University of Bath

UNIVERSITY OF
BATH

PHD

Propagation forecasting for EHF and SHF systems

Hodges, Duncan David

Award date:
2006

Awarding institution:
University of Bath

Link to publication

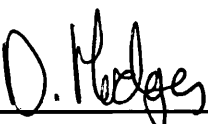# PROPAGATION FORECASTING FOR EHF AND SHF SYSTEMS

Duncan David Hodges

A thesis submitted for the degree of Doctor of Philosophy
University of Bath
Department of Electronic and Electrical Engineering

2006

COPYRIGHT

Duncan David Hodges

UMI Number: U223518

UMI

Dissertation Publishing

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

This thesis presents a unique method for estimating the attenuation on EHF and SHF communication links. The approach is to recreate the meteorological environment and from this, estimate the propagation characteristics. Numerical weather prediction (NWP) techniques are used to provide three-dimensional estimates of temperature, humidity and pressure, two-dimensional fields of large scale and convective rainfall rates and u- and v- winds are also used. The estimates of rainfall rate can be supplemented by the use of rainfall radar data. Downscaling techniques can be used to improve the grid scale of the rainfall rate estimates for use in simulation studies.

Historical time series can be created from archived meteorological data products in order to produce time series for network-centric simulation. Forecast meteorological data can also be used to create propagation forecasts in order to enable proactive fade mitigation.

The procedure is shown to produce time series with the correct long-term statistical distribution and the correct fade-slope characteristic. Comparisons between Ka- band beacon measurements and estimated time series have also been performed in both historical and forecast paradigms.

High-level network-centric simulations of a DVB-RCS system with time-diversity fade countermeasures are performed. These simulations provide example error-bounds for the forecast time series.

Acknowledgements

Contents

List of Figures

List of Tables

Acronyms

ACTS	Advanced Communications Technology Satellite
BADC	British Atmospheric Data Centre
BFWA	Broadband Fixed Wireless Access
CAMRa	Chilbolton Advanced Meteorological Radar
CAPE	Convective Available Potential Energy
CDF	Cumulative Density Function
CIN	Convective Inhibition
COST	Cooperation européenne dans le domaine recherche scientifique et technique
CSIP	Convective Storm Initiation Project
DLR	Deutschen Zentrum für Luft- und Raumfahrt
DVB	Digital Video Broadcasting
DVB-RCS	Digital Video Broadcasting - Return Channel via Satellite
DVB-S	Digital Video Broadcasting - Satellite
EDS	Extreme Dependence Score
EHF	Extra High Frequency
EIRP	Effective Isotropic Radiated Power
ENIAC	Electronic Numerical Integrator and Computer
EPSRC	Engineering and Physical Sciences Research Council
ESA	European Space Agency
ETS	Equitable Threat Score
ETSI	European Telecommunications Standards Institute
FCC	Federal Communications Commission
FFT	Fast Fourier Transform
FMT	Fade Mitigation Technique
GBS	Global Broadcasting Satellite
GFS	Global Forecast System
GNSS	Global Navigation Satellite System
GSM	Global System for Mobile communications
HAP	High Altitude Platform
HDTV	High-Definition Television
ITU	International Telecommunications Union
JPL	Jet Propulsion Laboratory
LBC	Lateral Boundary Condition
LWC	Liquid Water Content
MAC	Medium Access Control
MM5	Mesoscale Model - 5

MPM	Millimeter Propagation Model
MRO	Mars Reconnaissance Orbiter
NAE	North Atlantic and Europe Model
NASA	National Aeronautics and Space Administration
NCAR	National Centre for Atmospheric Research
NWP	Numerical Weather Prediction
OBP	On-Board Processing
ONERA	Office National d'Etudes et de Recherches Aérospatiales
PDF	Probability Density Function
PES	Packetised Elementary Stream
PFE	Propagation Forecast Engine
PSD	Power Spectral Density
QEF	Quasi-Error Free
OPEX	Olympus Propagation Experiment
QPSK	Quadrature Phase Shift Keying
RAL	Rutherford Appleton Laboratory
RS	Reed-Solomon
SHF	Super High Frequency
SOHO	Small Office - Home Office
SNR	Signal to Noise
TDMA	Time Domain Multiple Access
TS-MUX	Transport Multiplex
TSaR	Telecommunication Space and Radio
USGS	United States Geological Survey
VSAT	Very Small Aperture Terminals
UKMO	UK Met Office
UFAM	Universities Facility for Atmospheric Measurement
UM	Unified Model

Chapter 1

Introduction

1.1 The history of satellite communications

Communication by satellite was first considered in the autumn of 1945 when Arthur C Clarke wrote a short article in *'Wireless World'* that described the use of manned satellites in a geostationary orbit to distribute television broadcasts. The article provided a strong case for the use of orbiting systems for the provision of assorted services and the article was repeated in Clarke's later work entitled *'The Exploration of Space'*.

The first commercial investigations in to the technical challenges involved in the use of satellites for communications were performed John R Pierce of AT&T's Bell Laboratories. In a 1954 speech and 1955 article he elaborated upon the idea of a communications *'mirror'* in space, a medium orbit *'repeater'* and a geosynchronously orbiting *'repeater'*. Pierce estimated the communications capacity at 1000 simultaneous telephone calls which he compared to the first trans-Atlantic telephone cable which could only carry 36 simultaneous calls.

In 1957 Sputnik 1 was launched and the interest in space technologies grew, driven not only by the associated benefits and profits, but also by the prestige. In order to accelerate development and reduce overlap initial NASA experiments were involved with *'mirrors'* or *'passive'* communications satellites; whilst the Department of Defence investigated *'active'* or *'repeater'* satellites that amplified the received signal before re-transmission.

By late 1961 NASA had awarded a contract to RCA to build a medium-orbit active communication satellite (RELAY); AT&T was building its own medium-orbit satel-

lite (TELSTAR) and NASA had awarded a sole-source contract to Hughes Aircraft Company to build a 24 hour orbiting satellite (SYNCOM).

By 1964, two TELSTARs, two RELAYs and two SYNCOMs had successfully operated in space. Initial experiments with these systems involved broadcasting parts of the 1964 Olympics from Tokyo. During this time several earth stations across the world were built and on August 20th 1964 the International Telecommunications Satellite Organisation (INTELSAT) was created.

At this time Communications Satellite Corporation (COMSAT) was in the process of contracting for their first satellite. COMSAT's initial capitalisation of $200 million was considered to be sufficient to build a system involving dozens of medium-orbit satellites. COMSAT rejected a joint RCA and AT&T offer of a medium-orbit satellite incorporating the best of TELSTAR and RELAY and chose the geosynchronous satellite offered by the Hughes Aircraft Company. COMSAT's first satellite, named EARLY BIRD was launched on April 6th 1965 from Cape Canaveral.

By the end of 1965, EARLY BIRD had provided 150 telephone half circuits and 80 hours of television service. The INTELSAT II series was more capable and longer lived than EARLY BIRD, whilst INTELSAT III was the first to provide coverage over the Indian Ocean and complete global coverage.

1.2 Opportunities within the EHF and SHF bands

The frequency bands typically occupied by satellite systems are shown in figure 1.1. Historically satellites operated in the C- and Ku- bands, for example Eutelsat's HotBird 1 operated with an uplink contained in the band 17.3 GHz to 18.1 GHz, with a downlink between 11.45 GHz and 12.5 GHz.

The use of EHF (30 to 300 GHz) and the upper end of the SHF (3 to 30 GHz) frequency bands has many advantages and industry has been very quick to seize the opportunity to expand into these frequency bands. For example in the United States of America 14 V-band satellite systems were proposed to the FCC following a notice of enquiry offering access to the millimeter band, a summary of these is shown in table 1.1.

These advantages offered by these higher frequency bands include:

- Greater bandwidth,

300Hz ┌──────┐
      │ ULF │
3KHz ├──────┤
      │ VLF │
30KHz ├──────┤
      │ LF │   ┌─ L- Band (1.5 GHz)
300KHz├──────┤   │
      │ MF │   S- Band (2.5 GHz)
3MHz ├──────┤   │
      │ HF │   C- Band (4/6 GHz)
30MHz ├──────┤   │
      │ VHF │   Ku- Band (11/12/14 GHz)
300MHz├──────┤   │
      │ UHF │   ┐ K- Band (20 GHz)
3GHz ├──────┤   │
      │ SHF │   Ka- Band (30 GHz)
30GHz ├──────┤   │
      │ EHF │   ┘ V- Band (40 GHz)
300GHz└──────┘
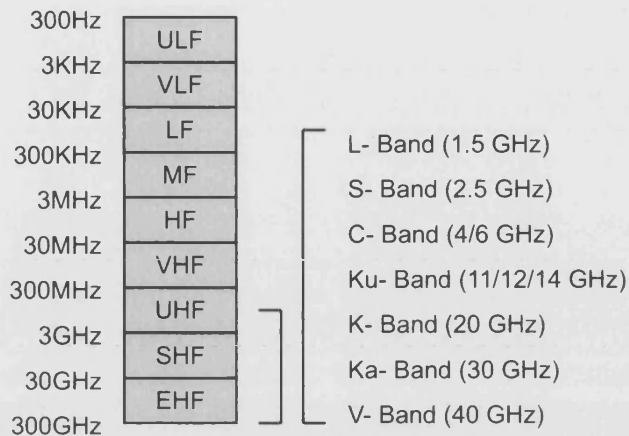
Figure 1.1: Frequency bands defined in ITU-R V.431-7.

- Less radio congestion,

- Increased antenna gain (for the same aperture size).

Each of these advantages is now covered in more detail.

1.2.1 Greater bandwidth

One of the greatest advantages of using EHF and SHF frequencies is the large band-widths that are available. As can be seen from table 1.1 bit-rates of the order of $Gb\,s^{-1}$ are achievable. Within broadcast networks these higher bandwidths are very important particularly with the introduction of broadcast high-definition television (HDTV).

The greater bandwidth also encourages the use of return channel links, the DVB-RCS standards are built around a MF-TDMA return link. These return links allow a far greater level of interactivity than ever before; interactivity is also a key objective of the ESA funded 'SatMode' concept. SatMode is expected to have a maximum upstream capacity of $64kb\,s^{-1}$ and will use SES Astra's Ka- band capacity for its return traffic.

Other recent return channel experiments include the EMS 'school-net' experiment which used the Ka- band capability of the Anik F1 satellite and DVB-RCS equipment to link schools around Canada.

Broadband internet connectivity via satellite is also becoming even more viable, and with large bit-rate payloads it will surely become more viable. The provision of broad-

3

Table 1.1: Proposed United States Q-/V-band satellites [Evans and Dissanayake, 1998]

Company	System	Investment ($B)	Capacity (Gb s^{-1})	EIRP (dBW)
Denali Telecom	Pentriad	1.9	≈ 36	73.0
GE Amercom	GEstar+	3.4	≈ 70	59.0
Hughes	Xpressway	3.9	≈ 65	55.0
	SpaceCast	1.7	≈ 64	72.0
	StarLynx	2.9	≈ 6	70.5
Lockheed Martin	Q/V band	4.75	≈ 45	64.0
Loral Space	Cyberpath	1.17	17.9	77.5
Motorola	Mstar	6.4	≈ 3.6	38.0
Orbital Sciences	Orblink	0.9	≈ 75	62.0
PanAmSat	VStream	12	≈ 3.2	60.0
Spectrum Astro	Aster	2.4	≈ 10	66.7
Teledesic	VBS	1.9	4	59.0
TRW	GESN	3.4	≈ 60	83.0

band internet services for the SOHO or corporate user via a satellite system has many advantages. In remote regions or disaster areas where local infrastructure is unable to support broadband access, users can still be provided with services at a suitably high bit-rate. Even within developed countries many people are unable to receive broadband access over terrestrial networks; [SpaceWay] suggests that in North America there are over 4 million businesses and 40 million households that are unable to receive broadband access over terrestrial networks.

1.2.2 Increased antenna gain

The Effective Isotropic Radiated Power (EIRP) of systems operating in the EHF and SHF band is greatly improved for a given physical antenna size, since the gain of a parabolic antenna is linked to the frequency through equation 1.1:

$$G = \eta \left(\frac{\pi D f}{c} \right)^2 \tag{1.1}$$

where G is the antenna gain as a ratio to an isotropic antenna, η is the antenna efficiency, D is the antenna diameter (m), f is the frequency (Hz) and c is the speed of light (m s^{-1}).

Since the antenna gain is increased, the EIRP is increased for a given physical antenna size, this allows either a reduction in RF power or greater link flux for the same transmit power. The physical size of antenna can be very important in both commercial and

military applications, in military communications the use of man-portable high-gain antennas can be a big advantage allowing a network infrastructure to be assembled (or disassembled) very quickly. Within the commercial environment, especially within land-mobile and nautical systems, a large antenna size can make a system unattractive to the end-user.

The directivity of a given antenna working at higher frequency will also be greater than an antenna of equal size working at a lower frequency, following equation 1.2:

$$G = kD_i \qquad (1.2)$$

where G is the antenna gain as a ratio to an isotropic antenna, k is the radiation efficiency and D_i is the directivity.

The increased directivity of the antenna can be used either to give a greater level of control over downlink spot beam coverage or alternatively to allow a reduction in the spatial separation of orbital platforms, both of which encourage frequency reuse and increased spectral efficiency.

1.2.3 Less congestion

The 'traditional' frequency allocations for satellite communications, L-, S-, C- and Ku-bands are rapidly reaching, or have reached, saturation. This now requires complex frequency allocations to 'in-fill' allocation gaps, for example the Hotbird 6 frequency allocation is a mixture of K-, Ku- and Ka- band frequencies, an overview of the allocation is shown in table 1.2.

Table 1.2: Hot Bird 6 frequency plan overview.

Uplink (GHz)			Downlink (GHz)		
13.89902	to	13.84656	10.81508	to	10.87262
14.01918	to	14.23016	11.56574	to	11.68082
			12.51984	to	12.59656
			19.72918	to	19.74800
14.25066	to	14.47891	10.97141	to	11.19966
29.52918	to	29.54836	19.82508	to	19.84426
29.62508	to	29.64426			

The relatively free allocations in the Ka- and V- bands are becoming increasingly attractive; indeed terrestrial links at 55 GHz are becoming increasingly popular and the 64 GHz band is the next to used. The 42 GHz band has also been proposed for the

use with high altitude platforms (HAPs) and tethered balloons.

In addition to frequency congestion the problems of orbit congestion are equally important. There is a limited amount of physical space within the geostationary orbit, especially in locations over heavily-populated developed areas such as Europe and the Continental USA. The increased directivity of uplink antennas and the increased capacity on each payload allows higher bit rate payloads to be placed closer together than ever before, hence increasing the potential data capacity of the orbit.

1.3 Technical challenges within the EHF and SHF frequency bands

The previous section discussed the possible advantages of using EHF and SHF based communication systems when compared to 'traditional' Ku- and C- band systems. The purpose of this section is to provide a brief discussion of the disadvantages of using Ka- and V- band systems, over lower frequency Ku- and C- band systems.

1.3.1 Established user base at Ku-band

The use of different frequency bands for communication systems requires new transceiver equipment. In the last decade throughout the commercial communication market it is apparent that the public are reluctant to adopt new technologies. Third generation mobile telephones can be used as an example; it can be seen that the offer of new services (such as streaming video) are not currently driving the public to adopt the new technologies, beyond the early adopters.

The recent upgrades to digital television are another good example, all satellite broadcast systems now use digital receivers, and in the UK the free-to-air terrestrial digital television, 'freeview', now has a receiver in a large number of homes. However even the relative success of 'freeview' must be compared to its predecessor 'onDigital'. 'onDigital' was a billed terrestrial service provided by ITV Digital from the late 1990's, until April 2002 when the company were forced to sell off assets and withdraw the pay TV service.

