climatological semivariogram parameters depending on the type of rainfall classified; the adopted values are given in Table 2.3. This idea is tested in the next sub-section.

2.3.6 Sensitivity Analysis of Climatological Semivariogram Parameters

To determine how sensitive final Kriged solutions are to the use of a fixed set of semivariogram parameters, a sensitivity analysis was undertaken. In the testing procedure, various 2D data sets of radar reflectivity values were selected that contained solely convective or stratiform rainfall. Random portions of the selected reflectivity data were then removed and Ordinary Kriging was used to estimate the missing data, but in this instance the Ordinary Kriging was carried out several times using different combinations of semivariogram parameters (a brief overview of Kriging methods will be given in the next section). Five different sets of parameters were used: firstly the centroid value of α and L for convective and stratiform rain and then the α and L values one standard deviation away from the centroid values, as computed by Eqns. (2.17) and (2.18) and illustrated in Figure 2.16 for the horizontal direction.

$$\alpha \pm \sigma_\alpha \quad , \quad 0 < \alpha \leq 2 \tag{2.17}$$

$$L \pm \sigma_L \quad , \quad 0 < L \tag{2.18}$$

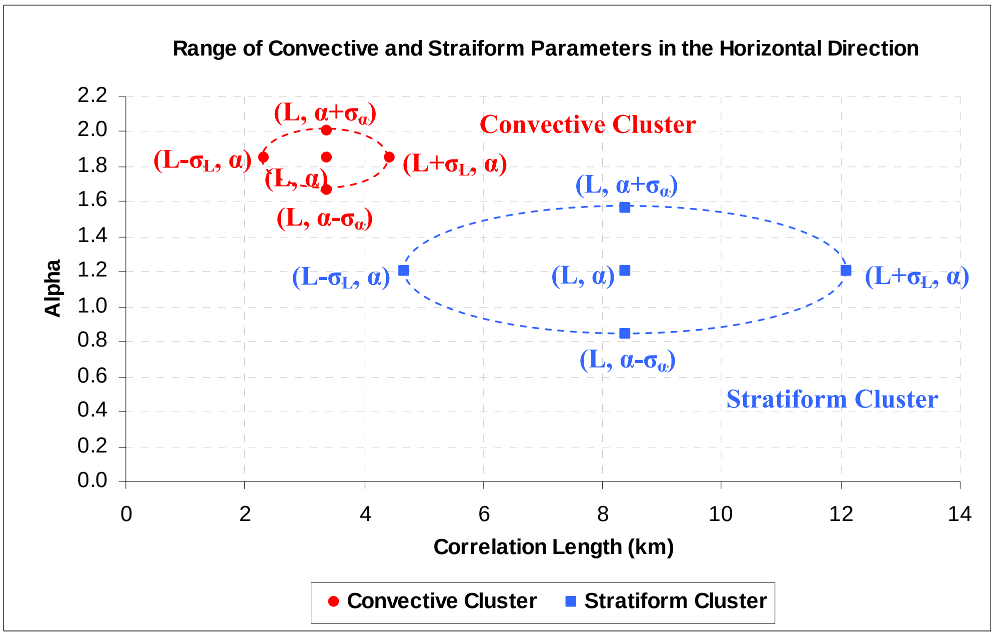

Figure 2.16 Schematic of the sensitivity analysis for the stratiform and convective clusters in the horizontal direction. The range of expected parameters for each cluster is indicated by the ellipses.

The infilled data were then compared to the original observed reflectivity data by computing the SSE, the means and standard deviations of the observed and estimated data. As indicated by Figures 17 and 18 there is no significant difference between the estimated values returned for the expected range of parameter values. The estimated Kriging values appear to be more sensitive to the shape parameter, α, than the correlation length, L.

The results indicate that there is no significant difference between the five sets of parameter values used in each instance, giving a clear indication that the final Kriged solution is reasonably insensitive to the range of the shape parameter, α, and correlation length, L, values that can possibly occur when a sample semivariogram is fitted to a specific rainfall type. Instead of solving for each set of semivariogram parameters for each neighbourhood one can simply use the centroid value computed for each rainfall type depending on the observed rainfall.

2.3.7 Application to 3D Reflectivity Data

When selecting control points for computing the estimated value of a target point in 3D space, three possible combinations of points can be selected: (1) all convective, (2) mixed, containing both stratiform and convective points and (3) all stratiform. For situations with either all convective or all stratiform, the appropriate subset of parameters in Table 2.4 are utilised for all estimations.

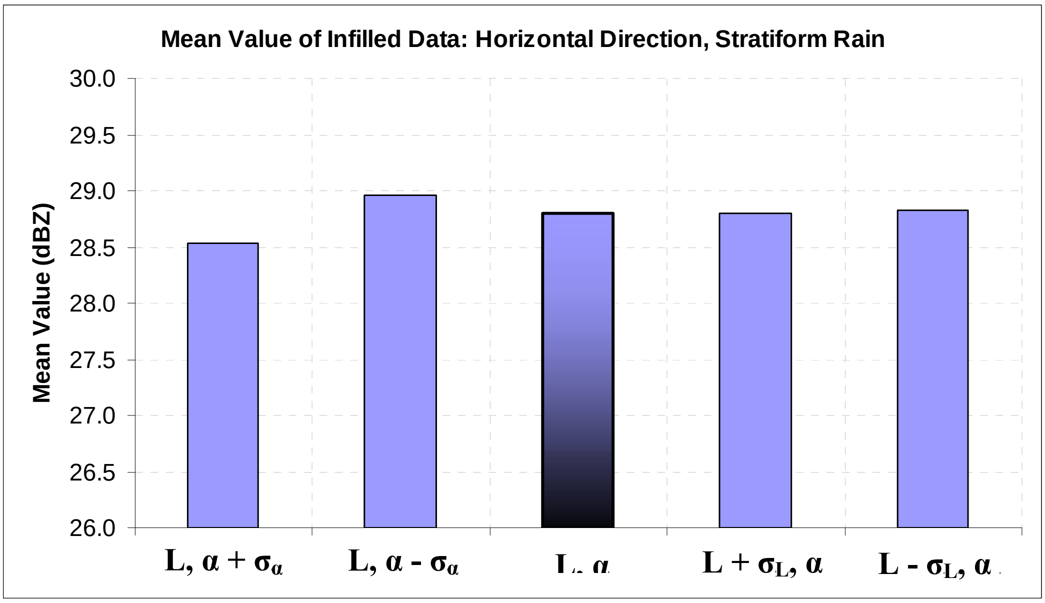

Figure 2.17 Comparison of the mean value for an infilled region with five different combinations of alpha (α) and correlation length (L) parameters as indicated on the horizontal axis on the graph.

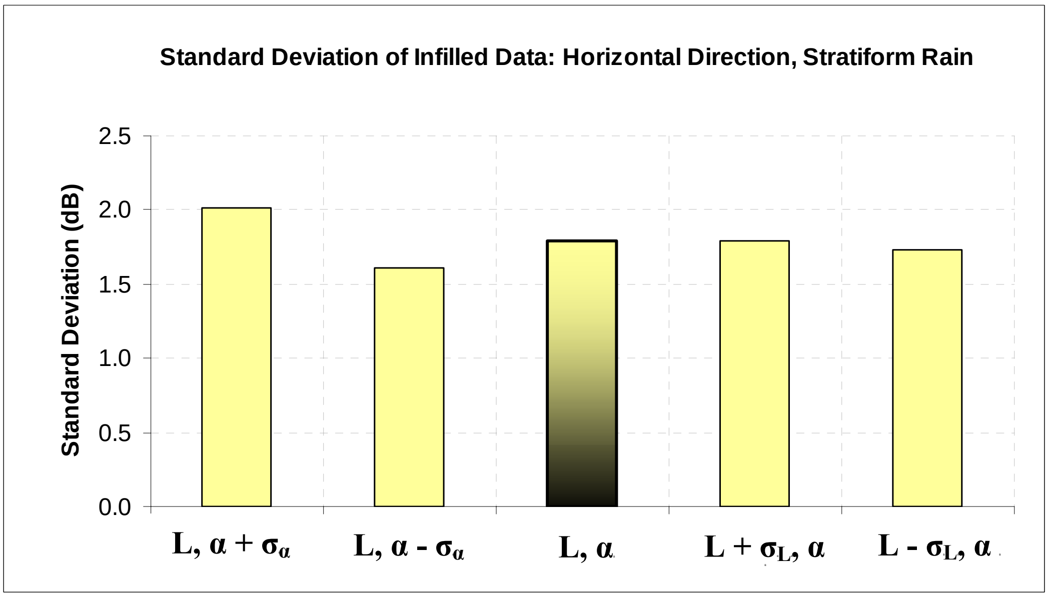

Figure 2.18 Comparison of the standard deviation for an infilled region with five different combinations of alpha (α) and correlation length (L) parameters as indicated on the horizontal axis on the graph.

Table 2.4: Semivariogram parameter values in the horizontal and vertical directions for stratiform and convective rainfall.

	HORIZONTAL		VERTICAL	
	α_H	L_H (km)	α_V	L_V (km)
STRATIFORM	1.53	8.40	1.33	2.56
CONVECTIVE	1.85	3.38	1.71	4.11

When it comes to the practical calculation of the semivariogram elements in the Kriging equations (Eqn. (2.21)) there are no numerical or technical difficulties if the control points are all of one type, either convective, or stratiform. A problem arises when the controls are of mixed type. It was found that if the individual semivariogram elements in the equation were set to their type, with a compromise value for a mixed pair, unpredictable instabilities occurred in the solution of the equation. The methodology adopted was to use a linear weighting of the α and L parameters, where the weights are the proportion of type. For example in a 2D situation in the horizontal direction if, 25 control points are selected of which 15 are of convective pixels and 10 stratiform pixels, the computed semivariogram parameters would be L_H = 5.39km and α_H = 1.72, and these values would be used for all calculations of $\gamma(h)$ in the Kriging equations (Eqn. (2.21)).

When Kriging with a 3D data set, the semivariogram model developed by Seed and Pegram (2001) was modified and is given by Eqn. (2.19):

$$\gamma(h) = \sigma^2 \cdot \left[1 - \exp\left(-h^\alpha\right)\right] \tag{2.19}$$

where: $$h^2 = \left(\frac{r}{r_0}\right)^2 + \left(\frac{z}{z_0}\right)^2$$

and $$r^2 = x^2 + y^2$$

here σ^2 is the field variance, h is the scaled distance in spherical co-ordinates, α is the scaling exponent, r is the distance in the horizontal plane, r_0 the horizontal correlation length, z is the distance in the vertical direction, z_0 the vertical correlation length and x and y are the distances in the Cartesian horizontal direction.

2.4 Kriging Methods

Kriging was chosen as the computational method to interpolate the observations missing within the CAPPI stack and to extrapolate the reflectivity values contained aloft to ground level. Kriging is considered to be the optimal technique for the spatial prediction of Gaussian data (Cressie, 1993: 106). In this application two types of Kriging were used, Ordinary and Universal Kriging (also called External Drift Kriging (Hengl et al, 2003)).

2.4.1 Ordinary and Universal Kriging

In Ordinary Kriging, the mean is assumed constant and unknown throughout the field whereas in Universal Kriging, the mean is also assumed to be unknown but varying (Chiles and Delfiner, 1999: 151). The Ordinary Kriging equations are given by Eqn. (2.20):

$$z(s_0) = \lambda^T(s_0) \cdot z \tag{2.20}$$

where $z(s_0)$ is the value to be estimated at the target point s_o, z is the vector of known reflectivity values, or control points, and the row vector $\lambda^T(s_0)$ contains the calculated weighting values which depend on s_0. The vector of weighting values is computed by the matrix Eqn. (2.21):

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} \cdot \begin{bmatrix} \lambda(s_0) \\ \mu(s_0) \end{bmatrix} = \begin{bmatrix} g(s_0) \\ 1 \end{bmatrix} \tag{2.21}$$

where G is a matrix of semivariogram values between the control points, u a unit vector of ones, $\lambda(s_0)$ the vector of weighting values, $\mu(s_0)$ a Lagrange multiplier and $g(s_0)$ a vector of semivariogram values between the target point at s_o and control point locations; note the dependence of the unknown and right hand side section on the position of the target point s_0. In Universal Kriging variables can be added to Eqn. (2.21) in order to model the variation of the mean throughout the field. This modification is indicated in Eqn. (2.22), by the addition of Q_c to the coefficient matrix and $q_t(s_o)$ to the right hand side and β to the vector of weighting values:

$$\begin{bmatrix} G & u & Q_c \\ u^T & 0 & 0 \\ Q^T{}_c & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} \lambda(s_0) \\ \mu(s_0) \\ \beta \end{bmatrix} = \begin{bmatrix} g(s_0) \\ 1 \\ q_t(s_0) \end{bmatrix} \qquad (2.22)$$

where the additional terms: $Q_c = \begin{bmatrix} q_1(s_1) & \cdots & q_p(s_1) \\ \vdots & & \vdots \\ q_1(s_n) & \cdots & q_p(s_n) \end{bmatrix}$

The selected variable $q_j(s_i)$ is associated with the point s and could be altitude, mean annual precipitation or some other classification variable. They range from $q_1(s_i)$ to $q_p(s_i)$ and are associated with the selected control point. The vector $q_t(s_0)$ contains the corresponding variables associated with the target point. Typical models for the variation of the mean are either linear or polynomial; alternatively variables associated with physical characteristics of the control points may be included, that influence the mean value of the field.

2.4.2 Computational Stability and Efficiency

In order to supply an estimate of the rainfall at ground level in a real time basis, computational efficiency is of the utmost importance. In South Africa, as volume scan data are to be processed at approximately 5 minute intervals, an estimated 120 000 target points need to be computed in this time. One of the disadvantages of the Kriging technique is that it relies on the solution of a linear system of equations whose size is proportional to the number of selected control points. For large systems of equations this can be time consuming, computationally burdensome and potentially unstable. Previous authors have also made reference to the computational burden associated with Kriging (e.g. Creutin and Obled, 1982).

There are several approaches that can be adopted for Kriging target data that can be adopted for Kriging target data, some of which are:

- compute all target values together.
- separate the target points into contiguous sets and identify their boundary control points.
- use neighbourhood Kriging to estimate each target value individually.

It was ascertained that the last of the three options is the most accurate, most stable and fastest by an order of magnitude (Wesson and Pegram, 2004) and is the method adopted here. One of the properties of Kriging that can be exploited to overcome the problem of dimension is the 'screening effect' (Chiles and Delfiner, 1999: 202-206). The screening effect refers to the observation that the significant weighting values are concentrated around the target point, with the weighting values rapidly decreasing with distance from the target point. For each target point only a cluster of the nearest control points needs to be selected. In this application the computed optimum number of control points based on an exhaustive comparison, is twenty-five. This has the effect of drastically reducing the computation time with little or no loss in accuracy of the final results.

Another unexpected problem associated with the Kriging technique, which is not so well known, is the ill-conditioning of the coefficient matrix. It was demonstrated by Wesson and Pegram (2004) that the coefficient matrix can be highly ill-conditioned dependent on a number of factors, but most notably the chosen parameterization of the semivariogram function. This is especially evident as the α value tends to become Gaussian ($\alpha \rightarrow 2$). The method of Singular Value Decomposition (SVD) is therefore used in conjunction with a trimming of the singular values to provide a computationally stable method to find the solution to the Kriging equations. Due to the fact the singular values are trimmed, matrix rank reduction techniques can then be exploited to dramatically decrease the time to find the inverse solution of the coefficient matrix with little or no loss in accuracy of the final Kriged solution (Wesson and Pegram, 2004).

2.4.3 Development of Ordinary and Universal Kriging Techniques

In the application of the Kriging technique the control data selected divides into three categories which are determined by the 25 control points in the selected neighbourhood and are defined as follows:

(1) Stratiform - control points consist entirely of stratiform rain.

(2) Mixed / Intermediate - control points consist of a combination of stratiform and convective rain.

(3) Convective - control points consist entirely of convective rain.

Although the rain is classified as purely stratiform or convective rainfall an intermediate rainfall zone is implied for computational purposes. Ordinary Kriging was used in the "pure" convective and stratiform zones while Universal Kriging was used in the mixed zone. The problem to be solved was how to get the best combination in the mixed cases.

In order to determine the most effective form of Kriging to use and the most appropriate variables to use in Universal Kriging a validation technique was devised. The CAPPI level 2km above ground level was selected as the target level for a variety of weather types. The data from aloft was then used to estimate all the reflectivity data on the 2km CAPPI level. The estimated data and observed data were then converted from reflectivity to rainrate values and the rain separated into the three categories listed above. These were compared in terms of SSE, means and standard deviations. In this exercise only the northern half of the image was considered in order to ensure there were no ground clutter reflectivity values in the data. Example of this procedure is illustrated in Figure 2.19. In this way the most appropriate form of Kriging and variables for Universal Kriging can be selected for each particular rainfall zone.

Stratiform zone

In the stratiform zone Ordinary Kriging is used to compute the estimate of the target point at the 2km level. Universal Kriging was investigated by modelling the vertical variation of the mean by use of either a second order polynomial or a linear relationship to take advantage of the trend in the vertical reflectivity profile. This did not result in any significant improvement in terms of the sum of the difference squared between the estimated and observed rainfall when the trend was ignored (Ordinary Kriging). A computational advantage of Kriging in the stratiform zone is that SVD does not need to be used in order to compute the inverse of the coefficient matrix.

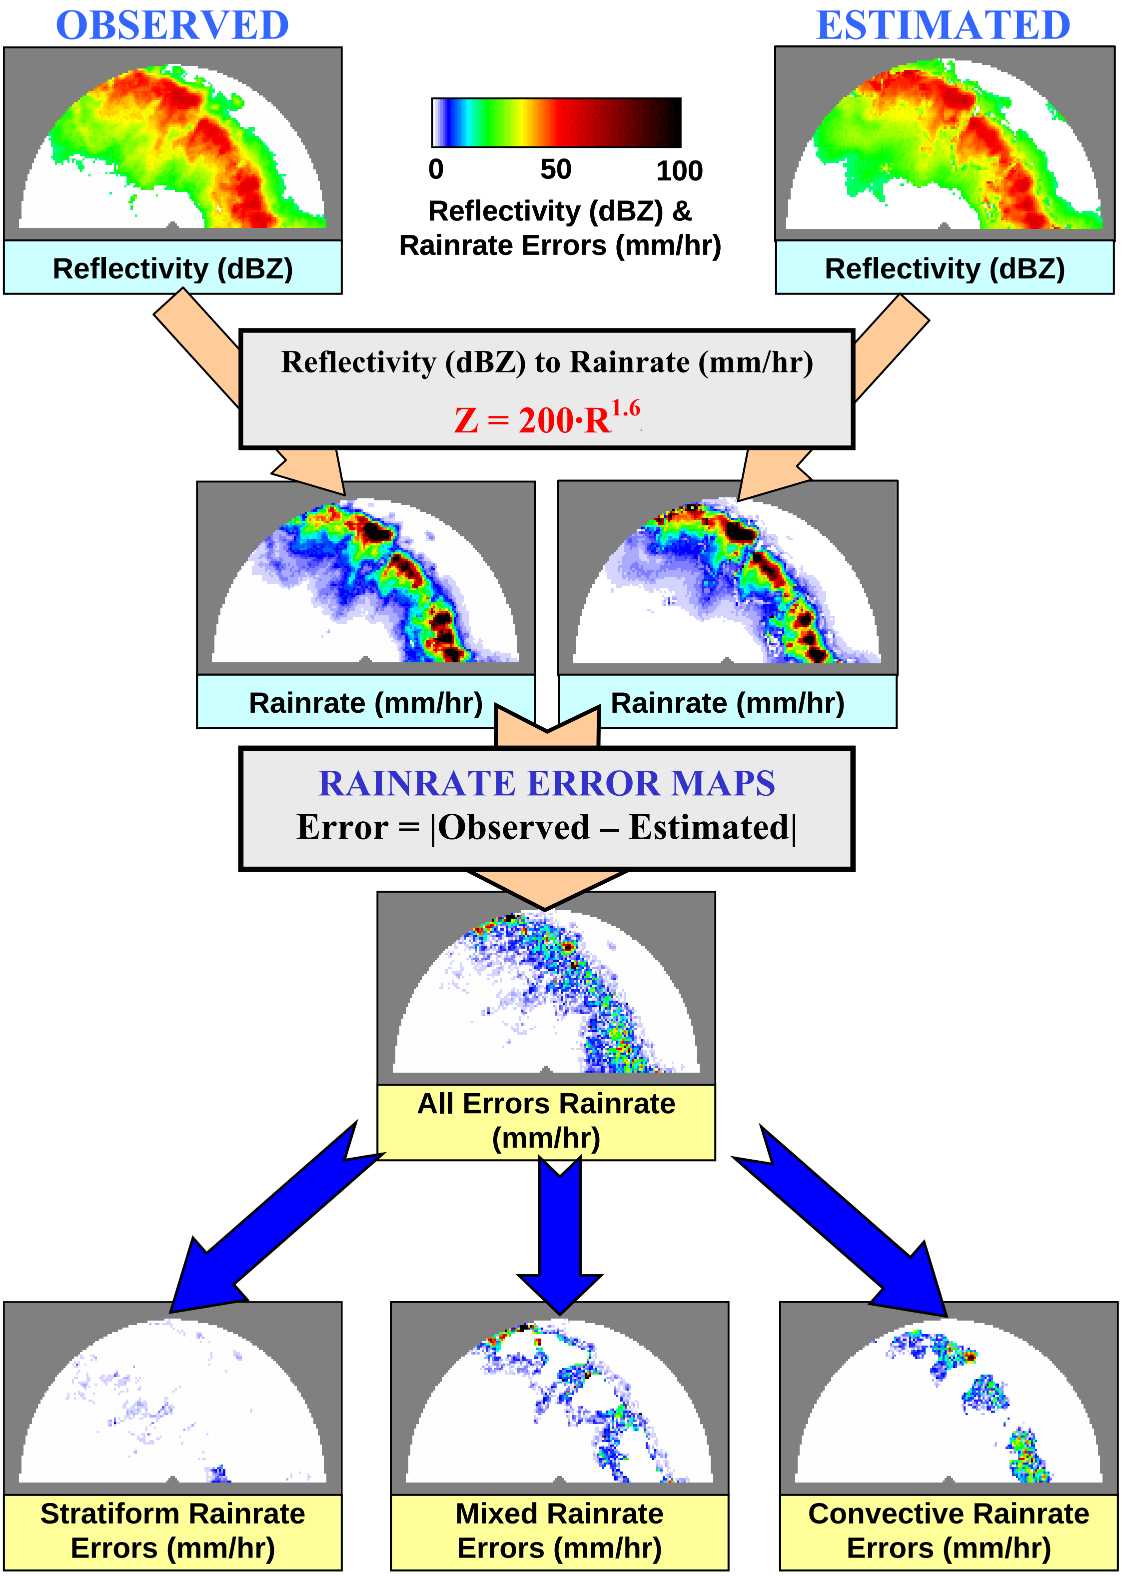

Figure 2.19 Separation of observed and estimated rainrates at CAPPI level 2 into stratiform, mixed and convective rainrate zones and the corresponding errors for each zone in order to determine effectiveness of Ordinary and Universal Kriging in each rainfall type. The estimated image at level 2 was completely filled in from the information above it.

This is due to the fact that the shape parameter, α, for stratiform rain is close to unity (exponential) which results in a stable coefficient matrix. By computing the condition number, which for a symmetric matrix is defined as the ratio of the largest to the smallest singular value (Wilkinson, 1988: 191), it was shown that the coefficient matrices for stratiform rain were much more stable than the coefficient matrices for mixed and convective rain. In this instance the quicker LU decomposition algorithm can be employed to find the solution to the Kriging equations.

Convective zone

The control points in this zone consist entirely of convective rainfall, so SVD is used to compute the inverse of the coefficient matrix because the shape parameter, α, is close to Gaussian. Universal Kriging was not found to provide a significant improvement in the estimates over Ordinary Kriging. As shown in Figure 2.6, as the convective rain falls towards ground level there is a continuous and significant increase in rain intensity. Based on this, it was decided to use the vertical distance from the control to the target points as the external variables in Universal Kriging. Initially a linear and second order polynomial relationship were investigated. However the linear relationship tended to overestimate the convective rainfall and the second order polynomial provided an underestimate at the target point. Based on the above, Ordinary Kriging is now used when the control points selected consist entirely of convective pixels.

Mixed / Intermediate zone

In the mixed zone, SVD is used to compute the inverse of the coefficient matrix because, for convective rain, the shape parameter α is close to Gaussian as seen in Table 2.4. This results in a numerical unstable coefficient matrix (Wesson and Pegram, 2004). In addition Universal Kriging in this instance provides an improved estimate over Ordinary Kriging. Testing the method on 10 different instantaneous images showed an improvement in estimates (in terms of SSE) on 8 of the images. Since the control points now consist of both stratiform and convective pixels a stratiform/convective binary switch was used as the two external variables in Universal Kriging. This is indicated as a specialization of Q_c in Eqn. (2.22) and is given as Eqn. (2.23):

$$Q_c = \begin{bmatrix} C_1 & S_1 \\ \vdots & \vdots \\ C_n & S_n \end{bmatrix} \tag{2.23}$$

Where C represents a convective and S a stratiform control point; so that a convective control point is assigned the values *[C S] = [1 0]*; and a stratiform control point is assigned the values *[C S] = [0 1]*.

A decision needs to be made as to which semivariogram is to be used in the coefficient matrix and which switch is to be set in Q_c (Eqn. (2.23)). This also affects the choice of q_t in the right hand side of Eq. (2.22). It was decided to let this be determined by examining the control points; if the majority of the control points are convective then the target point is assumed to be convective, if not then it is assumed to be stratiform. Including these external variables resulted in an improvement in the estimates. An example of this comparison is indicated by Figure 2.20 where first Ordinary then Universal Kriging were used to estimate the mixed rainfall pixels on an instantaneous image (24 January 2002) from the Bethlehem weather radar; the number of pixels to be estimated were 2628. As indicated in Figure 2.20, Universal Kriging provides an improved estimate over Ordinary Kriging.

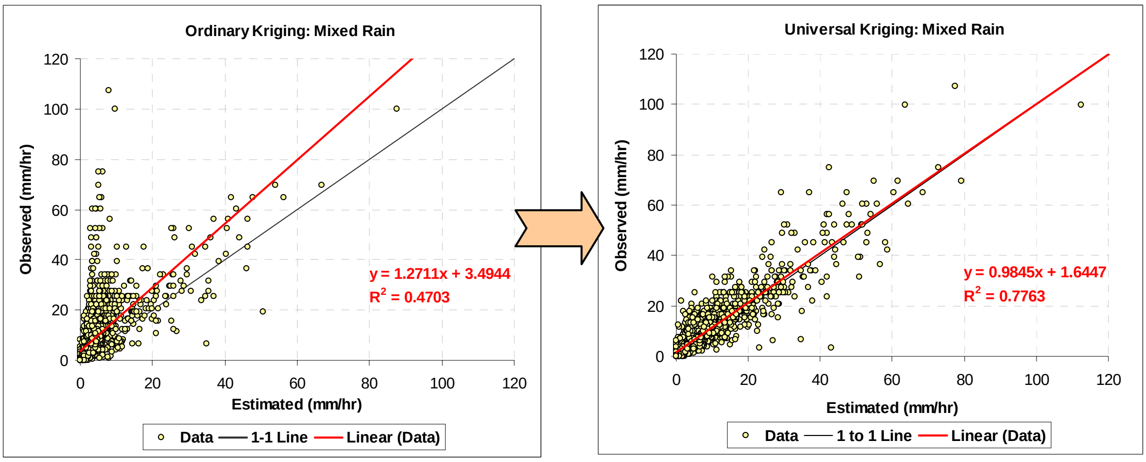

Figure 2.20 Scatter plot of observed and estimated rainrates at individual pixel points for mixed rainfall for a instantaneous image from the Bethlehem weather radar, 24 January 2002, the number of pixels estimated number 2628. Estimates are computed with both Ordinary and Universal Kriging with the Universal Kriging scatter plot indicating a substantial improvement in estimates; the bias is almost removed – there are actually two lines in the right-hand diagram.

2.4.4 Cascade Kriging

When Kriging directly to ground level using the full set of CAPPI data in a stack, there are unexpected problems which occur. A 24-hour accumulation of rainfall calculated from 5-minute images Kriged from the CAPPI volume down to ground level is shown in Figure 2.21. The data are from the Bethlehem weather radar (25 January 1996); the image is 200km square. As can be seen, Figure 2.21 shows serious discontinuities at the regions located directly under the edges of the CAPPI levels and also shows an inflation of the reflectivity estimates due to the numerical and geometric distribution of Kriging weights at the CAPPI edge locations.

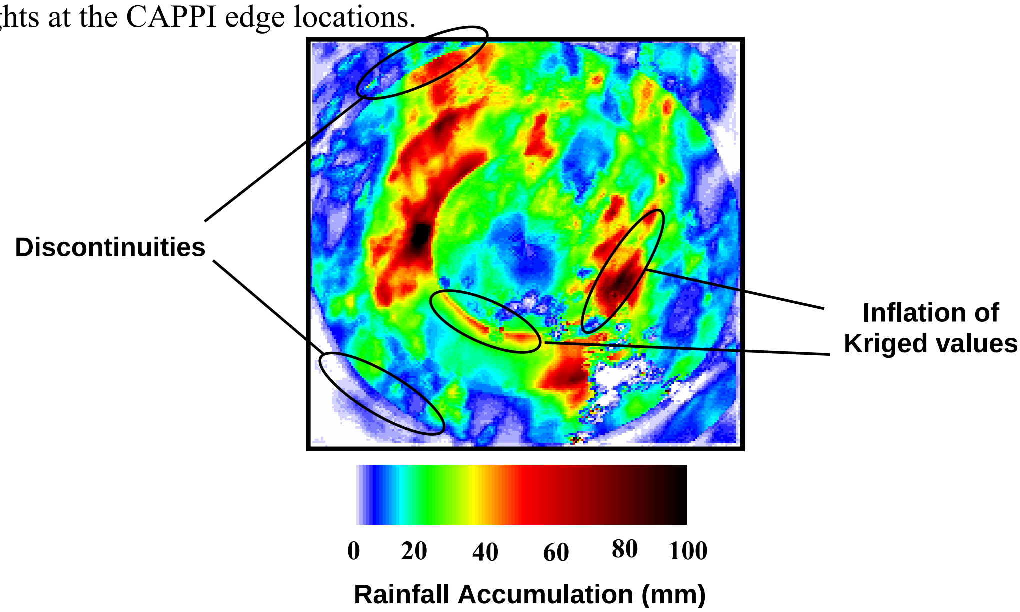

Figure 2.21 24-Hour rainrate accumulation for the Bethlehem weather radar (25 January 1996). Image dimensions are 200km square. Indicated on the image are the serious discontinuities associated with the CAPPI edges and also the inflation of the Kriged estimates at CAPPI edges.

A solution to the above problem is to use a Cascade Kriging approach that works as illustrated in Figure 2.22. The first step is infill the missing data in the CAPPI at 4km above ground level, pixel by pixel, using neighbourhood Kriging of the 25 control data vertically above and in the horizontal direction. Once all the unknown data on that CAPPI level have been estimated and infilled, the CAPPI level directly below is examined and any missing data here are estimated in the same manner; once again estimates from above (which now can include previously infilled data) and at the same level are used as control data. This is repeated until all the unknown data in each CAPPI level are estimated. The final step involves an estimate of the rainfall at ground level.

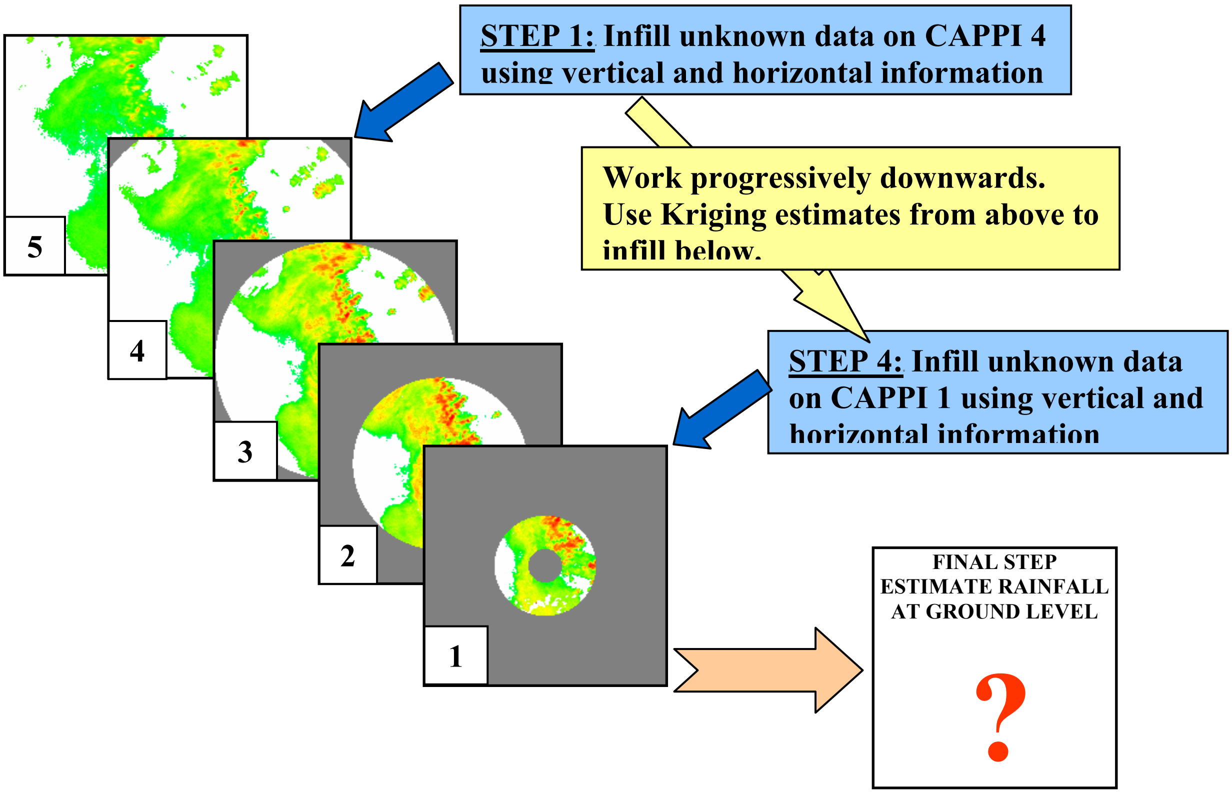

Figure 2.22 Cascade Kriging algorithm steps. Image taken from the Bethlehem weather radar (24 January 1996) where the images are 200km square and the grey portions indicate regions where the data are missing.

Figure 2.23 gives a 3D illustration of radar volume scan data before and after Cascade Kriging. The data are from the Bethlehem weather radar, 14 February 1996, where the image is 200km square, and the vertical extent is 3km. The levels above 3km have been ignored for clarity, but were used to estimate from level 4 downwards as indicated in Figure 2.22. The image on the left hand side indicates the volume scan data before Cascade Kriging and the right hand image represents the same volume scan data after Cascade Kriging has been used to estimate as much unknown data as possible and provide an estimate at ground level.

Cascade Kriging has the following influences on reflectivity data. There is increased smoothing of the reflectivity data as ground level is approached. There is also an increase in the reflectivity values as ground level is approached due to the pattern of Kriging Weights as illustrated in Figure 2.24. This is because, in the case where the reflectivity values increase with proximity to ground level, the negative weights on the upper level and the positive weights on the lower level create a gradient that results in

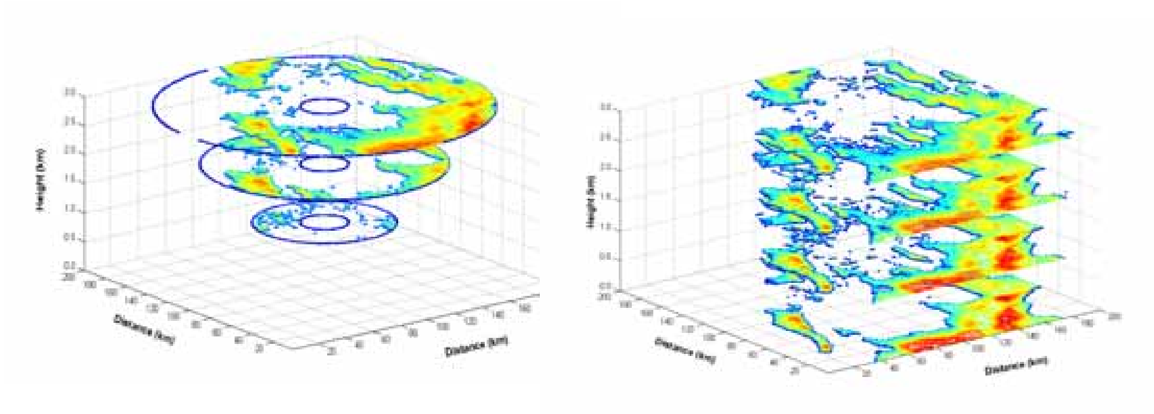

Figure 2.23 3D Graphical illustration of radar volume scan data up to 3km above ground level, before and after Cascade Kriging. Data are from the Bethlehem weather radar (14 February 1996). Data at levels 4, and higher are not shown, but were used to successively fill the stack from above.

an increase in the magnitude of the Kriged values, which corresponds with the climatological profiles shown in Figure 2.6.

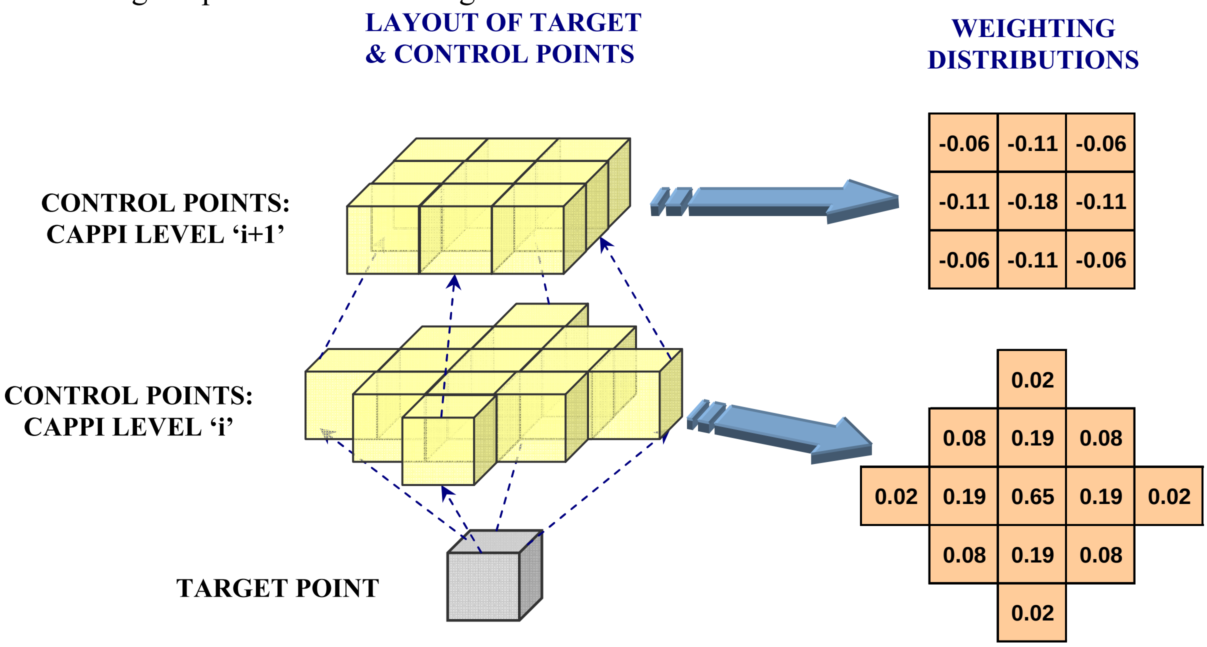

Figure 2.24 Layout of Kriging weightings, computed with convective semivariogram parameters, for 22 control points centred directly above a target point. The negative weights are situated on the upper CAPPI level and the positive weights on the lower CAPPI levels and maintain a gradient in the field if there is one above the target.

Figure 2.25 shows an example of the smoothing and the increase in magnitude of the reflectivity values as ground level is approached. The image on the left hand side is at the 4km CAPPI level and the right hand side image indicates the Cascade Kriging estimate at ground level. As can be seen the estimates are smoother and have increased in magnitude.

The next section details the validation testing procedure adopted to determine the efficiency and accuracy of the algorithm.

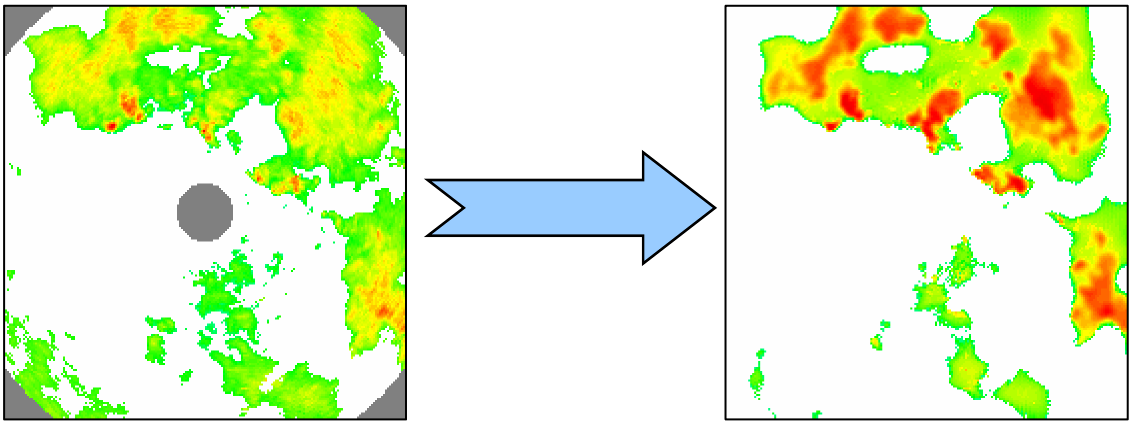

Observed reflectivity data 4km above ground level

Kriged estimate at ground level

Figure 2.25 Example of smoothing and increase in magnitude of the reflectivity values as ground level is approached. The image is from the Bethlehem weather radar from the 25 January 1996, dimensions of the images are 200km square.

2.5 Testing and Results

The testing and results section is divided into two sections (1) the effectiveness of the Cascade Kriging algorithm to infill ground clutter and (2) the accuracy of the rainfall estimates at ground level.

2.5.1 Ground Clutter Infilling

In order to determine how effectively the Cascade Kriging algorithm infills marked ground clutter contained in a radar volume scan, the following testing procedure was undertaken. Firstly a 3D ground clutter map from the Bethlehem weather radar was obtained from METSYS and is illustrated in Figure 2.26. Figure 2.26 indicates ground clutter contamination on CAPPI levels 1 to 4km above ground level occurring in the South East corner of the radar volume scan due to the Maluti and Rooiberge Mountains.