These examples have other influences that affect their ability to succeed, but a significant influence is the inability to attract sufficient numbers of users prepared to upgrade to a 'value-added-service'.

1.3.2 Implementation Issues

For a long time the implementation of payloads operating at frequencies above 20 GHz had been hampered by issues concerned with the practicalities of realising and building the transponder. Components such as TWTs and other high power amplifiers (HPAs) were not readily available, and the performance of HPAs that were available was poor.

However, as the frequency allocation begins to be populated more high-performance components are being made available. Currently high performance 40 GHz TWTs are available from Boeing, whilst smaller grid-amplifiers such as those from Wavetech can provide high-gain, low-noise, high-linearity amplifiers for terminal transceivers.

1.3.3 Propagation channel

At EHF and SHF frequencies the propagation channel can become very challenging. As the frequency of transmission is increased above 20 GHz the path attenuation caused by meteorological effects increases dramatically, the depolarisation effects also increase and scintillation caused by turbulence in the troposphere increases.

1.3.4 Structure of this thesis

In the following chapters a unique approach to this problem is discussed. Chapter 2 discusses the properties of the EHF and SHF propagation channel. Chapter 3 reviews a sample of the various approaches to propagation modelling. Chapter 4 describes the unique approach presented in this thesis. Chapter 5 describes downscaling techniques that can be used to improve the resolution of rainfall rate estimates for simulation purposes. Chapter 6 contains some example results from the propagation model and some verification of the time series. In chapter 7 the thesis continues with a brief description of a bespoke network simulator and some example results from a system that uses propagation forecasts to implement a time diversity system. Chapter 8 continues with a summary of the conclusions from the work and the thesis concludes with an assessment of the possible future direction of the work in chapter 9.

Chapter 2

Properties of the EHF and SHF communications channel

This aim of this chapter is to outline the various causes of the increased propagation impairments that effect communications at frequencies greater than 20 GHz. The emphasis is primarily on earth-space paths, however the impairments discussed are generic and can also be applied to other networks working at such frequencies, for example high bandwidth Broadband Fixed Wireless Access (BFWA) systems or overlay networks such as those within GSM or other mobile networks.

2.1 Attenuation due to gaseous absorption

The attenuation due to gaseous absorption is often considered insignificant upon a link path, however as the frequency increases the gaseous absorption can become comparatively large.

2.1.1 Attenuation due to oxygen

Attenuation by atmospheric oxygen is caused by the resonant absorption of radiowave energy by oxygen molecules. At sea-level pressures many absorption lines near 60 GHz merge together to form a single broad absorption band [Castanet, 2001]. As the altitude increases the pressure decreases and each of these absorption lines becomes visible. The attenuation due to oxygen remains relatively constant, with a very small variability

8

caused by large changes in temperature and pressure.

Maps of the annual average zenith attenuation[1] due to oxygen at 15 GHz and 50 GHz are shown in figure 2.1. These show the increase in attenuation due to the increase in frequency.

---

[1]The average values of temperature are used for each location as defined by ITU-R P1510. The global annual average pressure is used (1013.7 hPa, as derived from the hydrostatic equation, see appendix C).

Zenith attenuation due to oxygen at 15 GHz

(a) Oxygen: 15 GHz

Zenith attenuation due to oxygen at 50 GHz

(b) Oxygen: 50 GHz

Figure 2.1: Zenith attenuation (dB) due to absorption by atmospheric oxygen as predicted by the ITU-R recommendations.

2.1.2 Attenuation due to water vapour

In addition to the attenuation caused by absorption of radio power by oxygen the attenuation caused by water vapour is also significant at SHF and EHF frequencies. The water vapour molecule has a permanent electric dipole which causes this resonance. Nitrogen has no resonance within the EHF and SHF frequency bands, and hence despite being the most abundant gas in the atmosphere does not cause any significant attenuation [Ulaby et al., 1986].

There is a spot resonance at approximately 22.235 GHz, which can cause problems for lower frequency Ka- band systems where the specific attenuation reaches nearly $0.2\,\mathrm{dB\,km^{-1}}$. As the frequency is increased the absorption increases, until at about 100 GHz where there is a specific attenuation of nearly $0.5\,\mathrm{dB\,km^{-1}}$.

The variability of the attenuation due to water vapour is far greater than that of oxygen, since the attenuation depends on the water vapour density. This is typically greater during the summer months and in hot, humid conditions such as those experienced in tropical areas, typically for any location the greatest attenuation due to water vapour is during the season of the greatest rainfall [Castanet, 2001].

Maps of the annual average zenith attenuation[2] due to water vapour at 15 GHz and 50 GHz are shown in figure 2.2.

---

[2]The average values of temperature are used for each location as defined by ITU-R P.1510. The global annual average pressure is used (1013.7 hPa, as derived from the hydrostatic equation, see appendix C). The average values of water vapour are used as defined by ITU-R P836.

Zenith attenuation due to water vapour at 15 GHz

(a) Water Vapour: 15 GHz

Zenith attenuation due to water vapour at 50 GHz
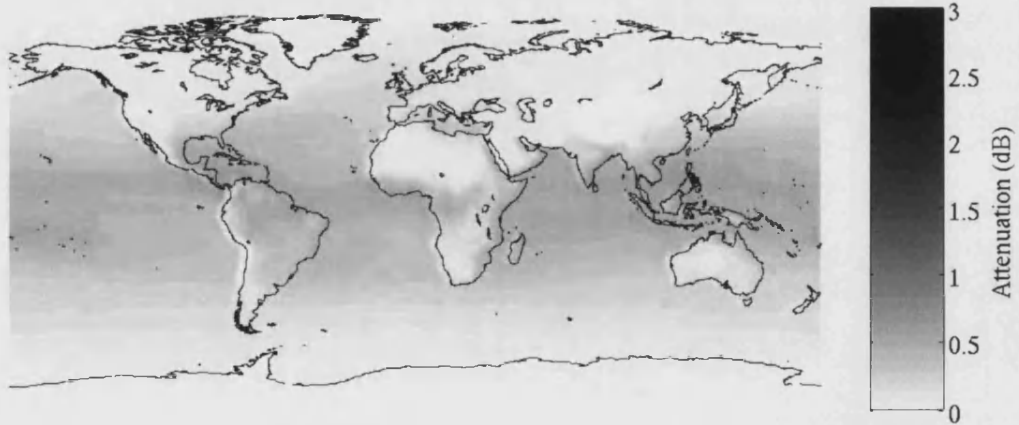
(b) Water Vapour: 50 GHz

Figure 2.2: Zenith attenuation (dB) due to absorption by atmospheric water vapour as predicted by the ITU-R recommendations.

2.1.3   Total gaseous attenuation

The specific attenuation due to oxygen and water vapour as a function of frequency is shown in figure 2.3. It can be seen that in addition to several spot resonances the general tendency is to increase with frequency.
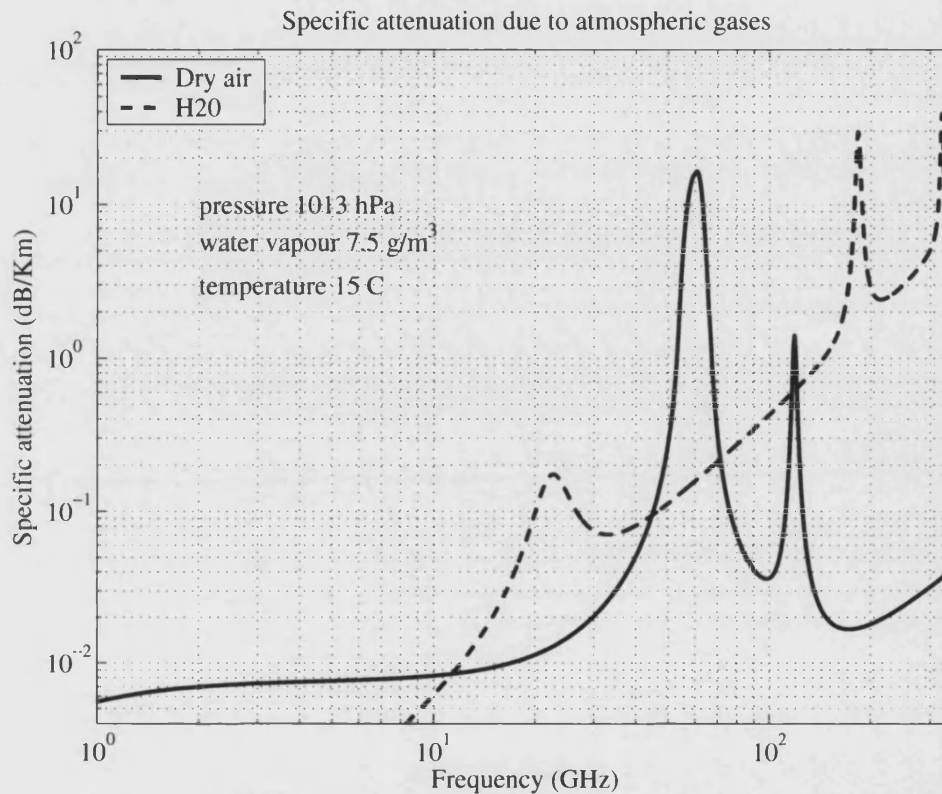
Figure 2.3: Specific attenuation due to gaseous absorption as predicted by the ITU-R recommendations.

The attenuation due to oxygen and water vapour both occur for all time (since the atmosphere contains both oxygen and water vapour for all time).

2.2   Effect of clouds

The attenuation due to clouds at EHF is very important especially in low availability systems. In the city of Bath, UK, the ITU-R recommended fade margin to overcome the attenuation due to cloud is equal or greater than that required to overcome the attenuation due to rain up to an exceedance of approximately 98%.

Clouds form in air which has become supersaturated, the most common way air becomes supersaturated is through ascent and adiabatic cooling. A more thorough discussion of the ascent process is covered in section 2.3.

Various types of ascent are possible each causing clouds of different types.

2.2.1 Convective clouds

Convective clouds are caused by the turbulent ascent of warm buoyant air parcels. Typically the bases of convective clouds are well-defined since the air at lower levels is well-mixed due to convective stirring, this results in a fairly uniform condensation level.

Whilst in the early stages of development the clouds show well-defined boundaries and consist mainly of liquid water even though they may extend above the zero-degree isotherm [Wallace and Hobbs, 1977]. As the cloud ages and the rising parcels of convective air lose their buoyancy the boundaries of the cloud become less defined. This reduction in definition is enhanced by the creation of ice particles which grow faster than water droplets and evaporate slower. Since these ice crystals survive longer at the cloud boundary they mix with the ambient air creating a diffuse cloud boundary.

The final stage of convective cloud growth occurs when the top of cloud begins to spread horizontally, causing what is commonly called an 'anvil'. Convective clouds which reach up to the tropopause are not uncommon, although less vigorous convection will be stopped by temperature inversions in the lower atmosphere.

These clouds have diameters ranging from a few hundred meters to tens of kilometers; within the cloud air ascends typically with vertical velocities of the order of a few meters per second. These ascents support a water content of generally around $1\,\mathrm{g\,m^{-3}}$, although water contents of $20\,\mathrm{g\,m^{-3}}$ have been measured in intense hailstorms. Typical lifetimes of convective clouds range from minutes to hours. Typical cloud droplet radii in convective clouds are around $40\,\mu\mathrm{m}$.

2.2.2 Layer clouds

Layer clouds are caused by the forced lifting of stable air. In temperate latitudes this is mostly associated with the development of a weather front. For example around a warm front the lifting of warm air over denser cold air produces layer clouds which decrease

in height closer to the front. Furthest from the front will be high clouds, typically consisting of ice particles several mm in diameter at low concentrations. Closer to the front the cold cloud ice layers thicken and descends taking on a uniform appearance. At this distance from the front several cloud layers may be present, and ice particles may fall from cold layers to other warmer layers.

In warm[3] layer clouds and fog the liquid water content is relatively uniform horizontally, and is typically of order of $0.1\,\mathrm{g\,m}^{-3}$. In warm non-precipitating layer clouds lifting rates range from approximately $2\,\mathrm{cm\,s}^{-1}$ to almost $10\,\mathrm{cm\,s}^{-1}$. These updrafts support cloud drops ranging in size from approximately $1\,\mu\mathrm{m}$ to $40\,\mu\mathrm{m}$, with the average cloud droplet size decreasing with height within the cloud layer.

2.2.3 Orographic clouds

These clouds are caused by the lifting of air above the condensation level as it passes over hills or mountains. Typically a mountain-wave cloud is created over terrain undulation. The liquid water in the cloud reaches a maximum over the windward slopes of the undulation whilst the cloud base is often higher on the windward side - since precipitation removes water from the air.

Under some conditions it is possible for the streamlines downwind of the undulation to undergo a series of oscillations creating 'lee-wave clouds' at intervals of the order of $5\,\mathrm{km}$, depending on the wind velocity, see figure 2.4. These are very common near long mountain ranges such as the Sierra and Cascade ranges in the USA or the Mendip Hills in the South West of the UK.
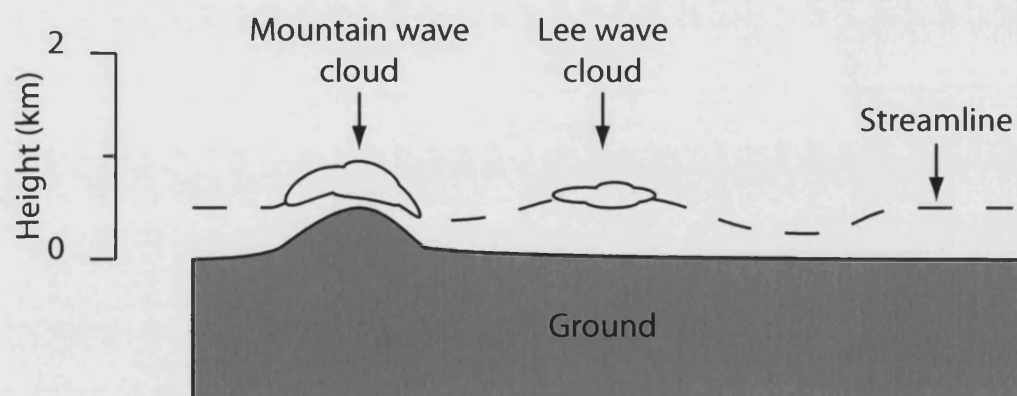
Figure 2.4: The creation of orographic clouds.

[3]In meteorological parlance 'warm' clouds refers to liquid clouds

Updrafts and lifetime will depend on the wind velocities and specific orography, however the water content is typically similar to that seen in layer clouds.

2.2.4 Attenuation due to clouds

Typically warm clouds below the zero degree isotherm clouds consist entirely of suspended liquid water with a drop size that rarely exceeds 0.1 mm, as discussed previously. This places the drop size within the Rayleigh scattering region at EHF and SHF frequencies. The Rayleigh region is valid for $|n|k_0a \ll 1$, where n is the refractive index of water, k_0 is the free space wave number and a is the radius of the drop. Even at the upper limit of the EHF and SHF frequency bands, the value k_0a only reaches values of the order of 0.25, and the Rayleigh approximation can been used up to frequencies of 100 GHz. In the Rayleigh region the attenuation is primarily due to absorption by the cloud droplets, scattering losses are less significant.

Maps of the zenith attenuation[4] due to cloud exceeded for 10% of the time at 15 GHz and 50 GHz are shown in figure 2.5.

---

[4]The average values of temperature are used for each location as defined by ITU-R P1510. The values of columnar content of liquid water for 10% exceedance are used as defined by ITU-R P840.