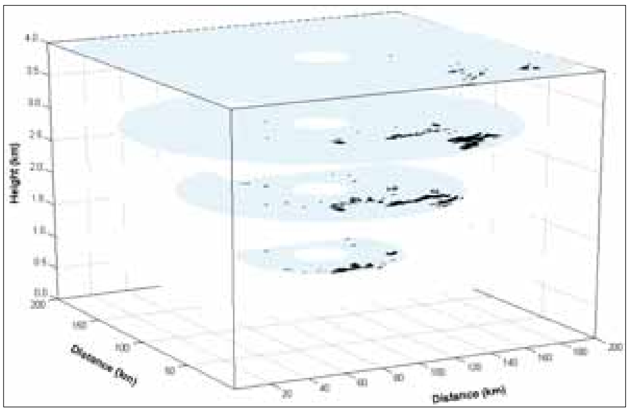

Figure 2.26 3D graphic of ground clutter contamination from the Bethlehem weather radar. Ground clutter contaminates the weather radar volume scan up to 5km in height.

The Bethlehem ground clutter map was then rearranged so as to lie in the North Eastern sector of the radar volume scan so as to be situated in regions where observed rainfall is now blocked out by the ground clutter on CAPPI levels 1 to 5. Three rainfall events were then selected that contained distinctly different rainfall patterns, all taken from the Bethlehem weather radar. The first rain event selected was from the 17 December 1995 which consisted of predominantly stratiform rainfall. The second rain event, 25 January 1996, consisted of a mixture stratiform and convective rainfall. The third rain event, 13 February 1996, also consisted of a mixture of stratiform and convective rainfall.

The testing procedure employed works as follows and is also illustrated in Figure 2.27. All three rain events described above were run over a 24 hour period. Each instantaneous image had the rearranged Bethlehem ground clutter map, now situated in the North-East corner of the image, superimposed over observed (and known) reflectivity values. The regions of marked ground clutter are then treated as unknown and contaminated data and the reflectivity estimated using the Cascade Kriging algorithm. All of the estimated and observed reflectivity values are then stored in a vector and converted to rainrate by the Marshall Palmer formula. This process is repeated for each instantaneous image throughout the rain event.

The accumulation values of the observed (real) and estimated reflectivity values over different time intervals, 6 hour, 12 hour and 24 hours, were then computed in order to determine how closely the estimated rainrate values matched the observed rainrate values. As indicated in Figure 2.28 the 24 hour accumulation scatter plots for the three selected rain events are displayed. The scatter plots indicate a high correspondence between observed and estimated reflectivity values. The majority of the data points are close to the 1 to 1 line and return a correlation coefficient (r^2) close to one.

Indicated in Figure 2.29 are the cumulative distribution functions (CDF's) of the observed and estimated rain rate accumulation values over a 24 hour period. For the three selected rain events the distributions returned for the estimated and observed accumulation values tend to exhibit a close correspondence, with the rider that (as with most extrapolation algorithms) the estimated values are biased downward, in this case by about 1 to 3 mm/hr, which features in Table 2.5.

A summary of the statistical results is contained in Table 2.5. A statistical comparison is carried out between the observed and estimated accumulation values over a 24 hour period where the mean, standard deviation, median, minimum, maximum and range of the values are evaluated. As indicated in Table 2.5 the returned statistical values are similar for the observed and estimated rain rate values.

Indicated in Table 2.6 is the percentage difference between the observed and estimated accumulation values. The percentage difference for the mean and standard deviation values is low with the greatest difference being 13.9% and the smallest 3.7%. There is also a similar percentage difference between the other summary statistics returned for the observed and estimated accumulation values.

The accumulation values for the 6 and 12-hour periods showed an equally high correspondence between the observed and estimated rainfall accumulation values. The scatter plots, cumulative distribution functions and summary statistics all indicate that the Cascade Kriging algorithm described in the previous sub-section estimates reflectivity values contaminated with ground clutter with a high degree of accuracy. The method proposed is computational fast and the results suggest that it provides an accurate estimate of missing or contaminated reflectivity values obtained in radar volume scan data.

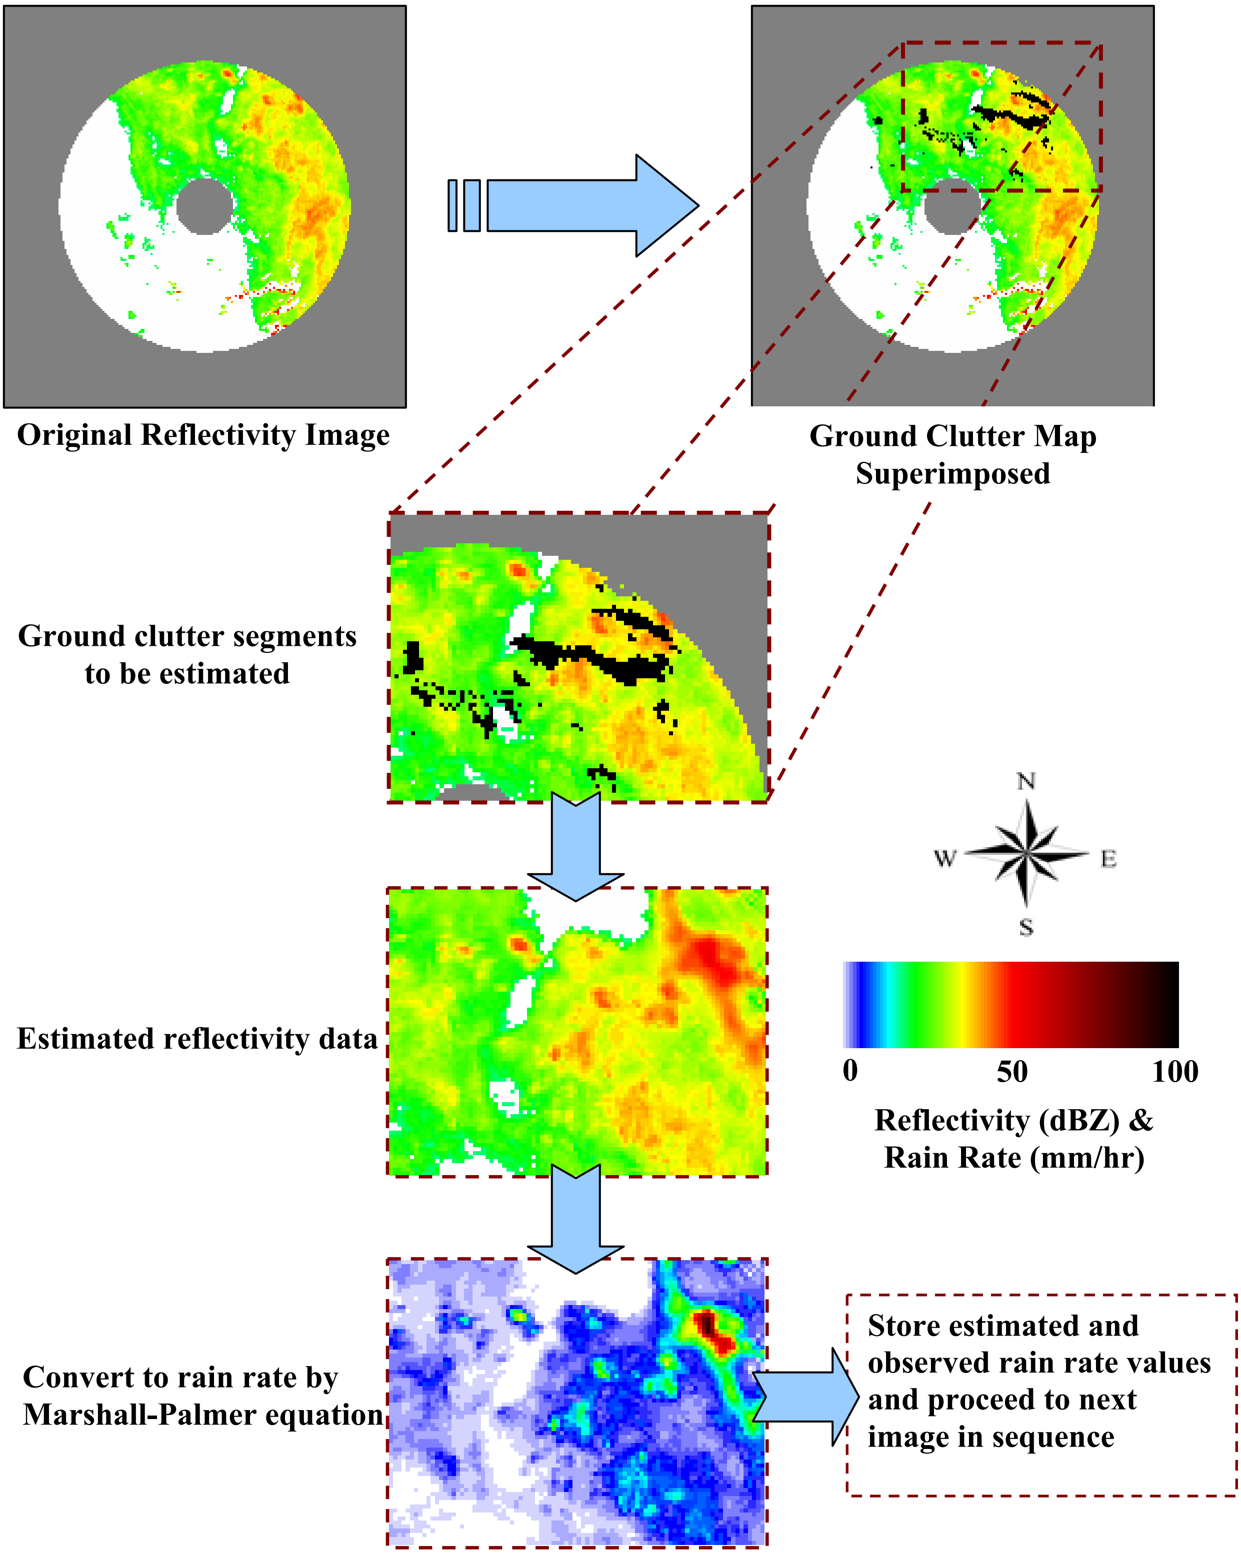

Figure 2.27 Testing and validation procedure to determine the effectiveness of the Cascade Kriging algorithm in infilling ground clutter. The instantaneous image is from the Bethlehem weather radar on the 13 February 1996.

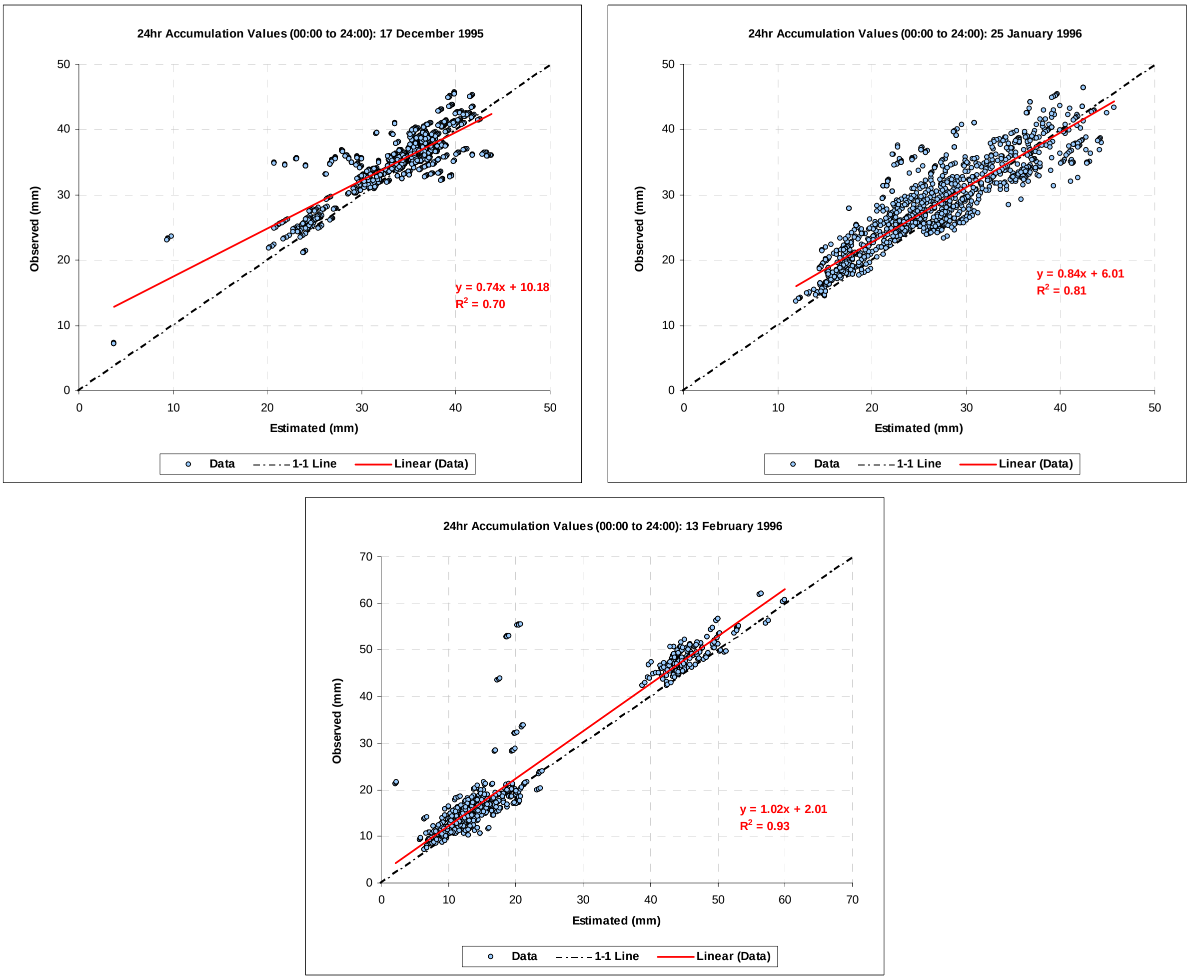

Figure 2.28 Scatter plots of 24 hour accumulation values for ground clutter infilling for three different rain events.

2.5.2 Rainfall Estimation at Ground Level

In order to determine the efficacy of the proposed technique in providing accurate rainfall accumulation values at ground level, a comparison was carried out between rain gauge and radar accumulation values for selected rain events; care was taken to compute values which were comparable at ground level. The Liebenbergsvlei catchment described in the Introduction and illustrated in Figure 2.2 was used as the test area for the comparison between radar and rain gauge accumulation estimates.

The radar rainfall estimates at ground level are estimated in 1km by 1km cells and the rain gauge provides a point estimate at ground level. In order to provide an appropriate way of comparing the two estimates, the following approach was taken. The average value of the 9 radar pixels (a 3km square area), with the centre square covering a gauge, was taken as the radar estimate of average rainfall at each gauge location. In computing the radar accumulation values one needs to take into account that the images are sampled at five-minute intervals, so a simple linear accumulation of the images may not provide an accurate accumulation, especially for fast moving rain events. The accumulations are therefore computed by a Morphing Algorithm that takes into account the motion of the rainfield between instantaneous images (Sinclair and Pegram, 2003).

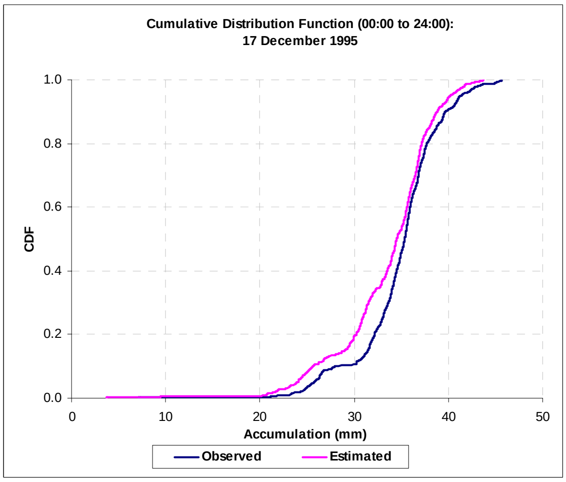

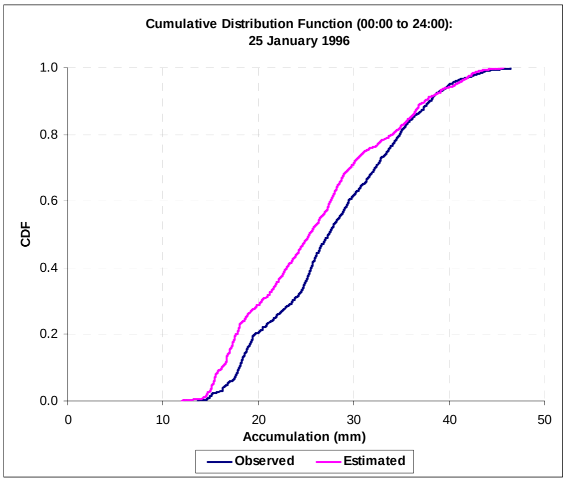

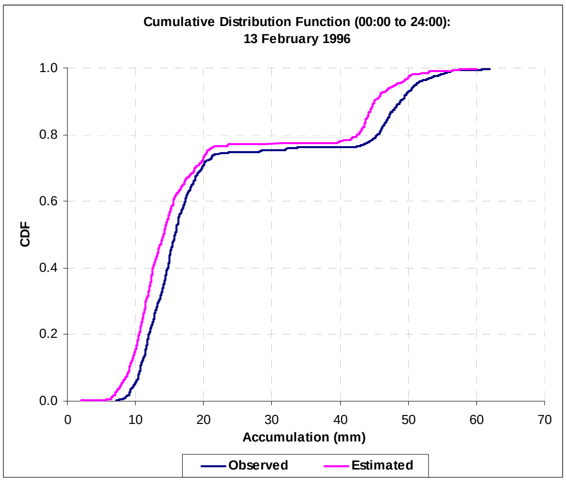

Figure 2.29 Cumulative distribution functions for the three selected rain events all indicating a close correspondence between the distributions returned for estimated and observed rain rate accumulation values for a 24 hour period.

Block Kriging was used to determine the average rainfall over the same $9km^2$ area by using all rain gauges within a range of 2 correlation lengths from the centre of the $9km^2$ area, at times corresponding to the radar estimates. The Block Kriging estimates are computed by Eqn. (2.24):

$$\hat{z}_D = \frac{1}{D}\int_D Z(x)\cdot dx \qquad (2.24)$$

where $\hat{z}_D$ is the average value over a selected region and D defines the region (Bras and Rodriguez-Iturbe, 1985: 402-404). Due to the wide spacing of the raingauges compared to the 3 × 3km area, the semivariogram parameters used were the same as those determined for the radar reflectivity data in the horizontal plane. However a nugget effect is included, which for stratiform rain has a value of 0.90 and convective rainfall 0.97 as was determined by Habib and Krajewski (2002). This allows for the random errors and uncertainties associated with rain gauge estimates. Illustrated in Figure 2.30 is the method of selection of the nearest rain gauges to be used in the estimation of the average rainfall over the 9-pixel region, where the rain gauges selected are within a range of 2 correlation lengths away from the centre of the 9-pixel region.

Table 2.5: Statistical summary of the observed and estimated accumulation values for a 24 hour period. The observed and estimated rainfall accumulation values return highly similar values.

	24 Hour Accumulation (mm)					
	Rain Event 1 (17 December 1995)		**Rain Event 2 (25 January 1996)**		**Rain Event 3 (13 February 1996)**	
	Observed	**Estimated**	**Observed**	**Estimated**	**Observed**	**Estimated**
Mean	34.8	33.5	27.8	26.0	22.8	20.5
Standard Deviation	4.54	5.17	7.41	7.93	15.02	14.23
Median	35.4	34.5	27.4	25.3	15.8	14.1
Minimum	7.1	3.7	13.7	11.9	7.2	2.1
Maximum	45.7	43.8	46.4	45.7	62.0	60.0
Range	38.5	40.1	32.7	33.8	54.8	57.9
r^2	0.70		0.81		0.93	

Table 2.6: Percentage difference between estimated and observed accumulation statistics for three rain events over a 24 hour accumulation period.

	Percentage Difference in Observed and Accumulation Estimates		
	Rain Event 1 (17 December 1995)	**Rain Event 2 (25 January 1996)**	**Rain Event 3 (13 February 1996)**
Mean	3.7 %	6.5 %	10.1 %
Standard Deviation	13.9 %	7.0 %	5.3 %
Median	2.5 %	7.7 %	10.8 %
Minimum	47.9 %	13.1 %	70.8 %
Maximum	4.2 %	1.5 %	3.2 %
Range	4.2 %	3.4 %	5.7 %

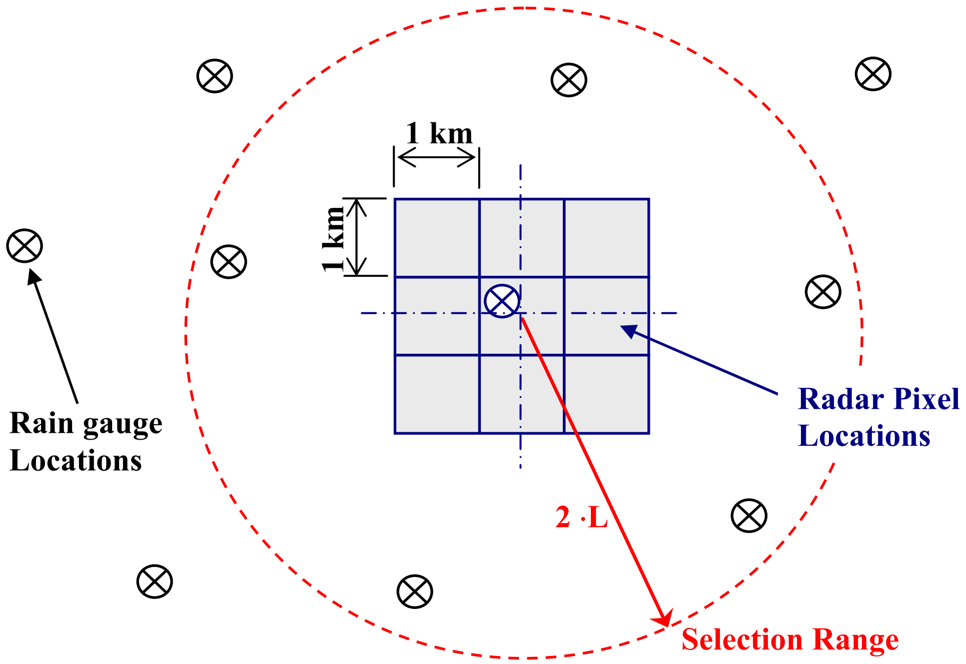

Figure 2.30 Method of selection of rain gauges to compute Block Kriging estimate over the 9 pixels centred at a rain gauge location. Rain gauges within a range of 2 correlation lengths are selected to compute the Block Kriging estimate over the 9 pixel region.

The following two rain events were selected for the validation tests: (1) 24 January 1996 - a rain event which consisted of extremely high rainfall intensities over an approximately 12 hour period, (2) 13 February 1996 - a rain event consisting of a combination of convective rain with periods of stratiform rain.

Statistical Comparison of Rain gauge and Radar Accumulations

In order to determine the quality of the radar estimates at ground level a statistical comparison of the rain gauge and radar accumulations over different periods of 24 hours, 12 hours and 6 hours was conducted. A scatter plot of the radar and rain gauge accumulations was done and the correlation coefficient computed. The Kolmogrov-Smirnov test was also used to determine if the cumulative distributions of the radar and rain gauge accumulations were significantly different at a significance level of 5%, where the hypothesis being tested is H_0: Rain gauge Distribution = Radar Distribution against H_1: Rain gauge Distribution ≠ Radar Distribution. The means and standard deviations (σ) of the accumulated values were also tested at a significance level of 5%.

RAINFALL EVENT 1: 24 JANUARY 1996

Severe convective rainfall was recorded in the late afternoon and evening of the 24 January 1996 (Mather et al, 1997). The analysis for this event was conducted over a 12-hour period from 12:00 to 24:00. For this rain event 29% of the images in the 12 hour accumulation period were identified as containing bright band and corrected. Figure 2.31 illustrates a 12 hour accumulation that gives an indication of the extreme rainfall intensity during the 12 hour period. Indicated on Figure 2.31 is an outline of the catchment area as well as the rain gauge locations. The region directly above the radar, a circle of 20km radius, has also not been infilled due to poor data quality close to the radar during this time period; two rain gauges within this region were excluded from the study.

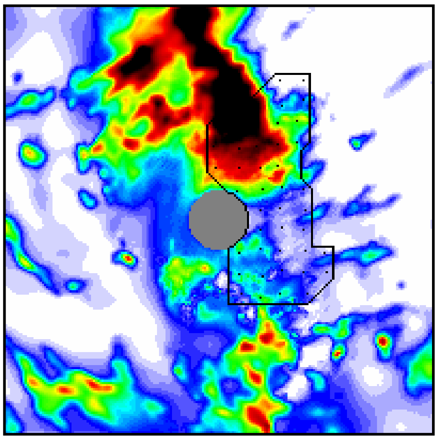

Figure 2.31 12-hour rainfall accumulation from the Bethlehem radar on the 24 January 1996. An outline of the test catchment is indicated on the image as well as the rain gauge locations. Dimensions of the image are 200km square.

A scatter plot of the rain gauge and radar accumulations after a 12 hour period at each of the 43 gauge sites computed by the appropriate Kriging method over each $9km^2$ region, is shown in Figure 2.32. There is a high correspondence for this rain event between radar and rain gauge accumulations with an r^2 value of 0.86 being returned. The highest pair of radar observations bias what is otherwise a reasonable fit, and also affects Figures 2.32 and 2.33. This is probably due to under estimation of the highly localised convective rainfall by the gauges concerned (Wilson and Brandes, 1979).

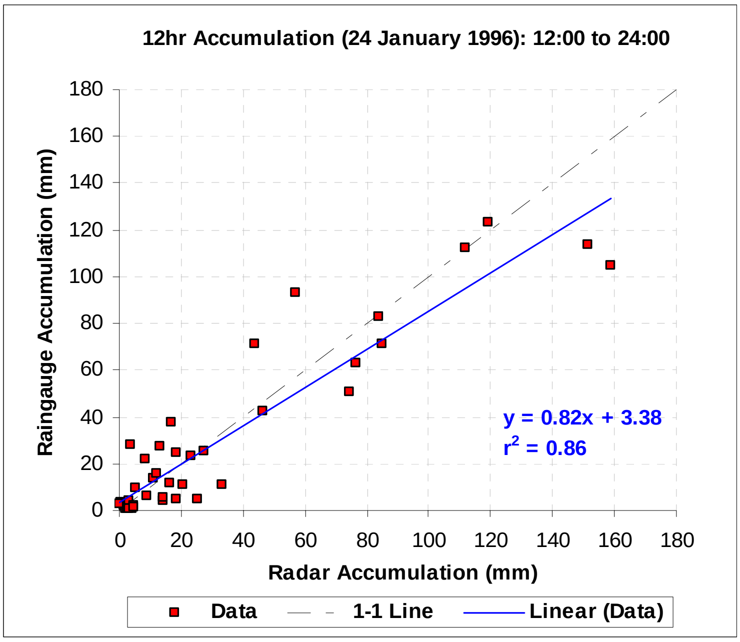

Figure 2.32 Scatter plot of radar and rain gauge accumulations for a 12-hour accumulation period for the 24 January 1996 rain event. A strong correspondence exists between the radar and rain gauge accumulation, especially for the high intensity rainfall.

Table 2.7 provides a summary of the statistical results returned for the rain gauge and radar accumulation over a 12 hour period and two 6 hour periods. For the 6 hour period (18:00 to 24:00) and the 12 hour accumulation period there is a close correspondence between the radar and rain gauge estimates. With the exception of one standard deviation, the means and standard deviations are statistically similar with the Kolmogrov-Smirnov test (K-S test) also indicating that the distributions of the accumulations for the radar and rain gauge are not dissimilar, except for the 6 hour accumulation period from 12:00 to 18:00. The word "Accept" in Table 2.5 refers to the null hypothesis H_0 being accepted and H_1 rejected, and "Reject" refers to H_0 being rejected and H_1 accepted. Figure 2.33 shows an example of the cumulative distribution functions for 12 hour rain gauge and radar accumulations where the distributions for this example were considered to be not dissimilar.

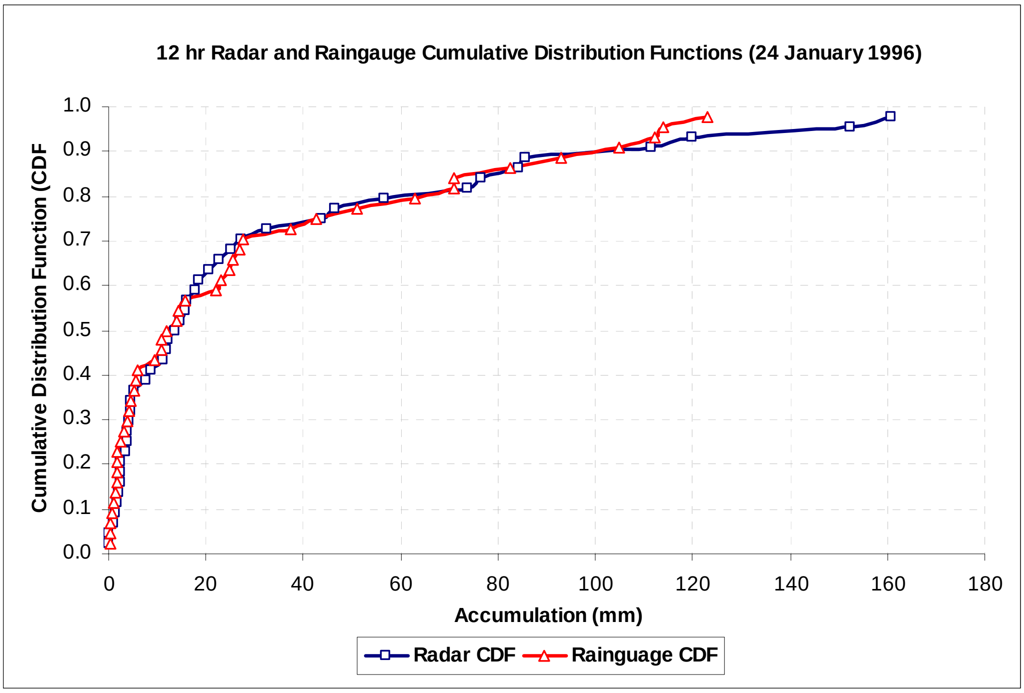

Figure 2.33 Cumulative Distribution Functions (CDF's) of the rain gauge and radar accumulations for a 12 hour period. Accumulations values are from the 24 January 1996 (12:00 to 24:00). Testing at a significance level of 5% indicate that the distributions of the accumulations for the radar and rain gauge are not dissimilar

Table 2.7: Summary of statistical results for different accumulation periods for 24 January 1996 rain event for radar and raingauge data.

	12hr accumulation (12:00 to 24:00)		**6 hr accumulation (12:00 to 18:00)**		**6 hr accumulation (18:00 to 24:00)**	
	Mean	**Stdev**	**Mean**	**Stdev**	**Mean**	**Stdev**
Radar	30.3	41.3	4.6	5.1	27.2	40.8
Raingauge	29.4	36.8	4.0	7.1	25.4	38.3
Accept / Reject H_0	Accept	Accept	Accept	Reject	Accept	Accept
r^2	0.86		0.08		0.88	
K-S TEST	Accept		Reject		Accept	

RAINFALL EVENT 2: 13 FEBRUARY 1996

The rain event on the 13 February 1996 consisted of a combination of convective and stratiform rain. The rainfall initially consisted of light stratiform rainfall but towards the evening and afternoon this area also experienced convective rainfall. In the 24 hour rain period 19% of the images were identified as containing bright band and corrected. The 24 hour accumulation for the 13 February 1996 is depicted in Figure 2.34 illustrating the combination of convective and stratiform rainfall that was recorded throughout the day. The exclusion of data directly above the radar applies in this instance as well. Two of the gauges were therefore excluded from the data analysis.

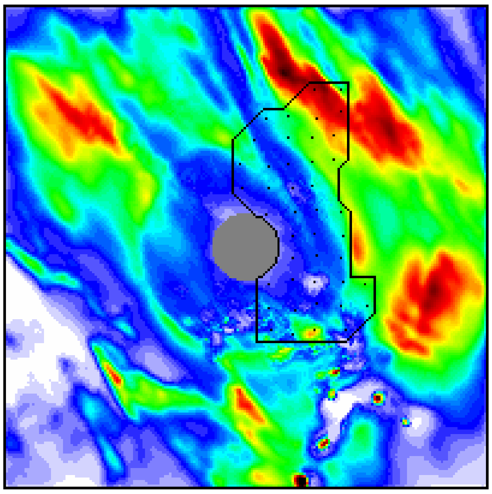

Figure 2.34 24 hour rainfall accumulation from the Bethlehem radar on the 13 February 1996. An outline of the test catchment is indicated on the image as well as the rain gauge locations. Dimensions of the image are 200km square.

Figure 2.35 illustrates a scatter plot of the radar and rain gauge areal average data for the 24 hour accumulation period of low intensity, wide spread, mostly stratiform rain, corresponding to Figure 2.34. The scatter plot indicates a strong correspondence between rain gauges and radar estimates with a correlation coefficient, r^2, of 0.78 being returned. The higher rainfall accumulation values show a strong correspondence, while the lower rainfall values exhibit a weaker relationship. Again, the rain gauges underestimate the highest two values of the rainfall when compared to the radar, inducing a bias because of the strong influence of the two largest values.

Table 2.8 provides a summary of the statistical analysis of the comparisons between rain gauges and radar accumulation for the full 24 hour period and two 12 hour periods. The mean and standard deviations for each of the accumulation periods closely correspond and are statistically similar, except for the mean value of the 24 hour accumulation. The correlation coefficient value, r^2, is also high except for the accumulation period from 12:00 to 18:00 which consisted of light stratiform rain. The cumulative distributions of accumulation values for each time period were also determined to be statistically similar via the K-S test.

The statistics of the comparatively high rainfall accumulations summarised in Table 2.7 and 2.8 indicate that the estimates of the rainfall at the ground from the radar images satisfactorily capture the amount and variability of the rainfall. The conclusion is that the methodology suggested in this report is operationally sound.

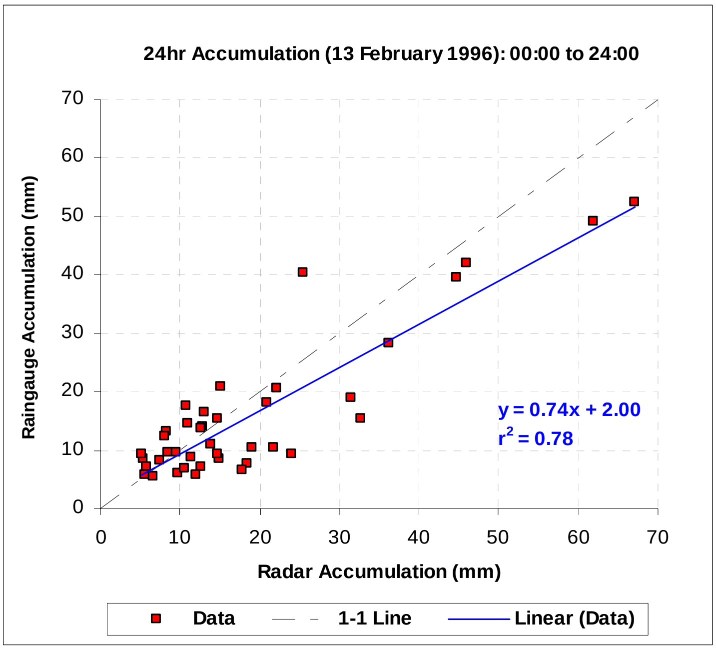

Figure 2.35 Scatter plot of radar and rain gauge accumulations for a 24-hour accumulation period for the 13 February 1996 rain event. A strong correspondence exists between the radar and rain gauge accumulation, especially for the high intensity rainfall.

Table 2.8: Summary of statistical results for different accumulation periods for 13 February 1996 rain event for radar and rain gauge data.

	24hr accumulation (12:00 to 24:00)		12 hr accumulation (12:00 to 18:00)		12 hr accumulation (18:00 to 24:00)	
	Mean	**Stdev**	**Mean**	**Stdev**	**Mean**	**Stdev**
Radar	18.1	14.7	1.1	1.4	17.0	14.5
Rain gauge	15.8	12.2	1.0	1.5	14.8	11.9
Accept / Reject H_0	Reject	Accept	Accept	Accept	Accept	Accept
r^2	0.78		0.14		0.76	
K-S TEST	Accept		Accept		Accept	

2.6 Summary

A technique has been presented that extrapolates radar volume scan data estimated at 1km intervals above ground level to the earth's surface, and in the process estimates missing data in the CAPPI stack.

A rainfall classification was firstly done to separate the rainfall into stratiform and convective types. By separating the rainfall, climatological semivariograms for each rainfall type were defined, a procedure which ensures high computational efficiency with little loss in accuracy. The 2km CAPPI level is then examined to determine if a bright band is present and if so the pixels are adjusted on an individual basis. Both 3D Universal and Ordinary Cascade Kriging were then used to infill missing data in the CAPPI stack and finally provide a rainfall estimate at ground level. In the Kriging computations care was taken to ensure computational efficiency and stability.

The effectiveness of the above algorithm was tested on two different rainfall events exhibiting distinctly different types of rainfall by comparing rain gauge and radar accumulation estimates at ground level. The technique demonstrated that both and high and moderate intensity rainfall exhibited similar estimates to rain gauges whereas stratiform rainfall showed a weaker correspondence to the rain gauge estimates; the research is ongoing.

3 Improving Satellite rainfall estimation

The satellite rainfall product is a key ingredient in the final merged rainfall estimate. In this section of the report we begin by outlining the MSRR algorithm detailed in Deyzel et al. (2004) used in the current SIMAR product to produce rainfall estimates from Meteosat-7 data. Following this is a description of the data reception mechanisms for Meteosat-8 data (EUMETCast). Finally examples of rainfall estimation using the MSRR algorithm as adapted to MET-8 are presented.

3.1 Introduction to satellite rainfall estimation

Satellite-based rainfall-estimations provide a cost-effective data source over temporal and spatial scales, not possible from any other in-situ or remote sensing system (Rozumalski, 2000). An important inherent advantage of remote sensing is that the view from space provides a means of extracting data from areas on the Earth's surface, regardless of country boundaries and inhospitable conditions. Unlike ground observations, which have varying accuracy, space borne sensors are a single source of data and therefore the errors of sampling and calibration are maintained at a known uniform level for every pixel (Cresswell, 1998). While space borne sensors cannot actually observe the precipitation reaching the Earth's surface, they can measure other variables that highly correlate with surface rainfall, such as cloud top temperature (Ebert and Le Marshall, 1995). Unfortunately this is also an inherent disadvantage in using satellite data, in that precipitation is only indirectly inferred from cloud observations as observed from space (Kurino, 1997). Non-precipitating high clouds may therefore obscure precipitating low clouds and the underlying dynamics of the precipitation process, thus leading to erroneous estimation of the precipitation. In the case of baroclinic systems the cloud canopy identified as precipitating may not necessarily correspond to the actual spatial distribution of the rainfall at the surface. These shortcomings must be kept in mind whilst developing a satellite rainfall estimation technique.

During the past three decades, there have been numerous attempts to utilize satellite measurements for precipitation estimations (Csiszar et al., 1997). One of the earliest satellite rainfall algorithms (Arkin, 1979; Arkin and Meisner, 1987) utilized the observed strong correlation between the frequency of cold cloud top temperatures below a regional-specific threshold temperature and rainfall rates observed at the surface in the tropics. For convective cells a correlation exists between the area occupied by a storm, its height, its life duration and the volume rain rate (Sauvageout, 1992). Cloud top temperature as inferred from infrared radiances relates directly to convective cloud height (Scherer and Hudlow, 1971, Scofield, 1987) and can therefore be used to estimate the precipitation falling at the surface. It has been found (Lovejoy and Austin, 1979; Tsonis and Isaac, 1985; Bellon and Austin, 1986) that combining visible (VIS) and infrared (IR) brightness temperatures, non-precipitating clouds can be screened out from the cloud top temperature field before applying the rainfall algorithm. However, due to the need to monitor precipitation at the highest possible temporal frequency, and in order to avoid deceptive biases in estimates of daily precipitation due to missing VIS data during nightfall (Porcu et al., 1999), it is commonly preferred to use IR data for its surpassed temporal resolution. In both the VIS and IR spectral bands, clouds are opaque and precipitation is inferred from cloud top structure, making it an indirect method at best.

3.1.1 Prelude to the development of a satellite rainfall estimation technique

The satellite rainfall estimation technique developed for the SIMAR project is a Multi Spectral Rain Rate (MSRR) estimation method. Active rain-cloud identification is accomplished by the discrimination of high-level cloud characteristics, associated with the rain/no-rain classes. Cloud textural patterns, intensity and high-level vapour content are used in this recognition process. The basis of this process is a simplified linear discriminant function, modelled on instantaneous radar spatial reflectivity fields. In addition to the satellite information, surface characteristic information is used for identifying possible warm orographic rainfall, which may be neglected by the satellite data alone.

The delay of the coming of age of the first user-operational data stream from Meteosat-8 (MET-8), necessitated the optimal utilization of the data sources from the then current Meteosat-7 system.

Data from all three channels are utilized in this switch-based rain area recognition rainfall estimation technique. However, only the IR channel data are used for the actual estimation of surface rainfall values. The IR information is separated into three classes, based on temperature thresholds, for improving quantitative rainfall estimation for cold convective clouds, middle layer clouds and warm coastal clouds.

3.1.2 METEOSAT-7 data sources

METEOSAT-7's primary mission is to observe the evolution of cloud systems. To do this it collects data from three different spectral channels at thirty minute intervals. The METEOSAT spin scan radiometer operates in three spectral bands:

0.5 - 0.9 μm (visible band - VIS)
5.7 - 7.1 μm (infra-red water vapour absorption band - WV)
10.5 - 12.5 μm (thermal infra-red band - IR)

During daylight hours ***visible imagery*** is recorded at half-hourly intervals, with one High Resolution Visible (HRV) measurement at 11:30 GMT daily. Although the VIS images have greater spatial resolution than IR images and are useful for visually identifying deep convection, their temporal availability is limited to daylight hours. IR imagery is preferred for precipitation estimation due to its continual temporal availability during 24-hours. Bi-spectral combinations of the IR and VIS channels for rainfall area masking are acceptable (Lovejoy and Austin, 1979; Tsonis and Isaac, 1985; Bellon and Austin, 1986) and form the foundation for the cloud classification scheme for SIMAR.

The amount of radiation absorbed by ***water vapour*** is dependent on the amount of moisture in the radiation's path and the wavelength of the radiation. Increased amounts of moisture, or water content, in the radiation's path lead to more absorption of the radiation emitted from lower layers. Therefore, if the air temperature decreases with height, higher moisture content results in colder brightness temperature. On a 6.7μm image the coldest temperatures correspond to high cloud tops, whilst the warmest are observed over lower altitude areas when the air is very dry through a deep layer in the atmosphere. For the 6.7μm water vapour channel, the radiation values may also be converted to brightness temperatures. A difference exists between WV (6.7μm) brightness temperature and that of the standard IR (11μm) channel. This is attributed to the absorption and re-radiation by water vapour above the earth's surface or clouds. It is this difference that allows a distinction to be drawn between cirrus and moist updraft regions.

Unfortunately, WV imagery is not available at the same temporal frequency as the 11μm IR data. It is available at half-hourly frequencies during late evening and early morning (22:00 - 05:30 UTC), hourly for early evening (19:00 - 21:30) and three hourly during the course of the day (06:00-18:00). The load on the data flow decreases during the evening when the demand for data decreases, thereby allowing WV to be subsequently available at the highest sampling frequency.