Zenith attenuation due to cloud at 15 GHz

(a) 15 GHz

Zenith attenuation due to cloud at 50 GHz
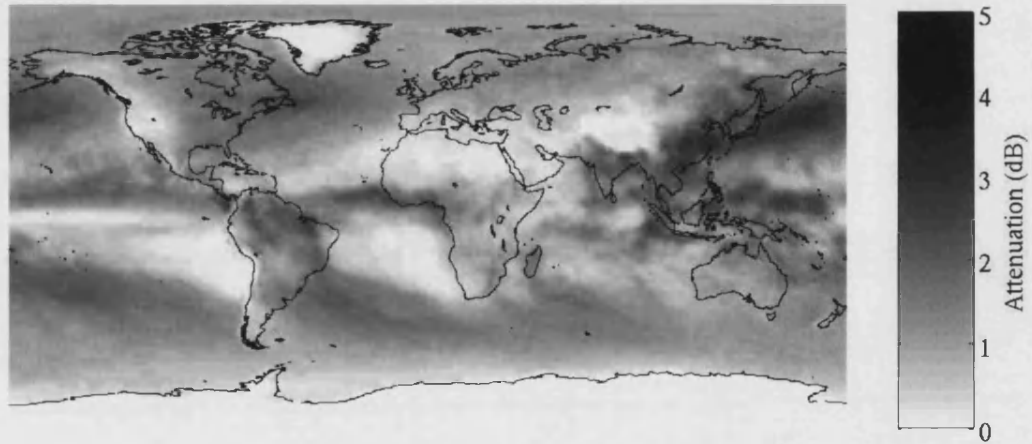
(b) 50 GHz

Figure 2.5: Zenith attenuation due to cloud attenuation exceeded for 10% of the time as predicted by the ITU-R recommendations.

Application of the ITU-R recommendations indicate that a 50 GHz link from Bath to a geostationary platform at 2W at 99.9 % availability requires an increase of 5 dB over the fade margin required for the same link at 20 GHz.

17

2.3 Hydrometeors

The attenuation due to hydrometeors is generally the source of the deepest fades at EHF and SHF frequencies. This section contains a discussion on the formation of hydrometeors and the effect they have on radiowave propagation.

The formation of rain in warm clouds is predominately due to two growth mechanisms: condensation, and collision and coalescence.

2.3.1 Condensation

The initial growth mechanism of water droplets is condensation. This process can be described by considering a parcel of air rising and expanding, this cools adiabatically and eventually reaches saturation. If the parcel of air continues to rise the parcel it becomes supersaturated; if this continues some droplets become activated (for discussion of this phenomena see appendix A). When the rate at which moisture is made available from adiabatic cooling equals the rate at which it is condensing into cloud droplets the supersaturation reaches a maximum value. Beyond this point the growing droplets consume more water than is being made available by the cooling of the air parcel and the supersaturation decreases.

It is possible to show that droplets growing by condensation initially grow very rapidly, but their rate of growth reduces with time. Consequently small drops grow quicker than larger drops and therefore over time the cloud droplets become an increasingly uniform size. Typically these processes can only produce droplets with radii up to about $10\,\mu$m yet drop sizes in clouds regularly exceed this size and rain drops are typically 1000 times larger.

2.3.2 Collision and coalescence

The process of collision and coalescence creates the larger droplets associated with rain from the smaller sized cloud droplets created by condensation. At sub-millimeter size the terminal fall velocity of cloud droplets increases with the droplet radius, hence larger droplets have a higher fall velocity and tend to collide with smaller droplets in their path.

Not all smaller droplets directly in a falling droplets path will collide with it, since very

small droplets will tend to follow the stream lines around the larger droplets. Larger droplets will have a smaller closure velocity so although the collision efficiency is higher there are fewer collisions.

If two drops collide it is not certain they will coalesce, it is possible for small droplets to bounce off a layer of a water surface. This occurs when air becomes trapped between the drops, this allows the drops to deform without touching. If this air is expelled before the drops rebound the droplets can touch and coalesce[5].

The 'continuous collision model' [Wallace and Hobbs, 1977] helps provide an understanding of the behaviour of the cloud drops as they grow to rain drops. The important conclusion of the model is the equation:

$$H = \frac{4\rho}{\omega_l} \left[\int_{r_0}^{r_H} \frac{w}{v_1 E} \, dr_1 - \int_{r_0}^{r_H} \frac{1}{E} \, dr_1 \right] \tag{2.1}$$

where, ρ represents the density of water, ω_l is the liquid water content of small drops that are coalesced, r_1 represents the radius of the collector drop, w is the updraft velocity, v_1 is the terminal fall velocity of a drop of radius r_1, E is the collection efficiency and r_0 and r_H represent the radius of the raindrop at the cloud base and a height H respectively.

Initially when the drop is small $w > v_1$ and the first integral in equation 2.1 is large hence the drop ascends. As the drop continues to grow v_1 now becomes greater than w and the first integral becomes less significant. At this point the drop falls through the updraft and, if it does not break up, will fall through the cloud base as rain.

Using equation 2.1 it can be seen that warm clouds with strong updrafts produce rain in a shorter time than clouds with weaker updrafts. However, clouds with stronger updrafts must be thick in order to produce raindrops. Raindrops that form in deep clouds with strong updrafts should be larger drops. A summary of typical precipitation characteristics is shown in table 2.1.

2.3.3 Rain formation in cold clouds

Clouds that exist above the zero-degree isotherm height are typically termed cold clouds. These clouds consist of a mixture of supercooled water droplets and ice particles. Ice crystals grow in cold clouds by three different mechanisms:

---

[5]It is worth noting that even if drops coalesce it is still possible for the new drop to break up if the motion is unstable.

Table 2.1: Typical precipitation characteristics, [Wiesner, 1970].

Meteorological situation	Average vertical velocity $(\mathrm{m\,s^{-1}})$	Average size of water droplet (mm)	Precipitation Type
Widespread uplift			
Depressions on a synaptic scale	0.026	0.1	Steady light rain
Depressions with layer clouds	≤ 0.05	0.1-1.0	Moderate rain
Convection			
Convective clouds	≤ 10.3	1.0-2.5	Showers
Cumulo nimbus clouds and thunderstorms	≤ 20.6	>3	Heavy showers

Growth from vapour. In a mixed cloud dominated by supercooled droplets, the air is close to saturated with respect to liquid and hence supersaturated with respect to ice. Air saturated with respect to liquid water at -10°C is supersaturated with respect to ice by 10 %, and at -20°C it is supersaturated by 21 % [Wallace and Hobbs, 1977]. These supersaturations are much higher than the supersaturations with respect to water and hence ice crystals grow much quicker than liquid droplets.

Growth by rimming. In a mixed cloud ice particles can increase in mass by the collision with supercooled droplets which freeze to the ice particles; this process is called rimming. When an ice particle becomes noticeably rimmed it is often termed 'graupel'.

Growth by aggregation. A further way in which ice crystals can grow in cold clouds is by aggregation. This is the process whereby ice crystals within the cloud collide and adhere to one another. This process is particularly significant in clouds with crystals that have undergone a large degree of rimming. This creates a large variation in the size (and terminal fall velocity) of the rimmed crystals[6], this increase in the variability of fall speeds causes an increase in the number of collisions and hence an increase in aggregation.

In common with warm clouds, deposition from vapour is not able to produce significantly large raindrops. For example, at -5°C in air saturated with respect to water an ice crystal can obtain a diameter of about 1 mm in about 30 minutes. Beyond this

---

[6]Un-rimmed particles tend to have a fall velocity independent of diameter [Wallace and Hobbs, 1977].

period growth decreases very rapidly. This crystal can form a rain drop of approximately 0.13 mm in diameter (assuming it does not evaporate as it leaves the cloud base and the updraft is less than the terminal fall velocity of the crystal), this drop size is consistent with very light rain.

Growth by rimming and aggregation increase with particle size, and it can be shown that the 1 mm crystal falling through a cloud with a liquid water content of $0.5\,\mathrm{g\,m^{-3}}$ could develop into a graupel particle with a density of $100\,\mathrm{kg\,m^{-3}}$ in about 10 minutes. This particle would create a drop around 0.25 mm in radius. Indeed a snow flake can increase from 1 mm to 1 cm by aggregation in about 30 minutes if the ice content of the cloud is about $1\,\mathrm{g\,m^{-3}}$, this would create a rain drop about 1 mm in radius.

2.3.4   Shape and fall velocity of raindrops

Once a raindrop has left a cloud base (either a warm or cold cloud) it starts to display different characteristics. The first to be considered is the terminal fall velocity of a raindrop outside the cloud. Some of the most extensive measurements were performed by Gunn and Kinzer [1949] who described how the fall velocity increased as the raindrop size increased. As the raindrop diameter exceeds 2 mm the rate of increase in the terminal velocity decreases and at a diameter of around 4 mm the terminal velocity achieves the maximum terminal velocity ($\approx 9\,\mathrm{m\,s^{-1}}$).

A fit to this data (and similar work by other authors) was presented by Atlas et al. [1973], the equation 2.2 describes this fit (assuming stagnant air at 1013 mb, 20°C and 50 % relative humidity).

$$w_t\,(D) = 9.65 - 10.3\exp\,(-600D) \qquad (2.2)$$

where w_t is the terminal velocity ($\mathrm{m\,s^{-1}}$) and D is the rain drop diameter.

By plotting equation 2.2 the variation of terminal fall speeds with rain drop size is perfectly clear, this is shown in figure 2.6 along with the measurements presented by Gunn and Kinzer [1949] and Best [1950].

It is also important to understand the shape of rain drops falling at the terminal fall velocity. The work by Pruppacher and Beard [1970] and Pruppacher and Pitter [1971] made several observations regarding the shape of rain drops for different drop radius:

- rain drops with a radius $\leq 0.17\,\mathrm{mm}$ are very slightly deformed and can be con-
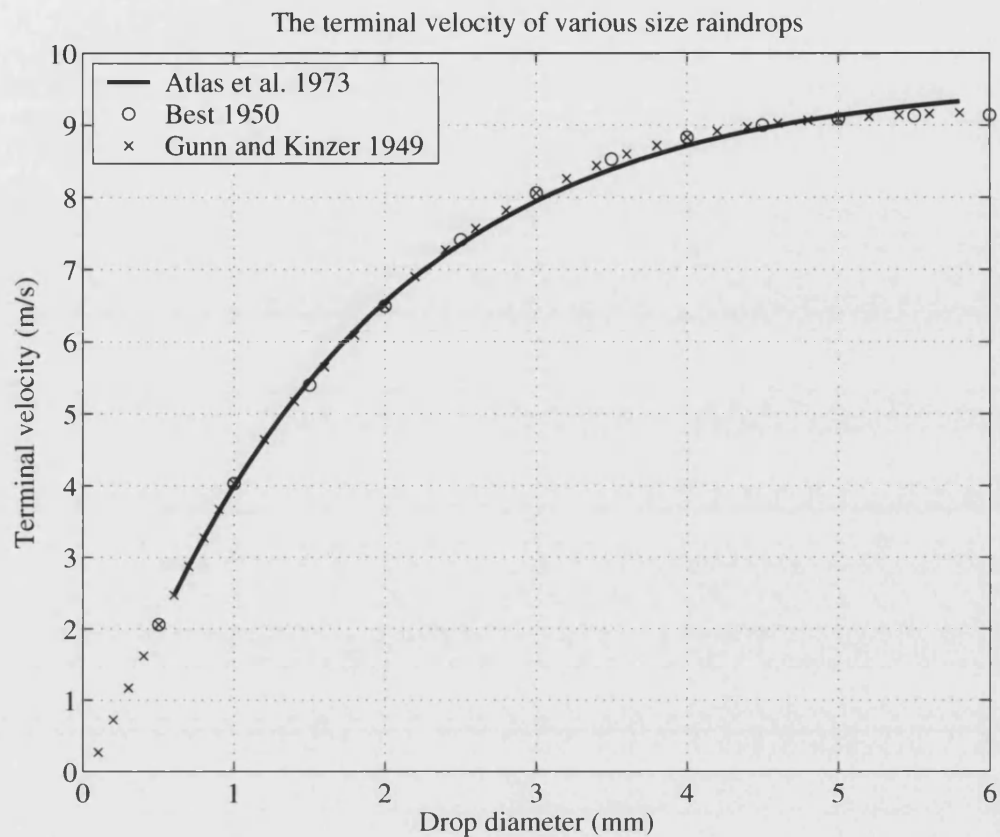
Figure 2.6: The variation of raindrop terminal velocity with rain drop diameter.

sidered spherical,

- rain drops with a $0.17\,\text{mm} \leq \text{radius} \leq 0.50\,\text{mm}$ can be closely approximated by an oblate spheroid,

- rain drops with a $0.50\,\text{mm} \leq \text{radius} \leq 2.00\,\text{mm}$ are deformed into an asymmetric oblate spheroid with an increasingly pronounced flat base,

- rain drops with a radius $\geq 2.00\,\text{mm}$ develop a concave depression in the base which becomes more pronounced.

Pruppacher and Pitter [1971] derived a semi-empirical relationship linking the rain drop shape and equivolumetric rain drop radius, a_o. Illustrations of rain drops of several equivolumetric rain drop radii have been calculated from this relationship and are included in figure 2.7.

From these shape measurements and formulations it is possible to calculate the vari-

a_o=1mm a_o=2mm

a_o=3mm a_o=4mm
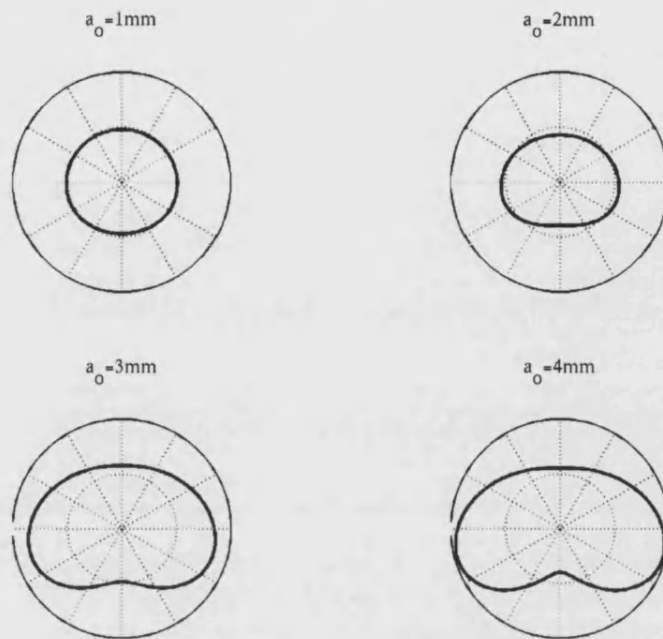
Figure 2.7: Typical rain drop shapes.

ation of the rain drop deformation[7] a/b and a_o, the variation is shown in figure 2.8.

From the discussion in the previous section it is apparent that since larger updrafts can support larger droplets before precipitation occurs there is a large range of possible raindrop size, ranging from a raindrop size comparable to those in some denser clouds up to about 4 mm in radius. Raindrops with a radius in excess of 4 mm are hydrodynamically unstable and will tend to break up as they descend [Oguchi, 1983].

2.3.5 Raindrop size distribution

The tools defined in section 2.3.4 allows the definition of the shape and speed of a falling rain drop given a certain rain drop size. To complete the characterisation information is required defining the distribution of rain drop sizes.

The most notable initial works on raindrop size distribution were by Laws and Parsons [1943] and the later work by Marshall and Palmer [1948]. In the former it was noted that for the same rainfall rate the size distributions varied greatly from rain event to rain event. Whilst the later work focused on providing a relationship between the

---

[7]where $2a$ is the height and $2b$ is the width of the drop
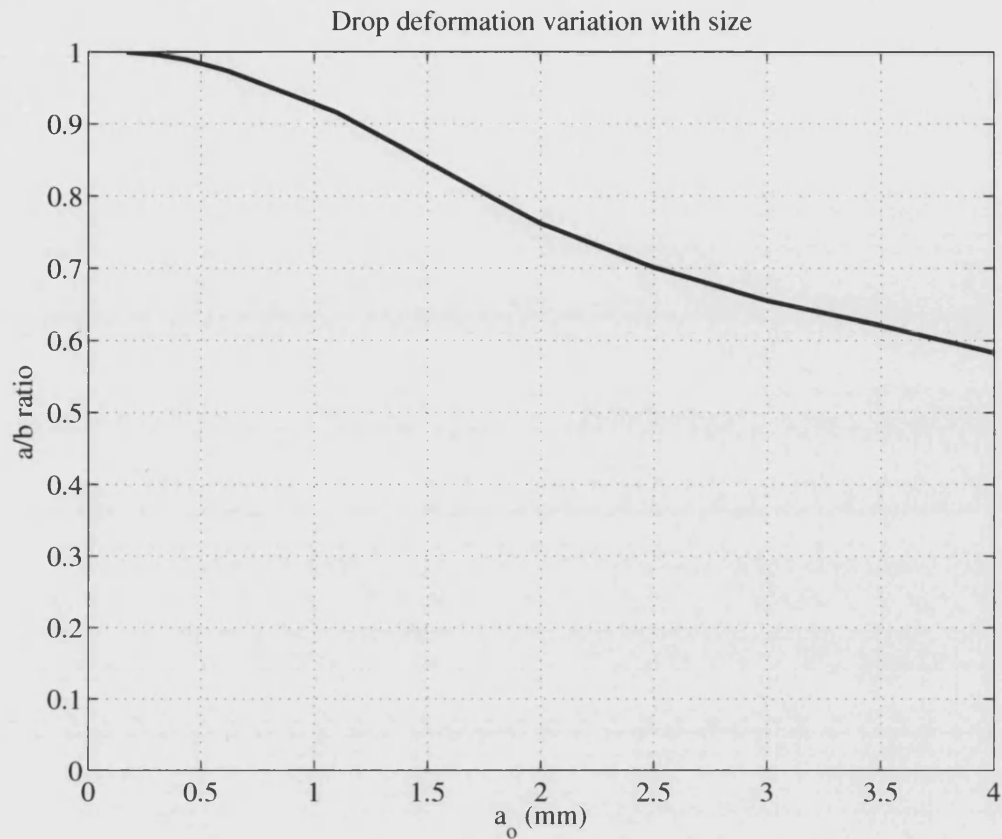
Figure 2.8: The variation of the raindrop deformation ratio.

rainfall rate and the distribution of raindrop size.

Marshall and Palmer using their own measurements in combination with those of Laws and Parsons found a relationship between the number of rain drops of per unit volume of space with radius a in a range[8] da and the rainfall rate, this is shown in equation 2.3.

$$n(a)\,da = N_o \exp(-\Lambda a) \tag{2.3}$$
$$\Lambda = 8.2R^{-0.21} \tag{2.4}$$

where N_o is defined as the intercept parameter ($N_o = 1.6 \times 10^4$), Λ is defined as the slope parameter, a is the rain drop radius (mm) and R is the rain fall rate (mm hr^{-1}).

The Marshall and Palmer rain drop size distributions for four different rainfall rates are shown in figure 2.9

---

[8]da denotes a size interval ranging from $a - \frac{da}{2}$ to $a + \frac{da}{2}$
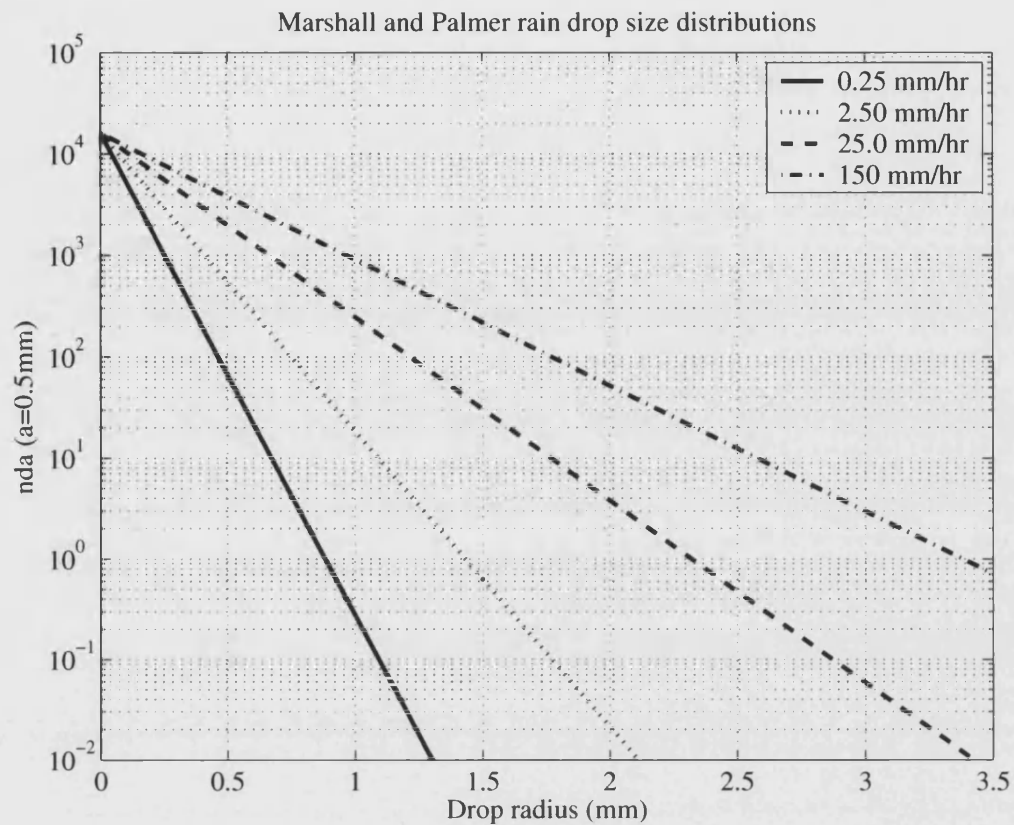
Figure 2.9: The Marshall and Palmer rain drop size distribution for several different rain rates.

The initial works by Laws and Parsons [1943] and Marshall and Palmer [1948], although performed over half a century ago, are still considered to be typical of the average rain drop size distribution both for both widespread and convective rain [Oguchi, 1983].

Marshall and Palmer, and others, suggested the use of the exponential distribution (for example equation 2.3), others have suggested the use of the gamma distribution [Ulbrich, 1983].

The gamma distribution has been introduced in order to account for the shape of the distribution at high rainfall rate. This is characterised by a reduction in small rain drops and by a convexity of the 'tail' for large drop diameters.

2.3.6   Canting angle

Non-spherical rain drops falling through a vertical airflow (i.e. through stagnant air or a vertical wind direction) will have the axis of rotational symmetry vertical, since the mean resultant drag force is vertical. It follows that the orientation of the axis of rotational symmetry for a rain drop is parallel to the direction of the airflow (and hence drag forces) around the drop [Brussaard, 1976].

Hence the inclination of the airflow with respect to the drop is equal to the canting angle of the drop as described in figure 2.10.
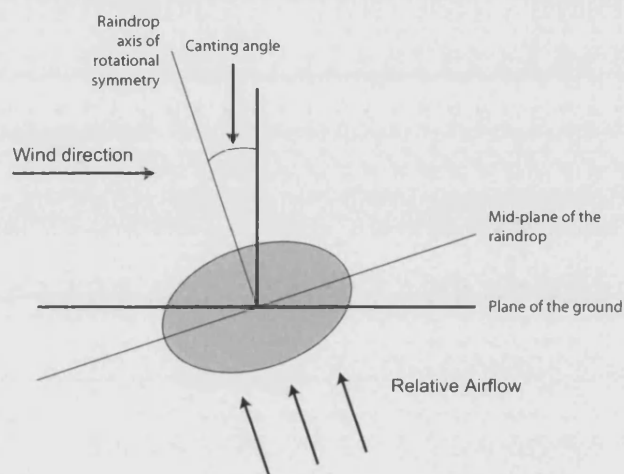
Figure 2.10: Raindrop canting.

In the simplest first approximation where the horizontal wind speed does not vary with height ($\frac{dV_{horiz}}{dz} = 0$) it is apparent that the raindrop is not canted. The raindrop is falling at the same horizontal velocity as the surrounding air particles so the airflow relative to the drop is vertical, although the drop trajectory is slanted. Essentially the drop is not being accelerated or retarded in the horizontal direction so there are no horizontal drag forces and the drop is stable with the axis of rotational symmetry vertical.

Consider a more realistic situation where the horizontal wind speed does vary with height ($\frac{dV_{horiz}}{dz} \neq 0$). It is apparent that the relative airflow is not vertical as the drop is accelerated or retarded in the horizontal direction, and the rain drop will be canted appropriately. Hence the canting angle is a function of the differential of the vertical wind profile, not the absolute value of the wind speed.

Raindrop canting is an important phenomena particularly when considering polarisation schemes. Orthogonal linear polarisations will suffer independent fading levels from a non-spherical canted drop, this will result in a reduced polarisation discrimination.

2.3.7 Rain attenuation

Rain attenuation is a very important source of attenuation at EHF and SHF frequencies since it causes the deepest fades. It is also one of the hardest to accurately quantify; as previously noted there is a large variability of rain drop size distribution from rain event to rain event, even if the events are of similar rain rates. This variation in rain drop size distribution is further compounded by the variability in vertical updraft profiles and micro-scale phenomena (such as embedded convection or supercooled liquids).

Statistics of the attenuation caused by rain can be generated using the ITU-R recommendation P.838-2; maps of the zenith attenuation due to rain are shown in figure 2.11.

Zenith attenuation due to rain exceeded for 1 % at 15 GHz
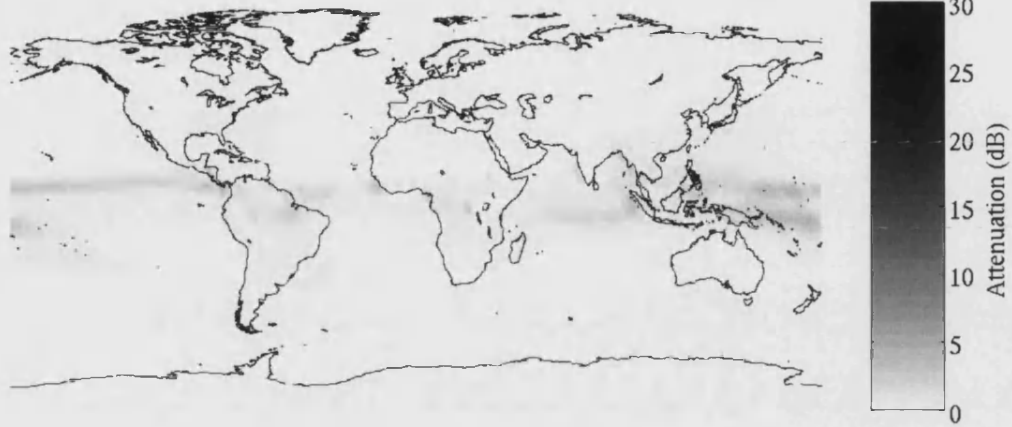
(a) 15 GHz

Zenith attenuation due to rain exceeded for 1 % at 50 GHz
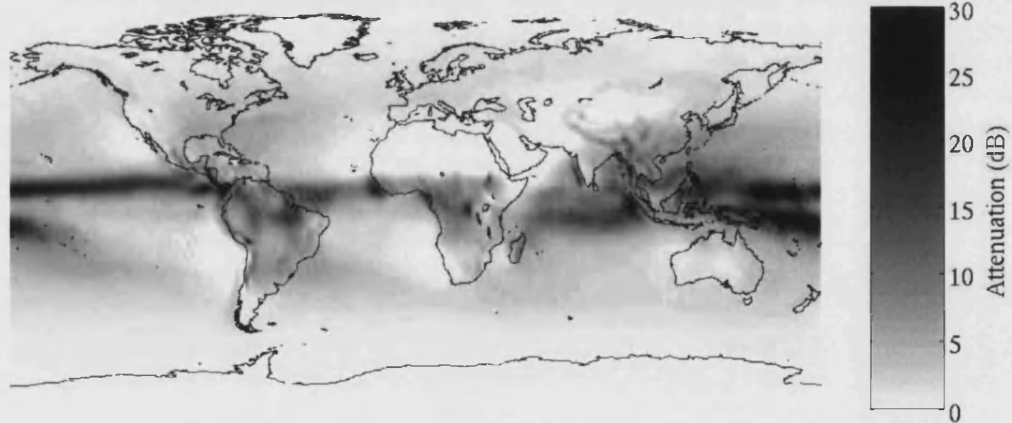
(b) 50 GHz

Figure 2.11: Zenith attenuation (dB) due to rain exceeded for 1% of the time as predicted by the ITU-R.

The same method can be used to calculate the expected attenuation for a given exceedance. This has been done at two frequencies (20 GHz and 50 GHz) for a link to a geosynchronous satellite at 2°West, this is shown in figure 2.12. As can be seen the increase in frequency causes a very high attenuation at high exceedances.
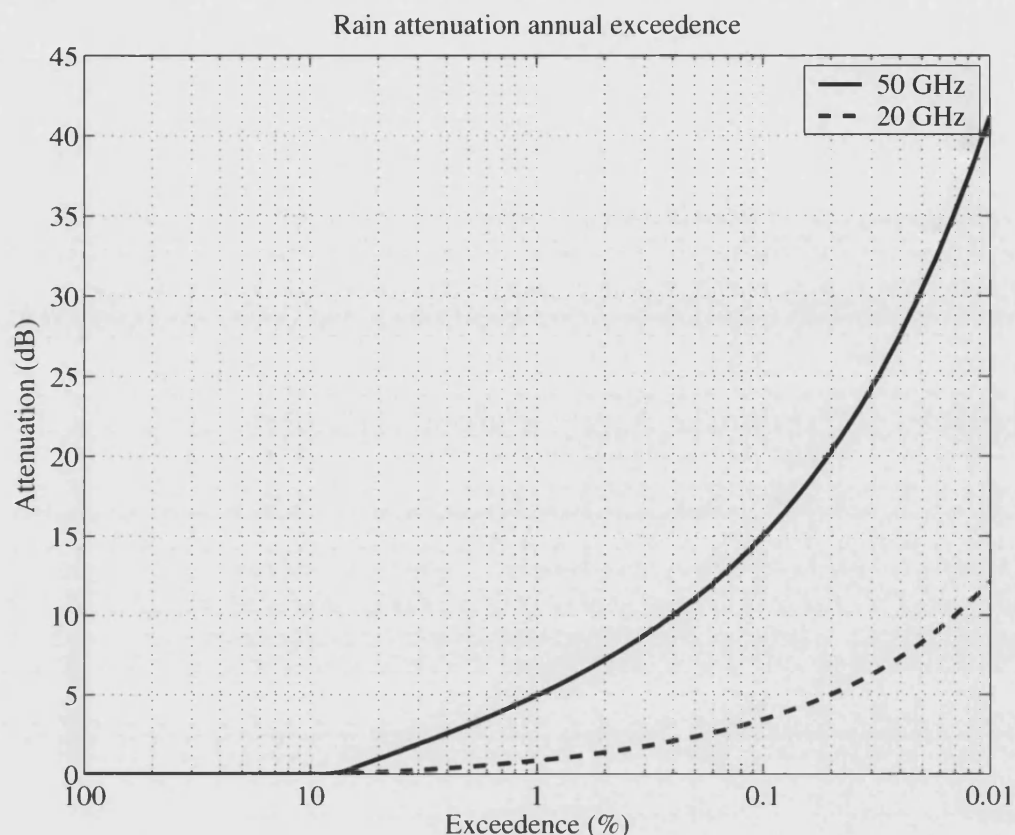
Figure 2.12: Annual rain attenuation exceeded at 20 GHz and 50 GHz for Bath.