A statistical relationship exists between the ***IR band*** measurement of cloud top temperature and surface rainfall (Arkin, 1979; Arkin and Meisner, 1987) and therefore analysis of cloud top temperature is used to estimate the rainfall falling at the surface. The thermal infrared data from METEOSAT-7 channel 2 are received every half-hour. This consistent sampling frequency is another reason why IR data are preferred for rainfall estimation.

All these data sets are received with the Tecnavia satellite data processing system in Pretoria, with an approximate lag-time of 15-30 minutes from the end of the scanning time. Since the satellite is not fixed absolutely in its orbit relative to the Earth the image needs to be standardized by the EUMETSAT ground station in Darmstadt before relaying it to the user via MET-7, resulting in a delay.

3.2 The Multi-Spectral Rain Rate algorithm

The two-stage process is: (i) to Mask out non-raining information and then (ii) produce the rainfall estimates from the IR map. The main difficulties are associated with cirrus clouds (high, cold and non-raining) and warm rain at the coast, due mainly to orographic forcing. The 'nuisance' data are thus cirrus clouds and streaks and speckles from sun activity and image enhancing operations, so image texture and enhancement techniques are used in developing the final satellite-rainfall product.

3.2.1 Masking the IR field to identify clouds

Various single channel infrared techniques were investigated for estimating surface rainfall for the Southern African region using geostationary MET-7 satellite data. This was done during the period Oct. 2000 - May 2001. In doing this, special attention was given to the following techniques: GOES precipitation index (Arkin and Meisner, 1987), Negri-Adler-Wetzel Technique (Negri et al., 1984), NESDIS Auto Estimator (Vicente et al, 1998) and infra-red Power law Rain rate (IPR) technique of Goodman et al. (1993). Ebert and Le Marshall (1995) successfully implemented and evaluated the GPI, IPR and NAWT technique for rainfall estimation in Australia. They found these techniques to perform reasonably well with good correspondence of their individual results with real observations.

Further investigation led to the development of the infrared Threshold Rainfall (ITR) technique, which is a fusion of the above-mentioned techniques. It is widely accepted that a specific satellite rainfall technique is not necessarily transferable to another climatic region (Adler and Negri, 1988), since dynamic rain processes differ from region to region. This was overcome by combining existing techniques. Evaluation studies indicated that the technique is reasonably well suited to the South African climate regimes, as a first guess field.

The ITR is an indirect infrared temperature based method. Precipitating areas are delineated by assuming all clouds colder than 253K to be active convection. This threshold signals a compromise between extensive overestimation and less significant underestimation (neglecting warm coastal rain) of spatial rainfall fields. It was found that even increasing the threshold to 263K led to extensive overestimation of rainfall, unless a more sophisticated cloud identification technique is used. The 253K threshold has been found to be optimal for cold cloud filtering, i.e. cold convective and frontal systems producing most of the rainfall in our region. Also, Negri et al.(1984), Goodman et al.(1993) and Porcu et al.(1999) found this threshold to be useful for identifying convective rainfall in the tropics.

The cloud mask is set up as follows: $IR_{mask} = 1$ if $IR \leq 253K$; $IR_{mask} = 0$ if $IR > 253K$

The Infrared Power Law Rain rate (IPR) technique of Goodman et al.(1993) is applied to the residual cloud filtered pixels, for estimating a half hourly rain rate. The IPR technique was originally developed for the Amazonian basin (Goodman et al., 1993) and was used to estimate three-hourly rainfall rates in the tropics by applying a power law relationship of surface rainfall and IR temperatures on the delineated precipitating pixels. This power law relationship was derived from a statistical non-linear regression of co-located surface rain gauges and 11.2 μm infrared cloud top temperatures from GOES imagery. This technique estimates rain-rate as a continuous function of cloud brightness temperature and can therefore reproduce the scale of features observed in the satellite imagery (Ebert and Le Marshall, 1995). By adapting the offset values it was found that the IPR technique led to promising and acceptable rainfall estimations for the South African region at half-hourly intervals.

Evaluation of the ITR technique on daily, monthly and seasonal rain gauge data indicated several limitations of this technique. These include:

- Warm coastal and orographic rain tends to go undetected. This is fundamentally due to the 253K threshold.
- Overestimation of spatial extent of convective rainfall as a result of cirrus contamination.
- Underestimation of severe convective rainfall. This indirect method cannot intelligently identify the true extent of convective dynamics.
- General underestimation of frontal rain systems. Again the method cannot identify the underlying dynamics relating to the rainfall processes.

This technique suffers from the inherent limitations of a single-channel IR-based technique. The main reason for the errors in an IR-based rainfall estimation scheme can be attributed to the fundamental assumption that precipitation at the surface is a function of cloud-top temperature. If this assumption is violated, errors in rainfall detection occur. Even if the clouds are correctly identified as precipitating, the estimated rain amounts may contain a large amount of random error, due to the different dynamical processes occurring inside the clouds.

These limitations warranted the redesign of the ITR satellite rainfall technique, leading to the development of the MSRR technique for implementation in SIMAR.

Infra-red Water Vapour Spectral Mask to refine the mask
Szejwach (1982) suggested a bi-spectral method to infer the cirrus cloud-top temperature. This technique relied on 6.5μm water vapour and 11.5μm standard infrared channel data. A bi-dimensional histogram, frequency analysis of grey level adjacency occurrence and determining the principal axis of the histogram formed part of the technique for identifying cirrus cloud. This technique indicates the use of a bi-spectral method in identifying cirrus. Inoue (1987) showed that it was feasible using a split window channel of the IR data, e.g. 11μm and 12μm for cirrus detection. Kurino (1996) applied this technique and found that an 11μm and 12μm spectral difference of larger than 3°K corresponded well with cirrus. In the use of the 6.7μm WV data, Kurino (1996) found that the areas with an 11μm and 6.7μm spectral difference less than or equal to 0°K corresponded to deep convection. The same principle is applied for the IR-WV MSRR Mask devised for SIMAR .

A definite contrast was observed along the coastal regions and NE inland regions of South Africa. In tropical convective systems associated with a surface and upper air trough, it was found that the mask threshold had to be lowered to counteract the spatial overestimation of rainfall over coastal and ocean areas. A higher threshold is used over the inland regions, the height taken from a Digital Elevation Model (DEM):

if WV - IR < -3°K and DEM-based elevation > 1000 m } IRmask = 1

if WV - IR < -5°K and DEM-based elevation < 1000 m } IRmask = 1

Convolution
Image processing suppresses information that is not relevant to the specific image processing or analysis task, enhances the image content and assists in identifying certain spatial features (e.g. Sonka et al., 1993). Processing methods used here are based on the Gaussian spatial convolution of a small neighborhood centered on a pixel, which produces a new brightness value in the output image by deriving the center value from some function of the neighborhood values. The Gaussian kernel has the effect of a weak high pass filter that enhances the features within the image and removes speckles which are considerably smaller than the kernel size.

Texture Analysis
A common characteristic of rainfield images is that neighbouring pixels are highly correlated (Csiszar et al., 1997). IR images of 11μm are no different and contain vast amounts of information on various spatial scales of cloud and surface features. One means of classifying a pattern is to use its texture characteristics (Pican et al., 1998). Texture is one of the most important defining characteristics of an image (Dulyyakarn et al., 2000). It is characterized by the spatial distribution of grey levels in a neighbourhood (Jain et al., 1985). For infrared images the classification problem is one of extracting (or discriminating between) textural features of clouds that are associated with significant surface precipitation.

Textures may be random, but with certain consistent properties, and can be described by their statistical properties. These include the mean intensity, variance, skewness and kurtosis computed from 1-D histograms (Morse, 2000). In order to capture the spatial dependence of grey-level values, which contribute to the perception of texture, some well know techniques exist, e.g. Markov Random Fields, Grey Level Co-occurrence Matrix (GLCM), self-organizing maps, fractal components and 2-dimensional FFT (Pecan et al., 1998, Morse, 2000, Dulyyakarn et al., 2000, Chan et al., 2000). The choice of a suitable method depends on the constraints of the application in terms of the nature of the texture and computation time (Pican et al., 1998).

The GLCM is a well-known statistical technique for feature extraction (Dulyyakarn et al., 2000) and has been successfully applied for seabed classification (Pican et al., 1998), sea ice identification (Kaleschke and Bochert, undated) and has been used widely for land cover classification (Dulyyakarn et al., 2000, Chan et al., 2000, Smits and Annoni, 1999). GLCM statistically samples the way certain grey-levels occur in relation to other grey-levels (Morse, 2000) and is based on the repeated occurrence of a given grey-level configuration. Haralick (1979) defined the GLCM as the relative number of occurrences of a pair of grey level values within a given pixel distance and in a given direction. The GLCM provides quantitative descriptions of the spatial and attribute relationship between cells within a gridded framework (Wood, 1996).

For the images processed in SIMAR, the decoded values (IR or VIS) are scaled and coded in grey-level before calculating their GLCMs. For IR images SIMAR research has indicated that a distance of 3-4 pixels is sufficient to retain most of the cloud characteristics useful for the feature classification. This relates to a neighbourhood matrix of approximately 9x9 pixels. This kernel approach to solving the GLCM improves the computational efficiency of the technique dramatically so that it can be useful for real-time processing; the standard method of calculating the GLCM on the whole data domain makes it far too cumbersome for real-time application.

For the classification of the cloud textural features in the rain and no-rain classes, a Linear Discriminant Analysis (LDA) was performed to produce a general Discriminant Function (DF) based on the GLCM texture images. This is a simplified approach, in that it is a two-class identification problem. A different LDA was performed for each data set, be they IR or VIS, used for texture analysis. In SIMAR, The LDA was trained on matching instantaneous radar reflectivity data and IR or VIS satellite images, pixel-wise. The radar data was preprocessed for quality purposes, removing ground clutter or other contamination and averaging the radar data to a coarser grid corresponding with that of the satellite. This was done to remove any biases, which may negatively affect the LDA coefficients. The resulting DFs for IR and VIS are given by:

$$IR_{(tex)} > -1.495(-115-IR) \} \ IRmask=1$$
$$VIS_{(tex)} < 6.08(VIS-42) \ \} \ IRmask=1$$

The algorithm to produce the rain/no-rain mask is summarised in Table 3.1 and proceeds in more detail as follows, depending on the availability of the IR, WV & VIS images:

IR & WV available (column 3 of Table 3.1)

1. Infrared Water Vapour Spectral Mask is produced, by subtracting the IR brightness temperature from the WV field. The potential of using WV data is indicated by the observation of relatively warm water vapour pixels over deep convective clouds (Tjemkes et al., 1997). The negative response of this spectral field relates to the high, moist convective cloud tops, which should produce the highest intensity of surface rainfall. This specifically masks the cold clouds (≤ 218K), where the WV information is more reliable in the vertical.

2. For the remainder of the cloud field (>218K) the textural features of the clouds, as derived from the GLCM, are used to distinguish between the rain and no-rain classes. A discriminant function (DF), trained with the aid of radar data, is used

for this classification process. The thermal IR is used as a baseline data set for the classification of cloud textures.

Table 3.1. Summary of thresholding and filtering activities to produce a “Possibly raining” mask, working down the columns

<table>
<tr><th>Data availability</th><th>IR</th><th>IR & WV</th><th>IR & VIS</th><th>IR, WV & VIS</th></tr>
<tr><td rowspan="2">Initial Screen</td><td>-</td><td>WV-IR<-3°
<-5° (coast)</td><td>-</td><td>WV-IR<-3°
<-5° (coast)</td></tr>
<tr><td colspan="2">-</td><td colspan="2">Sun Angle > Th’ld</td></tr>
<tr><td rowspan="3">Texture</td><td colspan="2">IR_{tex} GLCM > 180K</td><td colspan="2">VIS_{tex} GLCM > 180K</td></tr>
<tr><td colspan="2">IR Discriminant Fn
$IR_{(tex)} > -1.495(-115-IR)$</td><td colspan="2">VIS Discriminant Fn
$VIS_{(tex)} < 6.08(VIS-42)$</td></tr>
<tr><td colspan="2">-</td><td colspan="2">Qual Cont on IR & VIS</td></tr>
<tr><td rowspan="2">Filter</td><td>IR > Th’ld</td><td>WV>Th’ld</td><td>VIS warm</td><td>WV & VIS warm</td></tr>
<tr><td colspan="4">WAR Gaussian Speckle Filter</td></tr>
<tr><td colspan="5">IR mask: pass 1 – Go To Rain-Rate Estimation</td></tr>
</table>

VIS, IR & WV available (column 5 of Table 3.1)

1. During the day, when the sun angle is sufficient and the albedo data reliable, VIS images are preferred over IR. VIS data are obtained at a higher resolution than IR and the albedo relates to the optical depth of a cloud, which is a more direct measure of the clouds’ efficiency in producing rain. These allow the VIS images to produce higher quality texture images than IR. For warm orographic rainfall the results of King et al. (1995) showed that incorporating VIS data with IR produced higher correlation with ground truth data, than using IR data alone; for cold, bright thunderstorm clouds they found the correlation to be similar. This is why VIS data are utilised where possible. A unique DF is used for each data set, IR and VIS. When VIS is available the VIS texture mask is used to mask the IR field. This mask is used for rainfall estimation by empirical relationships between cloud top temperature and surface rainfall.
2. Since the DF is a simplified linear approximation to the classification process, post filtering must be applied to improve the resultant mask. The post-processing of the images entails two stages. Firstly, empirically derived WV and VIS thresholds are used to ensure that non-precipitating clouds, that is clouds with low mid/high atmospheric moisture and low albedo (low droplet concentrations), which have been incorrectly classified by the DF, are filtered.
3. Secondly, a convolution with a Wetted Area Ratio (WAR) Gaussian speckle filter diminishes the effect of small-scale rainfall areas. These outliers in the mask are either due to edges being classified as cloud areas or as a residual of the DF classification routine.
4. These produce the IR rainfall area mask which is used for rainfall estimation (Step 2).

IR and/or VIS available (columns 2 & 4 of Table 3.1)

1. Similar to the above, with the difference that the whole range of the IR or VIS field is used for the GLCM texture classification, and not only the mid to warm range.
2. The post-processing filter is similar. An empirically derived VIS threshold filter removes possible false alarms that have been incorrectly classified.
3. A post-processing WAR filter is applied, similar to the above after checking the threshold
4. These steps produce the masked IR rainfall area mask. This mask is used for the rainfall estimation.

The rainfall algorithm using the masked IR field

There are three temperature zones, each with its own equation, as shown in Figure 3.1.

$\boxed{IR \leq 218K}$

1. Identify possible cold convective cloud tops with temperatures below the empirically derived threshold. It is assumed that the cold rainfall areas classified by the WV and GLCM DF of the Masking process are associated with convection.

2. This cloud type produces the highest rainfall rates over SA and the Deep Convective Activity (DCA) index is used for estimating this type of rainfall. The DCA index was developed by Hendon and Woodberry (1993), based on the work of Fu et al., (1990) and is given by: $R_{hcs} = \alpha(230-T_B)$ where α = 0.45 mm^{-h}K^{-1} and T_B is the infrared brightness temperature. The coefficient α was adapted to produce more reliable convective rainfall for South African storms

3. The DCA index produces Rhcs, the cold portion of the half-hourly rainfall field.

$\boxed{218K < IR \leq 267K}$

1. Identify possible precipitating middle clouds, such as Nimbostratus or Altostratus. Rain from these types of clouds tends to persist at a low rate and for extended periods of time. A low rainfall rate generating technique is used for this class. The IPR technique (Goodman et al., 1993) was used as the rainfall estimation scheme with a modified temperature threshold of 253K. The offset value of the IPR technique was removed, resulting in increased sensitivity to warmer clouds. By doing so the IR threshold was increased to 267K. The performance of this technique is dependent on the correct delineation of raining clouds. The Adapted IPR (AIPR) estimations performed best for the mid range (218 - 267K) values in IR images. These values relate to moderate and fairly warm convection, stratiform middle cloud (Nimbostratus) and even significant low-level cloud (Stratocumulus). The AIPR index is: $R_{hms} = b(267-T_B)^{1.85}$, where b = 0.00303 and T_B is the IR brightness temperature. The original coefficients, developed for Amazonia (Goodman et al., 1993), have been adapted to suit the South African climate region, but the empirical nature of estimating rainfall was retained.

2. The AIPR technique handles this type of rainfall and produces Rhms, the middle range portion of the half-hourly rainfall field.

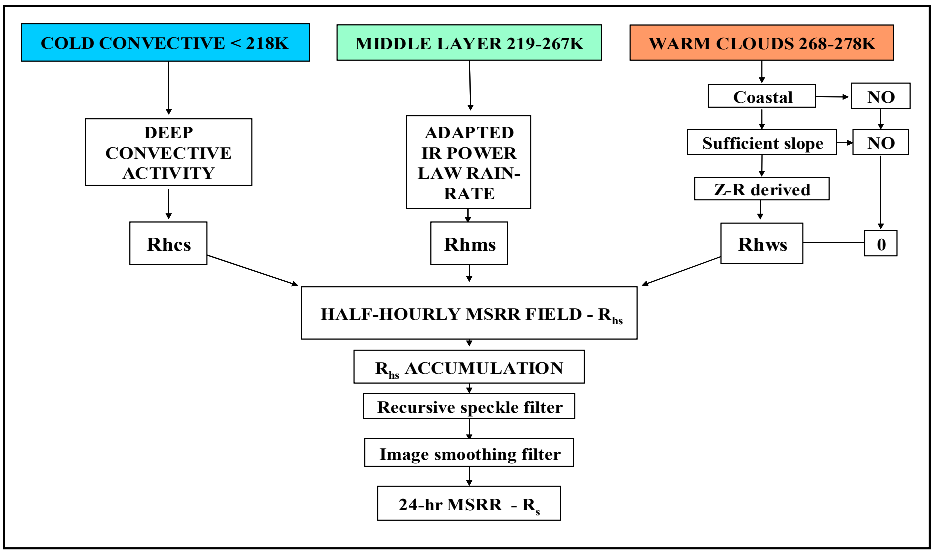

Figure 3.1 AIPR algorithm flow-chart for estimating rainfall from IR inside the 'possibly raining' mask produced from Table 3.1.

$267 < IR \leq 278K$

1. This algorithm aims to identify the rainfall from warm clouds, such as Stratocumulus.
2. Topographic information from a 1km DEM is used to constrain this index to the coastal areas, where these clouds usually produce rainfall.
3. The slope of the terrain surface features is utilized for finding possible orographic rainfall along this coastal plain.
4. A rainfall estimation technique, Warm Stratiform Rain Rate (WSRR), was developed specifically for the coastal zones with the aid of radar reflectivity fields. Investigation of the instantaneous IR, WV, radar and cloud data for a 3-month evaluation period, showed that significant amounts of rainfall originated from extremely warm stratiform clouds along the low lying regions of SA, especially along the Western Cape. The temperature of these clouds was found to be in the range 267 to 278K. The AIPR technique is not sensitive enough to identify such warm clouds. This warranted another technique to be used, which was developed locally due to the low thresholds that had to be used. The WSRR estimation technique was developed by linear regression of co-located IR and radar reflectivity values along the coastal region of South Africa for Spring/Summer 2001. This was done to provide some kind of memory of varying conditions (regional and synoptic related) that could produce warm rainfall. It was hoped that this technique would serve as a useful estimation method for warm rainfall, along the SA coastal region.

   The WSRR estimation formula is given by: $R_{hws} = [10\text{^}(73.32-0.1726T_B)/2000]\text{^}0.625$. Low lying areas are considered to be those below 1000 m above sea level. As an additional screen, all surface areas below a threshold elevation slope of 6 degrees are filtered out. This is done to counteract any spatial overestimation that may occur and to incorporate possible orographic effects. This Z-R derived technique produces Rhws, the portion of the warm rain in the half-hourly rainfall field.

Finally, all ½-hour rainfields are accumulated to produce the 24-hour satellite estimation field, S_R.

3.3 Meteosat-8 reception station

Radiance data collected in 12 spectral channels by the Spinning Enhanced Visible and Infrared Imager (SEVERI) instrument onboard MET-8 are transmitted from the satellite to the EUMETSAT ground station in Darmstadt, Germany. The data are processed to produce a rectified and calibrated 10-bit image for each of the 12 channels. Transmission of the data to users is achieved via standard Digital Video Broadcast (DVB) technology; the same system which allows many of us to view satellite television in our homes. Figure 3.2 gives an overview of the broadcast routing. The data products are transmitted to the HotBird-6 satellite from EUMETSAT's uplink station in Usigen, Germany; the HotBird-6 satellite re-broadcasts the data at a rate of 2 MBit/s at a frequency of 10.853 GHz in the Ku-band. European users receive the data using small (50cm) Ku-band antennas. The data stream is received in Fucino, Italy and uplinked to the AtlanticBird-3 satellite from where it is broadcast to African users at a rate of 2 MBit/s and frequency of 3.7317 GHz in the C-band. Figure 3.3 shows the C-band coverage over Africa. African users therefore need a larger (2.5m) C-band antenna to receive the data (Figure 3.4).

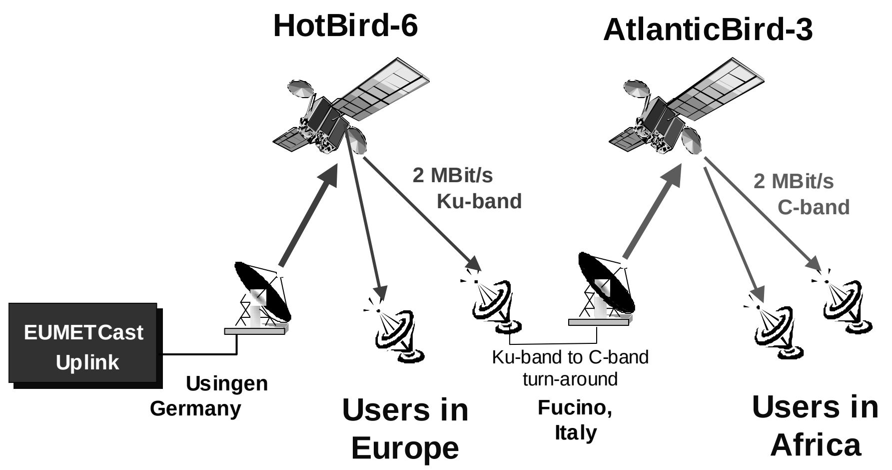

Figure 3.2 Overview of C-band data dissemination (after EUMETSAT, 2005).

In addition to the antenna, the receiving station consists of two PC's. One is the receiver PC and the second is a processing PC. In the configuration at UKZN, the second machine also doubles as a data server. The receiver PC is connected to the antenna by a DVB card and coaxial cable. While the data is being received by the DVB card it must be decrypted using the EUMETCast client software and Encryption Key Unit (EKU) in real time. Due to the high data volumes (1.6 MBit/s) this machine is reserved for decoding the data stream. The second machine (processing server) is connected by a direct GBit link to the receiving machine and runs software which manages the raw data files. The raw data are processed into 10-bit PNG images and archived in a three day rolling buffer. All further data processing is carried out using the 10-bit PNG image format. The processing server has a second network card installed and is accessible, to selected users, via the local network.

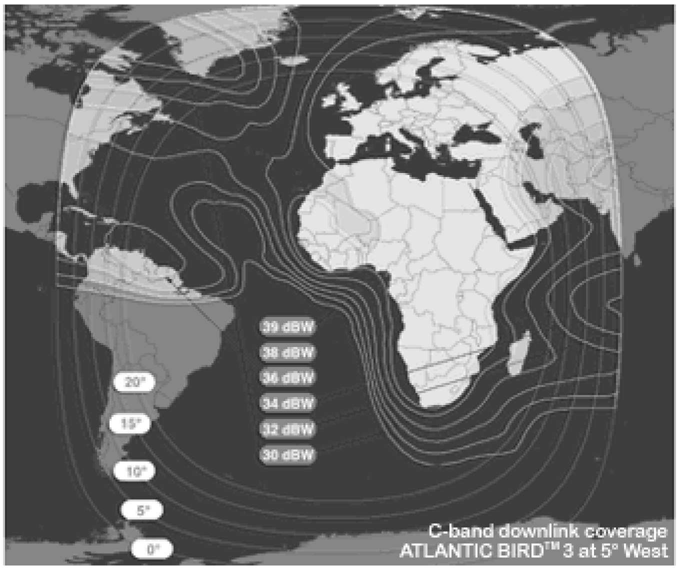

Figure 3.3 C-band data coverage from the Atlantic Bird satellite (after EUMETSAT, 2005).

Figure 3.4 The Meteosat-8 receiving station at UKZN. The left-hand panel shows the installed antenna on the roof of the Durban Centenary Building. The right-hand panel shows the receiving PC and Data Server.

3.4 Rainfall estimates using Meteosat-8

The MSRR algorithm described in section 3.2 resulted from a significant research effort which took place over a period of several years. Due to the long delay in MET-8 data becoming available, the current MET-8 rainfall product is based on the MSRR algorithm. Three spectral channels with wavelength bands similar to the existing MET-7 channels are used as input to the algorithm. The main improvement at this stage is in the increased spatial and temporal sampling frequency provided by MET-8. This is strikingly obvious when comparing the “before and after” images of daily accumulations in Figs. 3.5 and 3.6 which were processed at UKZN, without the added benefit of the new accumulation algorithm discussed in the next section. Note the different colour scales in the two images.

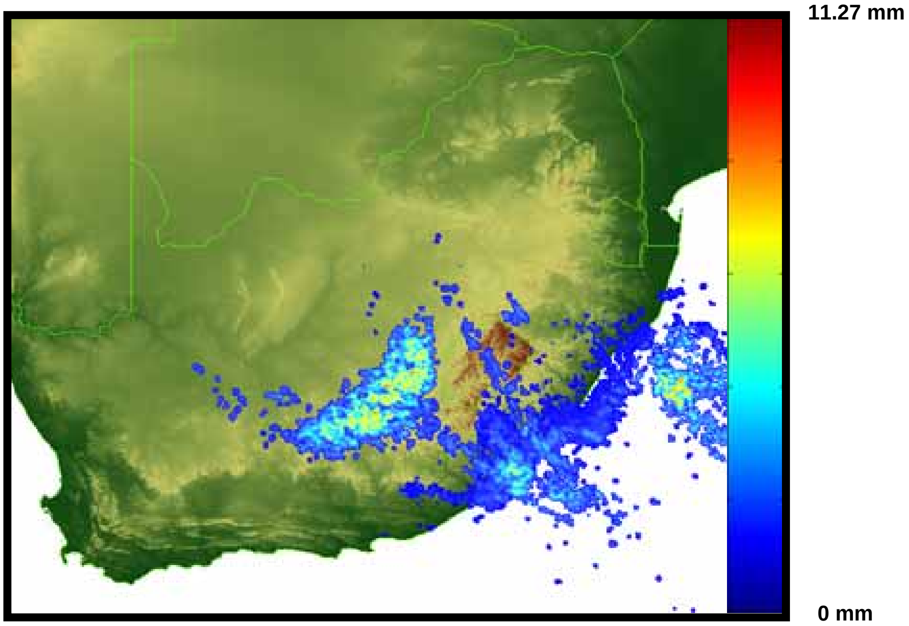

Figure 3.5 Daily rainfall accumulation. MET-7 06:00 GMT 22/06/2005.

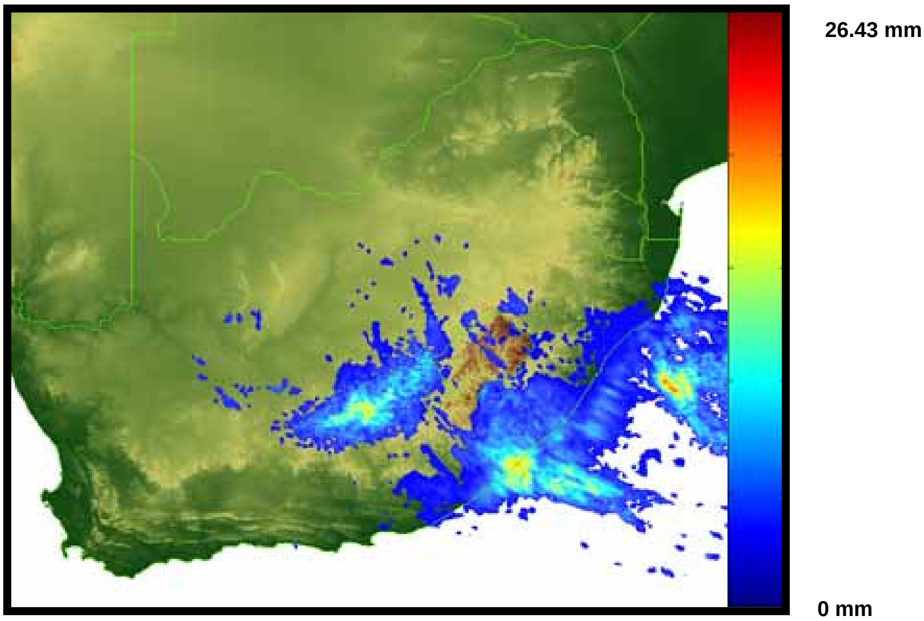

Figure 3.6 Daily rainfall accumulation. MET-8 06:00 GMT 22/06/2005.

4 Accumulating Spatial Rainfall

Rain gauges record accumulated rainfall depth and might be manually read or logging gauges. The accumulated totals are recorded either as the total depth in a given time period or as the times over which a given (small) depth is recorded, allowing rainfall rates to be estimated. Rainfall information derived from Radar and Satellite sources are instantaneous samples over large areas, from which the spatial rainfall rate is estimated.

Temporally accumulated rainfall is the basic input for hydrological models and water resource management applications while many operational spatial rainfall estimation schemes work with accumulated rainfall as a basis (e.g. Smith and Krajewski, 1991;

Seo 1998a; Seo 1998b; Anagnostou and Krajewski 1999a; Anagnostou and Krajewski 1999b).

The temporal accumulation of rainfall estimates needs to be handled carefully, in particular, the instantaneous fields from radar and satellite. A first possibility is a superimposition technique where the rainfall rates from each radar scan are appropriately scaled and added together. The method used by Fulton et al. (1998) for the NEXRAD radar network is to take an average rainfall rate between consecutive radar scans and apply this rate over the time interval between scans. The accumulations between scans are then summed over the required time-period. A more complex method is described by Hannesen (2002), who suggests a path integral accumulation scheme based on a field of derived motion vectors. Hannesen's (2002) scheme turns out to be in the same spirit as a technique developed by (Anagnostou and Krajewski, 1999a), with the difference being that Anagnostou and Krajewski (1999a) view the accumulation as a time integral which they approximate with a summation at fixed time-steps. The decision regarding which algorithm to adopt for operational purposes is determined by the trade off between computational speed and the value of information lost due to the sampling interval inadequately capturing changes in the rainfall intensity field. The accumulation method presented here uses the length of computed advection vectors to choose the required computational complexity.

4.1 Algorithm description

The accumulation technique developed here accounts for the spatial and temporal evolution of the rainfall field between successive radar (or satellite) scans. The temporal sampling frequency of these instruments is not always short enough to ensure that important evolutionary features of the field are adequately captured. The advection of the field between successive scans is computed using an optical flow (OF) algorithm developed by Bab-Hadiashar et al. (1996) and introduced in the context of short term rainfall forecasting by Seed (2001). The optical flow computation (as implemented here) produces a dense field of advection vectors each of which is associated with a particular pixel. Figure 4.1 shows a portion of the optical flow field computed between two consecutive rainfall intensity images from the SAWS weather radar in Durban, South Africa, during November 2000. The left hand side of Figure 4.1, shows two superimposed radar scans taken five minutes apart. The first scan is in greyscale while the second is in colour. A portion of the field of advection vectors, computed using an optical flow algorithm, is shown on the right of the figure. The complex nature of the advection is evident from the advection vectors shown.

The total depth of rainfall accumulating on any pixel between scan times is computed from Equation 4.1, over the path defined by the advection vector. The individual depths can then be summed to produce accumulations over longer timescales, depending on what is required.

$$A_i = \frac{1}{V_i \Delta s}\left[\int_{S_0}^{i}\left(q - S_0\right)R_0(s)dq + \int_{S_1}^{i}\left(r - S_1\right)R_1(s)dr\right] \tag{4.1}$$

The terms in Equation 1 represent the following quantities; A_i is the accumulated rainfall depth at pixel i between successive scans, V_i is the magnitude of the velocity vector associated with pixel i, R_0 and R_1 are the initial and final rainfall intensity fields, S_0 is the position that pixel i would have occupied on R_0 at time t_0 and S_1 is the position that pixel i would have occupied on R_1 at time t_0.

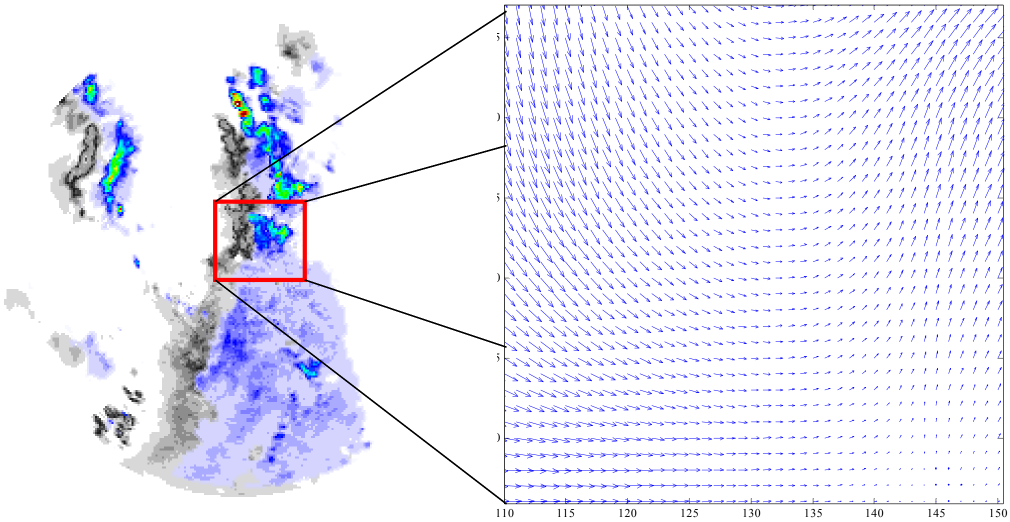

Figure 4.1 Advection field computed using the optical flow algorithm of Bab-Hadiashar et al. (1996).

The integrals are computed using a discrete Trapezoidal rule approximation. The computational effort of the integration is directly related to the number of interior points evaluated to compute the integral between consecutive scans, therefore the number of interior points used is computed as the ratio between the length of the advection vector and the spatial resolution of the data. Thus for a zero advection case the accumulation reduces to a simple average between the intensity values for successive scans and for large values of advection, many interior points are chosen with the same spatial resolution as the data.

Figure 4.2 shows a comparison between several accumulated daily rainfall fields using the operational simple averaging technique (top row) and the accumulations made using the OF based technique (bottom row). It is clear from qualitative observations that the precipitation swaths produced by the OF accumulation are smoother and more in keeping with what one would expect real precipitation swaths to look like.

4.2 Algorithm Performance

The performance of the newly proposed accumulation scheme was compared to the current operational scheme by generating an entire years worth of daily accumulations using each method and computing the total time taken to complete the computations in each case. This was done so that the performance of the algorithm could be evaluated for a wide variety of weather types. Figures 4.3 and 4.5 summarise the performance of the algorithm relative to the operational SIMAR implementation of the accumulation scheme. Figure 4.3 shows the total time taken to complete the daily accumulations for an entire year, for a number of interior points ranging between 0 and 8 as well as for the operational algorithm and scaled routine described in section 4.1.

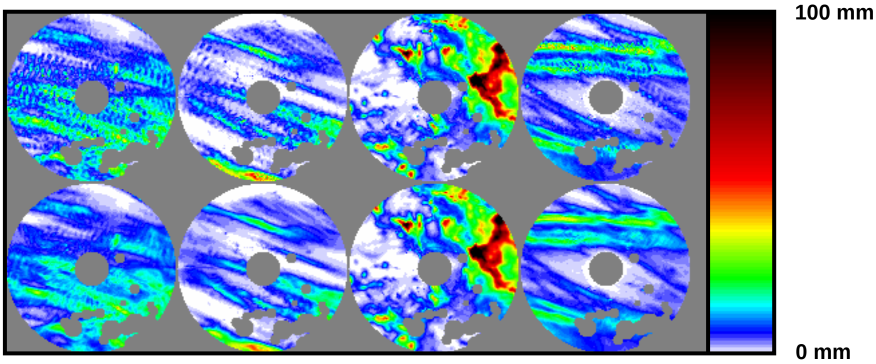

Figure 4.2 Comparison of accumulated daily rainfall fields, produced by the two different methods.

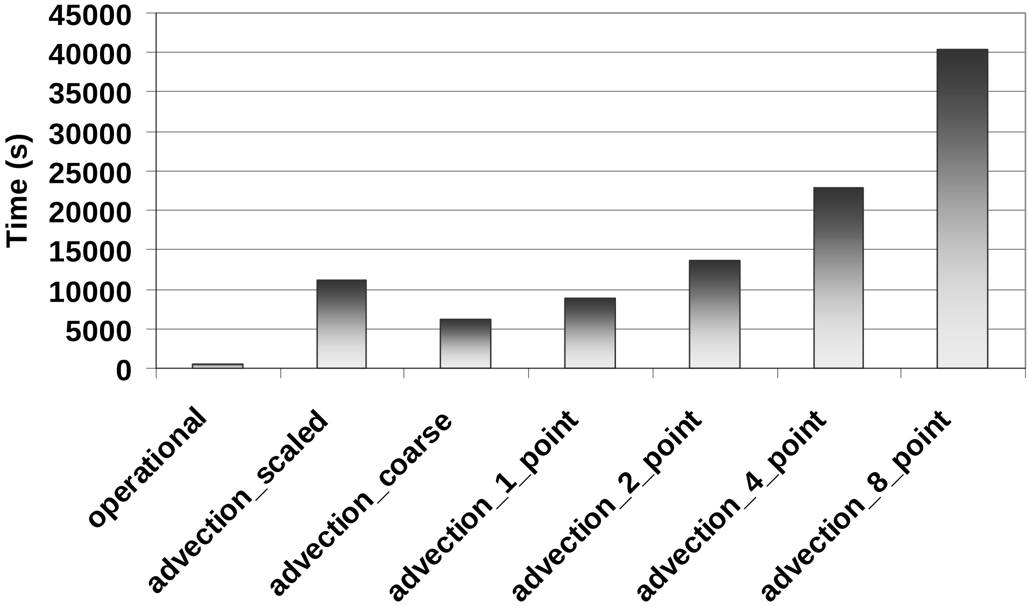

Figure 4.3 Comparison of timings for a variety of accumulation schemes. The times represent the total time taken to perform one year's worth of daily accumulations for the MRL5 radar in 1999.

Considerable effort was put into optimising the implementation of the advection based accumulation algorithms. The performance of the original implementation shown in Figure 4.3 was improved significantly as a result. The performance of the optimised code is presented in Figure 4.4. Although the scaled advection algorithm is 18.4 times slower than the lightning fast operational algorithm, its performance is still perfectly reasonable for operational purposes. The average time (during 1999) taken on currently available computers for a daily accumulation (the accumulation of approximately 288 radar scans) is 24.5 seconds.

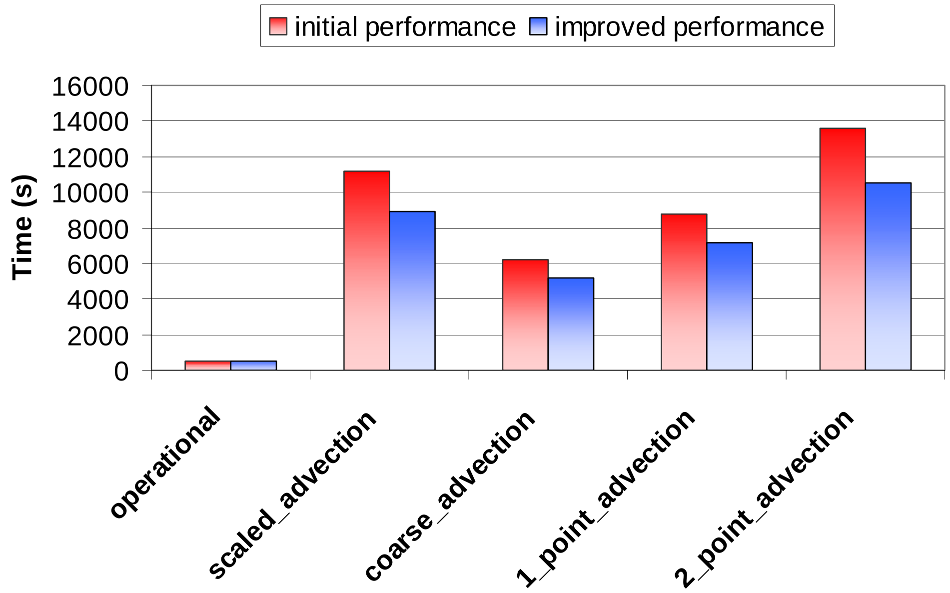

Figure 4.4 Changes in algorithm performance after careful optimisation of the code.

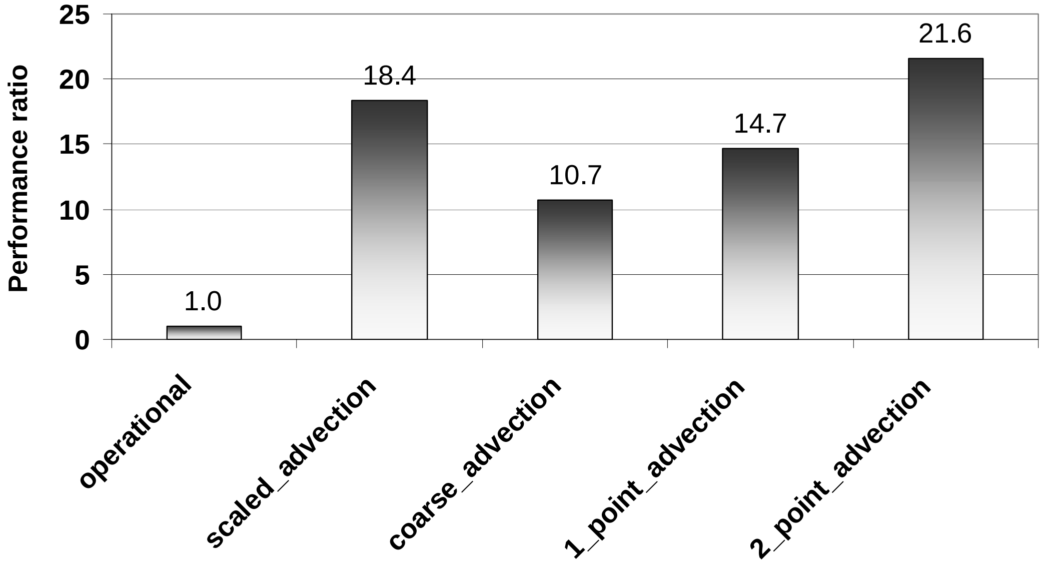

Figure 4.5 Performance of accumulation schemes relative to the operational algorithm.

5 Data fusion – Merging rainfall estimates

The Hydrologist's traditional tool for measuring rainfall is the rain gauge. Rain gauges are relatively cheap, easy to maintain and provide a direct and suitably accurate estimate of rainfall at a point. What rain gauges fail to capture well is the spatial variability of rainfall with time, an important aspect for the credible modelling of a catchment's response to rainfall. This spatial variability is particularly evident at short timescales of up to several days. As the period of accumulation increases the expected spatial variability is reduced and rain gauges provide improved spatial rainfall estimates. Due to the fractal variability of rainfall in space, simple interpolation between rain gauges does not provide an accurate estimate of the true spatial rainfall field, at short time scales.