2.4 Effects of ice, snow and the melting layer

The melting layer is the layer in a cold cloud below the zero-degree isotherm where there is the transition between snow and water, often appearing as a 'bright-band' on radar scans. This band corresponds to the area where ice particles are melting and the ice particles become coated in water greatly increasing the radar reflectivity. The melting layer is typically associated with stratiform rain events, since the increased turbulence in convective rain events tends to disrupt the forming of a distinct layer. It has been shown that the thickness of melting layer ranges from a few hundred meters to a maximum thickness of ≈ 1417 m dependent upon the rain drop size [Zhang, 1994].

Work presented by Zhang [1994] and Zhang et al. [1996] addressed the behaviour of the reflectivity, specific attenuation and specific phase shift of a melting layer consisting of oblate-spheroidal hydrometeors, it is accepted that the approximation of actual snow particles as oblate-spheroids is sound [Zhang et al., 1996]. The results show that at 35 GHz the specific attenuation of the melting layer was a maximum of $0.1\,\mathrm{dB\,km^{-1}}$

and $3\,\mathrm{dB\,km^{-1}}$ for rain rates of $0.25\,\mathrm{mm\,hr^{-1}}$ and $12.5\,\mathrm{mm\,hr^{-1}}$. A rainfall rate of $12.5\,\mathrm{mm\,hr^{-1}}$ is the maximum rain rate considered in order to ensure that the rain events were stratiform. There is also a phase shift of up to $50°\,\mathrm{km^{-1}}$ at $35\,\mathrm{GHz}$ for a rainfall rate of $12.5\,\mathrm{mm\,hr^{-1}}$, and a differential phase shift between H and V polarisations of around $3°\,\mathrm{km^{-1}}$ under the same conditions.

2.5 Increased noise temperature due to increased attenuation

During a fade due to rain or clouds there is an increase in the system noise temperature, caused by the effect of a warm attenuating media. This compounds the fade as not only is the signal degraded, but the noise temperature is also increased further reducing the signal-to-noise ratio.

As a first-order approximation the warm attenuating media can typically be assumed to be a black body radiator. This means that it absorbs all energy incident upon it, reflecting none, and it is also a perfect emitter. In this case the energy emitted (in the form of noise) from clouds, rain and other warm attenuators is equal to the energy absorbed.

The atmospheric contribution to the antenna noise in a ground station may estimated to a first order by equation 2.5

$$T_{eff} = T_s + T_m \left(1 - 10^{\frac{-A}{10}}\right)$$ (2.5)

where T_{eff} is the effective (or apparent) sky noise temperature (K), T_s is the clear-sky noise temperature as seen by the antenna (K), A is the path attenuation (dB) and T_m is the effective noise temperature of the attenuating media (K).

ITU-R recommendation P.618-8 suggests that the effective noise temperature is dependent on the attenuation due to scattering, the physical effect of clouds and rain cells on the vertical temperature profile and the antenna beamwidth. Below $60\,\mathrm{GHz}$ the upper limits of sky-noise temperature can be taken to be $280\,\mathrm{K}$ and $260\,\mathrm{K}$ for clouds and rain respectively.

The effect of this noise increase is shown in figure 2.13 for a $T_m = 260\,\mathrm{K}$, and a clear sky noise temperature $T_s = 60\,\mathrm{K}$.
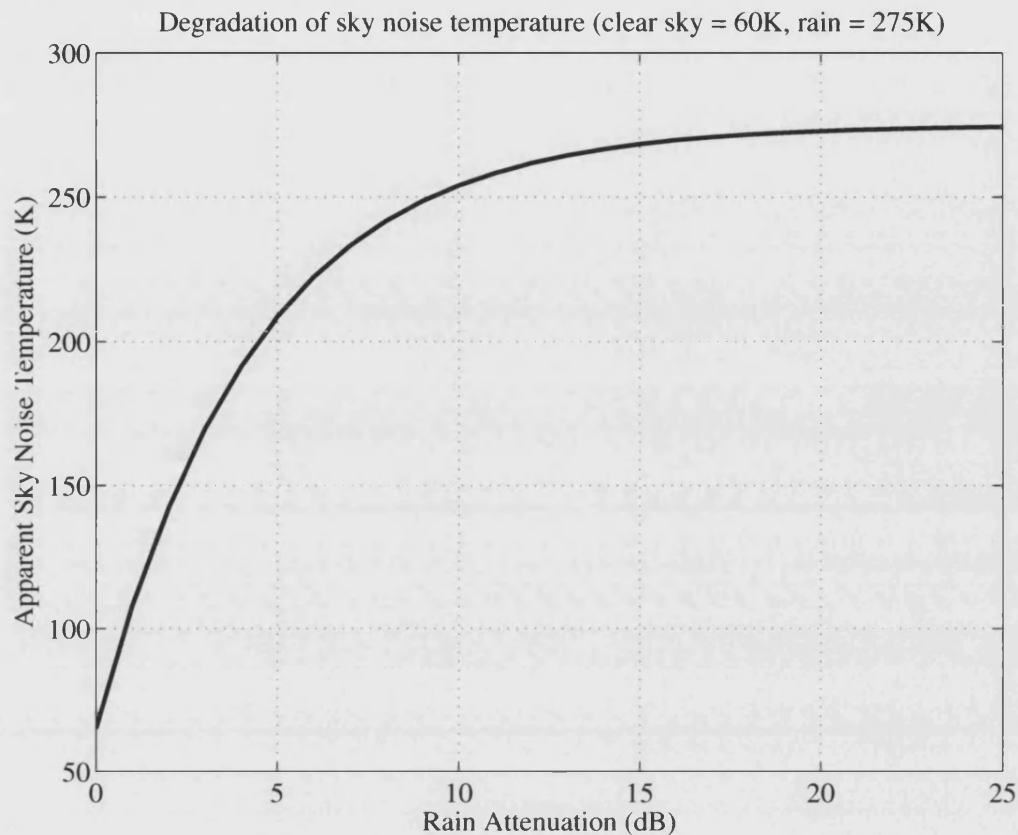
Figure 2.13: Atmospheric contribution to receiver noise temperature.

2.6 Scintillation

Scintillation is the rapid fluctuations of signal level caused by variations in the physical transmission media. This phenomena occurring at optical frequencies is what causes the stars to appear to 'twinkle'. Scintillation primarily occurs in two areas of the atmosphere: the ionosphere and the troposphere.

2.6.1 Ionospheric scintillation

The ionosphere is a region of ionised gas at an altitude peaking in concentration between 300 and 400 km. The level of ionisation fluctuates diurnally and from day to day with solar activity. Small-scale variations in the level of ionisation cause small scale irregularities in refractive index and hence cause scintillation. Typically the effect of the ionospheric scintillation can be ignored for frequencies above 20 GHz [Willis, 1991].

2.6.2 Tropospheric scintillation

Scintillation in the troposphere is caused by turbulent irregularities in temperature, humidity and pressure. These irregularities lead to small-scale variations in refractive index. The rapid fluctuations in link power caused by these irregularities are essentially random since the autocorrelation function of the scintillation component has been shown to fall away after a few seconds [Grémont et al., 1996]. The only way to describe the scintillation component is to define the distribution function, which is well described by a Gaussian distribution and therefore the defining factor is the scintillation variance (often called the scintillation intensity).

The tropospheric scintillation variance is understood to increase with temperature and humidity (often expressed by the 'wet term of refractive index', N_{wet}), and hence the scintillation component becomes more significant in equatorial climates and in temperate climates during the local summer.

Other work has shown that the variance of tropospheric scintillation is better correlated with the presence of 'heavy clouds' [van de Kamp et al., 1999]. In this context heavy clouds are defined as a cloud layer with an integrated water content in excess of $0.7 \, \text{kg} \, \text{m}^{-2}$ essentially this describes the presence of convective clouds. As discussed in section 2.2.1 these convective clouds show a large amount of turbulent mixing under the base. As a direct result of this turbulence it is apparent that the variance of the scintillation component will increase.

Scintillation effects are very important especially in low fade margin systems where the magnitude of scintillation fades can be comparable to the fade margin. The scintillation effects are also exaggerated on low elevation systems where the slant-path intersects with a considerable number of the turbulent irregularities.

To avoid scintillation fades causing burst errors it is important to have either a static fade margin large enough to encompass the scintillation component or an FMT with a fast enough update to be able to track the scintillation fades [Grémont et al., 1999].

2.7 Fade dynamics

The dynamic characteristics of the channel define how the propagation environment changes over time. Knowledge of the fade dynamics is incredibly important within an FMT environment as these ultimately influence the design of control systems.

2.7.1 Fade slope

Fade slope typically refers to the change in the slow varying components of the attenuation (typically attributed to the attenuation due to rain, gases and clouds) the fast fluctuations caused by effects such as tropospheric scintillation are not considered in the fade slope estimates. It is for this reason that experimental measures of fade slope are typically very difficult to obtain; all beacon measurements must have the scintillation component removed without removing any of the slow component.

Typically a low pass filter with a cut-off frequency of 20 mHz is used to remove the scintillation component [van de Kamp, 2003]. More recent work has shown that the use of the wavelet transform can be used with much greater success [Baxter and Garcia, 2004].

Beacon measurements from various Ka- band propagation experiments were used by van de Kamp [2003] to draw several conclusions about the behaviour of the slope of the slow varying components of the propagation environment.

The probability distribution of the fade slope (ζ) at a given attenuation level (A) was found to be a reasonable fit to the following model distribution, equation 2.6:

$$p\left(\zeta \mid A\right) = \frac{2}{\pi \sigma_\zeta \left(1 + \left(\frac{\zeta}{\sigma_\zeta}\right)^2\right)^2} \tag{2.6}$$

where ζ is the fade slope (dB s^{-1}), A is the fade level (dB) and σ_ζ is the standard deviation of the slope at that attenuation level.

If the attenuation quantization bins are suitably small it is possible to accurately estimate evaluate the value of σ_ζ using equation 2.7:

$$\sigma_\zeta\left(A\right) = \sqrt{\int_{-\infty}^{\infty} \zeta^2 p\left(\zeta \mid A\right) d\zeta} \tag{2.7}$$

The conditional probability distribution for several different values of σ_ζ is shown in figure 2.14. It is apparent that the distribution is symmetrical, i.e. at a given fade level the probability of a negative fade slope of a given magnitude is identical to that of a positive slope of the same magnitude.

It was also found that the standard deviation of the fade slope in each attenuation bin is approximately proportional to the attenuation level. The constant of proportionality
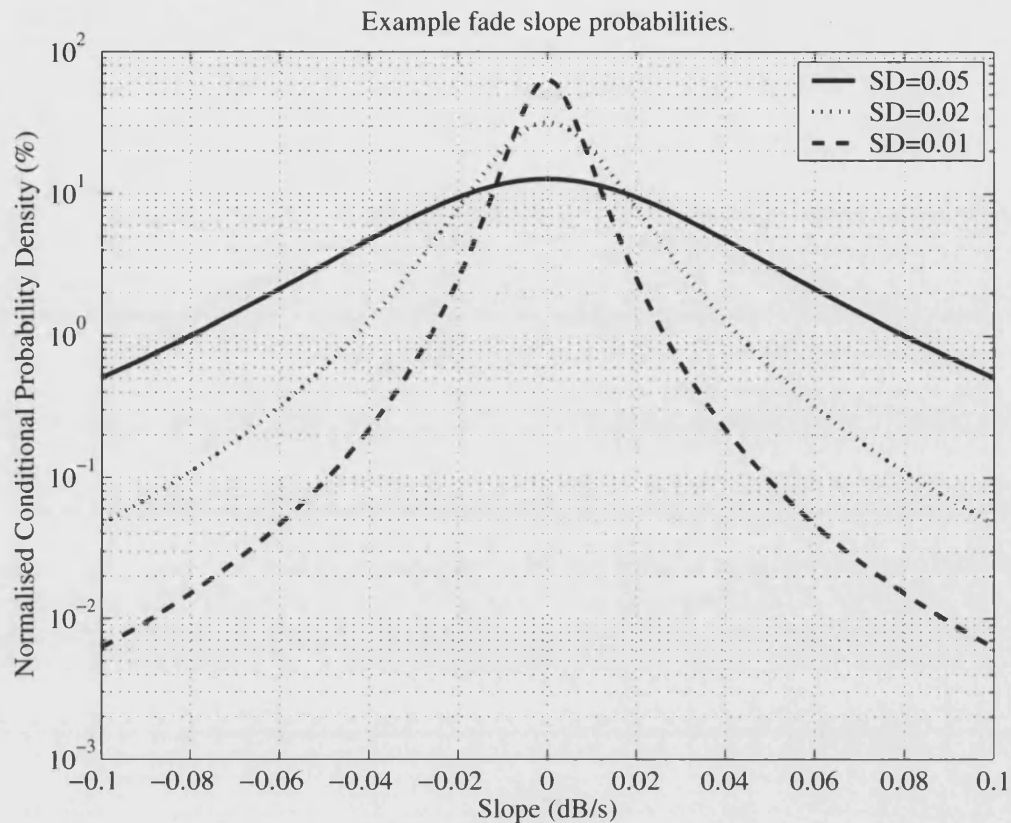
Figure 2.14: Example fade slope conditional probabilities.

is a function of the time constants of the receiver[9], climate zone and elevation angle [van de Kamp, 1999].

Further conclusions were also drawn that at a given attenuation level there is little correlation between fade slope and carrier frequency. However, fade slope standard-deviation does appear to be a function of the elevation angle. Low elevation links appear to exhibit weaker fade slopes, it is believed that this is due to an averaging effect. This averaging effect occurs since the probability of several rain cells or clouds occurring on the link is higher [Feil et al., 1997].

2.7.2 Fade duration

Fade duration is a metric that defines the length of time during which a given attenuation level is exceeded. This metric can be very important in defining the outage time

---

[9]This refers to the integration time of the receiver as well as the filter bandwidth used to remove the scintillation component

of communications system, it can also provide a statistical background within which resource allocation strategies or FMTs can be designed.

It has long been understood that fade durations should be divided into two time scales, long fade durations (longer than typically 30 seconds) caused by attenuation from meteorological conditions and shorter fade durations essentially due to tropospheric scintillation effects.

Fade duration statistics are typically presented as either:

1. Statistics of the number of fade events above a given attenuation threshold and longer than a given duration.

2. The total duration of fade events higher than a given attenuation threshold and longer than a given duration.

3. The statistics of fade duration with respect to the fraction of time during which attenuation is above a given threshold.

An example of the results from the OPEX experiment [Castanet, 2001] is shown in table 2.2. These show the percentage time that a fade level of 10 dB is exceeded; for example at 12.5 GHz, a 10 dB fade with a duration of 700 s or longer will occur for 1% of the time.

Table 2.2: Example fade durations from the OPEX experiment.

% of total outage time	Frequency		
	12.5 GHz	20 GHz	30 GHz
50 %	2 s	3.5 s	9 s
10 %	100 s	150 s	250 s
1 %	700 s	1000 s	1500 s

The most widely used model of fade duration in Europe is the model described by Paraboni and Riva [1994]. This model describes long duration fades above a transition time D_t using a log-normal model and shorter fades (shorter than the transition time) using a power law model. The transition time is computed from equations 2.8 to 2.11.

$$\ln\left(\frac{D_t}{D_o}\right) = p_1\sigma^2 + p_2\sigma - 0.39 \qquad (2.8)$$

$$p_1 = 0.885\gamma - 0.814$$

$$p_2 = -1.05\gamma^2 + 2.23\gamma - 1.61$$

The parameters D_o (conditional mean duration), σ (conditional standard deviation)

and γ (power-law exponent) have been evaluated for the OPEX data used in this particular study and have been evaluated as:

$$D_o = 3050\phi^{-0.4} \tag{2.9}$$

$$\gamma = 0.2\exp(-0.4A + 0.0125f + 0.15) \tag{2.10}$$

$$\sigma = 0.15 \tag{2.11}$$

where ϕ is the elevation angle, A is the attenuation (dB) and f is the frequency (GHz).