Weather radar provides (with a single instrument) a highly detailed representation of the spatial structure and temporal evolution of rainfall over a large area. Estimated rainfall rates are derived indirectly from measurements of reflectivity and are therefore subject to a combination of systematic and random errors. Satellite measurements provide an additional data source, although these are also indirect measurements with precipitation rates being inferred from cloud top temperatures. Sections 2 and 3 of this report contain further discussion on rainfall estimation using these remote sensing instruments.

Several significant sources of complexity arise when attempting to optimally blend rainfall estimates from a variety of different sensors. In this section the major issues pertaining to this blending are discussed and a technique to estimate the true rainfall field is presented. The key issue in obtaining and assessing the accuracy of spatial rainfall estimates is that the true spatial rainfall field cannot be directly measured using current technologies.

5.1 Problem Definition

Consider modelling the rainfall rate as a spatially correlated random process $\left\{R(s): s \in D \subset \mathbb{R}^d\right\}$ at any instant in time, where $R(s)$ is the value of the rainfall rate at location s, and $d = 3$ (2 space and 1 time dimension) if we restrict ourselves to a spatial rainfall field at ground level. For hydrological applications we require an estimate of the areal average rainfall rate $R(A)$ over an area of interest $A \in D$, where

$$R(A) \equiv \begin{cases} \frac{1}{|A|} \int_A R(u)du & A > 0 \\ ave\left\{R(u): u \in A\right\} & A = 0 \end{cases}$$

The marginal distribution of rainfall rate can be considered Log-normal (Kedem et al., 1991; Pegram and Clothier, 1999) therefore it is possible to define $Z(s) = \ln\left(R(s)\right)$ where $Z(s)$ is a normally distributed random variable .If A_i is an element of a regular grid with *m* rows and *n* columns, filling D, then the problem becomes one of estimating each $Z(A_i)$ from the available observations.

It is at this point that the difficulties in estimating spatial rainfall become apparent, each of the three sensors we are considering have a different spatial support and temporal sampling frequency. To confound things further, the estimates of rainfall to be combined are obtained from sensors measuring independent quantities at different temporal and spatial scales. In order for credible rainfall estimation and validation to be done, equivalent quantities must be obtained. This implies that the estimates must first be transformed to common temporal and spatial scales, or that the merging procedure must account for these factors. It is assumed in this section that the best available estimates of rainfall rate have already been derived from the remotely sensed quantities.

Rainfall is variable in space as well as in time, making it important to capture the spatial structure (in addition to temporal information) to provide a truly representative estimate. Since each instrument samples the rainfall field at a different spatial resolution, blending the data requires consideration of the change-of-support (COS) problem (Cressie, 1991; Gotway and Young, 2002). Cressie (1991, pp 284), in his treatment of the issue, points out that spatial (and temporal) averaging reduces the variance of the resulting averaged data compared to that of the underlying point process. The magnitude of the cross-correlation between the aggregated variables is also increased, however the sample mean is unaffected. The different statistical properties of the processes should therefore be taken into account in the estimation of the unknown process at the spatial and temporal scale of interest. The differing statistical properties are a direct result of the scales of measurement.

The spatial estimation of rainfall by combining information from multiple sensors has received considerable attention in the Hydro-meteorological literature. Early work (e.g. Brandes, 1975) focussed on the correction of bias in radar estimates of rainfall using an adjustment factor. Krajewski (1987) suggests a Cokriging procedure that he demonstrates by numerical experiment. A procedure that accounts for the fractional coverage of rainfall (spatial intermittency) is suggested by Seo (1998a, 1998b) using conditional expectations. Seo (1998b) computes the expectation of rainfall at an ungauged site conditional on the observed rain gauge and radar data. A Bayesian merging technique suggested by Todini (2001) relies on a Kalman filtering scheme to remove the error variance (after a separate bias reduction) by taking a Block-Kriged spatial estimate based on rain gauges as the observation vector. The Block-Kriged estimate is introduced to handle the spatial sampling differences. The novel techniques developed during the SIMAR project (Kroese, 2004; Deyzel et al., 2004; Pegram, 2004) combine the rainfall estimates from three separate data sources using an explained variance weighting technique. Recent work which isn't directly related to Hydro-Meteorology is presented by Wikle and Berliner (2005) who suggest a hierarchical Bayesian framework for combining information at different spatial scales and present an example of computing the stream function based on satellite and NWP model based wind observations at different spatial resolutions.

The greatest challenge in implementing the merging techniques discussed here is estimating the structure of the error covariances between the observations and the true rainfall process. Optimal blending of the information also requires knowledge of the errors associated with each of the measurements. The covariance relationships between the estimates themselves may either be estimated from historical observations on the assumption of stationarity or estimated (and updated) online using a parameter updating procedure similar to that employed by Anagnostou and Krajewski (1999a, 1999b) for adjusting radar rainfall estimates based on real-time rain gauge observations. It is also possible that remote sensing based estimates be compared to interpolated information derived from rain gauge measurements as is done by Todini (2001). However, the true rainfall field remains unknown and Hydro-meteorologists must rely on the relationships between the observational data, to produce space-time rainfall estimates. In SIMAR the contributions from each sensor are interpolated onto a common spatial grid. The estimates are then weighted by expected information content based on the measured covariance structure in South African radar data for both radar and gauge observations and a skill score approach is used to determine the influence of the satellite observations (Deyzel et al., 2004). Given the previous discussion, this approach leaves room for improvement.

Rain gauges measure "point" rainfall at the ground. Remote sensing techniques provide rainfall estimates based on information above the ground at the cloud top level (satellite) or at several levels above the surface (radar). Wind effects can result in spatial mismatches between the estimates derived from information aloft and the rainfall sampled by the gauges at the ground. The significance of these mismatches is difficult to quantify but one which we must remain aware of. A closely related problem is that of rainfall caused by low altitude weather systems. If the radar beam elevations are such that they overshoot these systems then any combined estimates that include the radar rainfall estimates will contain substantial errors. Cloud top heights estimated from Satellite data could provide a useful switch for determining when to ignore the radar information.

Before change-of-support issues can be addressed and blending of different data sources can be attempted, it is imperative that the data is referenced in a consistent co-ordinate system. The co-ordinate systems in which spatial rainfall data is collected may be different for each data source and may be different from the co-ordinate system in which combined estimates are required. Transformations between different co-ordinate systems may change the properties of the data and this aspect of the rainfall estimation problem needs to be accounted for.

The merging algorithm should account for availability of data. The question is: how does one handle the situation where one or all of the available data sources become unavailable? This is a particularly pressing problem when considering automated algorithms that operate in (near) real-time. The available options are i) produce no output when there is missing data ii) continue to produce a rainfall estimate that is based on the subset of data available iii) attempt to infill the missing data record based on surrounding observations iv) some combination of the preceding, depending on which data source is missing.

The key to producing a useful rainfall product is performing credible validation exercises that account for the difficulties already presented in this chapter. This validation process is non-trivial and needs to be carefully formulated and interpreted. Since rain gauges are the only direct measurement of rainfall available validation is usually only possible at points in space and time where gauge measurements are available. This issue remains to be adequately addressed and is the subject of ongoing worldwide research initiatives.

5.2 Conditional Merging

In previous work (Sinclair and Pegram, 2004b) we presented a comparison between the Bayesian Merging technique of Todini (2001) and the Conditional Merging technique presented below. Todini showed that the Bayesian Merging technique outperformed several others, using a small-scale numerical experiment. Sinclair and Pegram (2004b) showed that Conditional Merging performed better than Bayesian Merging in a comparison experiment. Further refinement of the technique and case studies using real-world data are presented here. Finally, a proposed algorithm, incorporating Conditional Merging, is presented as a candidate to replace the Explained Variance weighting algorithm currently used in SIMAR.

Radar produces an observation of the unknown true rainfall field that is subject to several well-known sources of error (e.g Wilson & Brandes, 1979; Habib & Krajewski, 2002) but retains the general covariance structure of the true precipitation field. The information from the radar can be used to condition the spatially limited information obtained by interpolating between rain gauges and produce an estimate of the rainfall field that contains the correct spatial structure while being constrained to the rain gauge data (where it is available). The conditional merging technique of Ehret (2002), suggested by Pegram (2002) does just this, making use of Ordinary Kriging (e.g Cressie, 1991) to extract the information content from the observed data. Figure 5.1 and the text following give an overview of the technique (for the one dimensional case) which is adapted from Ehret's (2002) work.

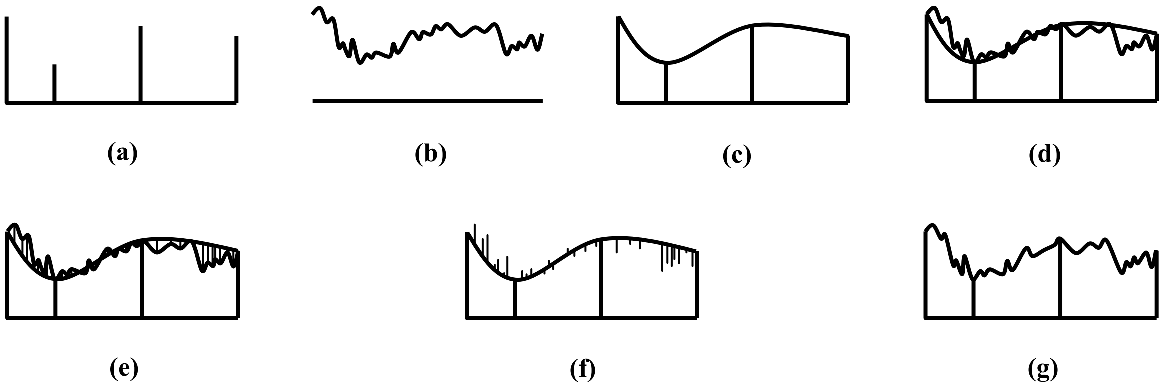

Figure 5.1 The conditional merging process

(a) The rainfall field is observed at discrete points by rain gauges.

(b) The rainfall field is also observed by radar on a regular, volume-integrated grid.

(c) Kriging of the rain gauge observations is used to obtain the best linear unbiased estimate of rainfall on the radar grid.

(d) The radar pixel values at the rain gauge locations are interpolated onto the radar grid using Kriging.

(e) At each grid point, the deviation **C** between the observed and interpolated radar value is computed.

(f) The field of deviations obtained from **(e)** is applied to the interpolated rainfall field obtained from Kriging the rain gauge observations.

(g) A rainfall field that follows the mean field of the rain gauge interpolation, while preserving the mean field deviations and the spatial structure of the radar field is obtained.

The spatial structure of the merged field is therefore obtained from the radar while the rainfall values are "stitched" to the gauge observations of the true rainfall field. The approach taken here has great similarities to the technique of Conditional Simulation by Kriging discussed by Chiles and Delfiner (1999, pp 452). However, the key difference in this case is that the radar rainfall estimate is *not* a simulation, unrelated (except by it's statistical properties) to the rainfall field we want to observe, but is in fact an observation of the true unknown field. This important link means that the radar data provides an estimate of the *actual* Kriging error and in particular it's spatial structure.

The error structure of the merged estimate can be examined by considering equations 5.1-5.7 below:

$$Z(s) = G_K(s) + \varepsilon_G(s) \tag{5.1}$$

$$R(s) = R_K(s) + \varepsilon_R(s) \tag{5.2}$$

$$M(s) = G_K(s) + \varepsilon_R(s) \tag{5.3}$$

$$E\left[Z(s) - M(s)\right] = E\left[\varepsilon_G(s) - \varepsilon_R(s)\right] \tag{5.4}$$

$$\operatorname{var}\left[Z(s) - M(s)\right] = \operatorname{var}\left[\varepsilon_G(s) - \varepsilon_R(s)\right] \tag{5.5}$$

$$= \sigma^2_{\varepsilon_G(s)} + \sigma^2_{\varepsilon_R(s)} - 2\operatorname{cov}\left[\varepsilon_G(s), \varepsilon_R(s)\right] \quad (5.6)$$
$$= \beta - 2\sigma_{\varepsilon_G(s)}\sigma_{\varepsilon_R(s)}\rho \quad (5.7)$$

$$\varepsilon_G(s) = Z(s) - G_K(s)$$
$$\varepsilon_R(s) = R(s) - R_K(s)$$
$$\beta = \sigma^2_{\varepsilon_G(s)} + \sigma^2_{\varepsilon_R(s)}$$

where $Z(s)$ is the true rainfall field at location s, $G_K(s)$ is the Kriged (mean field) estimate of $Z(s)$ from the rain gauge data, $R(s)$ is the radar rainfall estimate, $R_K(s)$ is the Kriged (mean field) estimate of $R(s)$ using the radar values at rain gauge locations and $M(s)$ is the merged estimate of $Z(s)$.

The term $\varepsilon_G(s)$ in equation 5.1 is unknown since $Z(s)$ is unknown. In equation 5.2, $\varepsilon_R(s)$ is known and, on the basis that $R(s)$ is a measurement of $Z(s)$, equation 5.3 can be used to estimate $Z(s)$. Equation 5.4 shows that the expected value of the error between the merged estimate and the true field is zero if the fields are Gaussian, since the Kriged estimates are unbiased in this case. The variance of the error estimate given by equation 5.5 can be decomposed as shown in equation 5.6. The variance of the error is (trivially) zero at the gauged points, while at any other position in the field it is bounded by a maximum value of β, for positive correlations ρ between $\varepsilon_G(s)$ and $\varepsilon_R(s)$. If $\varepsilon_G(s)$ and $\varepsilon_R(s)$ are strongly (positively) correlated as one would expect since both gauges and radar are measurements of $Z(s)$, then the variance of the error will be *significantly* less than β as suggested by equation 5.7.

5.3 An illustration based on the "String of Beads" model

An artificial experiment was carried out to test the efficiency of the technique at estimating the true rainfall field. Since the true rainfall is unknowable in reality, a sequence of 1000 independent 128x128 pixel rainfall fields was produced using the "String of Beads" simulation model (Pegram and Clothier, 2001). Each pixel is representative of a 1x1 km area and all 1000 simulated fields were generated with identical statistical properties. The simulated rainfall fields were treated as the "true" rainfall and "observed" radar estimates produced by adding systematic (bias) and random (noise) error components to the simulated fields. The "true" field was then sampled at 83 "rain gauge" locations distributed randomly on the pixel grid. This gives an approximate coverage of one gauge for every 200 km^2. A single realization of the "true" field as well as the corresponding "observed" radar estimate and the computed Kriged gauge field are shown in Figure 5.2, while a more detailed description of the experiment and the results obtained follows. Figure 5.3 shows the correlation structure in the simulated fields.

For each of the 1000 simulated radar rainfall fields and rain gauge observations the conditional merging technique was applied and the subsequent merged field compared to the true simulated rainfall field in terms of the mean and variance of it's deviations from the true field. All of the Kriging and merging computations were done on the logarithms of the variables in order to transform the log-normally distributed rainfall rate simulations to a Gaussian space. Comparisons, however, were done on the back transformed variables. The mean error of the merged field (over 1000 realisations) was computed for each of the 16384 pixels and plotted as a histogram (Figure 5.4). Also plotted in Figure 5.4 are histograms for the simulated radar deviations from the true field and the Kriged rain gauge deviations for comparison. The figure clearly shows the improvement of the merged fields in estimating the true field when compared to the simulated radar and Kriged gauge fields. Figure 5.5 shows similar plots for the variance of the errors, again the merged estimate performs well.

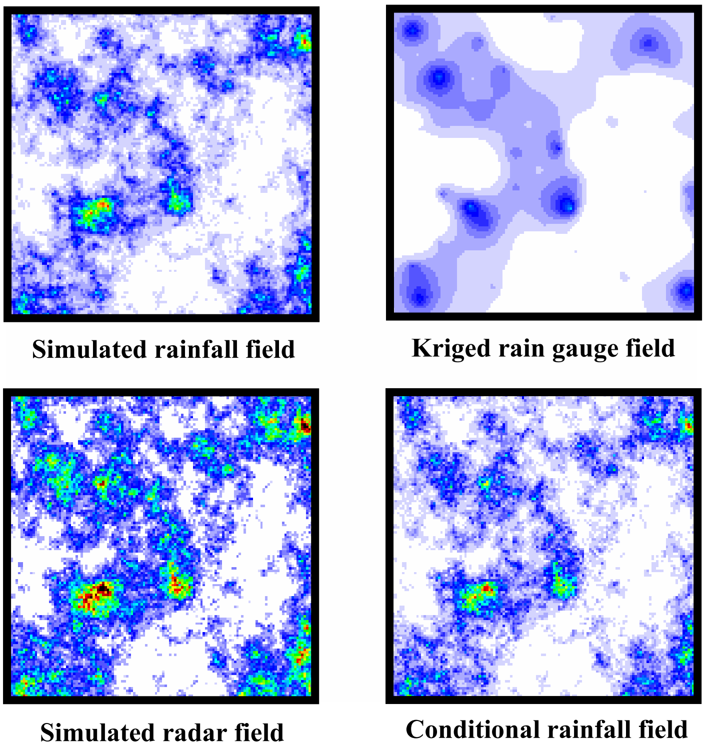

Figure 5.2 A single realization of the simulated and merged rainfall fields

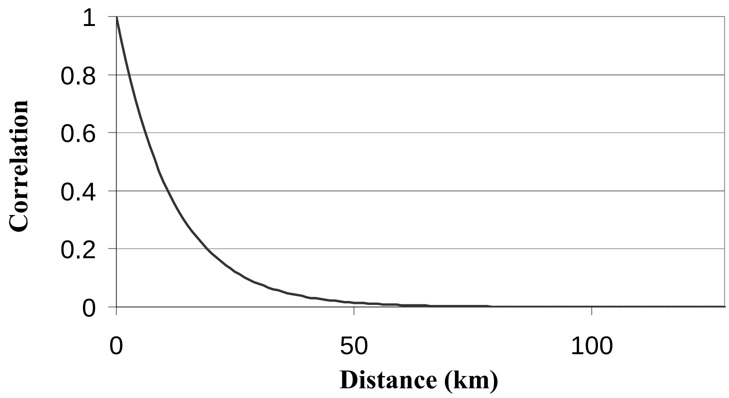

Figure 5.3 A plot of the exponential correlation function used in generating rainfields

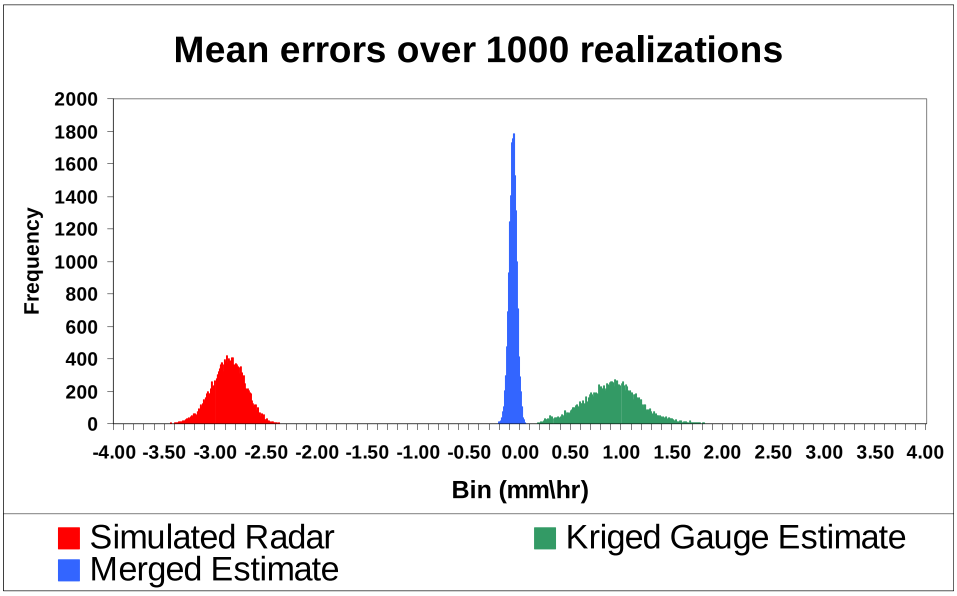

Figure 5.4 Comparison of the mean errors

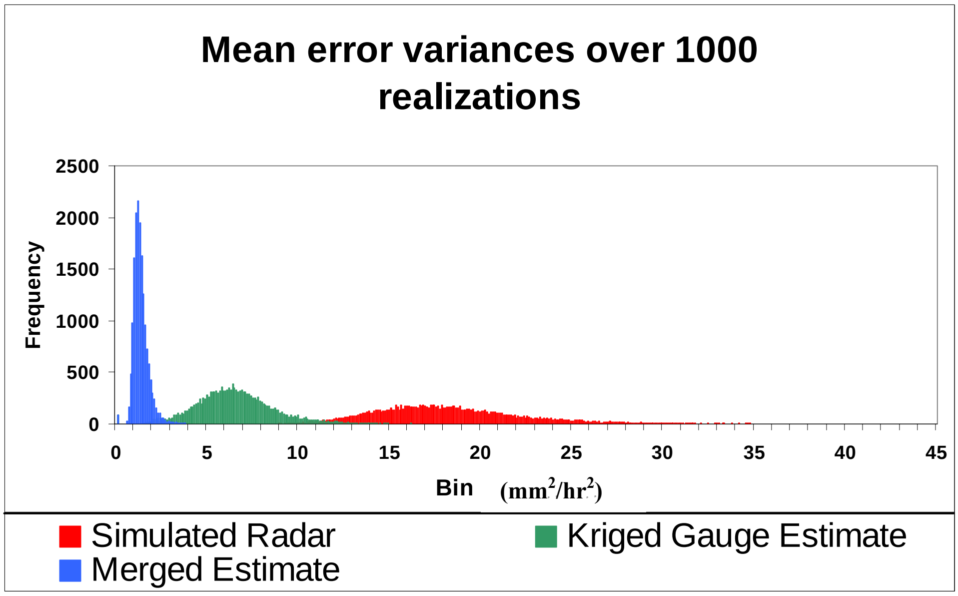

Figure 5.5 Comparison of the variance of the errors

5.4 A case study – Liebenbergsvlei, South Africa

The conditional merging algorithm was applied to radar and tipping bucket rain gauge data over the Liebenbergsvlei catchment in South Africa. The data analysed consisted of several daily accumulations from the South African Weather Services MRL5 radar (situated approximately 5 km from the catchment at it's closest point) and daily totals obtained from a network of 45 tipping bucket rain gauges shown (as red pins) in Figure 5.6.

A cross-validation study was carried out in the same way as described for the simulated rainfall fields and the mean error reduction at the gauges is illustrated for several days in Figure 5.7.

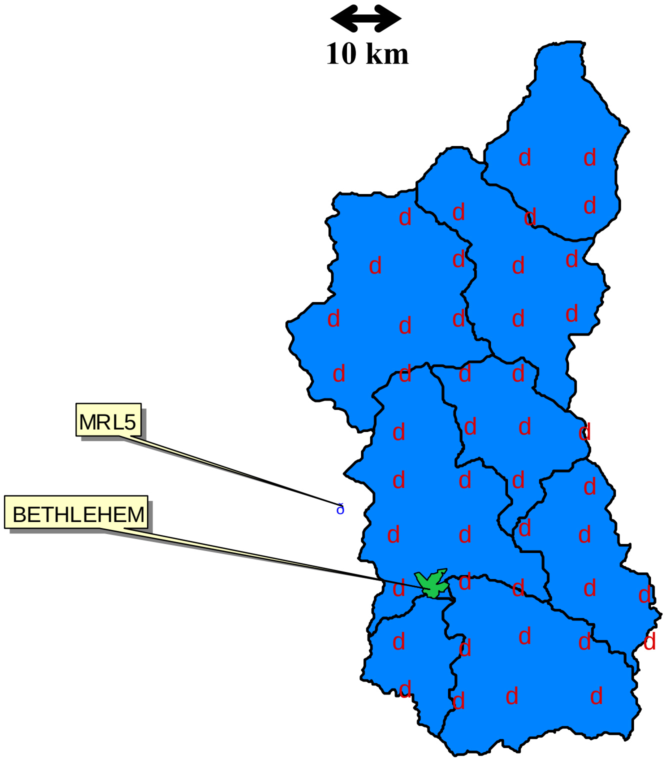

Figure 5.6 The Liebenbergsvlei catchment, showing the rain gauge network (red pins) and MRL5 radar

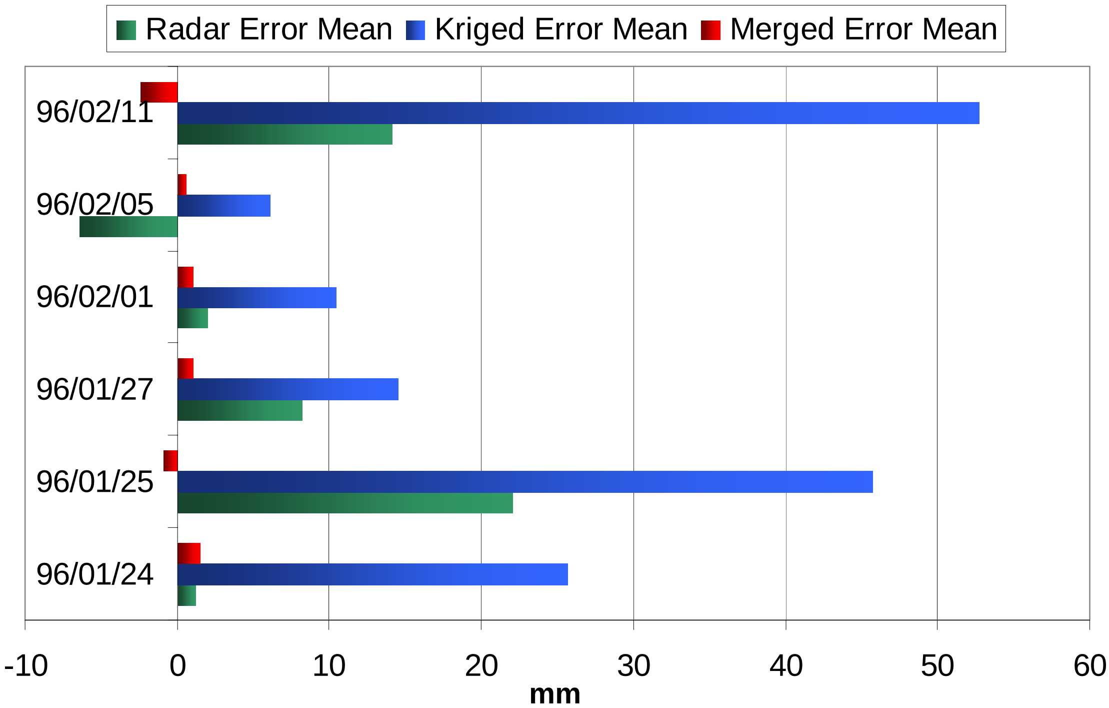

Figure 5.7 Error reduction achieved during cross-validation comparisons

5.5 Application of conditional merging to a Spanish Radar data set

The conditional merging technique described in Sinclair and Pegram (2005) was applied to a data set provided by Professor Daniel Sempere-Torres and Carlos Velasco from

GRAHI in Catalunya, Spain, an experiment which arose from a conversation between Professors Sempere-Torres and Pegram at the Symposium on Hydrological Applications of Weather Radar in Melbourne, 2004. The data set consists of twenty concurrent rain gauge and radar rainfall rates, measured at 10-minute intervals for the duration of a 13.5-hour rainfall event. The event began at 17:00 GMT on 28 September 2000 and ended at 06:30 GMT on 29 September 2000. The radar data consists of the 1kmx1km radar pixel above each rain gauge and its 8 surrounding neighbours as indicated schematically in Figure 5.8. The pixel labelled 0 is the radar pixel which has it's centre point located nearest to the rain gauge, the numbering scheme is relevant as it is used in the figures which follow.

6	5	4
7	0	3
8	1	2

Figure 5.8 Schematic local cut-out of the Spanish radar data, surrounding each of 20 rain gauges.

The first step in the data analysis was to compute accumulated event totals for each of the gauges and surrounding radar pixels. A simple accumulation was made which does not account for the effects of advection and sampling frequency in the radar estimations. The Optical Flow accumulation scheme described in section 5 could not be implemented due to the limited spatial extent of the radar data. Figure 5.9 shows a sample of the gauge and radar accumulations, for 4 of the 20 gauges examined. In most cases the radar produced an under-estimate relative to the gauge accumulation. There were only 5 cases in which the radar rainfall estimates were higher than the gauge estimate.

A cross-validation experiment was carried out to compare the rainfall estimates obtained by Kriging the rain gauge data, Conditional Merging and the raw radar estimates. Figures 5.10 and 5.11 show a sample of comparisons between gauge rainfall and the estimated rainfall at surrounding pixels computed by Kriging and Merging when excluding the gauge in question from the computation. Figures 5.12, 5.13 and 5.14 summarise the results of this comparison. The merged estimates clearly sit closer to a 1:1 relationship despite a slightly smaller R^2 value (indicating a wider scatter about their best fit line) than the raw radar estimates. The Kriged estimates do not perform as well as the Merged for this event and exhibit a wide scatter about their "best fit" line as indicated by the low R^2 value of 0.49.

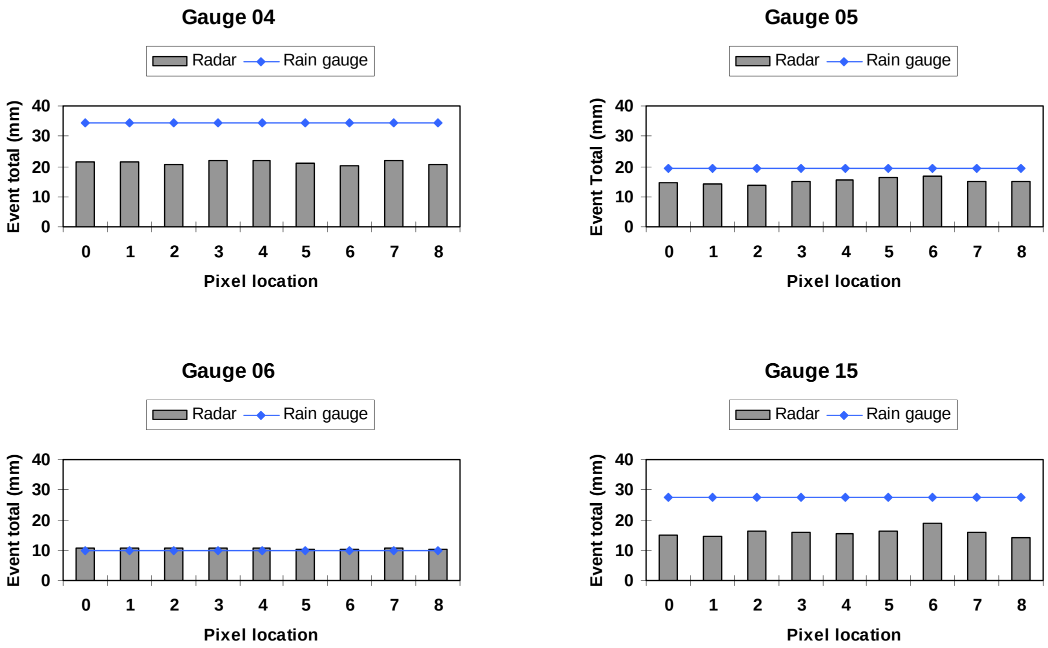

Figure 5.9 Event total accumulations

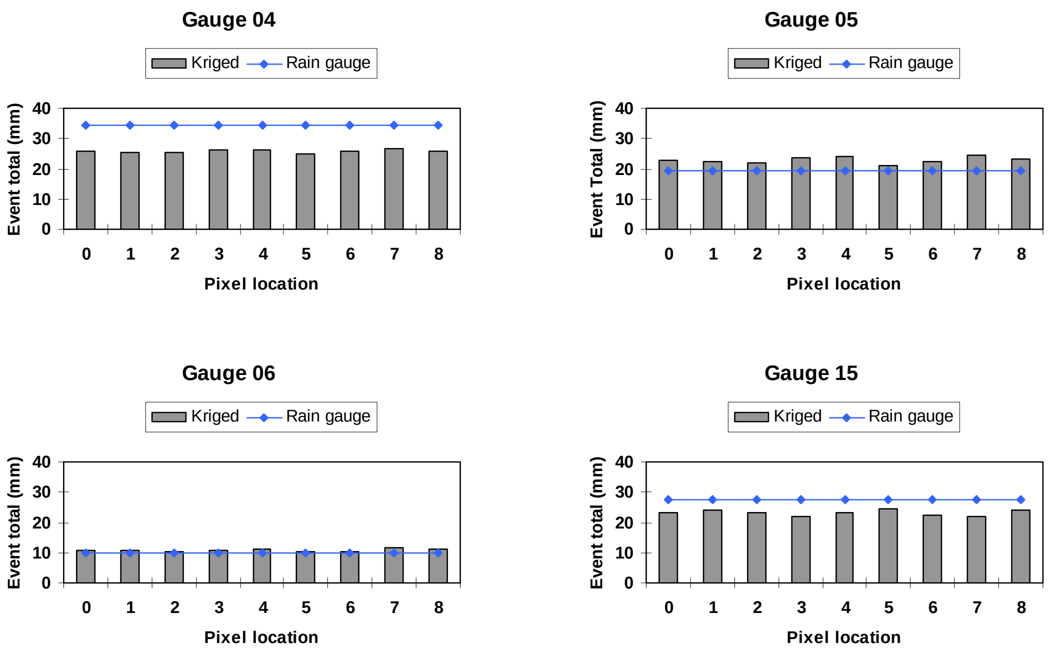

Figure 5.10 Cross-validation of Kriging

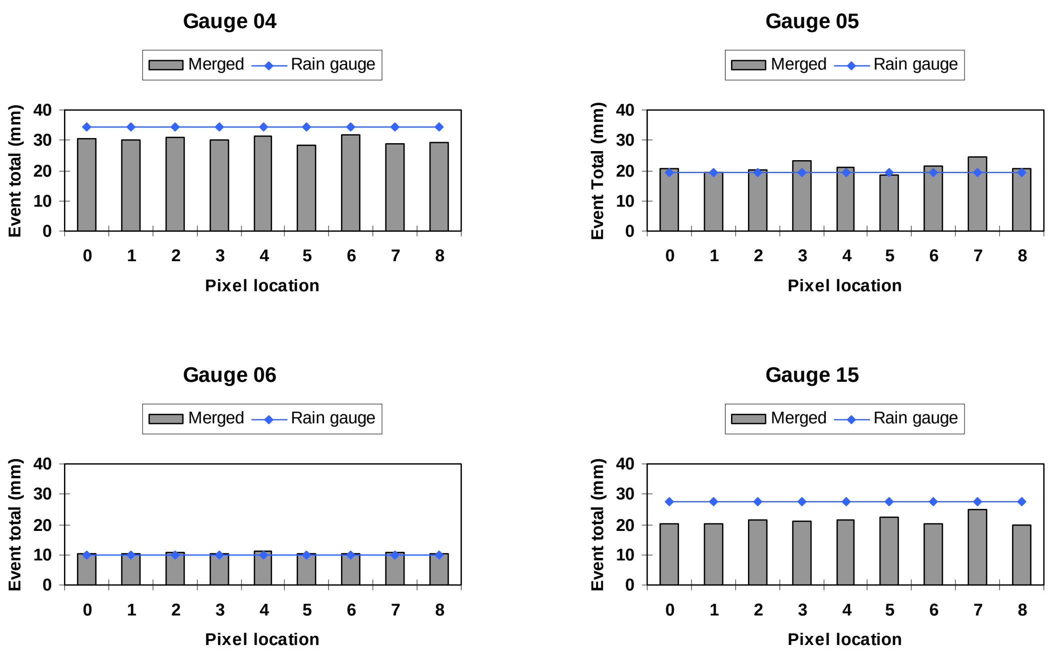

Figure 5.11 Cross-validation of Conditional Merging

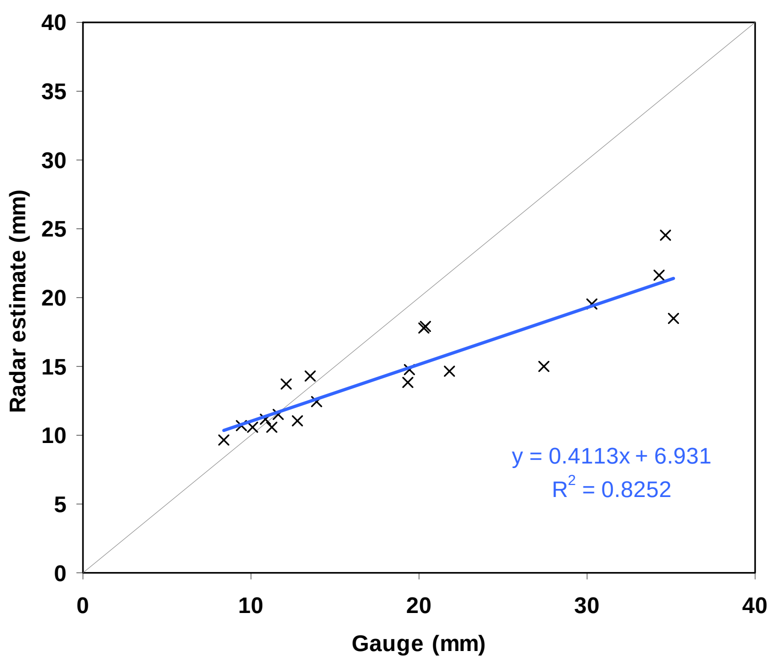

Figure 5.12 Comparison of gauge and Radar estimates

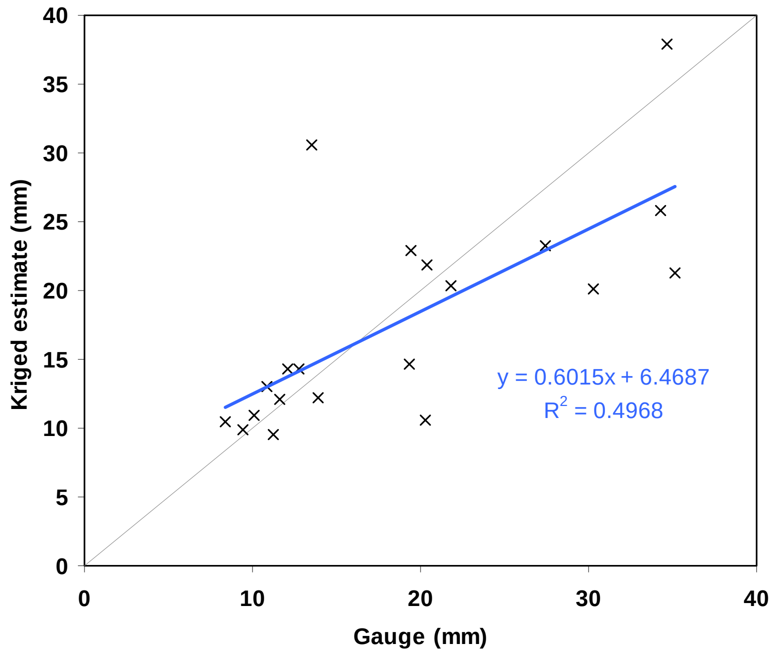

Figure 5.13 Comparison of gauge and cross-validated Kriged estimates

5.6 Comparing Explained Variance and Conditional Merging

The operational SIMAR algorithm (Deyzel et al., 2004; Pegram, 2004) uses explained variance weighting to combine estimates of rainfall. As discussed in section 5.3, the "true" rainfall field in unknown making a synthetic experiment necessary to compare the Explained Variance and Conditional Merging methods. A 25 by 25 spatially correlated random field was produced and subsequently sampled at 25 random locations. Additive noise was applied to the "true" field and the noisy field merged with the point samples using each of the two merging techniques. Figure 5.15 shows an example of the explained variance weighting technique applied to a synthetically generated spatially correlated field and Figure 5.16 shows the result of applying Conditional Merging to the same data. A simple visual inspection shows that Conditional Merging produces a resulting field that is more closely representative of the "true" field. The experiment was repeated many times and the sum of squared errors proved to be consistently lower for the Conditionally Merged fields.

5.7 Operational SIMAR products

A description of the operational SIMAR algorithm is repeated here for clarity, further detail is available in Deyzel et al. (2004). This operation is performed in three steps - suitably combine the gauge and radar information into a field R_G - suitably combine the gauge and satellite information into a field S_G - and finally combine R_G and S_G into the SIMAR field, which is published daily on the SAWS web-site with the others.