Chapter 3

Review of propagation channel modelling

In this chapter the aim is to give a brief overview of some EHF channel models. There are many different types of model but this chapter aims to cover a range of different approaches to the challenge of modelling the EHF and SHF propagation channel. The first method is a statistical method to generate the cumulative distribution function of attenuation using the ITU-R recommendations, there then follows several 'synthetic time series' techniques which attempt to generate time series of attenuation that have the same statistics as observed time series. Finally the approach presented in this Thesis is contrasted with these techniques.

3.1 ITU-R P.618-8: 'Propagation data and prediction methods required for the design of Earth-space telecommunications systems'

ITU-R recommendation P.618-8 can be used to provide important information for the planning of satellite links. This provides estimates of the statistics of various propagation effects that should be considered on earth-space links.

The recommendation considers the effects of the following:

- Absorption in atmospheric gases.

- Absorption, scattering and depolarization by hydrometeors and emission noise.

- Effects of beam-divergence due to refraction.

- The decrease in antenna gain due to phase decorrelation across the antenna aperture caused by irregularities in the refractive index.

- Slow fading from beam-bending caused by large-scale changes in refractive index; and scintillation and variations in angle-of-arrival, due to small-scale variations in refractive index.

- Multipath effects.

- Attenuation by the local environment of the ground terminal.

- Short term variations in the frequency scaling ratio.

- The effects of variable elevation angle for non-geostationary satellites.

- Ionospheric effects (typically only important for frequencies <1 GHz).

The statistics of these effects are provided as an attenuation cumulative density function. The recommendation P.618-8 also suggests using other ITU-R recommendations for some propagation impairments. Using the combination of recommendations shown in figure 3.1 it is possible to create an average annual attenuation CDF. This has been done for four example frequency bands for an uplink from the University of Bath to a geostationary satellite at 2°West; the results of which are shown in figure 3.2 in two different forms. One form is as a function of the availability and the other as a function of the available fade margin.
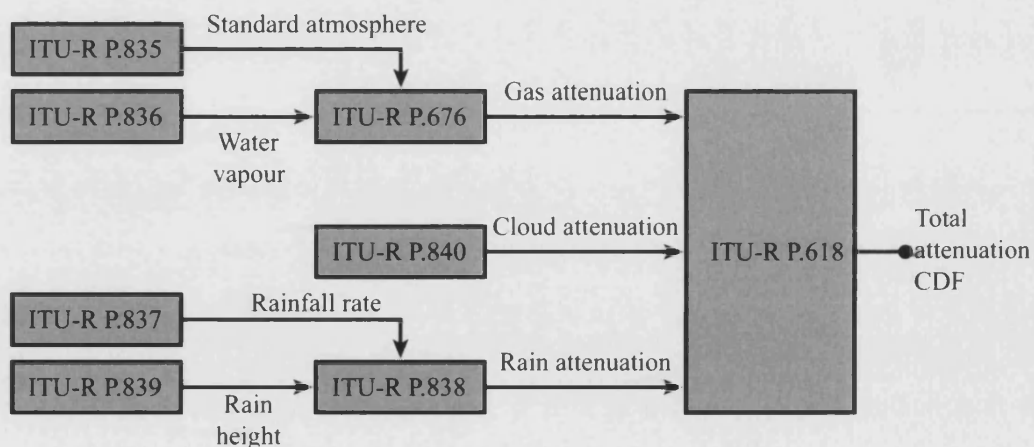
Figure 3.1: The procedure used to provide ITU-R recommended fade margins.
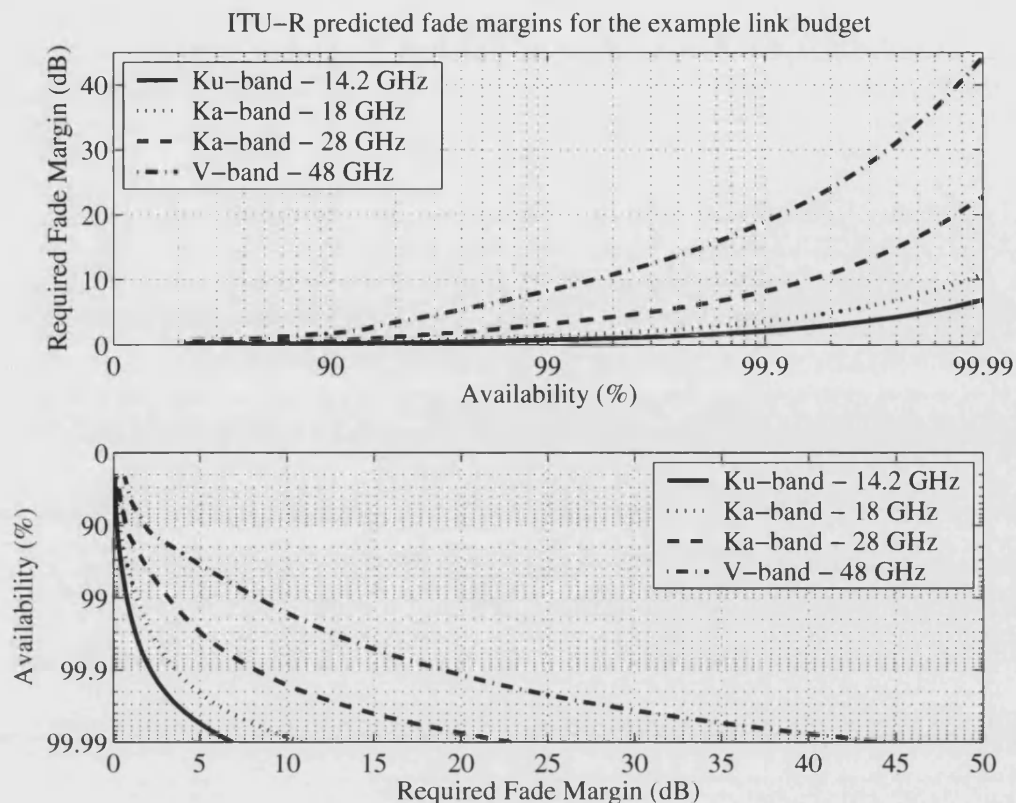
Figure 3.2: The ITU-R recommended fade margins for an average year for a link from the University of Bath to a geostationary satellite at 2°West.

3.2 Spectral approach - B. Audoire, Office National d'Etudes et de Recherches Aérospatiales (ONERA)

This is possibly the most simple single-site model available and was part of the PhD of Audoire [2001] with significant input from Boris Grémont at the University of Portsmouth. This time series generator consists of two parts: the first part creates the slow-varying component whilst the second part creates the fast-varying component.

The slow varying component is generated following the procedure shown in figure 3.3. The procedure starts with a Gaussian random variate which is then filtered by a first order low pass filter. This filter is defined by the theoretical power spectrum from the stochastic rain model of Maseng and Bakken [1981]. After the filter the time series is then normalised to a relevant CDF. This CDF can either be from ITU-R recommendations or from propagation experiments such as the Italsat or Olympus beacons measurements.
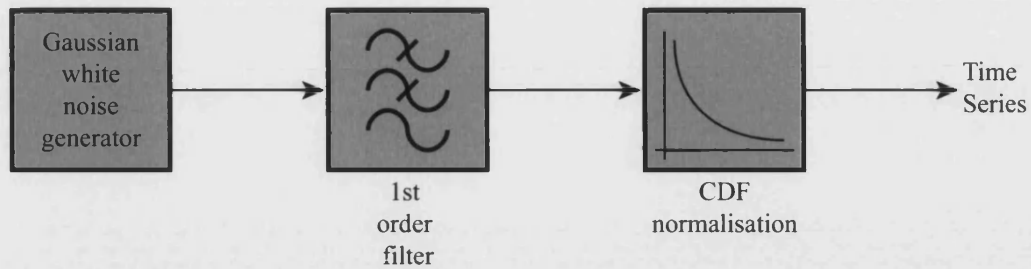
Figure 3.3: The procedure used to generate the slow varying component in the spectral approach.

The fast-varying component is created in a similar way. A Gaussian random variate (independent of the variable used to create the slow-moving component) is filtered with a Yule-Walker filter to create the correct power spectrum. Following this the component is normalised using the long term CDF of the scintillation variance, typically from the model of van de Kamp [1999].

3.3 N-state Markov chain models - L. Castanet, T. Deloues and J. Lemorten, Office National d'Etudes et de Recherches Aérospatiales (ONERA)

A more complex time series generator from ONERA is the 'N-state Markov model' [Castanet et al., 2003]. The model consists of two main processes, a macroscopic model that generates the rain/no-rain part of the attenuation time series and a microscopic model that fills in the rain portion of the attenuation time series.

The macroscopic model is a two-state model, the two states being 'rain' and 'clear-sky'. Therefore, the probability vector is:

$$\pi_0(t) = \left(\begin{array}{cc} p_0 & p_1 \end{array} \right) \tag{3.1}$$

where p_1 is the probability of rain, and $p_0 = 1 - p_1$.

From this it is possible to say that:

$$\pi(t+1) = P \cdot \pi_0(t) \tag{3.2}$$

where P is the transition matrix for this two state model which is simply:

$$P = \begin{pmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{pmatrix} = \begin{pmatrix} 1 - p_{01} & p_{01} \\ p_{10} & 1 - p_{10} \end{pmatrix} \tag{3.3}$$

It can be shown that:

$$p_{01} = \frac{p_1}{1 - p_1} p_{10} \tag{3.4}$$

where p_{01} is the probability of a transition from a state of rain to no-rain.

Thus the macroscopic model is a model of two parameters, the probability of rain, p_1, and the probability of a change from rain to clear sky, p_{10}. The authors recommend that the value of p_1 can be extracted from ITU-R statistics whilst the value of p_{10} can be extracted from fade duration models (such as the model of Paraboni and Riva [1994]).

Once the macroscopic model has been used to identify each rain event the microscopic model then 'fills-in' these rain events with a suitable structure. The microscopic model follows an N-state Markov model, where N states are used to define the attenuation levels, it is important to understand that these attenuation levels need not be uniformly distributed.

For the microscopic model the probability vector is similarly defined as:

$$\pi_1 (t) = \begin{pmatrix} p_1 & p_2 & \ldots & p_N \end{pmatrix} \tag{3.5}$$

The authors recommend that the initial probability vector is $\pi_1 (t) = \begin{pmatrix} 1 & 0 & \ldots & 0 \end{pmatrix}$.

In order to proceed to the next time step the transition matrix is required, this fully maps the transition probabilities from any state, i, in $\pi_1 (t)$ to any state, j, in $\pi_1 (t + \Delta t)$. This is defined as a matrix thus:

$$P = \begin{pmatrix} p_{11} & p_{12} & \cdots & & \cdots & p_{1N} \\ p_{21} & \ddots & & & & \vdots \\ \vdots & & p_{ij} & & & \vdots \\ \vdots & & & & \ddots & p_{N-1N} \\ p_{N1} & \cdots & & \cdots & p_{NN-1} & p_{NN} \end{pmatrix} \tag{3.6}$$

The method that the authors recommend for defining this model is to use a fade slope model. Fade slope models define the probability of a particular fade slope ζ given a particular attenuation level. If the states are defined at attenuation levels s_1 to s_N, it

follows that the transition matrix can be defined by:

$$p_{ij} = p\left(\zeta = \frac{s_j - s_i}{\delta t} \mid s(t) = s_j\right) \tag{3.7}$$

In order to verify the results a set of synthetic events were created and the cumulative density function of the results compared with those from the OPEX campaign.

This approach has several inherent problems. The first concerns are connected with the macroscopic model. The ITU-R models provide annual statistics of rain occurrence and hence shorter term periods demonstrate a considerable variation about these points. The Markov chain also relies on each event being independent, which is inherently not the case. Heuristically we understand that the meteorology 6 hours from now is not independent of the meteorology now.

The greatest concern is the procedure employed to 'fill-in' the macroscopic model. The microscopic and macroscopic models are independent, this does not allow the microscopic model any information regarding the length of the event. This results in the microscopic model producing rain events that individually do not show an expected time profile, although they demonstrate the correct first order statistics over a long enough period.

A final concern is the validation of the time series generator was performed by comparing the first order statistics of the generated time series with those of the OPEX campaign. However it should be noted that the two models that form the time series generator namely the fade duration model of Paraboni and Riva [1994] and the fade slope model of van de Kamp [2003], were also derived using these measurements.

3.4 Markov chain models - U.-C. Fiebig, Deutschen Zentrum für Luft- und Raumfahrt (DLR)

This work is performed primarily by Uwe-Carsten Fiebig at the Institute for Communications Technology at the German Aerospace Center (DLR). This model creates single site time series.

This model assumes that the received power is formed of three different types of 'short-time-segments':

1. Signal segments of essentially constant received power, denoted 'C-segments',

2. Signal segments of increasing received power, denoted 'U-segments',

3. Signal segments of decreasing received power, denoted 'D-segments'.

The assumption then is followed by assuming that the attenuation level at a given instant depends on the attenuation level some time, $\Delta\tau$, before and on the type of signal segment.

By analysing ITALSAT data and following the above assumptions, the conditional probability $p_i(x(t+\Delta t) \mid x(t))$ can be calculated, where x is the attenuation time series and i is the type of segment (C, U or D). The resulting conditional probabilities were found to be approximately Gaussian in shape. The attenuation time series $x(t)$ are quantised at typically 1 dB steps up to 40 dB to provide the time series $x'(t)$. Using this typical quantisation environment the required model parameters are the mean and standard deviation of the conditional probabilities at each of the quantisation levels: $\mu_i(x')$ and $\sigma_i(x')$ where $i \in C, U, D$ and $x' \in 1, 2, 3, \ldots, 40$. These parameters are evaluated from beacon measurements. After the model parameters are evaluated the time series can be created following the procedure in figure 3.4.
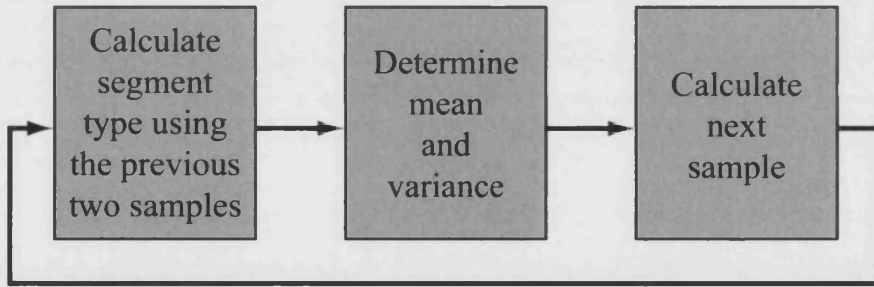
Figure 3.4: The DLR Markov chain model.

An output $y(j)$ at a time sample time j is created as follows. The first stage is to assess what 'segment' type the current output sample is in, this is done by examining the last two samples:

$$y(j-2) = y(j-1) \qquad \Rightarrow \qquad i = C \qquad (3.8)$$

$$y(j-2) > y(j-1) \qquad \Rightarrow \qquad i = D \qquad (3.9)$$

$$y(j-2) < y(j-1) \qquad \Rightarrow \qquad i = U \qquad (3.10)$$

The characteristic values, $\mu_i(y(j-1))$ and $\sigma_i(y(j-1))$ of the Gaussian distribution $p_i(y(j) \mid y(j-1))$, are then evaluated. The value $y(j)$ is then an independent experiment from this distribution.