5.7.1 Merging radar and rain gauge data to produce the field R_G

The purpose of merging the radar and gauge estimation of rainfall fields is to retain the superior sampling resolution of areal rainfall from radar and to remove any quantitative biases by bringing the values closer to those of the 'ground truth' rain gauges. Merging the two rainfall fields is accomplished in the following steps:

- Extrapolate radar data as far as possible
- Interpolate rain gauge point measurements
- Determine the explained variance fields for radars and gauges
- Merge the two rainfall fields

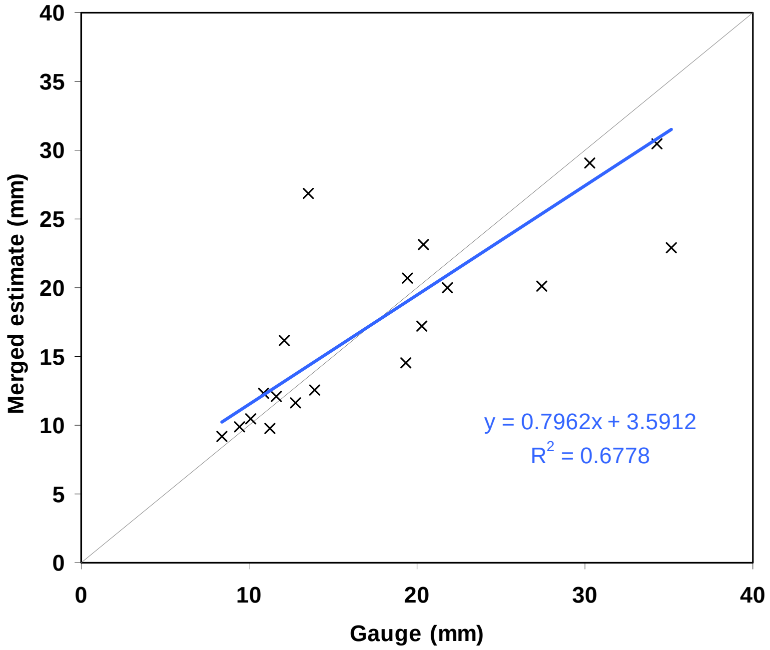

Figure 5.14 Comparison of gauge and cross-validated Conditionally Merged estimates

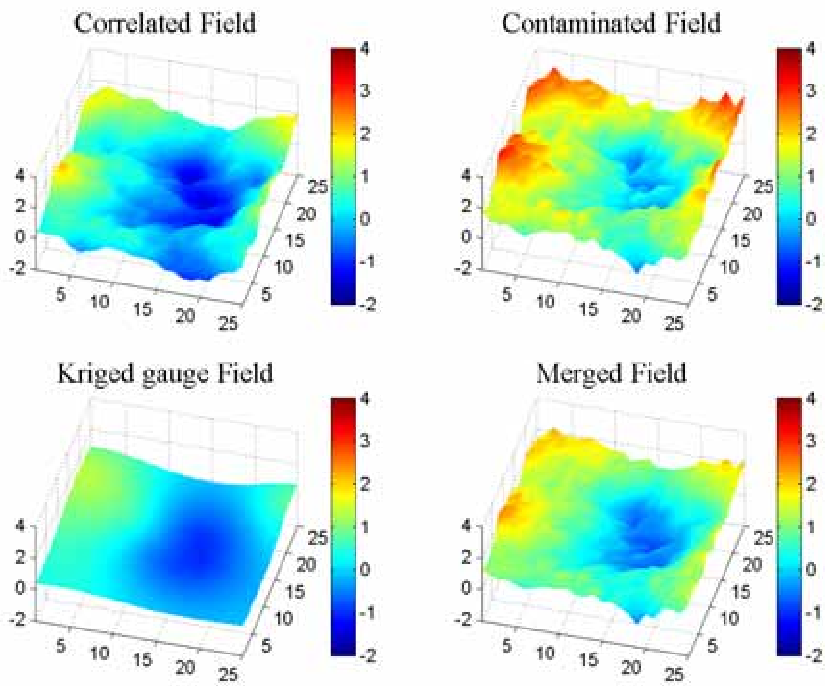

Figure 5.15 Synthetic Explained Variance example

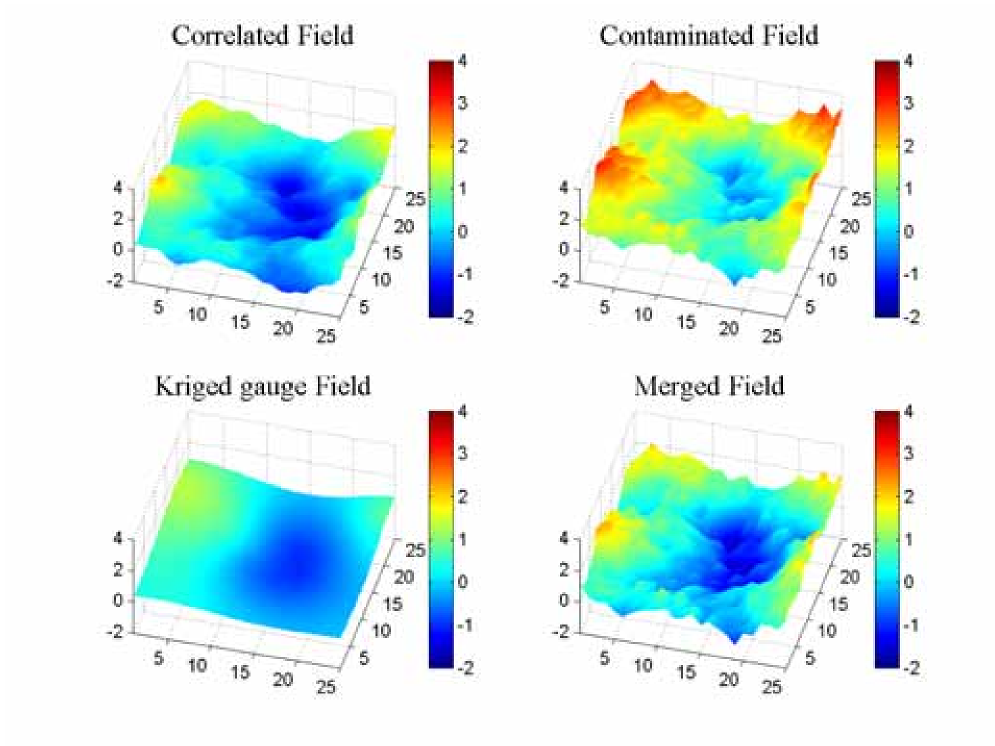

Figure 5.16 Synthetic Conditional Merging example

Extrapolate radar data

In SIMAR, it was decided to devise a method of Kriging over the 1024 by 1024 pixels (each of 1 minute of arc square) which would assume a constant covariogram over the region, estimated from work done on radar-rainfield modelling (Pegram and Clothier, 2001). The technique was to use the Fast Fourier Transform (FFT) in conjunction with Iterative Constrained Deconvolution (ICD) to perform the Kriging over the area. This novel technique requires too much space to describe in this context and is available in detail in Pegram (2004). Equivalently, Kriging could also be performed conventionally using nearest neighbours.

Kriging by FFT & ICD was applied to the radar rainfall field R_R to produce R_K. This field is a low-pass filtered version of the original field, in that the ICD and Kriging convolution act to blur the original structure of the spatial fields (Figure 5.17a). It has the beneficial effect of enhancing image quality, especially for areal rainfall. Erroneous high values, e.g. from ground clutter, are reduced during this process but, if not repaired, its effect is convolved with the surrounding area during Kriging, thereby degrading the qualitative fields. This has both a positive and negative impact on the resultant rainfall field. Other blemishes on the original radar rainfall field, such as anomalous propagation and bright band effects, are also lessened but their effect is spread with this technique. It is therefore imperative that the processed 24-hour radar rainfall field is of the highest quality before processing, this aspect is addressed in detail in section 3 of this report. In addition, the areas of missing data within the radar mosaic, such as occurs when radars are offline and where ground clutter information was removed, are masked out and thus the missing data are automatically interpolated.

Interpolating rain gauge measurements

The covariance function derived from historical radar fields was used in Kriging the rain gauge data between the gauge positions to produce the surface G_K, shown in Figure

5.17c. If the gauge density is sufficient and the values accurate, the resultant areal rainfall volume and structure more closely resembles that of the radar. When gauges are far apart and the rainfall heavy, interpolation rings are evident in the rainfall fields, which are typical of standard Kriging interpolation with an assumed field mean of zero. In such a case the areal rainfall differs widely from the radar. By increasing the rain gauge density and sampling resolution, the estimates of areal rainfall of the Kriged field and of the subsequent merged field will greatly improve.

Determine explained variance field for radars and gauges

To represent the accuracy of the data and the interpolated information where no observations are made in the radar fields and the rain gauge fields, each pixel in either field has a value of explained variance assigned to it. The explained variance fields, R_V and G_V respectively, are computed as a by-product of the Kriging used to interpolate the information through the fields. This was achieved in a novel way using FFT and ICD (Pegram, 2004). The whole radar domain, corresponding to a 200km range surrounding each radar, is accepted as being 100% accurate (assuming that all missing data areas were accurately interpolated) so the explained variance where there are observations is 100%. The explained variance diminishes to 0.1% if the point in question is more than 3 correlation lengths (about 60km) from observed data; this applies in both the radar and gauge explained variance fields R_V and G_V. The pixels corresponding to the daily reporting rain gauge positions are assumed to have 100% explained variance and because they are sparsely distributed, the G_V field takes on the appearance of a collection of small hyperbolic cones centred on the gauges. The R_V and G_V fields appear in Figures 5.17b & d. These fields will remain unchanged from day to day, unless a gauge or radar is not operating.

Merge the two rainfall fields

When it comes to merging interpolated radar and gauge rainfields, the algorithm chosen for SIMAR was to weight the values by their corresponding explained variances using the following equation: $R_G=[R_K R_V + G_K G_V]/[R_V + G_V]$. The merged field appears in Figure 5.18, which represents a rainfall field produced by integrating the quantitative values of the rain gauge (ground truth) with the superior spatial information of the distribution of areal rainfall of the radar. The spatial structure of the radar is maintained (Figure 5.17a), but the weight of the rain gauge information reduces the effects of bias or trend in the radar rainfall field.

5.7.2 Conditioning Satellite Data on Gauges to Produce the Field S_g

The satellite rainfall field is the alternative spatial data source in areas where there are no radars and the gauges are sparse. Because of the inherent limitation of determining rainfall from the opaque satellite channels, there is a consistent bias in the Multi Spectral Rain Rate (MSRR) satellite rainfall estimates. Conditioning the satellite rainfall field on the values at rain gauges, allows us to remove the bias in rainfall estimates at the rain gauges, and to retain the satellite's information with a reduction in the bias in between the gauges. The steps are:

- Interpolate an average satellite rainfall field S_Z from the best estimate S_R
- Interpolate the rain gauge data G_K
- Condition the satellite on the rain gauge information

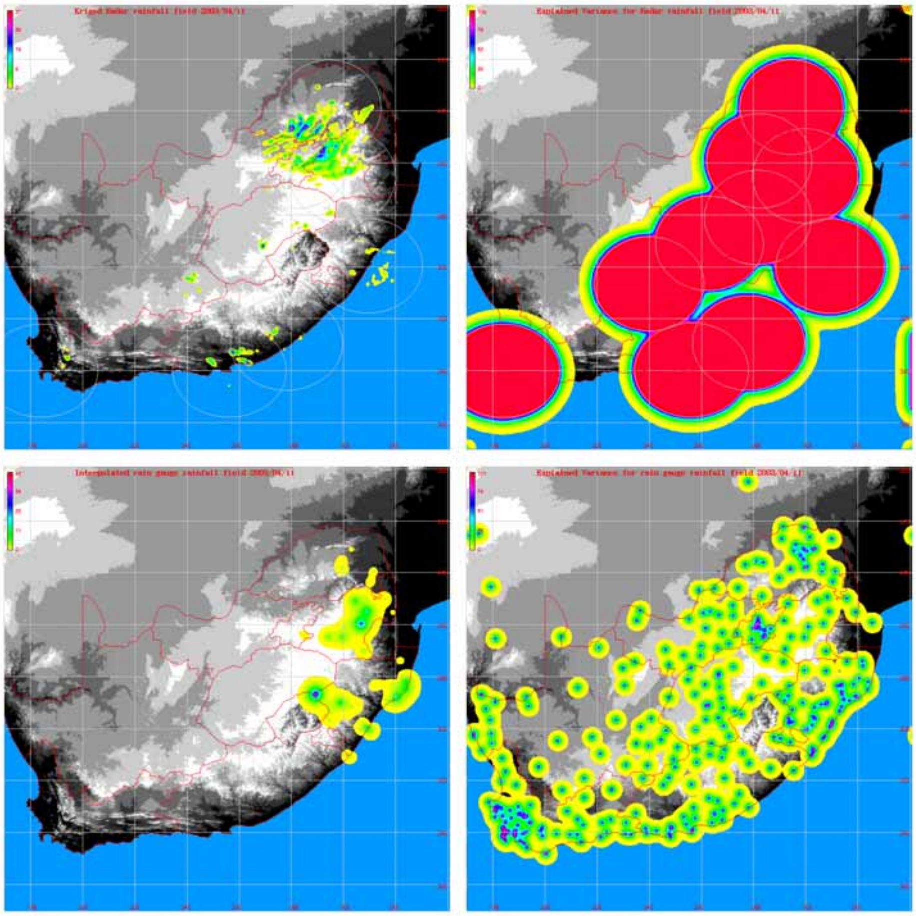

Figure 5.17 a) Kriged radar rainfall field, R_K; b). Explained variance field for Radar domain, R_V. c). Kriged rain gauge field, G_K; d). Explained variance field of rain gauge domain, G_V [from Deyzel et al. (2004)].

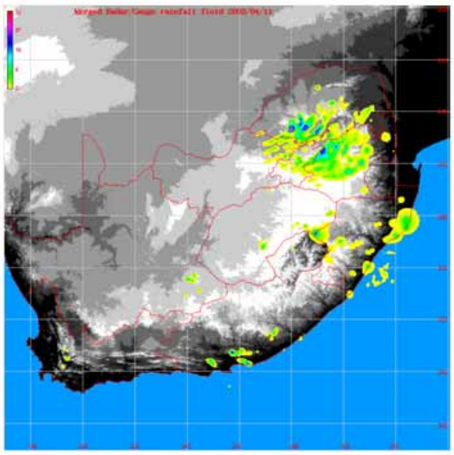

Figure 5.18 Combined radar and gauge field produced by merging 5.17 a) & c) using b) & d) [from Deyzel et al. (2004)].

Interpolate an average satellite rainfall field

The satellite rainfall pixels are sampled at the positions of the rain gauges, to produce a set of satellite rainfall point estimates: $S_{z,j} = S_{,j}$ wherever there is an observation G_j

Instead of using Kriging, the alternative of thin plate splines was employed to perform the interpolation. The $S_{z,j}$ data field is interpolated to a regular grid using completely regularized splines (CRS) with tension (Mitasova and Mitas, 1993) to recreate the satellite rainfall grid. CRS and spline interpolation have been widely applied on geographic (Mitasova and Mitas, 1993) and climatological (Hartkamp, et al., 1999) variables.

The spline interpolation cannot reconstruct the original satellite rainfall (Figure 5.19a) and at best produces a smoothed and homogeneous version of this rainfall field (Figure 5.19b). This field is taken to be the mean satellite rainfall field S_Z.

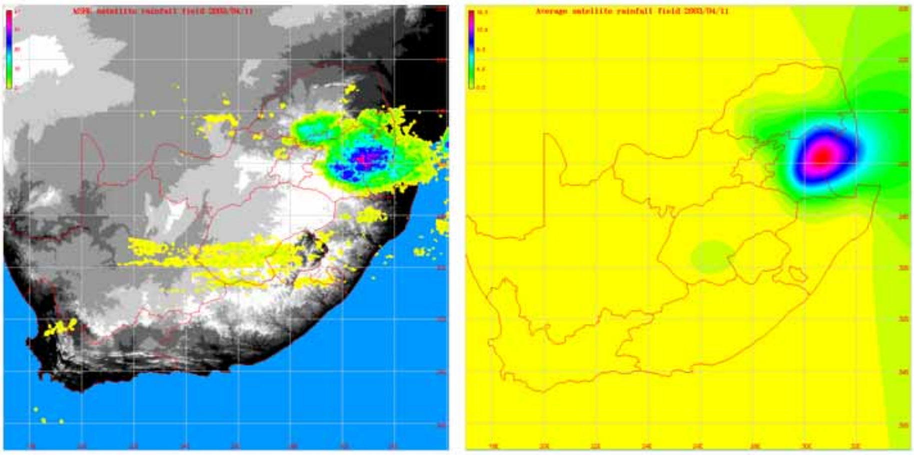

Figure 5.19 a) The 'best' 24-hour MSRR satellite rainfall field S_R and b) the corresponding average satellite rainfall field S_Z interplolated by splines from the gauge locations [from Deyzel et al. (2004)].

Interpolate the rain gauge data

The rain gauge point measurements were already interpolated to a rainfall surface with the Kriging algorithm, using the radar covariance information to spread the rain gauge data to produce G_K (Figure 5.17c).

Condition the satellite on the rain gauge information

The variability of the satellite rainfall field S_R about its mean surface S_Z is obtained by simple subtraction and the result added to the interpolated gauge field G_K, to produce a field that retains the signal describing the spatial distribution of the rainfall, as estimated from satellite, but with reduced bias: $S_G = S_R - S_z + G_K$. This field S_G in effect conditions the satellite data on the 'ground truth', by matching the satellite data at the rain gauge positions and accepting the conditioned field between the gauges, more so with increasing distance. The conditionally merged satellite field appears in Figure 5.20. It remains to combine this with the radar and gauge combined field of Figure 5.18.

5.7.3 Producing the combined SIMAR rainfall field

The methodology for the merging of the final rainfall field is the optimal combination of the fields R_G and S_G. Rather than merging these two fields as equals, an individual weight field is composed for each merged data set. For the radar and rain gauge merged field the explained variance masks, G_v and R_v, are integrated with the Boolean "or" statement to produce RG_v whose values range from 0% to 100% on each pixel. This

RG_V field appears in Figure 5.21a and remains fixed from day to day, only changing when a radar or gauge is not working.

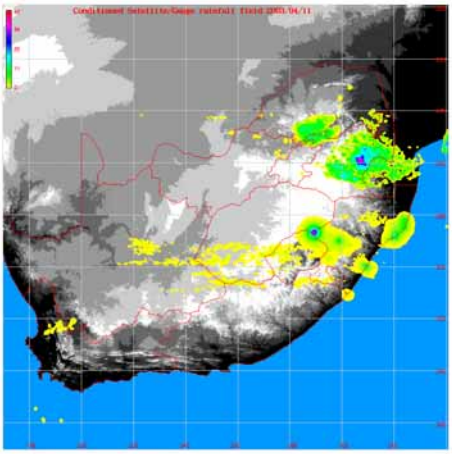

Figure 5.20 The field S_G, formed by conditionally merging the satellite rainfall field (Figure 5.19a) with the interpolated rain gauge information (Figure 5.17c) [from Deyzel et al. (2004)].

For the conditioned satellite rainfall field, a quite different and novel technique was developed for determining its accuracy. The satellite covers the whole domain and although measurement errors in the data set are quite consistent, the whole data set cannot be handled at the same accuracy level relative to the ground truth. For this purpose the bias skill scores of the satellite rainfall in a neighbourhood surrounding the rain gauges, are used as a weight field to accomplish the final merging. This field depends on the S_G field, so will change daily.

This is done by:

- Calculating the bias skill score S_{BSS} of the satellite rainfall field S_R relative to Kriged rain gauge values in a neighbourhood around rain gauge pixels G_K.
- Interpolating these skill score values expressed as percentages, to a regular grid S_{SS}, to be used as a weight field for combining the conditionally merged satellite field with the merged radar-gauge field.

The contingency table, formed by sampling each of the pixels in a 9x9 neighbourhood surrounding the rain gauge position, is used to calculate the bias skill score at the gauge location and is given by:

		Satellite	
Gauge		Rain	No Rain
	Rain	H	M
	No Rain	FA	CN

where H is correct hits, M misses, FA false alarms and CN correct nulls. Thus, in the vicinity of a gauge, S_{BSS} = Sum(from pixel j=1 to 81)[$H_j + CN_j$]/81. The bias skill score (Haralick, 1979) gives the fraction of correct rain/no-rain classifications made by the satellite. S_{BSS} indicates the bias of the MSRR technique in identifying the rainfall occurrences correctly. If $S_{BSS} = 1$ there is no bias in the technique's ability. With S_{BSS} decreasing to zero, the technique’s ability to classify the rainfall classes correctly, deteriorates.

Again splines are used to interpolate the discrete S_{BSS} data set to a gridded surface, to produce a continuous weight surface for the whole data domain of the satellite. This is expressed as a percentage between 100% and 0%, compatible (from the point of view of units and weighting) with the explained variance field.

The limitation of this technique of determining merging weights is that if the satellite rainfall field S_R has either a false alarm or miss when compared to an isolated rain gauge in data sparse areas, a low value of skill will be extended to an area where rainfall may have been correctly classified. This diminishes the impact of the satellite rainfall on the final merged field. To overcome this limitation, the average skill score of the whole data domain is used as a measure of accuracy of the satellite rainfall field in those areas far from radar or rain gauges, where the S_{BSS} may be unreliable, so the final skill-score field S_{SS} tends to the average S_{FSS} in sparse areas. A final S_{SS} field appears in Figure 5.21b and will change on a daily basis.

The final SIMAR merged rainfall field for 2003/04/11 appearing in Figure 5.22 is produce by integrating the radar-gauge and satellite-gauge rainfall fields in proportion to their weight fields RG_V and S_{SS}:

$$R_{MERGED} = \frac{[S_G \times S_{SS} + R_G \times RG_V]}{[S_{SS} + RG_V]} \qquad \} \; RG_V > 0$$

$$R_{MERGED} = \frac{[S_G \times S_{FSS}]}{100} \qquad \} \; RG_V = \text{NULL}$$

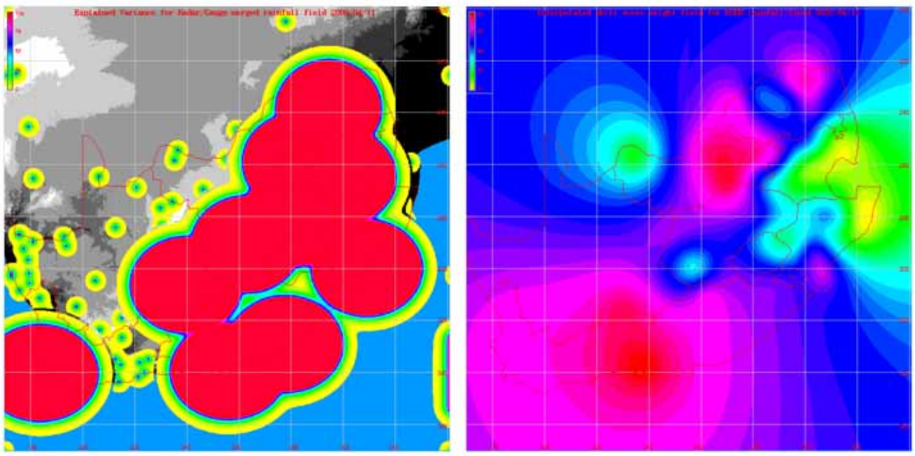

Figure 5.21 a) Radar and rain gauge combined explained variance field RG_V and b) scaled interpolated bias skill score field for satellite rainfall field S_{SS} [from Deyzel et al. (2004)].

5.8 An Improved Rainfield Merging Algorithm

In this section an algorithm for merging satellite, radar and rain gauge daily totals is proposed. The algorithm uses the decontaminated radar rainfall images extrapolated to ground level using a new technique called Cascade Kriging and the new Accumulation algorithm based on field advection measurements which gives smooth realistic daily totals of rainfall measured by radar and satellite. It combines these with the gauge data using the Conditional Merging technique as its basis and replaces the Explained Variance based algorithm currently used to produce the SIMAR product.

It is appropriate at this stage to restate the objectives driving the research into the merging of gauge, radar and satellite estimations of rainfall. They are two-fold:

- obtain the merged radar-gauge rainfield as the core of the best combined regional rainfield, but equally importantly, as a good ground-validation set of data for refining satellite estimation of rainfall. This latter contribution is in direct response to the request from the IPWG and GPM Ground Validation workshops attended (by invitation) by the project leader in October 2004 and September 2005.

- obtain the best satellite-gauge rainfield, because it is all we have where there are no radars (and when they don't operate) and where the gauge network is sparse.

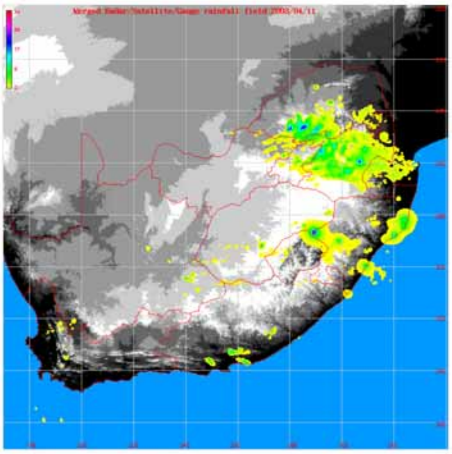

Figure 5.22 R_{MERGED}: The SIMAR merged rainfall field for 2003/04/11 produced by weighting Figures 5.18 & 5.19a using Figures 5.21 a & b [from Deyzel et al. (2004)].

Given these desiderata, recent experience with the SIMAR product shows that it is flawed in a number of areas, which can now be rectified. The first set of problems is infrastructure related:

- SIMAR estimates of rainfall are sometimes severely mismatched with the rain gauge readings in the sparse network. DWAF relies on some of these gauges for flood forecasting calculations, so their inaccuracy justifiably discredits the SIMAR field. To remedy this, SAWS is introducing a recapitalization program to introduce more gauges, as reported in the companion project K5/1426. They have also, as a result of this project and K5/1426, entered into co-operation agreements with other organisations (ARC, DWAF) which maintain their own rain gauge infrastructure. This increases the likelihood of including a greater number of gauges in future generations of the rainfall products.
- radars experience power outages, in spite of installed Uninterruptible Power Supplies (UPS), but SIMAR does not differentiate between 'no data' and 'no rain'; better maintenance and protection against wilful damage, besides modifying the algorithm to account for information absence, will help mitigate this effect.
- MET-7 is to be decommissioned on 01 February 2006 and data from its successor, MET-8, has to be used in it's place. Fortunately, although late, the data from MSG have been received since mid 2005 and are being incorporated into the algorithm.

The second set of problems derives from the setup of the computational algorithm producing the SIMAR field:

- the area (1024 x 1024 pixels) covered by the SIMAR field is too small, in that it does not cover parts of Angola and Mozambique. This is a consequence of using the FFT algorithm proposed in Pegram (2004), which was suitable for relatively limited computing power 4 years ago
- the relative accuracy of the radars was assumed to be 100% over their 200km range, a working hypothesis for the explained variance technique for combining their data with those of the gauges. The result was that, in areas further than 30km from the gauges but within the radar range, the radar was assumed to be the truth, ignoring any measurement bias which might exist.
- the satellite-gauge field (the first application of conditional merging using bicubic splines) when combined with the radar-gauge rainfield (formed by weighted explained variance) produced a smeared estimate of the rainfield which frequently ignored high rainfall at gauges. The SIMAR field is useful for getting a general idea of where it has rained in the last 24 hours but is inadequate for precise hydrological calculations.

How can combined daily rainfields be improved? The improvement to the infrastructure is covered in the companion study under contract K5/1426, which promises:

- an increasingly dense, more reliable network of gauges
- more reliable and better maintained radars
- improved satellite imagery

As far as improved algorithms are concerned, these fall into four groups, as detailed in this report:

- automatic online radar rainfield cleansing to remove contaminated data such as ground clutter and anomalous propagation once they have been identified
- extrapolation of radar estimates of rainfall to ground level
- conditional merging which has been proven in this study to be as good as any method currently available
- refined methods of accumulation of rainfields (which are instantaneous snapshots of estimated rain-rates) using smart field-advection algorithms

These four enhancements can be incorporated into the Improved Rainfield Merging Algorithm (IRMA), described below.

5.8.2 Improved Rainfield Merging Algorithm (IRMA)

Spatial coverage
There is now no limit to spatial coverage of the sub-continent. All Kriging will be performed on the gauge readings; even if their number increases ten-fold, the computational effort will not be a problem for current and future PCs and the algorithm.

Missing Data
Where gauge, radar or satellite measurements are incomplete, account will be taken of the missing information; a distinction will be made between ‘no data’ and ‘no rain’. SAWS are to provide guidance to define the criteria defining missing data when there have been partial observations.

Block Kriging over areas surrounding gauges

It is clear that gauges and radars/satellites have different measurement characteristics. As was introduced in Section 2.5.2 on ground-validation of extrapolated radar rainfall measurements, it is important to employ Block Kriging to estimate the appropriate areal average rainfall surrounding each operational gauge before combining gauge estimates with those from aloft.

Infer the spatial mean rainfall field

In conditional merging, a mean rainfall field between the gauges is routinely computed as the reference surface to which the radar (or satellite) rainfield is conditioned. In SIMAR, the basis function used in Kriging was the covariogram estimated from the radar rainfall observations. This resulted in a limited range of influence surrounding the gauges, limited to 30km. A possible alternative would be to use nearest neighbour ordinary Kriging, but this would result in discontinuities of the mean field in data sparse regions (Chiles & Delfiner, 1999, pp 201). The solution suggested here is to use ordinary Kriging of the complete gauge field with a distance basis function; effectively equivalent to the technique of Multiquadrics with a conical basis function.

Combination weighting

To combine the two individually merged fields: radar-gauge and satellite-gauge, the respective radar and satellite intrinsic errors relative to the gauges' measurements are to be used as relative weights. The satellite relative errors are assumed to be invariant with position, while the radar errors are assumed to vary with range from the radar. Thus at the locations of the gauges, the gauge readings will dominate (as is desirable), at intermediate distances from the radar (20 to 150 km) the radar-gauge field will dominate, just short of the 200km range of the radar, the radar-gauge and satellite-gauge fields will be equally weighted and beyond the range of the radar the satellite-gauge field will be the only estimate. Near the limit of range of the radar, it is expected that the relative errors of the radar and satellite estimates will be comparable: here, neither of them are good at estimating stratiform rainfall and both of them have skill at estimating convective rainfall – at this distance, beam-filling of the radar pixels degrades radar precision towards that of the geostationary satellite (a scale of about 2km).

In point form, the algorithm is as follows:

Error structure

- Compute the variance of the errors between satellite and gauge rainfall estimates over the entire data domain. The variance is assumed to be constant throughout the region. $S_{SG}^2 = \frac{1}{n}\sum_{i=1}^{n}\left[G_i - S_i\right]^2$, where G_i is the gauge estimate at point i and S_i is the associated satellite rainfall estimate. The number of gauges available is n. This is to be performed only where there is possible rain in the satellite field, as defined by the relevant IR mask.
- Fit a curve through $\varepsilon_r^2 = \left(G - R_r\right)^2$ out to the maximum range of the radars (200km). This is done to take account of the range dependant errors associated with the gauge radar rainfall estimation difference as described by Chumchean

et al. (2003). There is likely to be a nugget effect at the radar and the error function is likely to be monotonically increasing with range.

- Modify both of these error functions by the semivariogram computed from the radar rainfields as presented in Section 2; the relative error function of the combined rainfields should be zero at the gauge and should become the individual remote sensor's value at 30km. These two ideas appear in Figures 5.24 and 5.25 making use of the semivariogram depicted in Figure 5.23.

Figure 5.23 Average standardized semivariogram of Radar-rainfall in the Horizontal Plane Relative to a Gauge Position

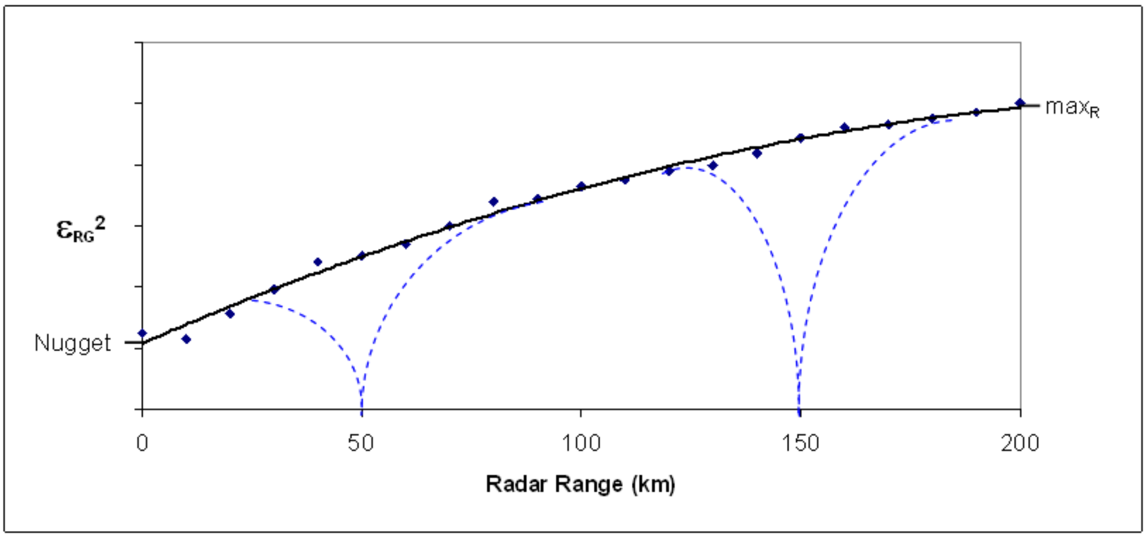

Figure 5.24 Variation of ε^2 with Range from Radar, where: ε_{RG} = [gauge – radar] estimate of daily rainfall. Superimposed is modification of error of conditionally merged radar-gauge field showing the effect of two gauges whose influence follows the curve of Figure 5.23.

Pairwise merging of fields

- Conditionally Merge gauge and radar using Multiquadrics. Multiquadrics are used to avoid the problem of having to assume a mean for the interpolated gauge field.
- Conditionally Merge gauge and satellite using Multiquadrics

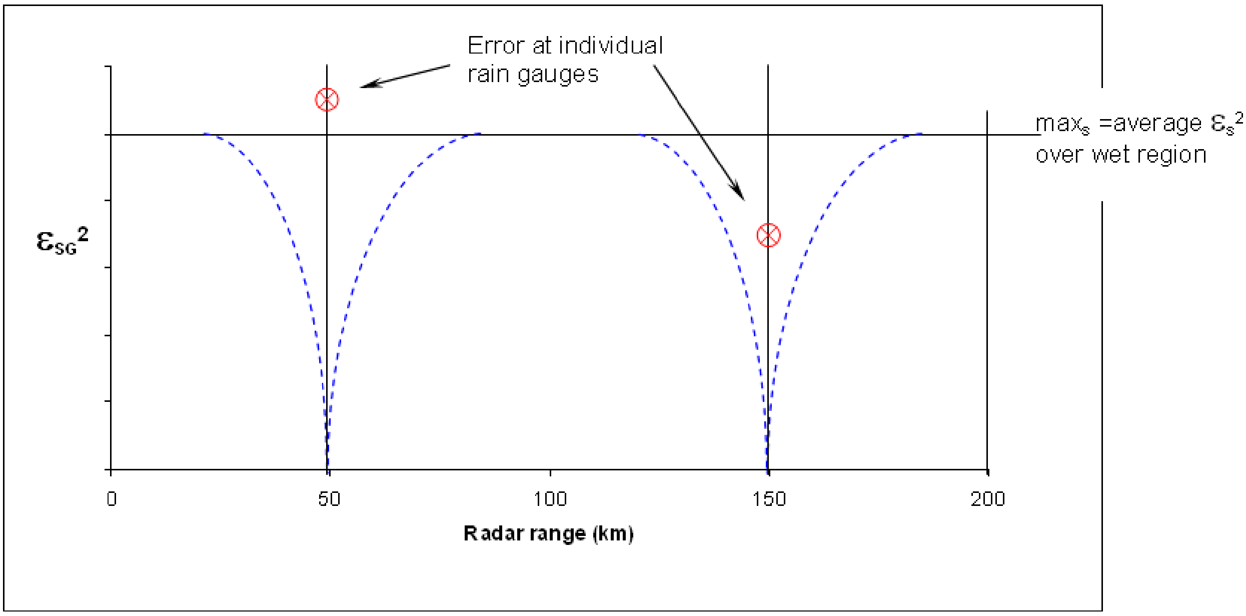

Figure 5.25 Estimation of max_s = mean of ε_s^2, where: ε_s = [gauge – satellite] estimates of daily rainfall when there is possible rainfall as defined by the mask. Superimposed is modification of error of conditionally merged satellite-gauge field in the same relative positions as in Figure 5.24.

Combination by error function weighting

- Scale the resulting curves and use as relative data influence weights. At the maximum radar range the influence of the merged radar and rain gauge field needs to tend to zero, making the only available estimate the satellite gauge combination. These two ideas appear in Figures 5.26 and 5.27, giving the weighting functions w_{RG} and w_{SG}.
- $R = \dfrac{w_{RG} R_{RG} + w_{SG} R_{SG}}{w_{RG} + w_{SG}}$

5.9 Conclusion

A conditional merging technique for combining rainfall information from radar and rain gauges has been presented and it's efficiency in terms of reducing the bias and variance of error estimates has been shown using an artificial simulation experiment and for daily data from the Liebenbergsvlei catchment in South Africa and on a Spanish data set. Rain gauge observations provide good point rainfall accuracy but poor spatial detail. The technique presented in this paper uses remotely sensed observations of rainfall fields to estimate the errors associated with using Kriging to interpolate between the rain gauge observations and condition the Kriged gauge field accordingly. Thus the spatial detail of the final merged field is improved while maintaining the mean field characteristics measured by the gauges. The operational SIMAR merging algorithm has been summarised and a candidate for an improved scheme based on the Conditional Merging technique proposed.

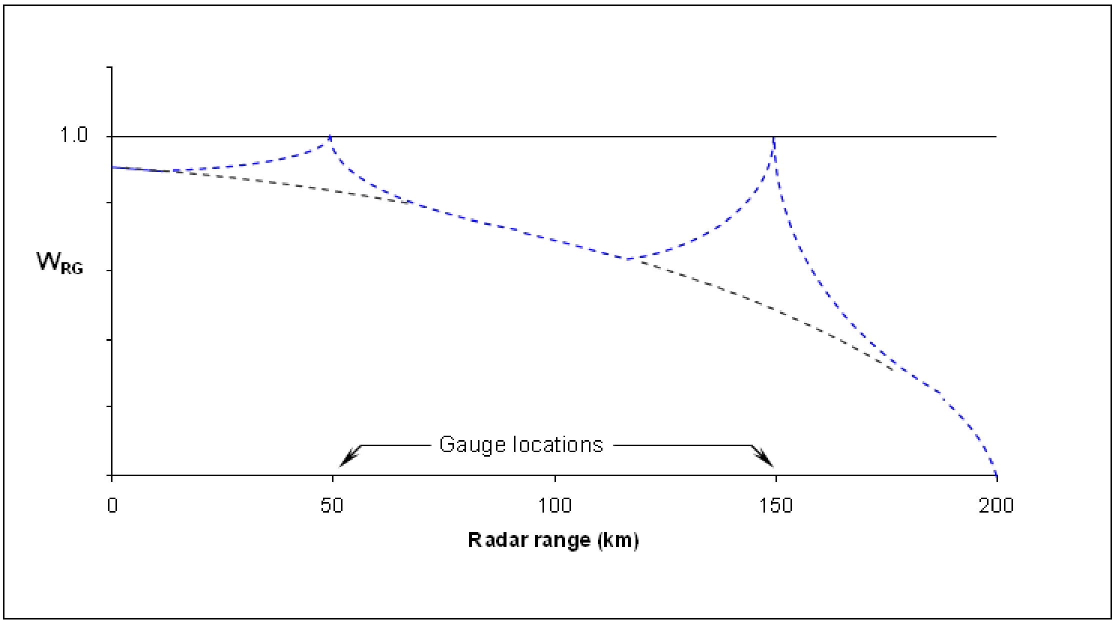

Figure 5.26 w_{RG}, the weighting function derived from Figure 5.24 as $w_{RG} = (max_R - \varepsilon_{RG}^2)/max_R$

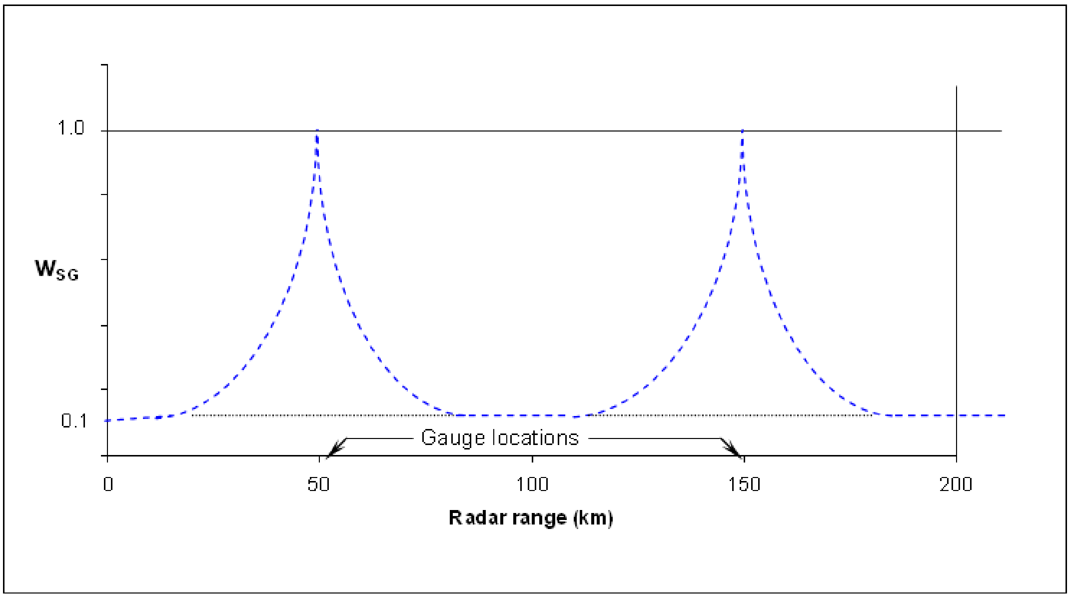

Figure 5.27 w_{SG}, the weighting function derived from Figure 5.25 as $w_{SG} = 0.1 + 0.9(max_s - \varepsilon_{SG}^2)$. Note: if the basic radar range error variation is linear, then w_{SG} and w_{RG} will cross at 180 km from the radar.

6 Rainfall nowcasting algorithm

The useful lead-time of stream flow forecasts can be significantly improved if reliable forecasts of the expected rainfall are available. This Chapter provides a discussion of several nowcasting techniques and presents a novel decomposition of two-dimensional rainfall data (based on 2D Empirical Mode Decomposition) into a range of spatial scales for use in a nowcasting context.

6.1 Overview of Short term rainfall forecasting

Forecasts of rainfall can be produced by applying one of two broad categories of model, these are; physically and statistically based models. The class of physically based models attempts to deterministically model the complex physics of the atmosphere, this relies on computing resources being available to make computations at the spatial and temporal resolution required to properly characterise the highly variable nature of rainfall. General circulation models (GCM's) are only able to perform operational simulations at synoptic scale spatial resolutions. Nesting limited area models (LAM) in a GCM can provide greater spatial detail up to resolutions of a few (e.g. 60x60 km) kilometres (Engelbrecht *et al.*, 2002), it is hoped that advances in computing power will allow such models to provide simulations and forecasts at higher spatial resolutions in the future.

Statistical models try to match certain key statistics of the rainfall as measured by a particular instrument; for example an application of the Neyman-Scott shot noise model presented by Cowpertwait (1991) tries to model the arrival times and duration of rainfall events as measured by rain gauges, while the "String of Beads" model (Pegram and Clothier, 2001) models the spatial and temporal characteristics of rainfall fields measured by weather radar. Here the discussion is limited to models capable of producing forecasts of spatial rainfall fields.

Rainfall forecasts applicable to flood forecasting (ideally) span a range of timescales from seasonal outlooks through several days ahead, right down to forecasts for the next few hours or minutes. In order that the longer-range forecasts are significant and useful, the uncertainties associated with them need to be reduced by decreasing the spatial resolution (increasing the basic area from pixel-scale) at which the forecasts can be made. This unfortunately has the effect of reducing their usefulness for flood forecasting, even for relatively large catchments. For this reason three statistically based models were investigated as possible candidates to provide short-term forecasts of spatial rainfall fields for the purposes of flood forecasting. The models considered are, the "String of Beads" Model (SBM) (Pegram and Clothier, 2001; Clothier and Pegram, 2002), the "Spectral prognosis" (S_PROG) model (Seed, 2001) and a new model based on Empirical Mode Decomposition (EMD) in two dimensions. The new model has been dubbed EMFOR (Empirical Mode FORecaster). Numerical weather prediction models are excluded here since they are currently unable to provide sufficient quantitative detail at appropriate spatial and temporal scales.

6.2 Nowcasts using Empirical Mode Decomposition

A data-driven method for extracting information, at temporally predictable scales, from spatial rainfall data (as measured by radar/satellite) is described, which extends the Empirical Mode Decomposition (EMD) algorithm into two dimensions. The EMD technique is used here to separate spatial rainfall data into a sequence of high through to low frequency components. This process is equivalent to a low-pass spatial filter, but based on the observed properties of the data rather than the predefined basis functions used in traditional Fourier or Wavelet decompositions. It has been suggested in the literature that the lower frequency components of spatial rainfall data exhibit greater temporal persistence than the higher frequency ones. This idea is explored here in the context of Empirical Mode Decomposition, to prepare rainfall data for nowcasts based on the temporal evolution of the lower frequency components. The paper focuses on the implementation and development of the two-dimensional extension to the EMD algorithm and it's application to radar rainfall data, as well as examining temporal persistence in the data at different spatial scales.