43

An example of the model parameters for the 40 GHz ITALSAT beacon measurements from Oberpfaffenhofen (elevation angle of 34.8°), were evaluated [Fiebig, 1999]. The time between samples was 64 seconds and a least mean squares fit of the beacon measurements provided the following model parameters:

$$\mu_C(x) = 0.12 + 0.98x - 0.001x^2 \qquad \sigma_C(x) = 0.16x - 0.0011x^2 \qquad (3.11)$$

$$\mu_D(x) = 0.12 + 0.89x - 0.00075x^2 \qquad \sigma_D(x) = 0.2 + 0.33x - 0.005x^2 \qquad (3.12)$$

$$\mu_U(x) = 0.12 + 1.1x - 0.0025x^2 \qquad \sigma_U(x) = 0.1 + 0.21x - 0.0025x^2 \qquad (3.13)$$

3.5 Synthetic storm technique - E. Matricciani, Politecnico di Milano

A similar technique to this synthetic storm technique was first proposed by Drufuca [1974] with the application in terrestrial links and has recently been re-engineered to become appropriate for earth-space links.

The technique takes a single-site rainfall rate time series, typically from rain gauge measurements, and converts them into a rain attenuation time series using a simple model of the vertical structure of precipitation.

The two layers of rain are shown in figure 3.5.
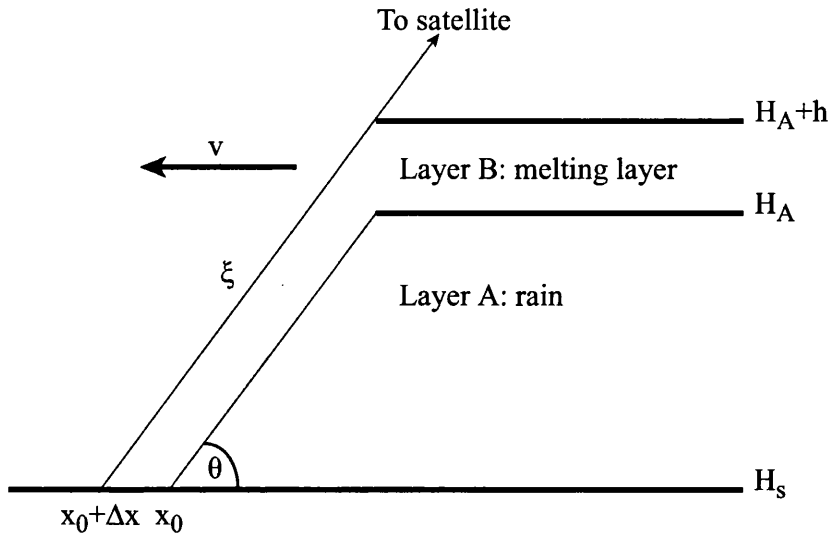
Figure 3.5: The modelling of precipitation in the synthetic storm technique.

Layer A consists of a uniform rain rate, R, as measured at the ground, whilst layer B

represents the melting layer with ice hydrometeors at the top which become raindrops at the top of layer A. Layer B is characterised as follows:

1. The mass of the hydrometeor is constant during melting and the density of the unmelted core is assumed to be $0.3\,\mathrm{g\,cm^{-1}}$.

2. There is a direct mapping between raindrops in layer A and melting hydrometeors in layer B.

3. The specific attenuation of layer B is calculated as if it were produced by raindrops with the same volume as the melting hydrometeors.

The 'apparent-rain-rate' in the melting layer is defined as a function of the rainfall rate in layer A so that $R_B = 3.134 R_A$ [Matricciani, 1996]. This melting layer model is only applicable to stratiform rain events, since the convective updrafts typically prohibits the formation of a well-defined melting layer. However, convective rain is typically formed at a higher altitude hence the 'apparent-rain-rate' creates a larger attenuation which qualitatively accounts for the increased rain height and the larger raindrop size distribution. The value of the zero-degree isotherm above sea-level $(H_A + h)$ is taken from ITU-R recommendations and the depth of the melting layer, h, is assumed to be constant at 400 m for all latitudes.

The specific attenuation, $\gamma(x)$, at a point x is calculated using:

$$\gamma(x) = kR(x)^\alpha \qquad (3.14)$$

Using the synthetic storm concepts the attenuation along the slant path of figure 3.5 becomes:

$$A(x_0) = k_A \int_0^{L_A} R(x_0 + \Delta x_0, \zeta)^{\alpha_A} d\zeta + k_B \int_{L_A}^{L_B} R(x_0, \zeta)^{\alpha_B} d\zeta \qquad (3.15)$$

where the path is given as:

$$L = \frac{H - H_s}{\sin(\theta)} \qquad (3.16)$$

and from simple geometry $\Delta x_0 = \frac{h}{\tan(\theta)}$ which is required since the path enters layer A at $x_0 + \Delta x_0$.

The Fourier transform of equation 3.15 results in the following equation within the domain of spatial frequency, $f_s = \frac{1}{x}$ in $\mathrm{km^{-1}}$:

$$S_A(f_s) = S_{\gamma,A}(f_s) L_A \mathrm{sinc}(f_s L_a) \exp(-j2\pi f_s \Delta x_0) + r^{\alpha_B} S_{\gamma,B}(f_s) \Delta L \mathrm{sinc}(f_s \Delta L) \qquad (3.17)$$

45

where :

$$S_A(f_s) = \mathcal{F}\{A(x_0)\} \tag{3.18}$$

$$S_\gamma(f_s) = \mathcal{F}\{kR(x)^\alpha\} \tag{3.19}$$

$$\text{sinc}(f_sL) = \frac{\sin(\pi f_s L)}{\pi f_s L} \tag{3.20}$$

The transform from spatial frequency to a temporal frequency, assuming Taylor's frozen field hypothesis, can be shown to be [Matricciani, 1996]:

$$f = \frac{v f_s}{\cos(\theta)} \tag{3.21}$$

and the slant-path speed is

$$v(\theta) = \frac{v}{\cos(\theta)} \tag{3.22}$$

Using this domain transform equation 3.17 represents the temporal frequency thus:

$$S'_A(f) = S'_{\gamma,A}(f) L_A \text{sinc}\left(\frac{f L_A}{v(\theta)}\right) \exp\left(\frac{-j2\pi f_s \Delta x_0}{v(\theta)}\right) +$$

$$r^{\alpha_B} S'_{\gamma,B}(f) \Delta L \text{sinc}\left(\frac{f \Delta L}{v(\theta)}\right) \tag{3.23}$$

where :

$$S'_\gamma(f) = v S_\gamma(v f_s) \tag{3.24}$$

and hence the attenuation time series can be considered to be:

$$A(t) = \mathcal{F}^{-1}(S'_A(f)) \tag{3.25}$$

3.6 Two sample model - M.M.J.L. van de Kamp, European Space Agency (ESA)

This model can be used for two purposes: the first is to generate attenuation time series for simulation purposes. The second implementation of this model is that it can be used to provide very short term ($\sim$ 10 seconds) propagation channel forecasts.

This model extends the hypothesis that the near-future attenuation level is a function of the previous samples [Grémont, 1997, Dossi, 1990]. This model aims to predict the probability distribution of the attenuation at time $t + \Delta t$, given the previous two samples and the value of Δt i.e. $p(A) = p(A \mid A_0, A_{-1}, \Delta t)$ where $A = A(t + \Delta t)$, $A_0 = A(t)$, $A_{-1} = A(t - \Delta t)$.

In order to evaluate the probability distribution several conditional statistical moments were extracted from beacon time series. The conditional standard deviation of the attenuation, $\sigma_A(A_0, A_{-1}, \Delta t)$was calculated and was seen to exhibit several important characteristics. The first characteristic is that if $A_0 = A_{-1}$ then the value of the standard deviation σ_A is approximately linear and proportional to Δt [van de Kamp, 2002a] these leads to the equation:

$$\sigma_A(A_0, \Delta t \mid A_{-1} = A_0) = A_0 \sqrt{\beta_2 \Delta t} \qquad (3.26)$$

For the case when $A_0 \neq A_{-1}$ a larger standard deviation is observed. This increase was found to demonstrate a linear dependence with the value $|A_0 - A_{-1}|$. Thus the conditional distribution of the value σ_A can be written as:

$$\begin{aligned}
\sigma_A(A_0, A_{-1}, \Delta t) &= \sigma_A(A_0, \Delta t \mid A_{-1} = A_0) + \gamma_2 |A_0 - A_{-1}| \\
&= A_0 \sqrt{\beta_2 \Delta t} + \gamma_2 |A_0 - A_{-1}|
\end{aligned} \qquad (3.27)$$

The author evaluated the constants from the 20 GHz beacon measurements at Eindhoven as $\beta_2 = 0.00034$ and $\gamma_2 = 0.11$ which results in the standard deviation characteristics shown in figure 3.6

In order to evaluate the mean of the probability distribution of the next sample, m_A, a similar process was followed as for that of the standard deviation. However for this statistical moment the value of m_A is approximately equal to A_0 plus a linear term which is a linear function of the difference between A_0 and A_{-1} thus:

$$m_A = A_0 + \alpha_2 (A_0 - A_{-1}) \qquad (3.28)$$

If this equation is evaluated with the constant α_2 as discovered for the Eindhoven link previously discussed the mean of the probability distribution of the next sample is as shown in figure 3.7.

The first two statistical moments have now been evaluated it is important to investigate the higher moments. The author found that the skew was approximately zero and the kurtosis was of the order of 5 [van de Kamp, 2002b]. This distribution is well described by the hyperbolic secant distribution. Hence the probability distribution of A can be described as:

$$p(A) = \frac{1}{2\sigma_A} \operatorname{sech}\left\{\frac{\pi(A - m_A)}{2\sigma_A}\right\} \qquad (3.29)$$
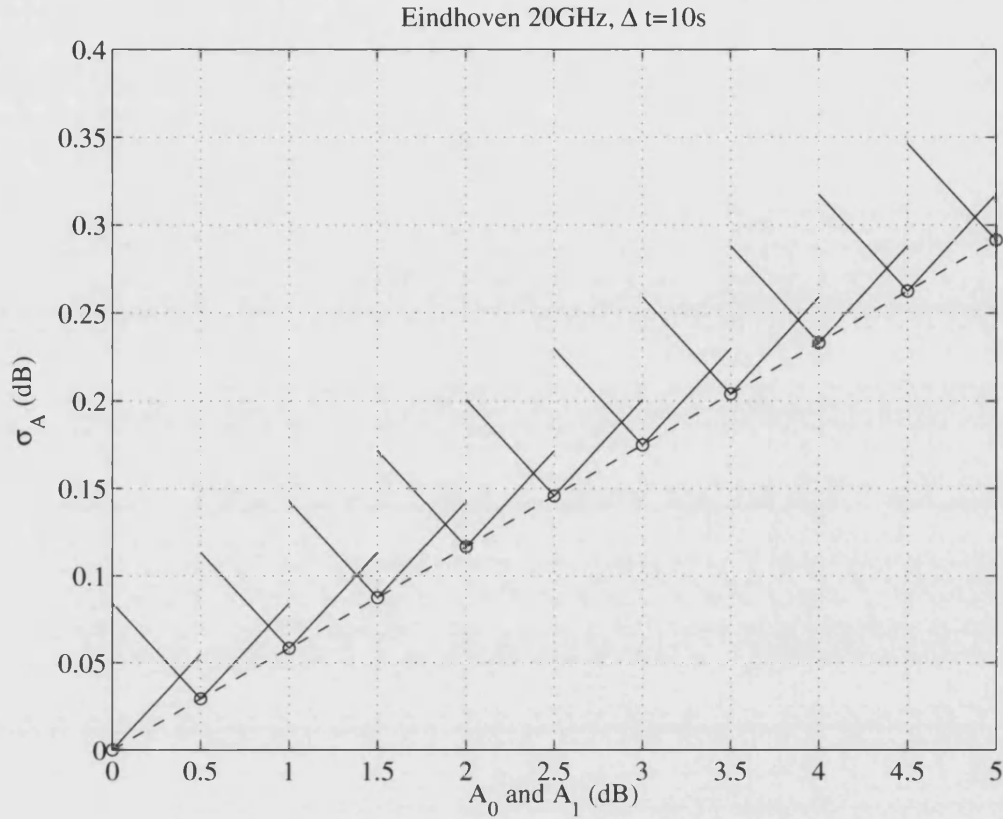
Figure 3.6: The standard deviation as a function of A_0 with $A_{-1} = A_0$ (dotted line) and as a function of A_{-1} with $A_0 = 0, 0.5, 1, \cdots, 5\,\text{dB}$ (solid line) with A_0 denoted as a circle.

3.7 A Spatio-temporal model - B. Grémont and M. Filip, University of Portsmouth

The model demonstrated by Grémont and Filip [2004] is a complex stochastic model, aimed primarily at site diversity experiments. The model assumes that point rainfall is distributed as a log-normal variable and hence the distribution of rainfall rate at two points on the horizontal plane is jointly log normal with a probability density function (PDF) as in equation 3.30.

$$f(R_1, R_2, r) = \frac{1}{2\pi\sigma_1\sigma_2 R_1 R_2 \sqrt{1-r^2}} \cdot$$
$$\exp\left[-\frac{1}{2(1-r^2)}\left(\frac{(\ln R_1 - m_1)^2}{\sigma_1^2} - \frac{2r(\ln R_1 - m_1)(\ln R_2 - m_2)}{\sigma_1\sigma_2} + \frac{(\ln R_2 - m_2)^2}{\sigma_2^2}\right)\right]$$

$$(3.30)$$
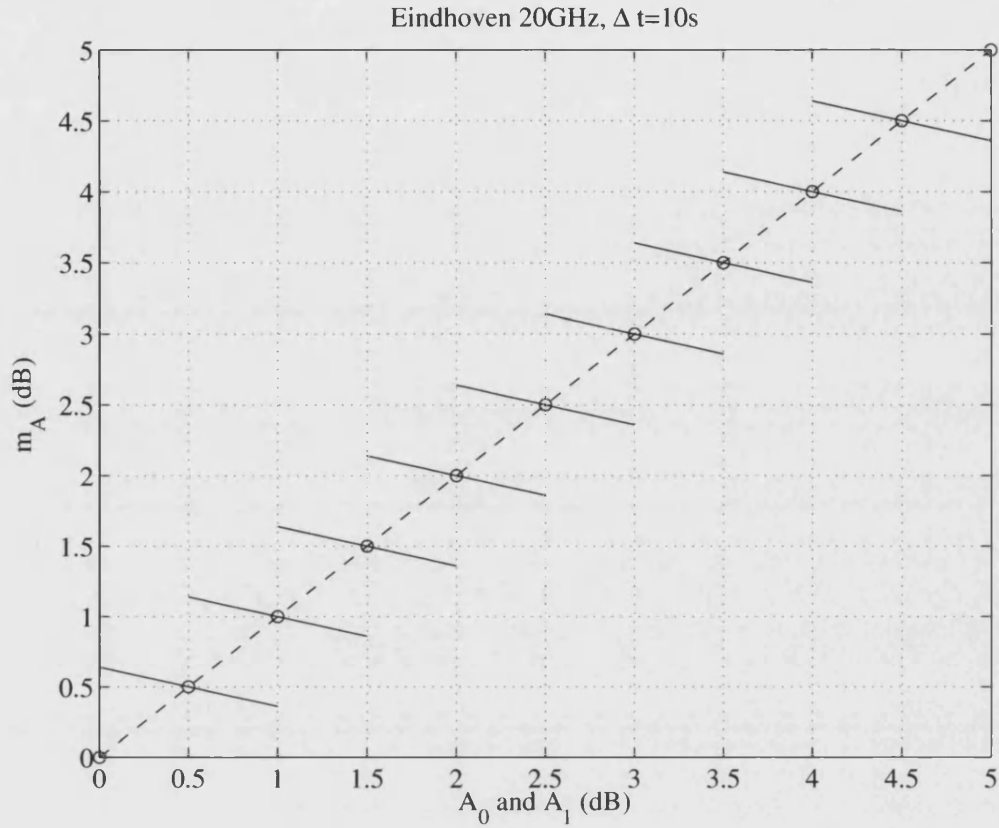
Figure 3.7: The mean as a function of A_0 with $A_{-1} = A_0$ (dotted line) and as a function of A_{-1} with $A_0 = 0, 0.5, 1, \cdots, 5\,\mathrm{dB}$ (solid line) with A_0 denoted as a circle.

or

$$f(R_1, R_2, r) = \Lambda_{R_1, R_2}(m_1, \sigma_1, r, m_2, \sigma_2) \qquad (3.31)$$

where r is the cross-correlation between the two locations and R_1 and R_2 represent the point rainfall rates at locations 1 and 2 respectively. The values of σ_i and m_i represent the log-normal statistics of the i-th location.