6.2.1 Introduction

Spatial rainfall data contain information at a broad range of spatial scales (Schertzer & Lovejoy, 1987; Harris et al., 2001; Pegram & Clothier, 2001). It has been suggested in the literature (Seed 2001; Turner et al., 2004) that the lower frequency components exhibit more temporal persistence than the higher ones; this premise is used here to prepare the data for nowcasts based on the evolution of the lower frequency components of space-time rainfall sequences. Examination of the (radially averaged) power spectrum (Figure 6.1) derived from a typical instantaneous estimate of rainfall rate obtained by weather radar (Figure 6.2) indicates that most of the power, hence potential for predictability in the context of nowcasting, is contained in the low frequency components. In this paper we focus on a data-driven technique to remove the high frequency (unpredictable) modes as the first step in a rainfall nowcasting scheme. The technique employed is a two-dimensional (2-D space) generalization of the one-dimensional Empirical Mode Decomposition (EMD) technique introduced by Huang et al. (1998). In a single dimension, EMD analysis produces a set of Intrinsic Mode Functions (IMFs) that are very nearly orthogonal; in two dimensions a set of Intrinsic Mode Surfaces (IMSs) is produced with similar quasi-orthogonal properties. Two-dimensional EMD appears to have been first introduced by Linderhed (2002) in the context of image compression; the key contribution in this paper is to introduce the concept of 2-D EMD to the Hydrometeorological literature as a tool for the analysis of space-time rainfall data. This paper focuses on the implementation and development of the two-dimensional extension of the EMD algorithm in this context, decomposing spatial rainfall data into its intrinsic spatial scale components.

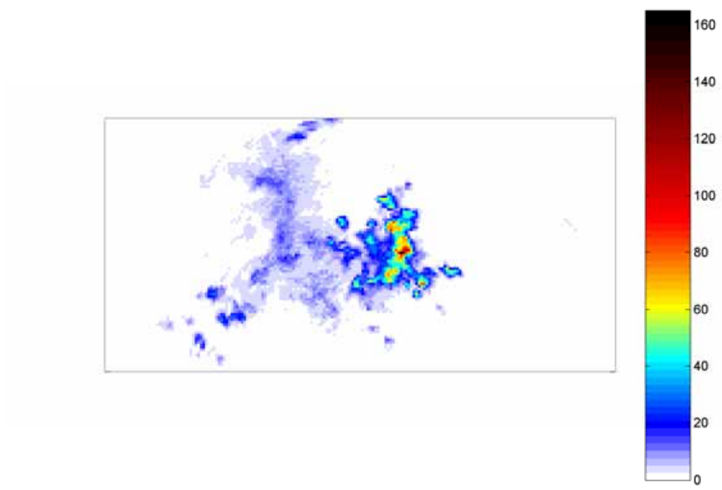

Figure 6.1 An observed convective rainfall field measured by S-Band weather radar at Bethlehem, South Africa (colour scale indicates instantaneous rain rate in mm/hr). The image is 100x200 km with $1km^2$ pixels.

In the application presented here, the least predictable IMS (exhibiting the highest local spatial frequency and least amount of spatial correlation – hence nearly white noise) is computed and removed from the raw rainfall data leaving a residual composed of the more predictable low frequency structural components in the data. This process is equivalent to using a low-pass spatial filter, based on the observed properties of the data rather than the predefined basis functions used in traditional Fourier or Wavelet scale decompositions. In Sections 6.2.2 and 6.2.3, simple theoretical examples, showing the power of EMD in one and two dimensions, are presented as a "proof of concept" before applying the procedure to observed radar rainfall data from Bethlehem, South Africa (Section 6.2.4). These complement and extend the original presentation by Huang et al. (1998) and Flandrin et al. (2004). Computational aspects relating to image processing and surface fitting are covered in detail (Section 6.2.4) and conclusions are drawn in Section 6.2.5.

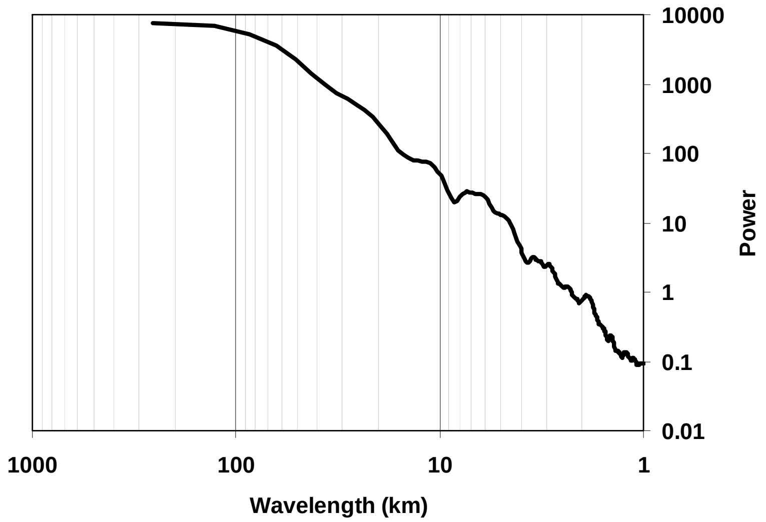

Figure 6.2 Radially averaged power spectrum of instantaneous rainfall rate from typical radar rainfall data shown in Figure 6.1.

6.2.2 Empirical Mode Decomposition in a single dimension

The basic idea embodied in the EMD analysis, as introduced by Huang et al. (1998), is to allow for an adaptive and unsupervised representation of the intrinsic components of linear and non-linear signals based purely on the properties observed in the data without appealing to the concept of stationarity. As Huang et al. (1998) point out in their abstract: "This decomposition method is adaptive and therefore highly efficient. Since the decomposition is based on the local characteristic time scale of the data, it is applicable to nonlinear and non-stationary processes."

To develop this idea further, the concept of stationarity can only be justified in the context of a theoretical hypothesis about the statistical properties of a population from which a sample is drawn. Few sequences of observations of natural phenomena are long enough to test the hypothesis of stationarity and frequently, the phenomena are patently non-stationary. This tacitly applies in the measurement of rainfall at a point or in space-time because sequences of rain are interspersed with dry periods and during the raining periods, the variability of the intensity due to mixtures of rainfall type (stratiform, convective, frontal) confound the stationarity definition. The EMD algorithm copes with stationarity (or the lack of it) by ignoring the concept, embracing non-stationarity as a practical reality. For a fuller discussion of the genesis of these ideas, see the

Introduction of Huang et al. (1998), who also heuristically demonstrate the implicit orthogonality of the sequences of Intrinsic Mode Functions (IMFs) defined by the EMD algorithm.

In the application of the EMD algorithm, the possibly non-linear signal, which may exhibit varying amplitude and local frequency modulation, is linearly decomposed into a finite number of (zero mean) frequency and amplitude modulated signals, as well as a residual function which exhibits a single extremum, is a monotonic trend or is simply a constant. Although EMD is a relatively new data analysis technique, its power and simplicity have encouraged its application in a myriad of fields. It is beyond the scope of this paper to give a complete review of the applications, however a few interesting examples are cited here to give the reader a feeling for the broad scope of applications. Chiew et al. (2005) examine the one-dimensional EMD of several annual streamflow time series to search for significant trends in the data, using bootstrapping to test the statistical significance of identified trends. The technique has been used extensively in the analysis of ocean wave data (Huang et al., 1999; Hwang et al., 2003) as well as in the analysis of polar ice cover (Gloersen & Huang, 2003). EMD has also been applied in the analysis of seismological data by Zhang et al. (2003) and has even been used to diagnose heart rate fluctuations (Balocchi et al., 2004).

Computing the one-dimensional EMD

The EMD algorithm extracts the oscillatory mode that exhibits the highest local frequency from the data ("detail" in the Wavelet context or the residue of a high-pass filter in Fourier analysis), leaving the remainder as a "residual" ("approximation" in Wavelet analysis). Successive applications of the algorithm on the sequence of residuals produce a complete decomposition of the data. The final residual is a constant, a monotone trend or a curve with a single extremum.

The EMD of a one-dimensional data set $z(k)$ is obtained using the following procedure:

1) Set $r_0(k) = z(k)$ and set $i = 1$.

2) Identify all of the extrema (maxima and minima) in $r_{i-1}(k)$.

3) Compute a maximal envelope, $\max_{i-1}(k)$, by interpolating between the maxima found in step 2. Similarly compute the minimal envelope, $\min_{i-1}(k)$. Cubic splines (as suggested by Huang et al., 1998) appear to be the most appropriate interpolation method for deriving these envelopes in one dimension (Flandrin et al., 2004).

4) Compute the mean value function of the maximal and minimal envelopes $m_{i-1}(k) = \dfrac{\left[\max_{i-1}(k) + \min_{i-1}(k)\right]}{2}$

5) The estimate of the IMF is computed from $IMF_i(k) = r_{i-1}(k) - m_{i-1}(k)$

Each IMF is supposed to oscillate about a zero mean and in practice it is necessary to perform a "sifting" process by iterating steps 2-5 (setting $r_{i-1} = IMF_i$ before each iteration) until this is achieved.

6) Once the IMF_i has a mean value that is sufficiently close to zero over the length of the data (defined by a stopping criterion within some predefined tolerance ε) the residual $r_i(k) = r_{i-1}(k) - IMF_i(k)$ is computed. Alternatively the sifting procedure can be

stopped when the difference in the standard deviation of successive estimates of IMF_i falls below a critical threshold (Huang et al., 1998).

7) If the residual $r_i(k)$ is a constant or trend then stop; else increment i and return to step 2.

Figure 6.3 shows the EMD of a composite data series (shown in the first panel) that is the summation of a sine wave, a triangular waveform and a slowly varying trend. The compact representation obtained by EMD extracts (almost perfectly – except at the ends) the three separate time series (shown in panels 2 to 4) that make up the composite signal, without resorting to Fourier or Wavelet techniques with restrictive assumptions about the form of the underlying oscillatory modes or basis functions. Figure 6.4 shows the analysis of the same data, using Wavelet decomposition. Here a fifth order Daubechies wavelet basis was (arbitrarily) chosen. Seven levels of decomposition were required before the trend became apparent; this decomposition is clearly far less compact and physically meaningful than the EMD results.

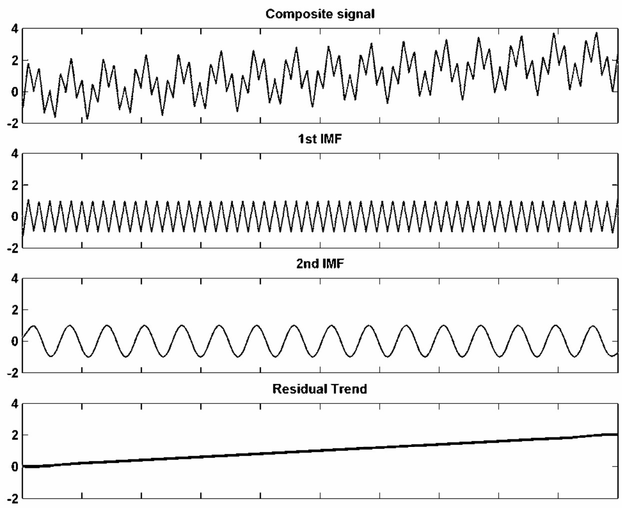

Figure 6.3 EMD based signal separation; all IMFs are plotted to the same vertical scale. Top panel is the combined signal; lower 3 panels are the decomposition which recaptures, almost exactly, the original components.

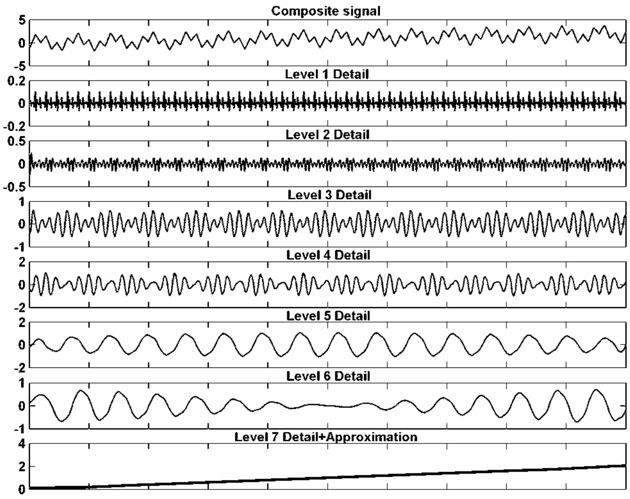

Figure 6.4 Wavelet based signal separation – The 'data' are the same as in Figure 6.3, the vertical scale has been compressed for a compact presentation. An arbitrarily chosen db5 wavelet basis has been used.

A similar decomposition analysis can be carried out using Fourier techniques. The Discrete Fourier approximation of a signal can be defined in terms of the Euler-Fourier coefficients (a_0, a_k, b_k) with $k = 1,2,...,m$ (Equation 6.1). The coefficients are all that are required to reconstruct the series and any signal can be well approximated (as long as it satisfies the Dirichlet conditions), provided *m* is sufficiently large. In equation 6.1, $F(x_j)$ is the Fourier approximation of the signal y_j at each of the *n* discrete (evenly spaced) data points x_j. L is the range of values x_j over which the data set is assumed periodic.

$$F(x_j) = \frac{a_0}{2} + \sum_{k=1}^{m} \{a_k \cos(2\pi k x_j / L) + b_k \sin(2\pi k x_j / L)\} \tag{6.1}$$

$$a_k = 2\sum_{j=1}^{n} \cos(2\pi k x_j / L) y_j / n \tag{6.2}$$

$$b_k = 2\sum_{j=1}^{n} \sin(2\pi k x_j / L) y_j / n \tag{6.3}$$

Figures 6.5 and 6.6 show the result of decomposing the data using a finite Fourier series. Figure 6.5 shows the first 5 harmonics; while Figure 6.6 shows the series reconstruction by accumulating the lower harmonics up to *m*. Computing the Euler-Fourier coefficients provides a compact approximation of the original signal but fails to extract physically meaningful information. The ability to determine meaningful structural information is clearly important in a nowcasting context, which cannot be bound by the periodicity assumption implicit in Fourier methods.

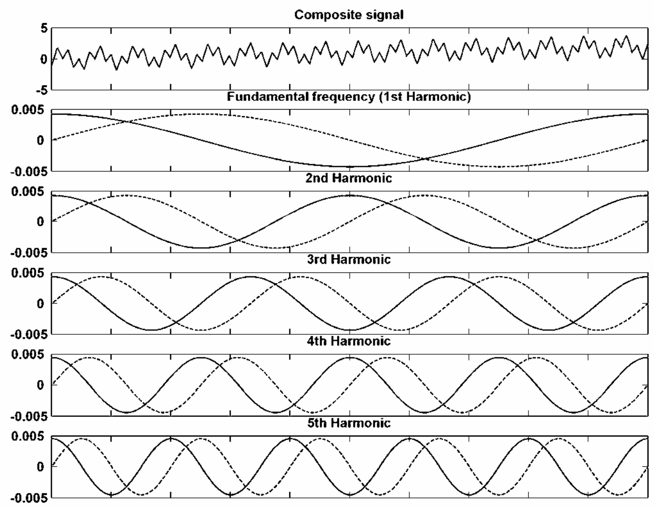

Figure 6.5 Fourier based signal separation, the first 5 of 75 components – The dashed lines show the sine component and the solid lines the cosine component.

6.2.3 Empirical Mode Decomposition in two dimensions

In two dimensions the EMD process is conceptually the same as for a single dimension, except that the curve fitting exercise becomes one of surface fitting and the identification of extrema becomes (a little) more complicated. Very little work appears to have been done which applies the EMD technique to two-dimensional data. Han et al. (2002) use EMDs in one dimension along four different directions to smooth Synthetic Aperture Radar (SAR) images and remove speckle. Nunes et al. (2003) develop a technique, which they term "Bidimensional Empirical Mode Decomposition" (BEMD) in the context of texture analysis in image data where they demonstrate several examples of intrinsic mode extraction from image data. Linderhed (2002; 2004) examined the use of EMD in two dimensions for image compression. Both of these applications are very similar to what we propose in this paper. The 2-D EMD analysis is a truly two-dimensional analysis of the intrinsic oscillatory modes inherent in the data. Fourier and Wavelet techniques are really applications of their one-dimensional counterparts in a number of principal directions, Fourier analysis concentrates on orthogonal "East-West" and "North-South" directions while 2-D Wavelet analysis examines an additional diagonal component. In contrast, EMD produces a fully two-dimensional decomposition of the data, based purely on spatial relationships between the extrema, independent of the orientation of the coordinate system in which the data are viewed.

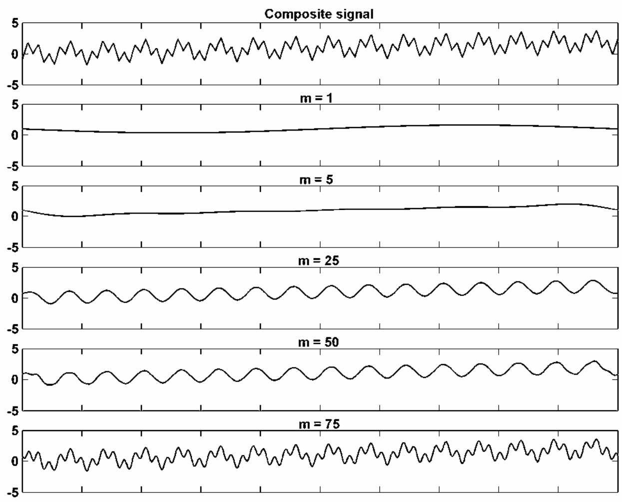

Figure 6.6 Reconstruction of the signal from the sine and cosine components, m represents the number of Euler-Fourier coefficients used in each reconstruction.

Description of the algorithm

The algorithm follows intuitively from the one-dimensional case and may be briefly summarised as follows:

1. Locate the extrema in the 2D space including maximal and minimal plateaus.

2. Generate the bounding envelopes using appropriate surface fitting techniques. We suggest conical Multiquadrics (for reasons explained in the description of surface fitting).

3. Compute the mean surface function as the average value of the upper and lower envelopes.

4. Determine the first estimate of an IMS by subtracting the mean surface from the data.

5. Iterate until the IMS mean surface function is close to zero everywhere.

6. Estimate the IMS and Residual.

7. If the Residual is a constant or a monotone trend, then stop; else return to step 2.

Surface fitting for extremal envelope generation

The generation of maximal and minimal envelopes is of key importance to a successful 2-D EMD implementation and is the most computationally intensive task. The problem is a familiar one of collocating a smooth surface to randomly scattered data points in two-dimensions. There are several options available to achieve this. Ultimately the fitting procedure reduces to computing the unknown value of the surface at a point $s_i = (x_i, y_i)$, by some linear (or nonlinear) weighting of the known data. In general, a basis function determines the influence of each known data point based on its spatial position relative to the unknown point s_i. Nunes et al. (2003) use radial basis functions while Linderhed uses bi-cubic splines (Linderhed, 2002) and later chooses the more suitable option of Thin Plate Splines (Linderhed, 2004). We use radial basis functions (technically, conical Multiquadrics), which are identical to Kriging (Cressie, 1991) with a purely linear semi-variogram model. It could perhaps be argued that it would be more appropriate to fit a semi-variogram model to the maxima and minima, but we feel this would be over-elaborate and presumptuous, as the extrema are only related by distance and cannot be considered drawn from a stationary correlated random field. Invoking Occam's razor in the spirit of Huang's original derivation of EMD, we wish to *let the data do the talking* and conical Multiquadrics assume the least structure of any linear surface fitting algorithm.

The Ordinary Kriging estimate $\hat{z}_i$ at any point i based on n observed data points is

$$\hat{z}_i = \sum_{k=1}^{n} \lambda_k z_k \tag{6.4}$$

where z_k are the observations and λ_k are weights associated with each observation and the target point. The mean is assumed unknown and the weights λ_k are constrained to sum to unity. The vector of weights $\underline{\lambda}$ is obtained by solving the linear system in Equation 6.5

$$\left[\begin{array}{c|c} \Gamma & u \\ \hline u^T & 0 \end{array}\right] \left[\begin{array}{c} \underline{\lambda} \\ \hline \mu \end{array}\right] = \left[\begin{array}{c} \gamma \\ \hline 1 \end{array}\right] \tag{6.5}$$

where γ is a vector of semivariogram values, in this application simply defined by the linear distance basis function $\gamma(s_{ij}) = |s_{ij}|$ with s_{ij} the distance between point i and the $j = 1, 2, ..., n$ observation locations. Γ is the matrix of distances between the observations, u is a vector of n ones and μ is a Lagrange multiplier ensuring that the Kriging weights λ_k sum to unity, as required. The solution of Equation 5 is obtained using Singular Value Decomposition (SVD) in this application to ensure that a stable solution is assured (when the matrix is ill conditioned). This is achieved by truncating singular and near-singular components. Although SVD is computationally less efficient than (for example) LU decomposition as a means of solving a dense linear system, it's use is preferred here because of it's robustness in the face of the near-singular Kriging systems which are frequently encountered in gridded data applications (Wesson & Pegram, 2004).

Simple two-dimensional EMDs

In this section, applications of the 2-D EMD technique are presented. As an artificially constructed example Figure 6.7 shows the successful removal of noise added to a synthetically generated two-dimensional sine signal. The noise (with it's high local spatial frequency) is almost completely described by the first IMS leaving a residual, which is closely representative of the underlying signal.

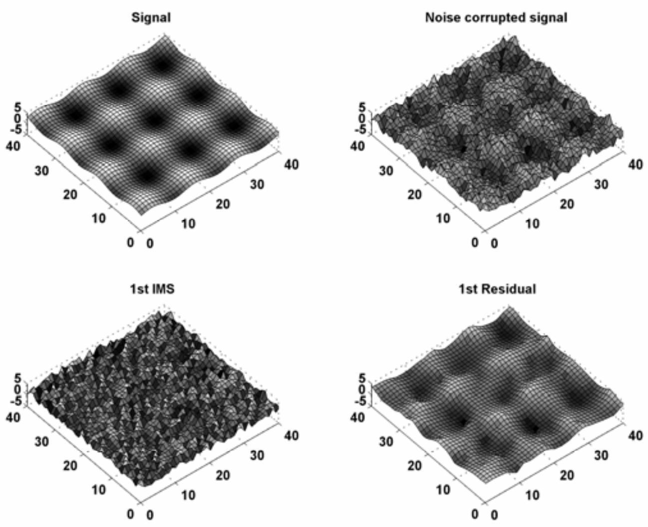

Figure 6.7 Example of EMD used for noise removal on a 2-D sine wave. The bulk of the additive white noise in the corrupted signal is well captured by the first IMS.

Turning to a realistic example of the type we have been aiming for, Figure 6.2 showed an instantaneous radar rainfall field with an area of 100x200 km. A complete EMD of this field is shown in Figure 6.8 using a direct application of the 2-D EMD process described in Section 6.3.1. The final residual (with a single extremum) gives a clear indication of the position of the largest convective raincell evident in the field.

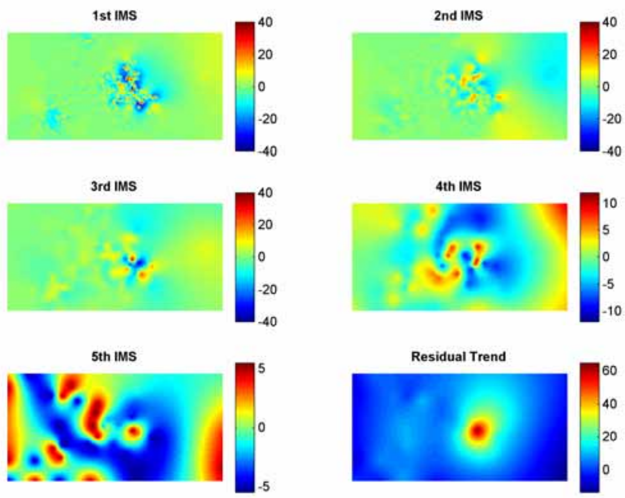

Figure 6.8 Naïve EMD of the observed rainfall field shown in Figure 6.1 – note the change in scale of the rain rates in the IMSs.

6.2.4 Application of 2-D Empirical Mode Decomposition to rainfall data

The simple 2-D EMD application presented in the previous section is computationally burdensome when applied to rainfall data. In this section, to overcome this drawback, a number of specific refinements are presented which combine to make EMD tractable in practical real-time situations.

Image processing techniques and optimisations

Since an application of 2-D EMD requires the use of surface fitting techniques, large linear systems must be solved. The size of a system is determined by the number of known data points which are to be used in combination to find the unknown values of the surface at each remaining position in the field. The highly variable nature of rainfall data means that the field contains a large number of extrema from which the bounding envelopes must be constructed. Additionally there are a large number of zero (no rain) data, which constitute minima. By only considering raining areas, the size of the linear systems requiring a solution are greatly reduced since each raining area (if more than one exists) will contain a considerably smaller number of extrema than the entire data region and each can be treated separately. Furthermore, it makes no sense to consider an EMD in areas where the variable of interest does not exist, in this case the areas that are not raining.

A number of well-known image processing techniques are implemented to isolate and process each raining area. Figure 6.9 summarises the steps taken in processing the data with the boxes numbered 1-7 indicating different steps in the process. First a mask is generated to separate the raining and non-raining pixels (Figure 6.9, Box 1); pixels below a threshold of 1 mm/hr are considered as non-raining and the remaining pixels are marked as raining.

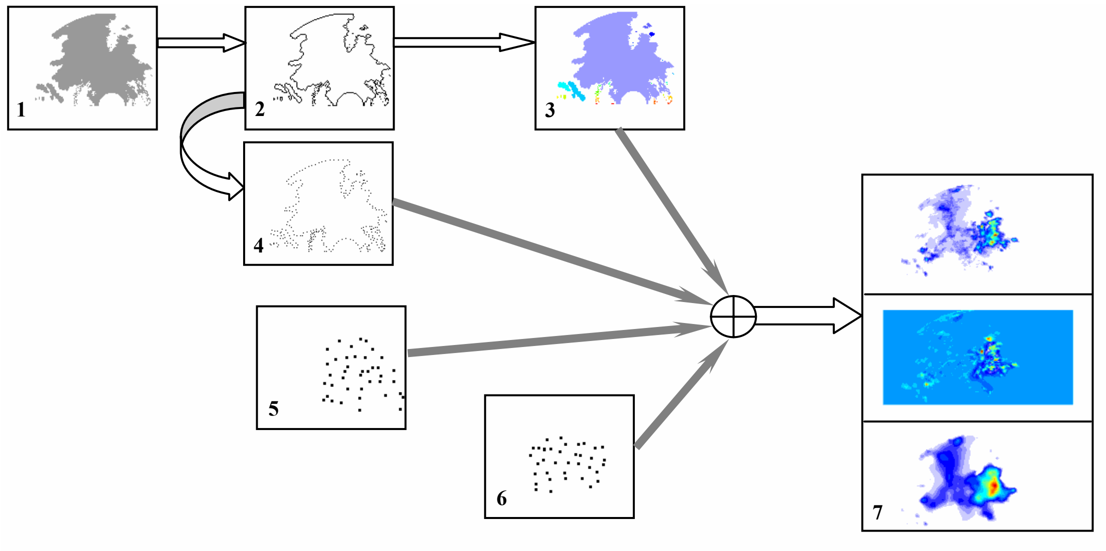

Figure 6.9 Summary of data processing; 1. Mask the wet and dry areas; 2. Trace the boundary of each wet region; 3. Separately label each wet region; 4. Decimate the fence by a factor of 5, isolating the ‘fence posts’; 5. Isolate the maxima in each sub-region; 6. Isolate the minima in each sub-region; 7. EMD analysis decomposes the data into the first IMS and the first residual using the maximal and minimal envelopes defined using the points in 4, 5 & 6.

An outer boundary border-tracing algorithm (Sonka et al., 1999) is used to establish a boundary “fence” around each raining area (Figure 6.9, Box 2) and a flood-fill procedure is then used to fill each raining area with a unique identifier, resulting in separately labelled raining regions (Figure 6.9, Box 3). To reduce the computational burden of the algorithm even further, the boundary “fence” is decimated by a factor of 5 to reduce the continuous string of border points to “fence posts” while retaining the gross shape of the raining areas (Figure 6.9, Box 4). The next step in the processing of the data is to isolate the extrema in the rainfall field (Figure 6.9, Boxes 5 & 6). Finally, the EMD analysis is carried out using the extrema within each raining area and the zeros at the "fence posts" of non-raining border pixels to specify the extremal envelopes (Figure 6.9, Box 7). Only one step of decomposition is shown here – the data is decomposed into the noisy first IMS and the first residual.

Results

An analysis of over 800 individual radar scans, embodying mixtures of various ratios of Stratiform and Convective rainfall types, was carried out to determine the effectiveness of the 2-D EMD algorithm in separating the high frequency spatial components from the original rainfall data. Working on the basis that the average characteristics of the data over a range of spatial scales summarised by the power spectrum is intuitively useful, the (radially averaged) power spectra of (i) the original data, (ii) the first IMS and (iii) the first residual of each image were examined and compared. Figure 6.10 shows a typical result; the power spectrum of the residual shows a very close correspondence with that of the original data at low wave numbers while it contains far less power at the higher wave numbers. In contrast, the spectrum of the first (noisy) IMS has very little power relative to the data's spectrum at low wave numbers but shows a strong correspondence at the highest wave numbers. Figure 6.10 clearly indicates how the 2-D EMD technique moves the bulk of the high frequency components in the original data into the first IMS and leaves the high power, lower frequencies in the residual. Figure 6.11 shows a time average of this behaviour by plotting the mean values at each wave number of the three spectra over five consecutive radar scans (beginning with the data used to produce Figure 6.10). The radar scans are captured at approximately five-minute intervals.

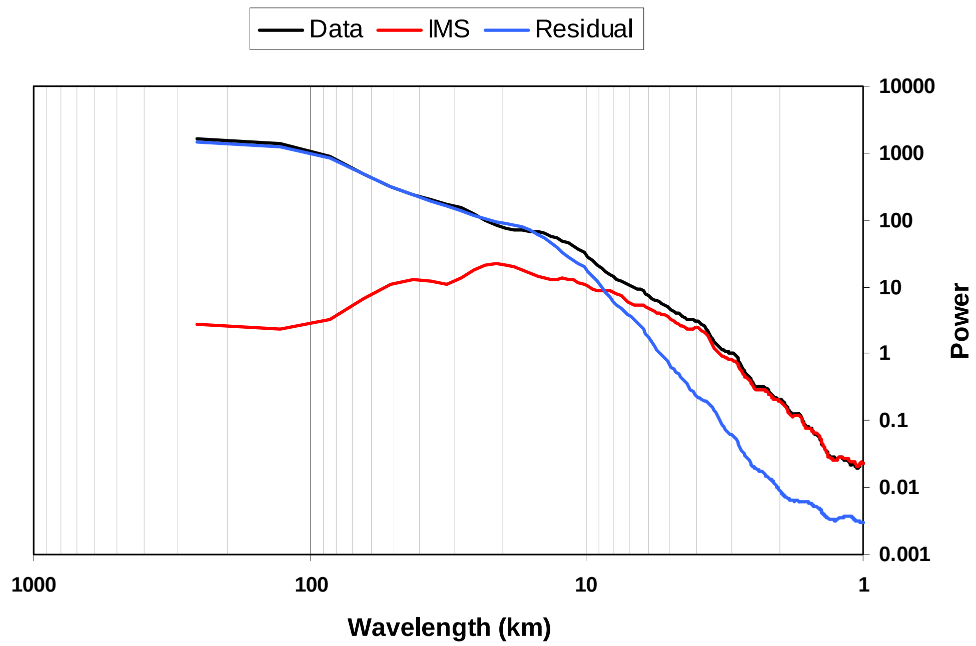

Figure 6.10 Comparison of individual radially averaged power spectra of the radar rainfall data (of Figure 6.1) with its EMD components: the first IMS and the first residual

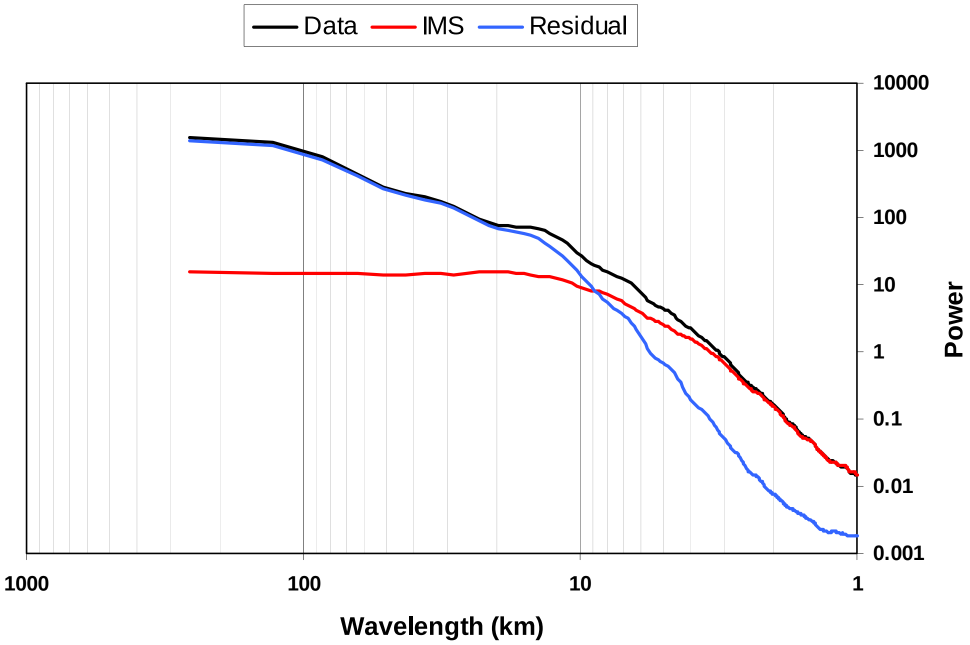

Figure 6.11 The same as Figure 6.10 but for the mean of individual power spectra for five consecutive, radar scans – Beginning with the spectra shown in Figure 6.10.

The temporal persistence exhibited at the spatial scales represented in each of the three sequences of: (i) the data, (ii) the first IMS and (iii) the first residual was examined by considering their temporally consecutive power spectra. The notion of "spectral persistence" was used to determine how variable the spatial structure (at a particular spatial scale) is in time and hence to give an indication of the temporal predictive capability at each spatial scale. A summarised example of the analysis of a sequence of 5 radar rainfall images is presented in Figures 6.12, 6.13 and 6.14 where a matrix of scatter plots is shown in each case. Scatter-plots of the power at each wave number for five consecutive spectra (with the 1:1 line indicated) are shown for, the original data (Figure 6.12), the first IMS (Figure 6.13) and the first residual (Figure 6.14). The rows and columns of the scatter-plot matrices are labelled from T_0 to T_4 and indicate separate radar scans between time $T=0$ and time $T=4$. Each block in the scatter-plot matrix represents a scatter-plot of the power at each wave number for the spectrum computed at T_i versus that of the spectrum computed at T_j. Clearly the plots on the diagonal each compare a spectrum to itself and a perfect 1:1 relationship is observed in this case. For the off-diagonal plots, the degree of scatter amongst the data points indicates the degree of similarity between the spectra at different time lags with a large scatter indicating a weak similarity. The trends shown here are typical of the data analysed and show how the first (high average frequency) IMS has a temporally incoherent spatial structure, while the first (low average frequency) residual shows a temporally consistent structure. The behaviour shown in Figures 6.10 - 6.14 suggests that the high frequency IMS components in spatial rainfall data do not contain much predictive capability, supporting the suggestions of Seed (2001) and Turner et al. (2004) to rely on the information contained in the lower frequency components as forecast lead time increases.

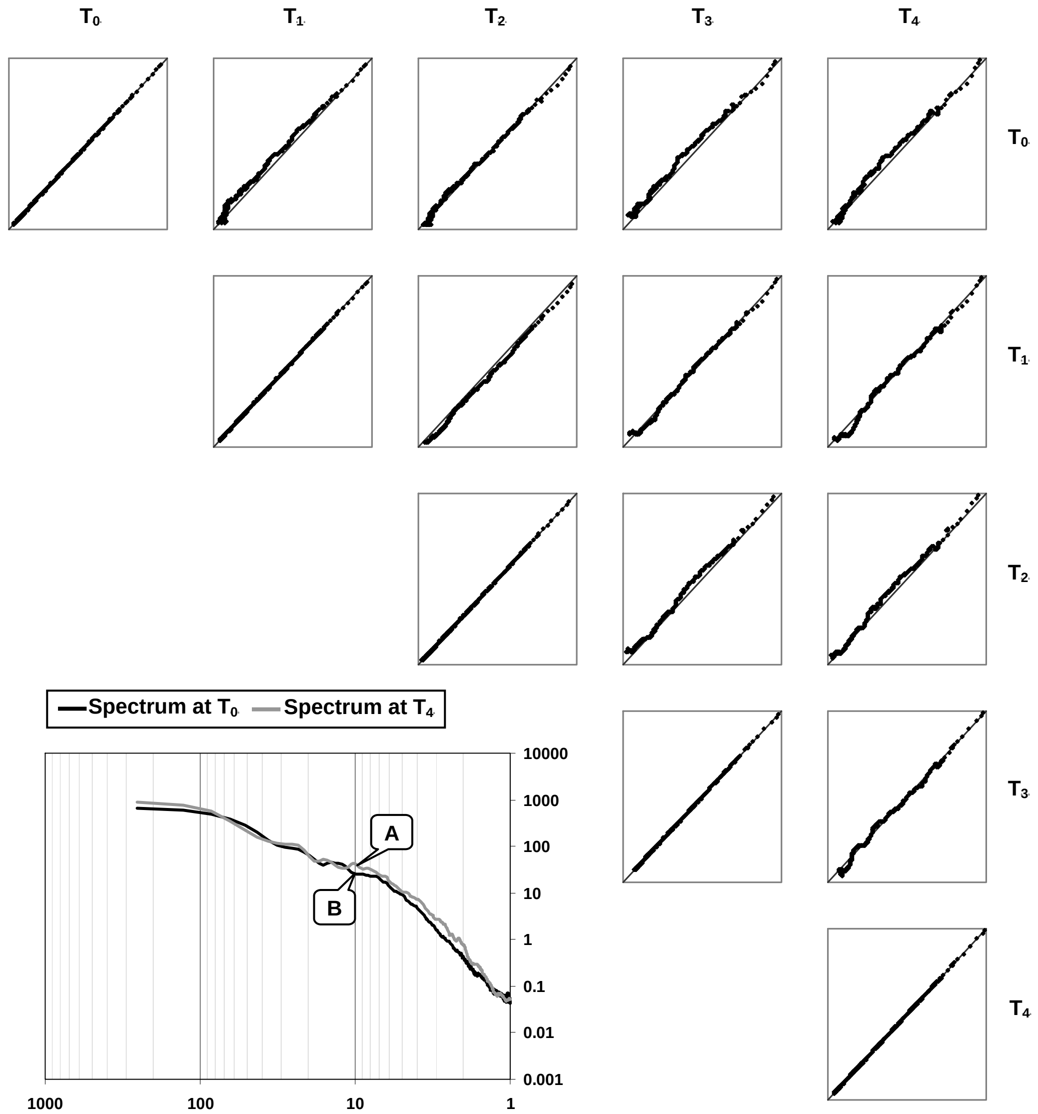

Figure 6.12 Spectral persistence scatter plots of the original data for a sequence of rainfall fields and those at successive intervals. This is constructed by plotting the values of power for each field at corresponding wavelengths coaxially. For example points A and B at the 10km wavelength are plotted against each other and appear ringed in the upper right diagram.

6.2.5 Conclusion

A new technique for analysing the spatial scaling structure of rainfall fields has been presented. The technique is a two dimensional extension of Empirical Mode Decomposition for the analysis of non-linear and non-stationary time series. An EMD analysis in two dimensions linearly decomposes the spatially distributed rainfall data into a set of Intrinsic Mode Surfaces, which are approximately mutually orthogonal and sum back to the original data. Each IMS contains an oscillatory mode inherent in the data at a different (narrow) range of spatial frequencies. The EMD analysis successively extracts the IMS with the highest local spatial frequencies in a recursive way, which is effectively a set of successive low-pass spatial filters based entirely on the properties exhibited by the data. The utility of the EMD technique for signal separation has been demonstrated in both one and two dimensions and applied to the analysis of a large set of 800 radar rainfall images in South Africa. The 2-D EMD technique is proposed here in the context of rainfall nowcasting to separate the less persistent high frequency components from the more persistent low frequency ones in the data. The aim is to remove the noisy high frequency components, which do not exhibit a strong temporal correlation and add little structural information to nowcasting algorithms. The scale separation achieved by 2-D EMD has been analysed using radially averaged power spectra to summarise the spatial structure of the data and filter outputs. In addition these power spectra have also been used to examine the temporal persistence of the spatial structure exhibited by the first IMS and it's residual. The results presented in this paper agree with other work in the Hydrometeorological literature, which suggests that the low frequency spatial components in rainfall data are most useful in a nowcasting context. This methodology is being exploited in ongoing research into rainfall nowcasting.

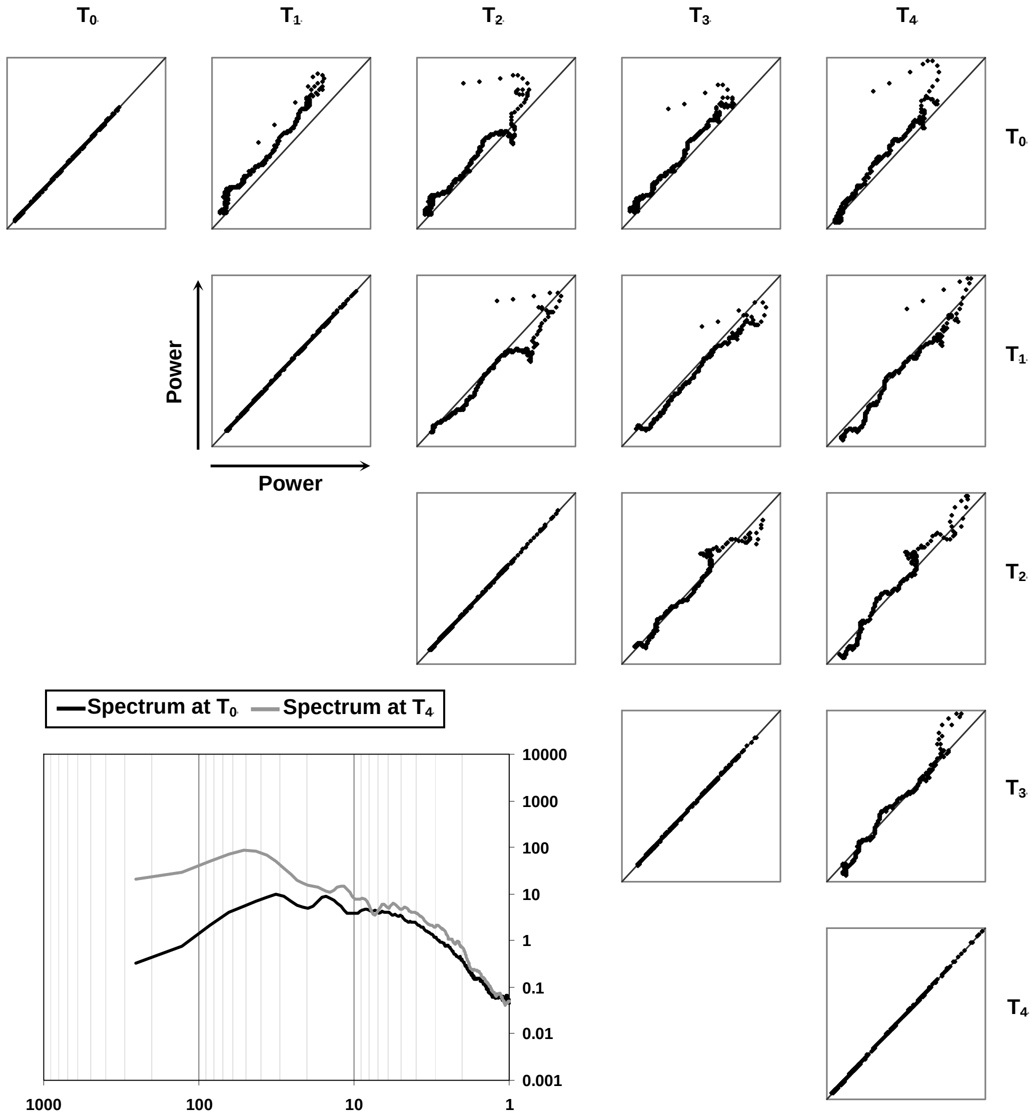

Figure 6.13 Spectral persistence scatter plots of the sequence of 1st IMSs of each pair of rainfall fields $T_0,...,T_4$.

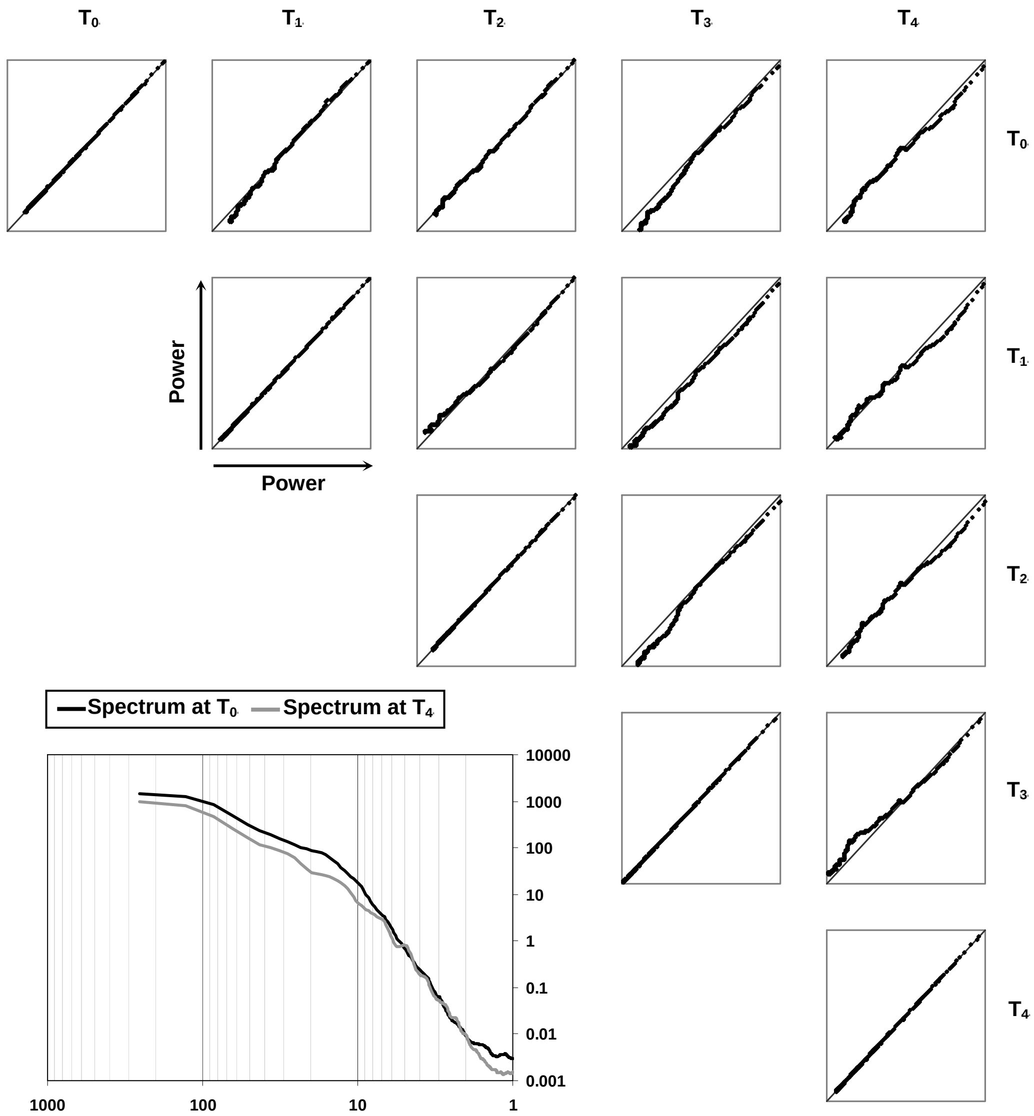

Figure 6.14 Spectral persistence scatter plots of the 1st Residual for each pair of rainfall fields $T_0,...,T_4$.

7 Conclusion

The triple umbrella WRC project called SIMAR (Kroese, 2004; Deyzel et al., 2004; Pegram, 2004) developed a computational procedure to provide a map of the previous 24 hour rainfall over the Southern African region, on a daily basis, on the SAWS website. This product has been operating routinely since early 2003, providing detailed rainfall information over the subcontinent. In its conception, it was a brave and ambitious project, enabling those involved to learn much from its routine operation, particularly concerning its shortcomings.

The work presented in this report is aimed at improving the daily mapping of rainfall, in particular through the development of improved algorithms, designed to make the product more accurate and trustworthy. To this end, the algorithms and techniques developed have been designed to build on SIMAR's strengths and minimise its weaknesses. This was done in the following ways:

- Some of the significant difficulties affecting rainfall estimation by radar have been addressed. These include, infilling of data contaminated by ground clutter, extrapolation of data (horizontally and from above) into "no data" areas and the extrapolation of radar data down to ground level. The result of the algorithms described in section 2 is an uncontaminated estimate of rainfall at ground level.
- The MSRR satellite rainfall estimation algorithm developed for SIMAR has been adapted to incorporate MET-8 data as described in section 3. Although this is a relatively new development (MET-8 only became available in early 2005) it is an important one for two reasons. Firstly, the current MET-7 data transmission will cease from February 2006 and second the increased spatial, temporal and spectral resolution of the data collected by the MET-8 satellite offer many opportunities to improve the satellite rainfall product even further. The work reported here has been an important first step in that direction.
- The credible accumulation of spatial rainfall fields (of the type produced by radar and satellite) is of great significance when these fields must be compared to rain gauge totals. Section 4 describes an accumulation algorithm which we believe to be novel in its approach to controlling the complexity of the integration scheme, based on advection vectors measured from the data.
- In section 5, the major issues pertaining to the blending of rainfall estimate from a variety of sensors are discussed and the Conditional Merging technique adopted as a basis for IRMA, a new data merging algorithm. IRMA goes some way to addressing the difficulties associated with data blending.
- In the final section (section 6), nowcasting of spatial rainfields was discussed and an extension of the EMD analysis technique presented which has great promise for rainfall nowcasting and analysis.

There was considerable spin-off and cross-pollination with other initiatives, especially the companion project K5/1429: a National Flood Nowcasting System (which benefits directly from the work presented in this report). There is a natural progression from daily rainfall maps at 1km resolution over the whole country to smaller areas of concern at higher sampling frequency and even rainfield nowcasting. Beneficiaries of these products will be everyone reliant on sound rainfall information for decision-making.

It is thus envisaged that catchment management authorities, local government disaster managers doing flood forecasting, agriculturists estimating drought and crop water demand, transport operators fearful of bad weather, eco-tourism operators checking the safety of their schemes to raft rivers and sports administrators dependent on the weather for their scheduling, will all benefit from the output from this project.

The stated aims of this research project were:

1. Combine satellite, radar and rain gauge estimates of rainfall into a meaningful composite to provide an improved optimum rainfield to users via the Internet.
2. To improve the estimation methods initiated in SIMAR by using MSG (METEOSAT Second Generation) and novel methods of field merging, notably Kriging with the Kalman Filter and conditional Geostatistical merging.
3. To continue to seek numerical techniques to reduce the huge computational load and accelerate convergence in producing large maps.
4. To collect, combine, archive and disseminate the raw data and the product fields through the efficient use of the IT division of NDMC/DWAF:PSU (National Disaster Management Centre/Department of Water Affairs & Forestry: Public Safety Unit).

The first three have been achieved successfully; the fourth was not completed because of difficulties in establishing communication links, as well as the responsibility for maintaining the information flow.

7.1 Recommendations

There was concern expressed at the final steering committee meeting that the access to the SIMAR web-site was limited to pictorial information, while numerical data of rainfalls on quaternary catchments was sought by some; indeed, gridded data at 1 km resolution would be preferred.

Two commitments rose from this concern: the first was to install the new software developed in this study; the second was to provide access to the data to those with a legitimate need, on approach to SAWS.

In more detail, to implement the first commitment, Stephen Wesson with the help of Scott Sinclair will adapt the research software to on-line code. Stephen will then spend as much time in Bethlehem as is necessary, before the end of 2005, to transfer the code and ensure with the help of Pieter Visser and Karel de Waal that it operates satisfactorily in on-line mode in real time.

Finally, it follows naturally from this study that there is a need for validation of the rainfall fields produced by the project: Daily Rainfall Mapping over South Africa: Modelling (DARAM). Towards this end, an invitation was sent to Professor Geoff Pegram by Drs Joe Turk and Peter Bauer, co-chairpersons of the International Precipitation Working Group (IPWG), to participate in the Pilot Evaluation of High Resolution Precipitation Products (PEHRPP), sponsored by the IPWG. From their letter of invitation: “PEHRPP is an effort to characterise as clearly as possible the errors in various high resolution precipitation products on many spatial and temporal scales over varying surfaces and climate regimes.”

The work started in SIMAR and continued in this study will provide a platform to make a valuable contribution to the global effort initiated by IPWG, because it will force those of us responsible for the rainfall estimation to closely evaluate the errors in our own products. This can only be good for making future improvements to this important work.

References:

Aitken A.C., Determinants and Matrices, *Oliver and Boyd Ltd.*, London, 1967.

Anagnostou E.M. and Krajewski W.F., Real-Time Radar Rainfall Estimation. Part I: Algorithm Formulation, *Journal of Atmospheric and Oceanic Technology*, **16**, 189-197, 1999a.

Anagnostou E.M. and Krajewski W.F., Real-Time Radar Rainfall Estimation. Part II: Case Study, *Journal of Atmospheric and Oceanic Technology*, **16**, 198-205, 1999b.

Bab-Hadiashar A., Suter D. and Jarvis R., 2-D Motion extraction using an image interpolation technique, *SPIE*, **2564**, 271-281, 1996.

Balocchi R., Menicucci D., Santarcangelo E., Sebastiani L., Gemignani A., Gellarducci B. and Varanini M., Deriving the respiratory sinus arrhythmia from the heartbeat time series using empirical mode decomposition, *Chaos, Solitons and Fractals*, **20(1)**, 171 – 177, 2004.

Bell T.L., A Space-Time Stochastic Model of Rainfall for Satellite Remote Sensing Studies, *J. Geophys. Res.,* **92(D8)**, p9631, 1987.

Brandes E.A., Optimizing rainfall estimates with the aid of Radar, *Journal of Applied Meteorology*, **14**, 1339 – 1345, 1975.

Bras R.L. and Rodriguez-Iturbe I.,. Random Functions and Hydrology, *Addison-Wesley*, Ontario, 1985.

Ciach G. J., Local Random Errors in Tipping-Bucket Rain Gauge Measurements. *Journal of Atmospheric and Oceanic Technology*, Notes and Correspondence, **20**, 752 – 759, 2002.

Chiew F.H.S., Peel M.C., Amirthanathan G.E. and Pegram G.G.S., Identification of oscillations in historical global streamflow data using empirical mode decomposition, *Regional Hydrological Impacts of Climatic Change - Hydroclimatic Variability, Proceedings of symposium S6 held during the Seventh IAHS Scientific Assembly at Foz do Iguaçu, Brazil*, IAHS Publ. 296, 53 – 62, 2005.

Chiles J.P. and Delfiner P., Geostatistics: Modelling spatial uncertainty, *Wiley*, New York, 1999.

Chumchean S., Sharma A. and Seed A., Radar rainfall error variance and it's impact on radar rainfall calibration, *Physics and Chemistry of the Earth*, **28**, 27 – 39, 2003.

Clothier A.N. and Pegram G.G.S., Space-time modeling of Rainfall using the String of Beads model: Integration of Radar and Raingauge data, *WRC report No. 1010/1/02 to the Water Research Commission*, Pretoria, South Africa, 2002.

Clothier A.N., High Resolution Space-time modeling of Rainfall: The String of Beads Model, *PhD Thesis*, University of Natal, Durban, South Africa, 2003.

Cowpertwait P.S.P., Further developments of the Neyman-Scott clustered point process for modeling rainfall, *Water Resources Research*, **27(7)**, 1431 – 1438, 1991.

Cressie N. and Hawkins D.M., Robust Estimation of the Variogram, 1. *Journal of the International Association for Mathematical Geology*, **12**, 115 – 125, 1980.

Cressie N., Statistics for Spatial Data, *Wiley*, New York, 1991.

Creutin J.D. and Obled C., Objective Analysis and Mapping Techniques for Rainfall Fields: An Objective Comparison, *Water Resources Research*, **18(2)**, 413 – 431, 1982.

Daubechies I., Ten Lectures on Wavelets, *Society for Industrial and Applied Mathematics*, Philadelphia, Pennsylvania, 1992.

Deyzel I.TH., Pegram G.G.S., Visser P.J.M. and Dicks D., Spatial Interpolation and Mapping of Rainfall (SIMAR) Volume 2: Radar and satellite products, *WRC report No. 1152/1/04 to the Water Research Commission*, Pretoria, South Africa, 2004.

Ehret U., Rainfall and flood nowcasting in small catchments using weather radar, *PhD Thesis*, University of Stuttgart, Stuttgart, Germany, 2002.

Engelbrecht F.A., deW Rautenbach C.J., McGregor J.L. and Katzfey J.J., January and July climate simulations over the SADC region using the limited-area DARLAM, *Water SA*, **28(2)**, 361 – 374, 2002.

Flandrin P., Rilling G. and Goncalves P., Empirical Mode Decomposition as a filter bank, *IEEE Signal Processing Letters*, **11(2)**, 112 – 114, 2004.

Franco M., Sempere-Torres D., Sanchez-Diezma R. and Andrieu H., A methodology to identify the vertical profile of reflectivity from radar scans and to estimate the rainrate at ground at different distances, *Proceedings of ERAD (2002)*, 299 – 304, 2002.

Fulton R.A., Breidenbach J.P., Seo D.J., Miller D.A. and O'Bannon T., The WSR-88D rainfall algorithm, *Weather and Forecasting*, **13**, 377 – 395, 1998.

Gloersen P. and Huang N., Comparison of Interannual Intrinsic Modes in Hemispheric Sea Ice Covers and Other Geophysical Parameters, *IEEE Transactions on Geoscience and Remote Sensing*, **41(5)**, 1062 – 1074, 2003.

Gordon A.D., Classification: Methods for the exploratory analysis of multivariate data, *Chapman and Hall*, New York, USA, 1981.

Gotway C.A and Young L.J., Combining incompatible spatial data, *Journal of the American Statistical Association*, **97(458)**, 632 – 648, 2002.

Habib E. and Krajewski W.F., Uncertainty Analysis of the TRMM Ground-Validation Radar-Rainfall Products: Application to the TEFLUN-B Field Campaign, *Journal of Applied Meteorology*, **41**, 558 – 572, 2002.

Han C., Guo H., Wang C. and Fan D., A novel method to reduce speckle in SAR images, *International Journal of Remote Sensing*, **23(23)**, 5095 – 5101, 2002.

Hannesen R., An enhanced surface rainfall algorithm for Radar Data, *Progress report for MUSIC*, European Commission contract No. EVK1-CT-2000-00058, 2002.

Haralick R.M., Statistical and structural approaches to texture, *Proceedings IEEE* **67(5)**, 786 – 803, 1979.

Harris D., Foufoula-Georgiou E., Droegemeier K.K. and Levit J.J., Multiscale Statistical Properties of a High-Resolution Precipitation Forecast, *Journal of Hydrometeorology*, **2**, 406 – 418, 2001.

Hartkamp A.D., De Beurs K., Stein A. and White J.W., Interpolation techniques for Climate variables. *NRG-GIS Series 99-01Mexico, D.F.:* CIMMYT, 1999.

Hengl T., Geuvelink G.B.M. and Stein, A., Comparison of Kriging with external drift and regression-kriging, *Technical note, ITC*, Available online at http://www.itc.nl/library/Academic_output/, 2003.

Huang N.E., Shen Z., Long S.R., Wu M.C., Shih H.H., Zheng Q., Yen N., Tung C.C. and Liu H.H., The Empirical Mode Decomposition and the Hilbert Spectrum for Nonlinear and Non-Stationary Time Series Analysis, *Proceedings: Mathematical, Physical and Engineering Sciences*, Royal Society London, **454(1971)**, 903 – 995, 1998.

Huang N.E., Shen Z. and Long S.R., A new view of Nonlinear water waves: the Hilbert Spectrum, *Annual Review of Fluid Mechanics*, **31**, 417 – 457, 1999.

Hwang P.A., Huang N.E. and Wang D.W., A note on analyzing nonlinear and nonstationary ocean wave data, *Applied Ocean Research*, **25(4)**, 187 – 193, 1999.

Jordan P., Seed A. and Austin, G., Sampling errors in radar estimates of rainfall. *Journal of Geophysical Research*, **105(D2)**, 2247 – 2257, 2000.

Journel A.G. and Huijbrechts Ch. J., Mining Geostatistics, *Academic Press,* London, UK, 1978.

Kedem B., Pfeiffer R. and Short D.A., Variability of Space-Time mean rain rate, *Journal of Applied Meteorology*, **36**, 443 – 451, 1997.

Krajewski W.F., Cokriging Radar-Rainfall and Rain Gauge Data, *Journal of Geophysical Research*, **92**, 9571 – 9580, 2002.

Kroese N.J., Spatial Interpolation and Mapping of Rainfall (SIMAR) Volume 1: Maintenance and upgrading of radar and rain gauge infrastructure, *WRC report No. 1151/1/04 to the Water Research Commission*, Pretoria, South Africa, 2004.

Lack S.A. and Fox N.I., Errors in Surface Rainfall Rates Retrieved from Radar due to Wind Drift, *Sixth International Symposium on Hydrological Applications of Weather Radar*, Melbourne, Australia, 2004.

Lebel T., Bastin G., Obled C. and Creutin J.D., On the Accuracy of Areal Rainfall Estimation: A Case Study, *Water Resources Research*, **23(11)**, 2123 – 2134, 1987.

Linderhed A., 2-D empirical mode decompositions – in the spirit of image compression, *Proceedings of SPIE, Wavelet and Independent Component Analysis Applications IXI*, Volume 4738, 2002.

Linderhed A., Variable Sampling of the Empirical Mode Decomposition of Two-Dimensional Signals, Preprint, Special issue on Sampling and frames in Wavelet Theory and Time-frequency analysis, *International Journal of Wavelets, Multiresolution and Image Processing*, 2004.

Marshall J.S. and Palmer W.M., The distribution of raindrops with size, *Journal of Meteorology*, **5**, 165 – 166, 1948.

Mather G.K., Terblanche D.E. and Steffens F.E., National Precipitation Research Programme: Final Report for the Period 1993 – 1996, *WRC report No.* 726/1/97 *to the Water Research Commission*, Pretoria, South Africa, 1997.

Matheron G., Traite de Geostatistique Appliquee, Tome I. *Memories du Bureau de Recherches Geologiques et Minieres*, **No. 14**. Editions Technip, Paris, 1962.

Mitasova H. and Mitas L., Interpolation by Regularized spline with Tension: I. Theory and Implementation, *Mathematical Geology*, **25**, 641 – 655, 1993.

Mittermaier M.P., Investigating the Characteristics of the Radar Reflectivity Profile, *MSc Eng dissertation*, Civil Engineering, University of KwaZulu-Natal, Durban, Howard College, 1999.

Mittermaier M.P., Investigating synergies between weather radar data and mesoscale model output. *PhD Thesis*, Department of Meteorology, University of Reading, 2003.

Nunes J.C., Bouaoune Y., Delechelle E., Niang O. and Bunel P., Image analysis by bidimensional empirical mode decomposition, *Image and Vision Computing*, **21**, 1019 – 1026, 2003.

Pegram G.G.S. and Clothier A.N., High resolution space-time modeling of rainfall: the “String of Beads” model, *WRC report No. 752/1/99 to the Water Research Commission*, Pretoria, South Africa, 1999.

Pegram G.G.S. and Clothier A.N., High Resolution Space-Time Modelling of Rainfall: The “String of Beads” Model, *Journal of Hydrology*, **241**, 26 – 41, 2001.

Pegram G.G.S and Sinclair S., A Linear Catchment Model for Real Time Flood Forecasting, *WRC report No. 1005/1/02 to the Water Research Commission*, Pretoria, South Africa, 2002.

Pegram G.G.S., Spatial interpolation and mapping of rainfall: 3. Optimal integration of rain gauge, radar & satellite-derived data in the production of daily rainfall maps, *Progress report to the Water Research Commission*, for the period April 2001 to March 2002 on contract K5/1153, April 2002.

Pegram G.G.S., Spatial Interpolation and Mapping of Rainfall (SIMAR) Volume 3: Data Merging for rainfall map production, *WRC report No. 1153/1/04 to the Water Research Commission*, Pretoria, South Africa, 2004.

Press W.H., Teukolsky S.A., Vetterling W.T. and Flannery B.P., Numerical Recipes in C –The Art of Scientific Computing, 2nd Edition, *Cambridge University Press*, Cambridge, 1992.

Sabyasachi B., Gunst R.F., Guertal E.A. and Hartfield M.I., The Effects of Influential Observations on Sample Semivariograms, *Journal of agriculture, biological, and environmental statistics*, **2**, 490 – 512, 1997.

Sanchez-Diezma R., Zawadzki I. and Sempere-Torres D., Identification of the bright band through the analysis of volumetric radar data, *Journal of Geophysical Research,* **105(D2)**, 2225 – 2236, 2000.

Schertzer D. and Lovejoy S., Physical Modelling and Analysis of Rain and Clouds by Anisotropic Scaling Multiplicative Processes, *Journal of Geophysical Research*, **92(D8)**, 9693 – 9714, 1987.

Seed A.W., Srikanthan R. and Menabde M., A space and time model for design storm rainfall, *Journal of Geophysical Research*, **100(D21)**, 31623 – 31630, 1999.

Seed A.W., A dynamic and spatial scaling approach to advection forecasting, *Proceedings of the Fifth International Symposium on Hydrological Applications of Weather Radar – Radar Hydrology*, Kyoto - Japan, 311 – 316, 2001.

Seed A.W. and Pegram G.G.S., Using Kriging to infill gaps in radar data due to ground clutter in real time,. *Proceedings Fifth International Symposium on Hydrological Applications of Weather Radar – Radar Hydrology*, Kyoto, Japan, 73 – 78, 2001.

Sempere-Torres D., Sanchez-Diezma R., Zawadzki I. and Creutin J.D., Identification of Stratiform and Convective Areas Using Radar Data with Application to the Improvement of DSD Analysis and Z-R Relationships, *Phys. Chem. Earth*, **25**, 985 – 990, 2000.

Seo D.J., Real-time estimation of rainfall fields using rain gage data under fractional coverage conditions, *Journal of Hydrology*, **208**, 25 – 36, 1998a.

Seo D.J., Real-time estimation of rainfall fields using radar rainfall and rain gage data, *Journal of Hydrology*, **208**, 37 – 52, 1998b.

Sinclair D.S., Ehret U., Bardossy A. and Pegram G.G.S., Comparison of Conditional and Bayesian Methods of Merging Radar & Raingauge Estimates of Rainfields, *Oral Presentation EGS/AGU/EUG Joint Assembly*, Nice, France, April, 2003.

Sinclair S. and Pegram G.G.S, Combining Traditional and remote sensing techniques as a tool for Hydrology, Agriculture and Water Resources Management, *Proceedings of the 11th South African National Hydrology Symposium*, Port Elizabeth, South Africa, 2003a.

Sinclair S. and Pegram G.G.S, Combining Radar and Rain Gauge Rainfall Estimates for Flood Forecasting in South Africa, *Proceedings of the 6th International Symposium on Hydrological Applications of Weather Radar,* Melbourne, Australia, 2-4 February 2004a.

Sinclair S. and Pegram G.G.S, A Flood Nowcasting System for the eThekwini Metro. Volume 1: Umgeni Nowcasting using Radar – An integrated pilot study, *WRC report No. 1217/1/04 to the Water Research Commission*, Pretoria, South Africa, 2004b.

Sinclair S. and Pegram G., Combining radar and rain gauge estimates using conditional merging, Atmos. Sci. Let., **6**, 19 – 22, 2005.

Smith J.A. and Krajewski W.F., Estimation of the Mean Field Bias of Radar Rainfall Estimates, *Journal of Applied Meteorology*, **30**, 397 – 412, 1991.

Sonka M., Hlavac V. and Boyle R., Image Processing, Analysis, and Machine Vision, 2nd Edition, *PWS Publishing*, Pacific Grove, 1999.

Steiner M., Houze R.A. and Yuter S.E., Climatological Characterization of Three-Dimensional Storm Structure from Operational Radar and Rain Gauge Data, *Journal of Applied Meteorology*, **34**, 1978 – 2007, 1995.

Terblanche D.E., Pegram G.G.S. and Mittermaier M.P., The development of weather radar as a research and operational tool for hydrology in South Africa, *Journal of Hydrology*, **241**, 3 – 25, 2001.

Todini E., A Bayesian technique for conditioning radar precipitation estimates to rain-gauge measurements, *Hydrology and Earth System Sciences*, **5(2)**, 187 – 199, 2001.

Turner B.J., Zawadzki I. and Germann U., Predictability of precipitation from continental radar images. Part III: Operational nowcasting implementation (MAPLE), *Journal of Applied Meteorology*, **43**, 231 – 248, 2004.

Vezzani G., Ground Clutter Suppression in Atmospheric Phenomena Observation, *COST 75 Weather Radar Systems International Seminar*, Brussels, Belgium, 281 – 295, 1994.

Visser P., Spatial interpolation and Mapping of rainfall: 2. Radar and Satellite products, *Draft final report to the Water Research Commission for the period March 2000 to March 2003*, 2003.

Wesson S. and Pegram G., Efficient cleansing of Radar Rainfall images, *Proceedings of the 11th South African National Hydrology Symposium,* Port Elizabeth, South Africa, 2003.

Wesson S.M. and Pegram G.G.S., Radar rainfall image repair techniques, *Hydrology and Earth System Sciences*, **8(2)**, 220 – 234, 2004.

Wikle C.K. and Berliner L.M., Combing information across spatial scales, *Technometrics*, **47**, 80 – 91, 2005.

Wilkinson J.H., The Algebraic Eigenvalue Problem, *Oxford Science Publication*, 1988.

Wilson J.W. and Brandes E.A., Radar measurement of rainfall – A summary, *Bulletin of the American Meteorological Society*, **60**(9), 1048 – 1058, 1979.

Zhang R.R., Ma S. and Hartzell S., Signatures of the seismic source in EMD-based characterization of the 1994 Northridge, California, earthquake recordings, *Bulletin of the Seismological Society of America*, **93(1)**, 501 – 518, 2003.

8 Appendix: Adaptation of Kriging to Image Infilling & extrapolation

To remove ground clutter and estimate the missing rainfall data in real time, a computationally fast and accurate technique is required. Kriging was chosen as the method for estimating the missing data. However for Kriging to be used in a real time application one needs to take advantage of various factors such as the pattern of the Kriging Weights in the "good" or control data pixels surrounding the Ground clutter. By taking these patterns into account, an effective and efficient method for estimating the missing target rainfall values can be provided.

Standard techniques to remove ground clutter include subtracting a "clutter map", derived from a period of no precipitation, from the observed reflectivity values. Unfortunately ground clutter is not totally stationary and its location can vary depending on atmospheric conditions (Vezzani, 1994).

The Kriging technique mentioned briefly above, and described in more detail later, is proposed as an alternative. The methodology is presented in this paper along with relevant examples of its application.

There are various quality problems associated with radar rainfall data viewed in images that include ground clutter, Beam Blocking and Anomalous Propagation, to name a few. To obtain the best rainfall estimate possible, techniques for removing ground clutter (non-meteorological echoes that influence radar data quality) on 2D radar rainfall image data sets have been developed and are presented here. These techniques concentrate on repairing the images in both a computationally fast and accurate manner, and are nearest neighbour techniques of two sub-types:

- Individual Target
- Border Tracing

The chosen method for estimating the contaminated data is Kriging, which is considered to be the optimal technique for the spatial interpolation of Gaussian data. Ordinary Kriging has various advantages and disadvantages, which need to be taken into consideration in this application.

For the radar rainfall images to be repaired in real time one needs a computationally fast and efficient method of estimating the missing contaminated data. This can be achieved by exploiting the "screening effect" that occurs with the Kriging weighting distribution around target points. Matrix rank reduction techniques in combination with Singular Value Decomposition (SVD) are also suggested for finding an efficient solution to the Kriging Equations which can cope with near singular systems. Nearest neighbours strategies and SVD both lead to considerable reductions in computation time but with no significant loss in accuracy of the estimated values.

8.1 Brief Overview of Ordinary Kriging Theory

Kriging was chosen as the computational method for infilling the contaminated rainfall data. Kriging is used to estimate the missing data as it is the optimal technique for the spatial prediction of Gaussian data (Cressie, 1993: 106) if the sum of estimation errors squared is to be minimized. Reflectivity values of rainfall have been shown to be approximately normally distributed, because the rain rate on an image follows a truncated lognormal distribution (Bell, 1987), (Pegram and Clothier, 2001). Thus, in this application, the estimation of the missing data by Kriging is carried out in the reflectivity rather than the rain rate domain. Kriging has several advantages over other spatial prediction techniques. The basis function used in Kriging is determined by the data set and its spatial structure; in this way the actual nature of the data set is taken into account. In other well known interpolation techniques such as Multiquadrics, Fourier series or splines the basis function is merely chosen for computational convenience.

Ordinary Kriging is used here where the mean is assumed unknown and is implicitly estimated from the data. The Ordinary Kriging equations used herein are given in Eqn. (8.1):

$$z(s_0) = \lambda^T(s_0) \cdot z \tag{8.1}$$

where $z(s_0)$ is the estimated data value at the target location $s_0 = (x_0, y_0)$, $\lambda\,(s_0)$ is the vector of Kriging Weights associated with the location s_0 and z is the vector of known or control data. This computation is carried out to estimate the reflectivity at each target pixel in turn. The weights are computed from Eqn (8.2):

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} \cdot \begin{bmatrix} \lambda(s_0) \\ \mu(s_0) \end{bmatrix} = \begin{bmatrix} g(s_0) \\ 1 \end{bmatrix} \tag{8.2}$$

where u is a unit vector of ones and μ a Lagrange multiplier. In Eqn. (8.2) $g(s_0)$ is the vector of covariances or semivariogram values between the target location (data point to be estimated) and the control points (known data values z) and G is the matrix of covariances or semivariogram values between the control points (the covariance or semivariogram function is estimated from the data). Eqn. (8.2) is derived by minimizing the estimation error variance and the Lagrange multiplier is included to constrain the sum of the weights to unity. The variogram is defined as the expectation of the square of the differences of the field variables as shown in Eqn. (8.3):

$$2 \cdot \gamma(x, h) = E\left\{ [Z(x) - Z(x+h)]^2 \right\} \tag{8.3}$$

the semivariogram is defined as γ (x, h) (Journel and Huijbregts, 1978: 11).

In a stationary random field Ordinary Kriging can be carried out by using either a covariance or semivariogram function, where one is the complement of the other with reference to the field variance. Mathematically they are equivalent in this context, however in some applications the variogram is more robust (Cressie, 1993: 70-73), so was used throughout in preference to the covariance. The isotropic semivariogram model chosen is the 2-parameter exponential model specified in Eqn. (8.4) and is used in all Kriging computations herein.

$$g(h) = 1 - \exp[-(h/L)^{\alpha}] \tag{8.4}$$

where s is the Euclidian distance between data points, L is a scale parameter we call the correlation length and α is the exponent parameter, which lies in the range $0 < \alpha \le 2$; the model is thus Gaussian when $\alpha = 2$.

The advantages of the Kriging technique can be listed as follows: it is considered to be the best linear unbiased spatial estimator (Journel and Huijbregts, 1978: 57, 304); in extrapolation the values converge to the mean of the field whereas other techniques do not necessarily exhibit this desirable behaviour; the technique can be easily used on 1D, 2D and 3D data sets; the precision of the estimated data can be obtained. The precision is termed as the Kriging Variance, which when using a semivariogram function, is computed by Eqn. (8.5) (Cressie, 1993: 122-123):

$$\sigma_k^2(s_0) = \lambda^T(s_0) \cdot g(s_0) + \mu(s_0) \tag{8.5}$$

The main disadvantage of the Kriging technique is that it relies on finding the solution to a set of linear equations so that in large data sets the matrix algebra can become computationally burdensome and time consuming. Not as well recognised, and one of the main thrusts of this paper is that the coefficient matrix can be highly ill-conditioned depending on the chosen parameters in the semivariogram function and the data configuration. This can lead to inaccurate solutions and numerical instability. These matters will be dealt with in more detail in the sequel.

8.2 Methods for Improving Computational Efficiency

In weather radar applications where data sets are processed at a frequency of 5 minutes in South Africa, the missing rainfall data need to be estimated in real time, thus a computationally fast and accurate technique is needed. The classical Kriging technique requires the solution of a linear system of equations whose size is determined by the number of control points. This can be time consuming, as the coefficient matrix becomes large. A radar image typically contains 120 000 pixels, leading to a coefficient matrix with dimensions greater then 100 000 by 100 000. Currently it is impractical to work with systems this size online.

There is an unexpected saving which can be exploited to ease the computational burden in a dense data set, as in radar data which typically occur on a lattice, not all of the control points need to be used in estimating the missing data points. A "screening effect" occurs where the significant Kriging Weights associated with the controls are concentrated around the target points (Chiles and Delfiner, 1999: 202-206). This can be used advantageously to significantly reduce computation time. For each individual target point, or set of target points in a cluster, a Kriging neighbourhood in its vicinity can be selected to reduce the dimensionality of the problem.

8.2.1 Theoretical Justification of the Screening Effect in One Dimension

To give some theoretical justification to the 'screening effect' a one-dimensional analysis is presented which shows this very nicely. This does not easily extend itself to two dimensions or more, but intuitively and computationally the idea is justified by extension from a single dimension. It turns out that, in one dimension, the screening effect is limited by the order of the model in cases where the co-variogram is that of an autoregressive process of order p. In fact only the first p intact (control) data nearest the point to be infilled (target) will have non-zero weights λ in the Kriging equation $z(s_0) = \lambda^T(s_0)\cdot\underline{z}$. A small one-dimensional example will suffice to explain the ideas.

One-dimensional AR(1) and AR(2) models

Suppose we have a sample $\{z_i\ , i = 1, 2, \ldots, 6\}$ with 2 missing values at positions $i = 3$ & 4 which we wish to estimate. Assume the sample is a realization of a zero-mean unit-variance Gaussian AR(1) process with parameter ϕ : $z_i = \phi z_{i-1} + a_i$ and (because of the reversibility of the AR(1) process) $z_i = \phi z_{i+1} + e_i$ where $\{a_t\}$ and $\{e_t\}$ are independent white noise processes with the same variance. If we were to attempt to infill the missing data using a forecast from z_2, we'd get $\{\hat{z}_3, \hat{z}_4\} = \{\phi z_2, \phi^2 z_2\}$ and hindcasting from z_5, we'd get $\{\hat{z}_3, \hat{z}_4\} = \{\phi^2 z_5, \phi z_5\}$, with no prospect of agreement. Any compromise would be subjective and *ad hoc*, unless one were to employ an iterative scheme embodying the EM algorithm for example.

As an optimal alternative, use Kriging to estimate the missing values. The correlation matrix between the known values is then formed from the usual 6x6 matrix and removing the two central rows and columns:

$$G_{11} = \begin{bmatrix} 1 & \phi & \phi^4 & \phi^5 \\ \phi & 1 & \phi^3 & \phi^4 \\ \phi^4 & \phi^3 & 1 & \phi \\ \phi^5 & \phi^4 & \phi & 1 \end{bmatrix}$$

whose inverse is:

$$G_{11}^{-1} = \begin{bmatrix} (1-\phi^6) & -\phi(1-\phi^6) & 0 & 0 \\ -\phi(1-\phi^6) & (1-\phi^8) & -\phi^3(1-\phi^2) & 0 \\ 0 & -\phi^3(1-\phi^2) & (1-\phi^8) & -\phi(1-\phi^6) \\ 0 & 0 & -\phi(1-\phi^6) & (1-\phi^6) \end{bmatrix} \div [(1-\phi^2)(1-\phi^6)]$$

$$= \begin{bmatrix} c & a & & \\ a & \beta & \alpha & \\ & \alpha & \beta & a \\ & & a & c \end{bmatrix}$$

To get the matrix of Kriging weights Λ^T we need we need the correlation matrix between the two missing and four known values:

$$G_{21} = \begin{bmatrix} \phi^2 & \phi & \phi^2 & \phi^3 \\ \phi^3 & \phi^2 & \phi & \phi^2 \end{bmatrix}$$

so that because of the symmetry in this case

$$\Lambda^T = \begin{bmatrix} 0 & \phi(1-\phi^4) & \phi^2(1-\phi^2) & 0 \\ 0 & \phi^2(1-\phi^2) & \phi(1-\phi^4) & 0 \end{bmatrix} \div [1-\phi^6]$$

The first thing to note is that only the observations contiguous to the gap have non-zero weights, so are the only observed data involved in the estimation. The second thing to note is that the inverse of G_{11} is a tridiagonal matrix, whose form can be exploited to find explicit expressions for its elements in the general 1-D case.

In general, for a sequence of n complete observations from an AR(1) process, the correlation matrix is $R_n = \{\phi^{|i-j|}, i, j = 1, 2, \ldots, n\}$. Its inverse is a tridiagonal matrix whose upper and lower diagonals are filled with equal elements $a = -\phi/(1-\phi^2)$, the main diagonal has equal elements $b = (1+\phi^2)/(1-\phi^2)$, except for the first and last which equal $c = 1/(1-\phi^2)$.

It is not difficult to show that when there is a block of data missing of width w after the m^{th} observed value, then the correlation matrix of the surviving variables becomes

$$R^*(n,m,w) = \begin{bmatrix} R_m & Q \\ Q^T & R_{n-m-w} \end{bmatrix}$$

where

$$Q = \begin{bmatrix} \phi^{m+w} & \phi^{m+w+1} & - & \phi^{n-1} \\ \phi^{m+w-1} & \phi^{m+w} & - & \phi^{n-2} \\ - & - & - & - \\ \phi^{w+1} & \phi^{w+2} & - & \phi^{n-m} \end{bmatrix}$$

and where R_m has the same form as R_n. The inverse of $R^*(n,m,w)$ will be a tridiagonal matrix of the same form as R_n^{-1} with identical diagonal elements except for the four elements either side of the gap, where instead of a and b, the elements will be α and β, which after some tedious but straightforward algebra are found to be equal to:

$$\alpha = (1-\phi^{2w+4})/[(1-\phi^2)(1-\phi^{2w+2})]$$

$$\beta = -\phi^{w+1}/(1-\phi^{2w+2})$$

which are seen to collapse to a and b when $w = 0$ (no gap).

When G_{11}^{-1} is pre-multiplied by G_{21}^T to give Λ^T, only the weights of the elements either side of the gaps will be non-zero. This happens in general and not only in the AR(1) example above. The behaviour extends to AR(p) models, where now the size of the screen is p elements deep either side of the gaps.

Other combinations of gaps in the data still produce a tridiagonal inverse whose elements either side of the gaps depend on the width of the gaps as well as the width of the intact elements between the gaps. A small-dimensional numerical example will demonstrate the ideas. In a sequence of 11 AR(1) values with $\phi = 0.5$, the targets are at i = 3, 4, 6 & 7. The G_{11} matrix, its inverse, G_{21} and the resulting matrix of Kriging weights Λ^T (where the solid lines indicate the gaps) is:

G_{11} =

element	**1**	**2**	**5**	**7**	**8**	**10**	**11**
1	1	0.5	0.0625	0.015625	0.007813	0.001953	0.000977
2	0.5	1	0.125	0.03125	0.015625	0.003906	0.001953
5	0.0625	0.125	1	0.25	0.125	0.03125	0.015625
7	0.015625	0.03125	0.25	1	0.5	0.125	0.0625
8	0.007813	0.015625	0.125	0.5	1	0.25	0.125
10	0.001953	0.003906	0.03125	0.125	0.25	1	0.5
11	0.000977	0.001953	0.015625	0.0625	0.125	0.5	1

$$G_{11}^{-1} = \begin{bmatrix} 1.3333 & -0.6666 & & & & & \\ -0.6666 & 1.3492 & -0.1269 & & & & \\ & -0.1269 & 1.082 & -0.2666 & & & \\ & & -0.2666 & 1.4 & -0.6666 & & \\ & & & -0.6666 & 1.4 & -0.2666 & \\ & & & & -0.2666 & 1.4 & -0.6666 \\ & & & & & -0.6666 & 1.3333 \end{bmatrix}$$

G_{21} =

element	**1**	**2**	**5**	**7**	**8**	**10**	**11**
3	0.25	0.5	0.25	0.0625	0.0313	0.0078	0.0039
4	0.125	0.25	0.5	0.125	0.0625	0.0156	0.0078
6	0.0312	0.0625	0.5	0.5	0.25	0.0625	0.0312
9	0.0039	0.0078	0.0625	0.25	0.5	0.5	0.25

Λ^T =

element	**1**	**2**	**5**	**7**	**8**	**10**	**11**
3		0.4761	0.1904				
4		0.1904	0.4761				
6			0.4	0.4			
9					0.4	0.4	

The only non-zero weights are those corresponding to the observations contiguous to the gaps, as asserted.

The ideas carry over to an AR(p) model as expected. This is useful because an AR(2) model is quite flexible for the purposes of describing the correlation models encountered in some random fields, particularly rainfields measured by radar. For example, given 14 values (with gaps at i = 2, 8 & 9) sampled from an AR(2) model: $z_i = \phi_1 z_{i-1} + \phi_2 z_{i-2} + a_i$, ϕ_1 = 1.10 & ϕ_2 = -0.32, the inverse of G_{11} is tridiagonal as asserted and the Kriging weight matrix Λ^T is:

element	**1**	**3**	**4**	**5**	**6**	**7**	**10**	**11**	**12-14**
2	0.4977	0.6570	-0.1448						
8					-0.2285	0.8932	0.4225	-0.1435	
9					-0.1435	0.4225	0.8932	-0.2285	

It will be noted that the weights in the second spaces from the gap are much more negative for the AR(2) case than the AR(1) case with these parameters. It is also noted that this behaviour (negative weights) is found when Kriging with Gaussian shaped correlograms. Nevertheless, even in the Gaussian case, the non-zero weights are confined to a zone close to the gaps.

An interim conclusion is that in the particular case of equally spaced data in one dimension, *only the observations contiguous to the gaps influence the infilling of the missing data*. This is likely to have greater economies in higher dimensions as will now be demonstrated computationally.

8.2.2 Demonstrating the Two Dimensional Screening Effect by Computation

An example of the "screening effect" is illustrated in Figure 8.1, which depicts a small dimensional example on a set of 8×9 pixels with 4 (marked by an asterisks) missing in the interior. The semivariogram model used had an $\alpha = 1{,}5$ and a correlation length of L = 11 pixels. (We note in passing that the value of L = 11 pixels, of dimension 1.5 kilometres, is an average value determined from observed radar rainfield intensities for a wide range of weather types and this value will be used in the sequel). The weights associated with the 68 control points for each of the 4 target points were computed one set at a time. These weights were then summed for all four targets and their totals appear in the shaded pixels. It is seen that the significant Kriging Weights are concentrated in the immediate vicinity of the target points and rapidly diminish in magnitude with distance. At a range of greater than two cells from a target point the Kriging Weights are nearly zero. As can also be seen in Figure 8.1, negative weighting values do occur some distance from the target. The Kriging Weights could be forced to be non-negative but a disadvantage in doing this is that the Kriged estimates are then constrained to lie within the minimum and maximum values of the selected control data (Chiles & Delfiner, 1999: 224), (Cressie, 1993: 143). In the rainfield infilling application it is more appropriate to allow for the Kriged estimates to range outside of the minimum and maximum values of the control data used.

0	0	0	0	0	0	0	0	0
0	0	-0.03	-0.07	-0.01	0.02	0.01	0	0
0	-0.04	-0.02	0.38	-0.17	-0.18	-0.03	0	0
0	-0.14	0.54	*	1.33	0.40	0.00	-0.03	0
0	-0.15	0.62	*	*	*	0.39	-0.08	0
0	-0.04	-0.04	0.61	0.75	0.49	-0.01	-0.03	0
0	0	-0.04	-0.16	-0.20	-0.14	-0.04	0	0
0	0	0.01	0.02	0.02	0.02	0.01	0	0
0	0	0	0	0	0	0	0	0

Figure 8.1 Sum of all Kriging Weights concentrated around four unknown data points. The variogram model was exponential with L = 11 pixels.

By considering only the control points with significant Kriging Weights a considerable reduction in size of the coefficient matrix can be made in cases where the target points are in clusters contained within an otherwise complete set of controls. With this reduction in size of the control set a concomitant decrease in computation time can be achieved with little or no loss in the accuracy of the final results.

8.2.3 Ill-Conditioning of the Coefficient Matrix

In the Ordinary Kriging computational procedure, a set of linear equations must be solved in order to determine the weights $\lambda(s_o)$. This solution is given by Eqn. (8.6), derived from Eqn. (8.2):

$$\left[\lambda^T(s_0) \quad \mu(s_0)\right] = \left[g^T(s_0) \quad 1\right] \cdot \begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix}^{-1} \tag{8.6}$$

Solving for $\lambda(s_0)$ requires inversion of the coefficient matrix (usually the equations are solved using an efficient method based on LU decomposition). We found that the coefficient matrix can be highly ill-conditioned depending on (i) its size, (ii) the chosen parameterisation in the semivariogram function (iii) the layout of the data set on a lattice and (iv) the ratio (s/L) as used in the semivariogram model, Eqn. (4), where a decrease in the ratio results in the coefficient matrix becoming more ill-conditioned. The ill-conditioning is most sensitive to the α-parameter, the exponent in the semivariogram function defined in Eqn. (8.4). As α increases from a value close to zero towards a value of two (converging towards a Gaussian semivariogram function) the coefficient matrix becomes increasingly ill-conditioned.

When $\alpha = 2$ in the semivariogram function, and the coefficient matrix is above a certain size, (as small as 68×68, from the 8×9 region of points shown in Figure 8.1) the coefficient matrix is essentially singular. As indicated in Figure 8.2 the determinant of the coefficient matrix (plotted on a log scale) decreases steadily as $\alpha \rightarrow 2$. The determinant was computed for a relatively small 50 by 50 coefficient matrix, with the semivariogram function having a correlation length of L = 11 pixels.

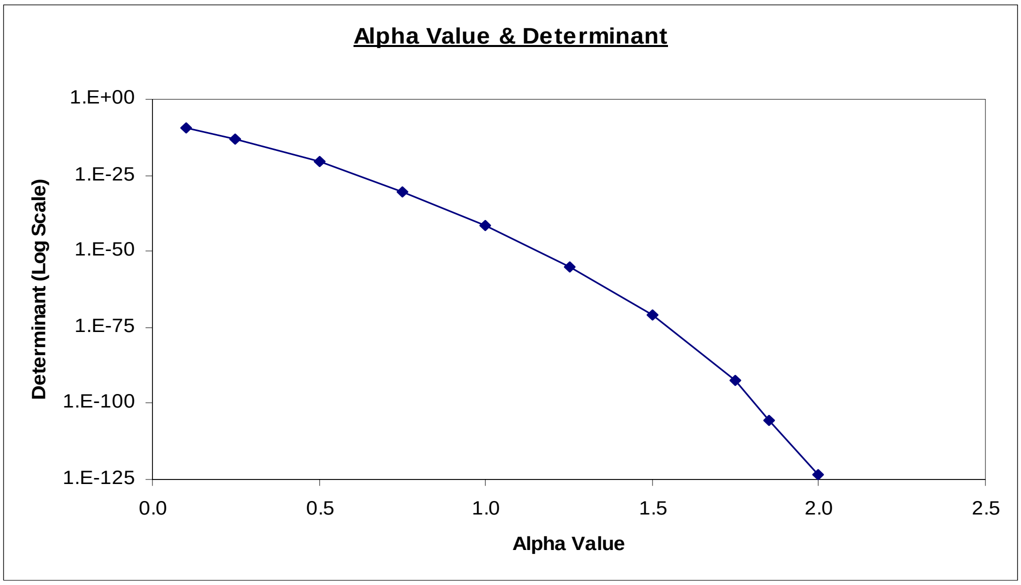

Figure 8.2 Semivariogram exponent α and corresponding determinant of coefficient matrix (size 50x50) for L = 11 pixels.

Although the Gaussian semivariogram results in a near singular coefficient matrix, it is still commonly used in practical Kriging applications (e.g. Todini, 2001) albeit for sparsely distributed control points. It is difficult to determine at what size the coefficient matrix does become noticeably unstable since this appears to be largely determined by factors such as the layout pattern of the control data points. However for a coefficient matrix in excess of 40 by 40 in size, and with $\alpha = 2$, the matrix can be ill-conditioned enough to return nonsensical weights. In such cases, conventional methods of computing the inverse coefficient matrix such as Gauss-Jordan and its derivatives cannot be used. The method of Singular Value Decomposition (SVD) can however be employed with advantage in order to determine an accurate and meaningful solution for the weights, despite the ill-conditioning problem described above.

8.2.4 Theoretical Justification of Matrix Rank Reduction Technique

Singular Value Decomposition involves the decomposition of a rectangular matrix into unique matrix "factors"; an overview of the method is readily available in Press et al. (1992: 59-67) which is briefly summarised here. The matrix is decomposed into a column orthogonal matrix *U*, a diagonal matrix *W* that contains the non-negative singular values along its diagonal and the transpose of an orthogonal matrix *V*. The decomposition of the Ordinary Kriging coefficient matrix is given in Eqn. (8.7).

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} = \begin{pmatrix} & & \\ & U & \\ & & \end{pmatrix} \cdot \begin{pmatrix} w_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & w_n \end{pmatrix} \cdot \begin{pmatrix} & & \\ & V^T & \\ & & \end{pmatrix} \tag{8.7}$$

To compute the inverse one simply needs to invert the diagonal elements of W and multiply out the matrices in the order indicated by Eqn. (8.8):

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix}^{-1} = V \cdot \left\{ diagonal\left(\frac{1}{w_j} \right) \right\} \cdot U^T \tag{8.8}$$

When the coefficient matrix is ill-conditioned, the diagonal values along the *W* matrix need to be carefully considered because the singular values which are near zero define the (near) null space of the matrix. To obtain a meaningful pseudo-inverse, the diagonal elements of W^{-1} are set to $1/w_j$ where w_j is above some chosen near-zero threshold and to zero otherwise. This reduces the rank of the matrix by the size of the (near) null space.

To demonstrate the ideas, a coefficient matrix of size 150 by 150 was computed from a sample reflectivity field. The coefficient matrix was then decomposed as shown in Eqn. (8.7). This was done several times for α values ranging from one to two. The singular values obtained in the W-matrix were ranked then plotted in Figure 8.3, where it is seen that the w_j values rapidly diminish in magnitude along the diagonal and become close to zero as α goes from one to two (Exponential to Gaussian semivariogram); the threshold of 10^{-15} reflects machine precision.

Since many of the singular values are significantly close to zero relative to w_1 which is approximately 60, the most appropriate action is to eliminate these values and replace them (and their inverses) with zeros instead. This solves the problem of ill-conditioning in the inverted coefficient matrix, yields useful accurate solutions, and incidentally reduces computation time.

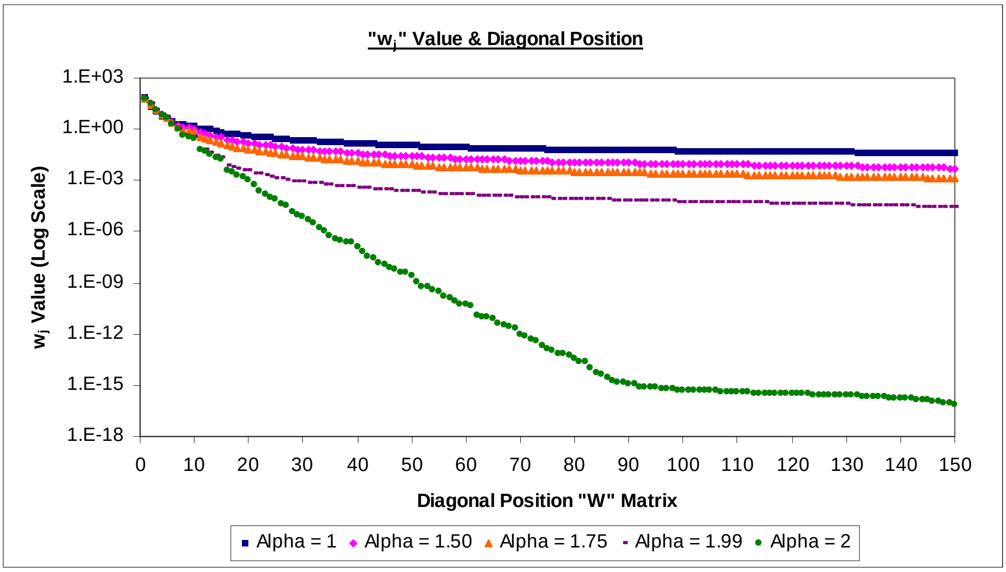

Figure 8.3 Singular values in W matrix for different values of semivariogram exponent α; note the logarithmic scale. The coefficient matrix was 150 square and L was kept constant at 11 pixels.

The Euclidian norm of the coefficient matrix is equal to the sum of the squares of the singular values; if the coefficient matrix were a covariance matrix this sum would also be equal to the variance, as indicated by Eqn. (8.9):

$$\sum_{j=1}^{n} w_j^2 = \sigma_x^2 \tag{8.9}$$

The variance is predominantly contained in the first few terms of the W matrix with the others contributing very little. Only the largest of the singular values need to be considered in most cases since their sum closely approximates the variance.

Once the appropriate number of w_j values has been trimmed and replaced with zero values only the corresponding rows and columns need to be retained in the U and V matrices; as shown in Eqn. (8.10) the columns of the U matrix and the rows of the V^T matrix can be trimmed. This speeds up the computation dramatically with little or no change in the accuracy of the final results.

$$\begin{bmatrix} G & u \\ u^T & 0 \end{bmatrix} = \begin{pmatrix} & U & \end{pmatrix} \cdot \begin{pmatrix} w_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & w_n \end{pmatrix} \cdot \begin{pmatrix} & V^T & \end{pmatrix} \tag{8.10}$$

8.2.5 Demonstration of Advantages of Matrix Rank Reduction by Computation

As an example, a sample reflectivity field was selected and portions removed to simulate a ground clutter scenario. To infill the missing 57 target points a 140 by 140 coefficient matrix had to be inverted. The semivariogram parameters used were α = 1.5 and L = 11 pixels. Initially, with none of the w_j values removed, the Kriging process (calculation of weights and estimation for each of the 57 target values) took approximately 5.6 seconds. The W matrix was then trimmed incrementally along with the corresponding rows and columns of the U and V matrix. It was found that up to 70% of the w_j values could be removed with no significant change in the final Kriged results (accords to the sum of the difference squared of hidden and computed targets) with a consequent decrease in time to 0.6 seconds as shown in Figures 8.4 and 8.5. This corresponds to a nine-fold decrease in computation time with no concomitant change in accuracy achieved with the matrix rank reduction technique. The results would be more dramatic for a larger α.

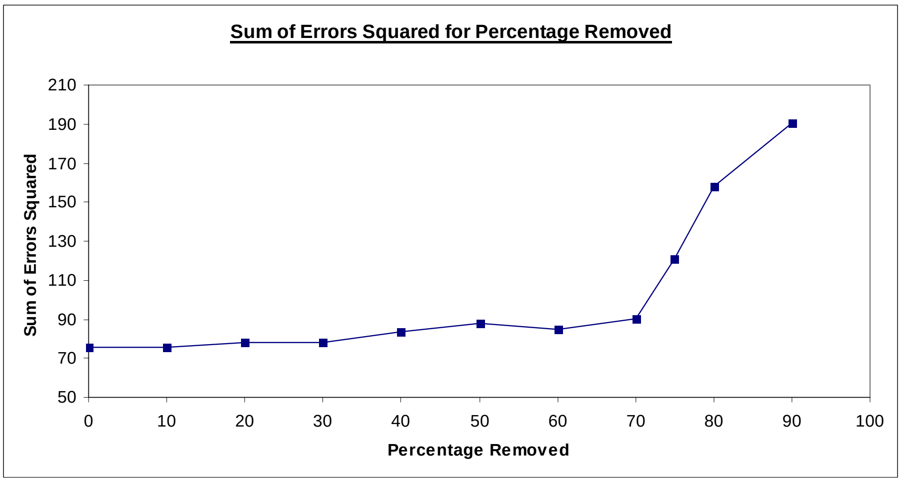

Figure 8.4 Sum of errors squared and corresponding percentage w_j removal, L = 11 pixels, α = 1.5, 140 controls (140 × 140 coefficient matrix).

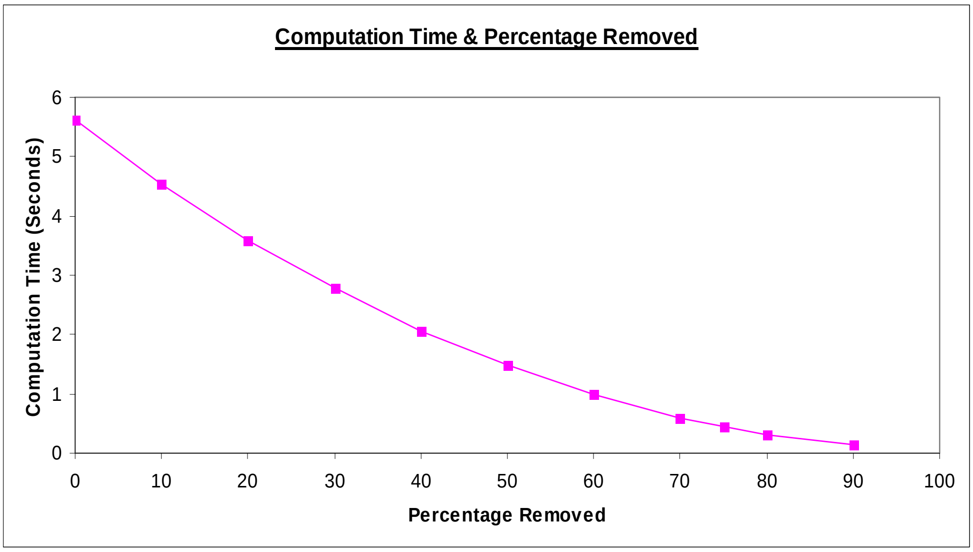

Figure 8.5 Computation time and corresponding percentage removal of singular values in a 140 × 140 coefficient matrix with L =11 pixels and α = 1.5.

An example of the numerical instabilities that can occur in the Kriging process is given in Figure 8.6. A spatially correlated Gaussian random field (the upper part of Figure 8.6) was generated with a known semivariogram function on a 25 by 25 grid field. The semivariogram parameters used were α = 1.5, sill (ω) = 1, L = 1000 pixels (the one exception from L = 11, to emphasize the instability), nugget (p) = 0 and a variance (σ^2) = 1, as shown in Eqn. (8.11):

$$g(s_0) = \sigma^2 - \left\{ p + \omega\left[1 - \exp(-(s/L)^{\alpha})\right] \right\} \tag{8.11}$$

Forty randomly scattered points were then selected from the correlated field as control points. Simple Kriging was then used to estimate the remaining points of the field from the control points. As shown in comparing the two parts of Figure 8.6, the Kriged field does not resemble the original correlated field at all due to the highly ill-conditioned coefficient matrix.

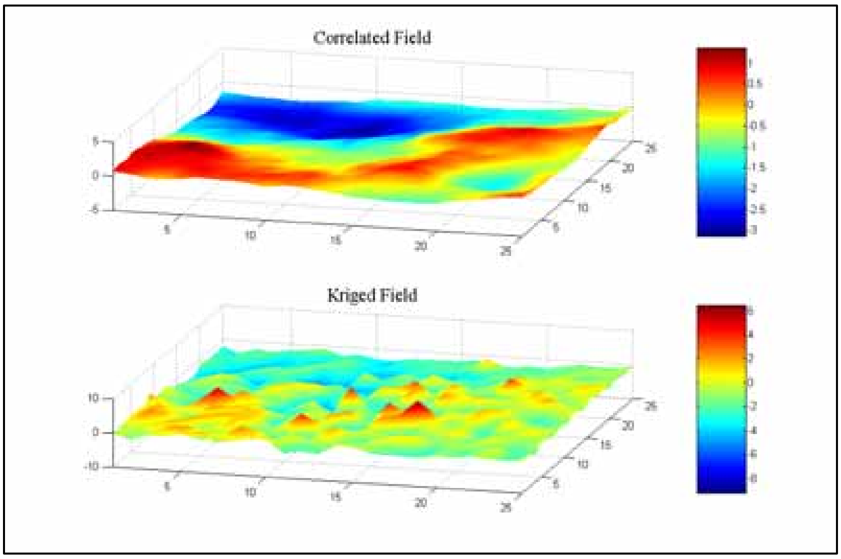

Figure 8.6 Simple Kriging without removal of near zero singular values. For definition of the variogram model, see text.

Figure 8.7 shows exactly the same scenario as in Figure 8.6, however in this instance the trimming of the near zero singular values from the W matrix was carried out, as illustrated in Eqn. (8.10). The Kriged field is now a smoothed replica of the correlated field as desired. In Figures 8.6 and 8.7 the 40 out of 250 points selected to perform the Kriging were randomly chosen. The α-value was midway between exponential and Gaussian, yet still the numerical instabilities occurred. We found this result surprising and thought it important enough to publish as a warning to the unwary and to provide an antidote.

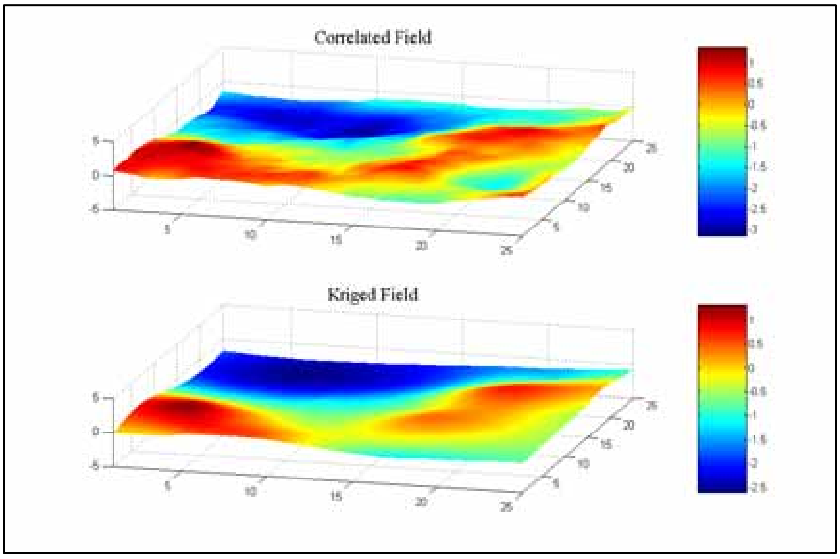

Figure 8.7 Simple Kriging with removal of near zero singular values.

For typical radar reflectivity fields of diameter 250 pixels (about 50 000 data points) the above two techniques, nearest neighbours and matrix rank reduction, need to be used in conjunction in order to achieve efficiency and stability. Image processing techniques also need to be taken advantage of so as to be able to manage the large data sets. Two techniques for ground clutter removal exploiting nearest neighbours will be compared in the next two sections; one works with individual target points in turn, the other by border tracing around a cluster of target points and solving the Kriging equations simultaneously.

8.3 Individual Target's Nearest Neighbours (ITNN) Technique

In the nearest neighbours technique applied to individual target points in turn (ITNN) the weighting distribution associated with Ordinary Kriging is taken into account in the sense that the "screening effect" described in the previous section (where the significant weights at control points are concentrated in the immediate vicinity of the selected target data point) is exploited. In the algorithm suggested here, only the nearest neighbours to the target point are identified. This results in a considerable decrease in computation time with no significant loss in accuracy. The SVD matrix rank reduction process is also carried out to reduce the computation time in finding the Kriging Weights.

The ITNN algorithm can be described as follows for a 2D data set. The radar reflectivity image is searched, pixel by pixel, starting at the top left hand corner and working across each row then down the image, one row at a time, the ground clutter having been already flagged as –10dBZ values. Once a ground clutter point is located, a search outwards from that point is initiated to locate the 25 nearest valid control points – the maximum would be 24 in a 5×5 square region if the target was an isolated pixel. Thus control values within 2 or so pixels of the target are selected. The search strategy proceeds as follows. Initially the rows directly above and below the point are searched, and then the columns on either side of that point. The search then moves incrementally outwards in a spherical pattern until 25 uncontaminated neighbours are located. This strategy is illustrated in Figure 8.8. The pixel marked with a red dot is the target point initiating the first search on the image; other targets (yet to be infilled) are masked in grey.

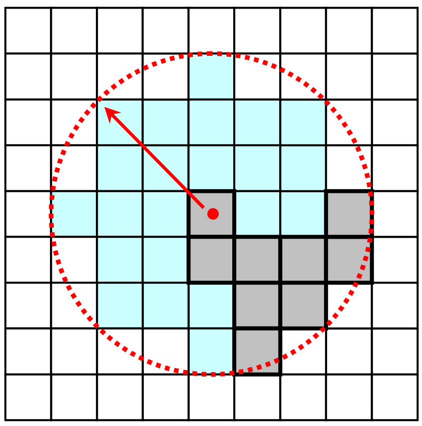

Figure 8.8 Nearest neighbours search strategy. Ground clutter pixels marked as grey squares.

Once the nearest neighbour controls are located they are tested to determine whether they are all zero, if they are, the target is set to zero. If the controls are all non-zero, Ordinary Kriging is carried out to estimate a value for the target.

Once the target point has been identified, an estimate of its value is calculated and then stored in a separate vector, the search resumes where it left off and continues searching across each row and then down the image, one row at a time, until the next ground clutter pixel is located. In this way the rainfall at each ground clutter pixel is estimated in a point by point manner unaffected by previous estimates until all the ground clutter pixels have been infilled. After the entire data set has been scanned the estimated reflectivity values, stored sequentially in a vector, are then inserted in place of the flagged ground clutter pixels.

As examples two radar-rainfall reflectivity images in South Africa have been selected for repair, Figure 8.9 from the Durban weather radar and Figure 8.10 from the Bethlehem weather radar. On each of them the ground clutter pixels have been flagged with -10dBZ values, which appear as black segments in Figures 8.9 and 8.10. The number of Ground clutter targets is approximately 1400 (for both the Bethlehem and Durban data) and the time taken to estimate the missing data for one 400 by 400 pixel image is typically 0.5 to 2 seconds computed on a Pentium(R) 4 with a 2,40GHz CPU and 512MB of RAM. The images on the right of each pair are those which have been "repaired", i.e. where the ground clutter values, or contaminated pixels, have been replaced with reflectivity values estimated by Kriging.

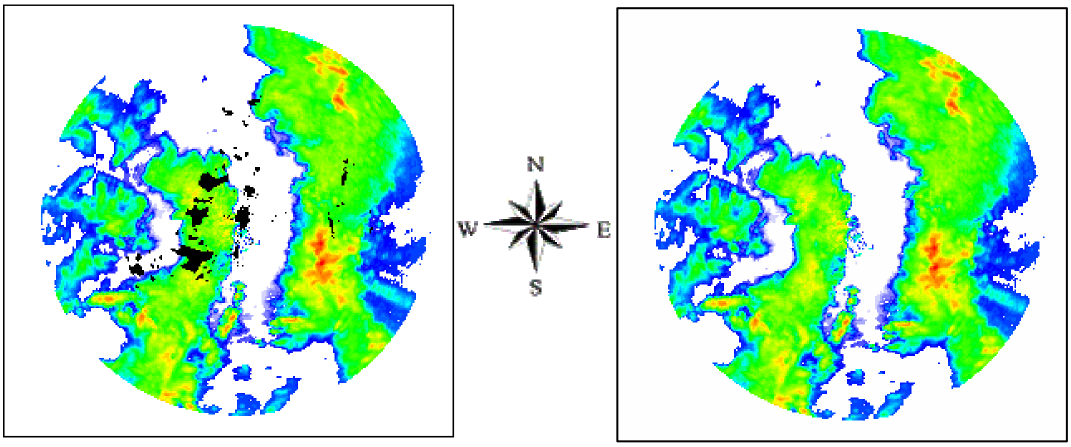

Figure 8.9 Durban radar rainfall reflectivity data before and after ground clutter removal (18 November 2000). Black areas indicate contaminated data. The square surrounding each radar image has sides of length approximately 270 by 270 kilometres.

8.4 Border Tracing Nearest Neighbours (BTNN) Technique

An alternative for identifying segments which have to be infilled is the Border Tracing with Nearest Neighbours (BTNN) technique, where advantage is again taken of the distribution of the Kriging Weights in order to reduce the computational load and improve computational efficiency by concentrating on the nearest neighbours around a cluster. In this method image processing algorithms were investigated to provide an alternative approach to infilling the missing data. In the ITNN technique the clutter was infilled point by point whereas in the BTNN technique the infilling is done a whole segment at a time, where a segment is one or more target pixels that are grouped together in a connected region.

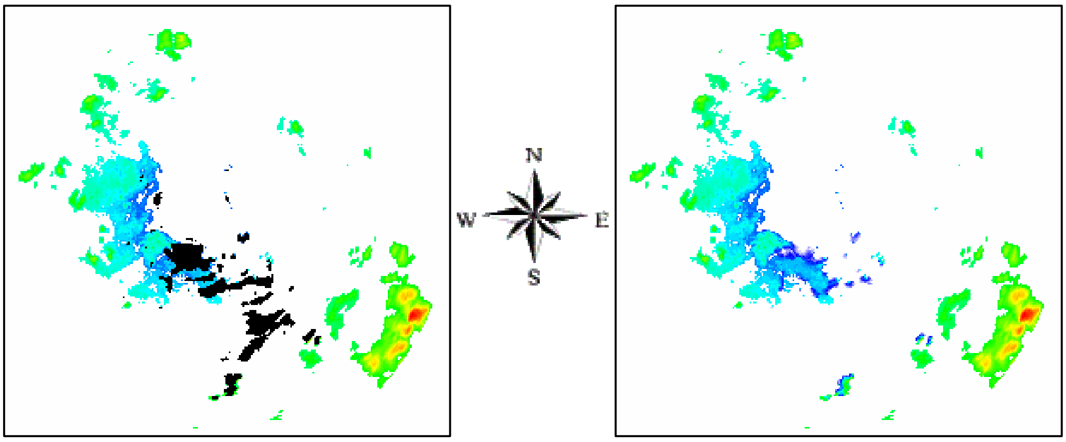

Figure 8.10 Bethlehem radar rainfall reflectivity data before and after ground clutter removal (25 February 2003). Black areas indicate contaminated data. The square surrounding each radar image has sides of length approximately 270 by 270 kilometres.

Once again, the image is searched from the top left hand corner, row by row, moving downwards until a flagged ground clutter pixel is located. An eight-connectivity border tracing technique (Sonka et al. 1999: 142-145) is then implemented from that local origin. From that pixel the eight immediate surrounding neighbours are searched in an anticlockwise direction until another new ground clutter pixel is located. This becomes the new local origin and the process is repeated. The border is then traced in this manner until the original starting point is reached. Figure 8.11a illustrates how the border is traced in an anticlockwise direction around the Ground clutter segment, with the dashed lines indicating pixels tested and identified as controls in the process.

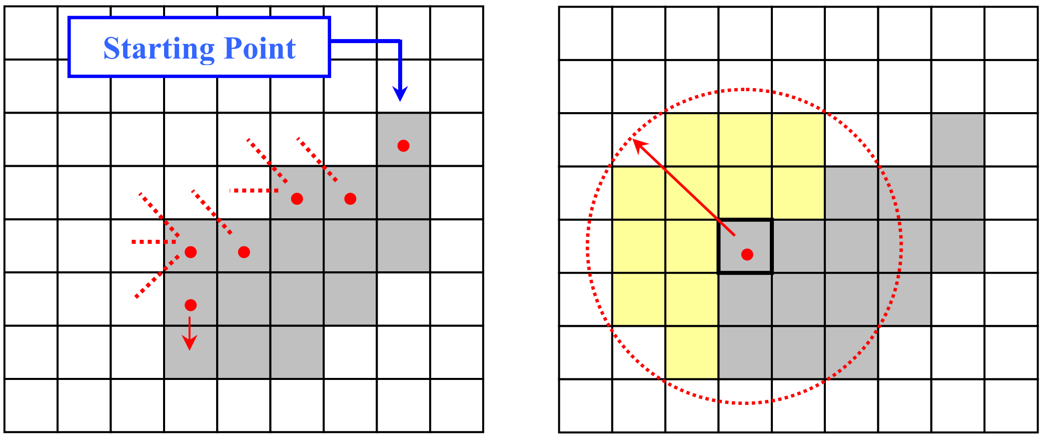

Figure 8.11 a) Border Tracing technique indicating trace direction and pixels tested in process (left image). b) Selection of Kriging neighbourhood for infilling of missing data (right image) (Sonka et al., 1999: 143).

While the border is being traced, a Kriging neighbourhood two pixels deep from each border target is simultaneously selected. As indicated in Figure 8.11b the control points within a range of 2.5 pixels are selected – this strategy neglects the corner pixels on a 5x5 square grid surrounding the target. All of the pixels selected in this manner are then identified as the Kriging control neighbourhood for the estimation of the cluster of target points in the particular segment. Once the border tracing has been completed, and the Kriging neighbourhood has been selected, an algorithm identifies the positions of the remainder of the target points that are located in the interior of the traced contour.

The Kriging control neighbourhood is then tested to determine whether the controls are all zero. If this is the case, all of the target points in the segment are set to zero. If the data points are not all zero, Ordinary Kriging is used to estimate values for all of the target points that make up that segment.

The solution to the Kriging equations is then computed by use of an Inversion by Partitioning Technique (Aitken. 1967: 148) in conjunction with the SVD matrix rank reduction technique in an effort to decrease the computation time. The Inversion by Partitioning Technique as employed here splits the coefficient matrix into 4 equal sub matrices, given by Eqns. (8.12):

$$A = \begin{bmatrix} P & Q \\ Q^T & S \end{bmatrix}$$

$$A^{-1} = \begin{bmatrix} \widetilde{P} & \widetilde{Q} \\ \widetilde{Q}^T & \widetilde{S} \end{bmatrix} \qquad \text{where:}$$

$$\begin{aligned}
\widetilde{P} &= (P - Q \cdot S^{-1} \cdot Q^T)^{-1} \\
\widetilde{Q} &= -(P - Q \cdot S^{-1} \cdot Q^T)^{-1} \cdot (Q \cdot S^{-1}) \\
\widetilde{Q}^T &= -(S^{-1} \cdot Q^T) \cdot (P - Q \cdot S^{-1} \cdot Q^T)^{-1} \\
\widetilde{S} &= S^{-1} + (S^{-1} \cdot Q^T) \cdot (P - Q \cdot S^{-1} \cdot Q^T)^{-1} \cdot (Q \cdot S^{-1})
\end{aligned} \qquad (8.12)$$

Further partitioning of the matrix into 9 sub-matrices was tested but this did not result in a decrease in computation time. The Inversion by Partitioning Technique was also found to be most effective for coefficient matrices with a size in excess of 30 by 30; for matrices below this size the SVD matrix rank reduction technique proved to be computationally faster.

Once the target values for a cluster have been estimated they are immediately inserted into the working data set and the pixels flagged as known. The algorithm then carries on searching from the initial local origin that identified the now infilled target segment, row-by-row, moving across then down the image until the next ground clutter pixel is located, at which point the process is repeated. Figure 8.12 is an example where the BTNN technique has been used to remove ground clutter, as before the left image contains the marked ground clutter and the right image the estimated reflectivity data.

To estimate the missing values in a typical radar rainfall image the BTNN technique takes approximately 5 to 20 seconds. However the time is highly dependent on the range used for selecting the Kriging neighbourhood for each segment. The greater the number of control points the greater the size of the coefficient matrix that needs to be inverted; where the computation time is dominated by the matrix inversion.

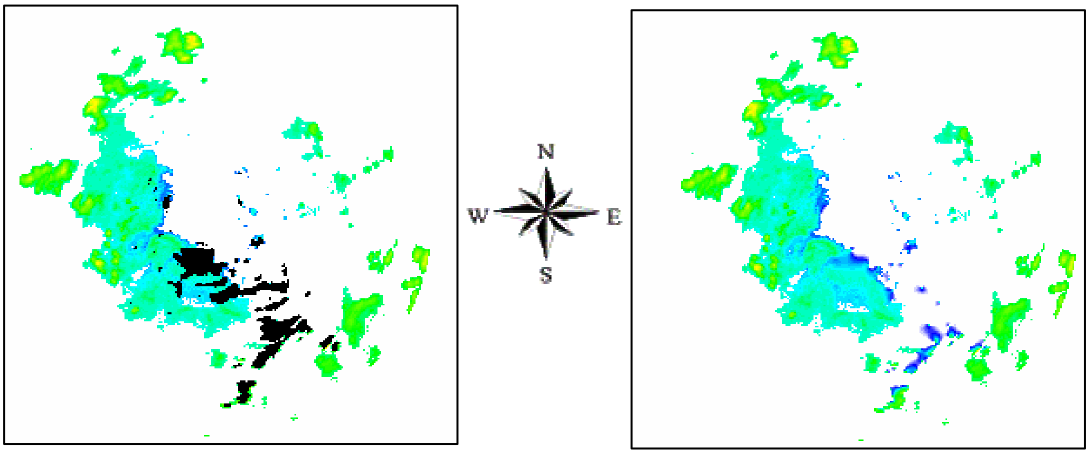

Figure 8.12 Bethlehem radar rainfall reflectivity image before and after ground clutter removal (25 February 2003). Black areas indicate contaminated data. The square surrounding each radar image has sides of length approximately 270 by 270 kilometres.

8.5 Comparison of the Two Nearest Neighbour Infilling Techniques

An example comparing the ITNN and BTNN techniques is now presented where a radar reflectivity image from the Durban weather radar containing no ground clutter on the 18 November 2000 was selected that contained convective rainfall. The Bethlehem ground clutter map (not coinciding with the Durban one) was then superimposed on this image, as indicated in Figure 8.13. The missing reflectivity data were then estimated using both the ITNN and BTNN techniques. In this application of the ITNN technique the twenty nearest neighbours for each target were used and for the BTNN technique pixels within a range of 2.5 from the ground clutter border were selected. Ordinary Kriging was used to compute the missing reflectivity data with a semivariogram function having L = 11 pixels and $\alpha = 1.5$ being used in both cases. The estimated values were compared to the real reflectivity values and the square of the errors computed at each target point; these sets are shown in the lower pairs of diagrams in Figure 8.13. The computation time was recorded for each technique; Table 1 gives a summary of the results.

Table 8.1. Errors and computation time for techniques tested in the example in Figure 8.13.

	Sum of Squared Errors (dB^2)	Computation Time (seconds)
ITNN	14 700	1.66
BTNN	19 000	16.8

For this particular radar reflectivity image the ITNN technique performed better than the BTNN technique. Not only was the sum of the errors squared lower than the BTNN technique but the computation time was faster by a factor of 10.

Limitations may exist for very large areas of ground clutter where the best estimate of the rainfall near the centre of the ground clutter will be the mean value of the control points (if the distance exceeds the Correlation Length) which may result in errors of high magnitude. Large errors may also occur where ground clutter completely or partially hides highly convective rainfall which is surrounded by stratiform rain, this phenomenon occurs in one of the images shown in Figure 8.13.

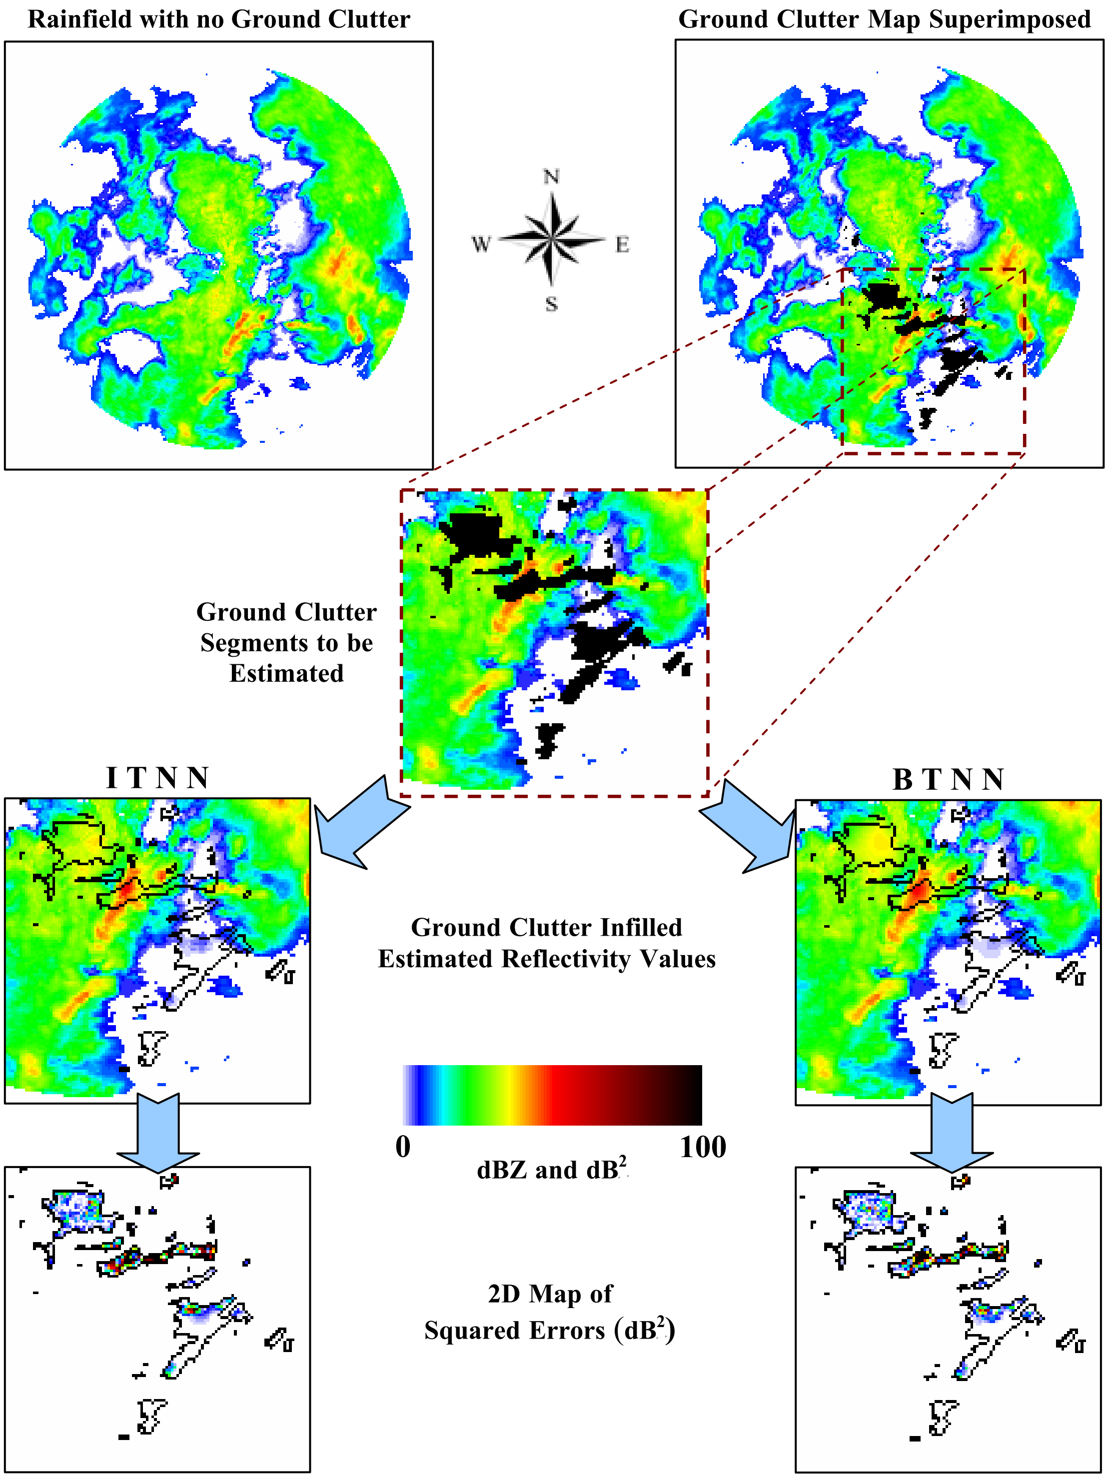

Figure 8.13 Example of ITNN and BTNN techniques.

8.6 Summary

Radar rainfall data have importance in many fields of endeavour. However the fact that errors persist in these data means that the rainfall estimates are not as accurate as they might be. These errors, noticeably ground clutter, can be detected and identified in the radar images. The technique proposed for estimating the missing data in this paper was Ordinary Kriging. Various methods to improve computational efficiency such as using nearest neighbours to a target point, matrix rank reduction techniques and inversion by partitioning have been implemented in order to make the Ordinary Kriging process computationally efficient in real time applications.

Two techniques for removing ground clutter contamination have been developed which exploit the screening effect in Kriging. These are the Individual Target Nearest Neighbour (ITNN) and Border Tracing Nearest Neighbour (BTNN) techniques. Testing on typical sample sizes indicates that ITNN is 10 times faster than BTNN and is more accurate. Both algorithms exhibit sufficient computational speed for online data cleansing, when supported by computationally efficient methods: (i) rank reduction using Singular Value Decomposition and (ii) matrix partitioning