It can also be shown that the correlation of the two point rainfall rates is:

$$r = \frac{e^{r'\sigma_1\sigma_2} - 1}{\sqrt{e^{\sigma_1^2} - 1}\sqrt{e^{\sigma_2^2} - 1}} \; ; \; r' \in [0, 1] \qquad (3.32)$$

where r' is the cross-correlation of the reduced variable:

$$R_i' = \frac{\ln R_i - m_i}{\sigma_i} \qquad (3.33)$$

By assuming the rainfall field cross-correlation can be described as an exponential

49

function it can be written that:

$$r' = \exp\left(-\frac{\rho}{L} - \beta_2 |t|\right) \qquad (3.34)$$

where ρ represents the separation distance between locations, L is the characteristic size of the rain cells, and β_2 accounts for the birth and decay of rain cells.

Taylor's frozen field hypothesis [Taylor, 1938] suggests that if the velocity of the air stream that carries turbulence is much greater than the turbulent velocities, it is possible to assume that the variation at a fixed point is due to the passage of an unchanging pattern of turbulent motion over that point. In this example this implies that $\rho = \overline{V} |t|$, where $\overline{V}$ represents the advection velocity of the rainfall field.

This allows equation 3.34 to be simplified to:

$$r' = \exp\left(-\beta |t|\right) \qquad (3.35)$$

where $\beta = \frac{\overline{V}}{L} + \beta_2$ and is called a 'compound time constant' representing the spatial properties and birth and decay of the rainfall cells. Assuming a mean storm velocity of the order of $10\,\mathrm{m\,s^{-1}}$ and a mean rain cell diameter of the order of $5\,\mathrm{km}$ and little birth or decay (i.e. $\beta_2 \to 0$) β is of the order of $0.002\ \mathrm{s^{-1}}$.

Using Taylor's frozen field hypothesis allows r' and hence r to be calculated spatially or temporally. This spatio-temporal analogy comes from an Eulerian view-point where the observer is fixed in space whilst the rainfall field passes over. The time pattern of the rainfall is a combination of the spatial structure of the rainfall (i.e. size, development, decay, micro-structure) and the velocity of the storm movement. In simple terms it implies it is impossible to distinguish between a small cell moving slowly or a large cell moving very quickly. This creates what the authors call a 'spatio-temporal duality', essentially equation 3.30 can be evaluated either spatially as $f(R_1, R_2, r(\rho))$ or temporally as $f(R_1, R_2, r(t))$.

It can be seen that the horizontal projection of a slant path, D_i, is

$$D_i = \frac{H_i - h_i}{\tan \theta_i} \qquad (3.36)$$

where H_i is the rain height at location i, h_i is the altitude of point i and θ_i is elevation angle of the slant path from location i. Assuming the rainfall is homogeneous along the slant path the specific attenuation along that path, Γ_i, can be expressed as a simple power law relationship:

$$\Gamma_i = a_i R_i^{b_i} \qquad (3.37)$$

From this it is apparent that equation 3.30 can be extended to give the joint PDFs of attenuation, of the form:

$$f(A_1, A_2, r_A) = \Lambda_{A_1, A_2}(M_1, \Sigma_1, r_A, M_2, \Sigma_2) \tag{3.38}$$

where,

$$M_1 = b_1 m_1 + \ln a_1 + \ln D_1 \; ; \; \Sigma_1 = b_1 \sigma_1 \tag{3.39}$$

$$M_2 = b_2 m_2 + \ln a_2 + \ln D_2 \; ; \; \Sigma_2 = b_2 \sigma_2 \tag{3.40}$$

$$r_A = \frac{\exp(r' \Sigma_1 \Sigma_2) - 1}{\sqrt{\exp(\Sigma_1^2) - 1}\sqrt{\exp(\Sigma_2^2) - 1}} \tag{3.41}$$

In conclusion the model relies on a particular exponential assumption for the space-time cross-correlation function, and that Taylor's hypothesis is valid. The model requires six parameters; m_1, m_2, σ_1 and σ_2 which characterise the log-normal distributions at the two locations. The space-time properties are defined by the remaining two parameters; β is the spatio-temporal parameter and $\overline{V}$ represents the velocity of rain cells.

3.8 University of Bath Propagation Forecast Engine

The approach presented in this Thesis differs significantly from the typical techniques for propagation modelling. The techniques mentioned use data from propagation experiments in order to generate statistics regarding the channels characteristics (fade slope, fade level distribution etc.). The models then attempt to recreate these characteristics in a synthetic time series.

The approach presented here attempts to recreate the meteorological environment, and from this, the propagation channel. This process was first considered by Page et al. [2005] who outlined the possibility of using meteorological data to estimate the attenuation on a EHF/SHF communication link. This work proved that the approach worked on several example days, however a slightly different approach was required to be able to process many days of example time series over large areas. The approach presented here also differs in it's approach to scintillation modelling since the model used by Page et al. [2005] uses the Warnock model [Warnock et al., 1985] to generate a mean C_n^2 profile which is then used to generate the scintillation variance, since the initial paper was written better approaches to the modelling of scintillation have been proposed. Another difference is the approach to sub-grid-resolution rainfall, the grid-scale rainfall was supplemented by a filtered white noise source to simulate the effects of small-scale structure within the rainfall. A more physical approach to this phenomena,

a procedure known as downscaling, has been considered in this Thesis.

The use of the meteorological environment in modelling link attenuation has several advantages over 'statistical propagation modelling'.

The first advantage is that the system is not fixed to a single link; even large scale propagation experiments such as Olympus or planned Stentor satellite experiments have a small number of links in similar climates with similar elevation angles etc. This results in techniques that are not able to model large variations in link characteristics.

The beacon time series also display several characteristics due to the nature of their collection. Typically the receiver dynamic range is limited, for propagation experiments this can be relatively high when compared to typically satellite receivers however it will still be limited by the component technology and design issues. This has the effect of the receiver losing phase lock and not being able to track the faded signal. In effect this results in time series that do not exhibit large attenuation levels. Since the receiver can not regain lock in zero-time the return from large fades is often not well recorded. This is confirmed by the 'semi-physical' SST model which displays some differences from measured cumulative density functions, especially at low exceedances ($< 0.1\%$) [Matricciani, 2004].

Our approach neatly avoids this problem since the meteorological data is, relatively, easy to collect. This allows the system to be deployed in any location in the world (or indeed off it !). Numerical Weather Prediction (NWP) techniques are becoming increasingly elegant and sophisticated and able to model very subtle meteorological phenomena. Meteorological radars are also becoming cheaper and more capable and with the recent American adoption of next generation Doppler radars these very sophisticated radars are dramatically dropping in price.

Of course, the use of meteorological estimators such as NWP or radars introduce errors inherent to these techniques. However, the errors associated with these techniques typically are well understood.

The second advantage is that the system is able to provide time-coincident multiple site time series. In essence this means that the model is able to provide time series for any number of spatially unique locations, and these time series display the correct spatial correlation. This is very important for network-centric simulations. Network-centric simulations are especially critical in systems that have dynamic link budgeting (such as techniques using FMTs). In these networks there will be the dynamic reconfiguration of any number of links. Thus the network load will fluctuate whilst the control centre is attempting to deal with dropping some users and initialising other users. It is essential

that robust protocol and network management techniques are in place to be able to manage the system.

However, the greatest advantage of using meteorological data to estimate the radiowave propagation condition is a forecast ability. If forecast data is used instead of archived historical data it is possible to create radiowave propagation forecasts. This flexibility allows long-term forecasts (or link-timetabling) and will be a significant ability for several key types of network.

Chapter 4

The University of Bath
Propagation Forecast Engine

As discussed previously the propagation model developed in this section adopts a different approach to that of other investigators. In essence, estimates of the meteorological conditions are used to recreate the propagation conditions. This has several distinct advantages:

- The spatial variation of the propagation condition is inherited from the spatial variation of the meteorological conditions, this allows any number of spatially-consistent time series to be created.

- The system can be easily deployed over several different coverage areas / climates without the need for expensive propagation experiments.

- The system is not biased by inaccuracies in beacon measurements or errors in empirical models.

As part of this study meteorological estimates are taken from two sources: NWP systems and weather radar.

4.1   Numerical Weather Prediction Systems

Numerical weather prediction (NWP) dates back to 1904 when Norwegian scientist Vilhem Bjerknes postulated that since air is a fluid it obeys the fundamental physical

54

laws of fluids and hence the weather could be quantitatively predicted by applying these hydrodynamic and thermodynamic equations to an initial atmosphere state. However this hypothesis could not be tested until World War 2 when neutral Norway was cut-off from weather information. In response to this Norway established the first dense network of weather stations and was able to provide highly accurate, highly dense initial conditions.

After the end of World War 2 a meteorologist named Jule Charney began working at the Princeton Institute for Advanced Study with John von Neumann. Using one of the first computers ENIAC (Electronic Numerical Integrator and Computer)[1], shown in figure 4.1, the first successful numerical weather prediction was produced in April 1950.
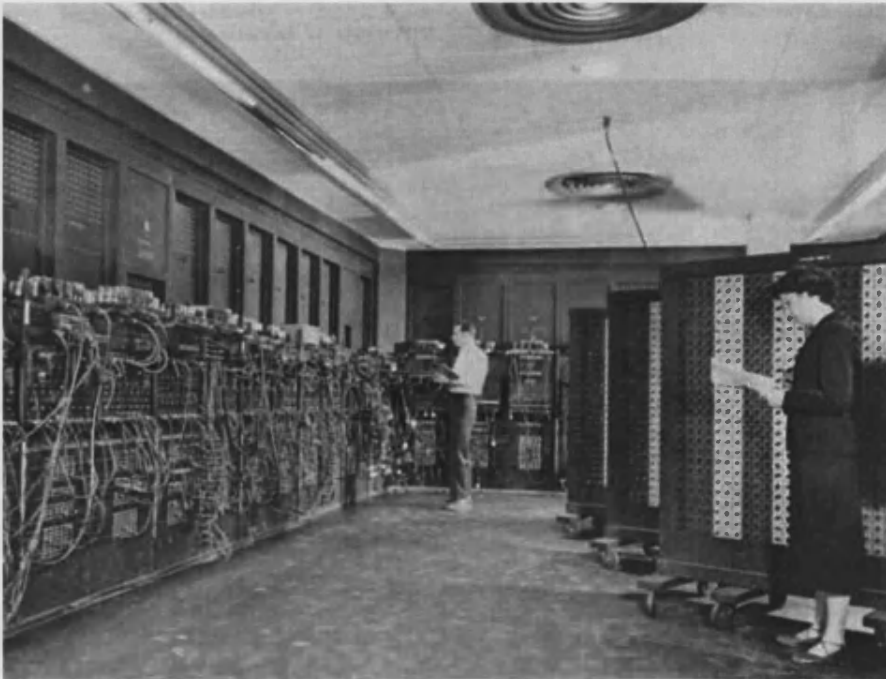
Figure 4.1: The ENIAC computer.

As NWP techniques improve, model resolutions decrease and coverage areas increase and the computational requirement increases. Indeed, the UK Met Office currently has several NEC supercomputers the most powerful of which is currently a 120 processor NEC SX-8 capable of 927.6 GFlops, estimated 960 GFlops peak (see figure 4.2).

There are many different mesoscale NWP models, with the UK Met Office (Unified

---

[1] ENIAC had 20 bits of memory, and could perform an addition in $200\,\mu s$, a multiplication in 3 ms and a division in 30 ms.

55

Figure 4.2: The UK Met Office NEC SX-8 supercomputer.

Model) and the American groups (MM5 and WRF) widely regarded as having the most sophisticated techniques.

The output of NWP models are a set of prognostic variables, typically temperature, pressure, humidity, u and v winds and geopotential height (see appendix E) and a set of diagnostic variables, such as rain accumulation.

4.1.1 Pennsylvania State University and National Centre for Atmospheric Research mesoscale model - MM5

The MM5 model is the fifth generation mesoscale model from NCAR and Penn State, this model is run as several distinct sub-models which makes it *relatively* easy to understand, install/compile and maintain.

A general execution procedure of the MM5 model is outlined in figure 4.3. Each of the boxes represent a separate sub-program. However, on the downside this system is harder to run as an automated un-supervised system (this is still possible, see later chapters).

The first part of the MM5 model is the 'terrain' module; this is compulsory and initialises the constant fields for the simulation. The most important role of this module is to initialise the domains in terms of coverage area and grid size. At this stage the grid is completely defined in terms of primary parameters (it is possible to analytically move from model co-ordinates to 'real-world' co-ordinate systems such as Lat-Lon). A
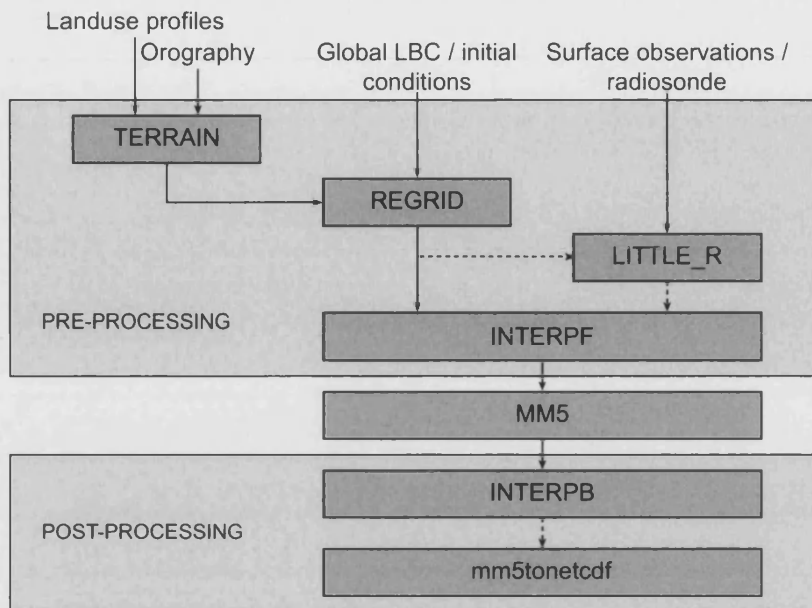
Figure 4.3: A general MM5 execution procedure.

very important part of this model is the ability to nest small-area low grid-size domains within larger coverage higher grid size domains, this allows for long robust simulation on small grid scales. There are some restrictions for the domains but in essence the domains can be deployed over any locations on the globe[2].

The other important role of the 'terrain' module is to create the orography and land use fields that are required within the MM5 simulation engine. Typically this information is extracted from USGS data at a spatial resolution of 30 second. The orography has a very important role in NWP systems especially within cloud and precipitation techniques as well as surface wind vectors. Land use has a more subtle effect within NWP techniques; the MM5 system models five characteristics that vary depending on the land use: albedo, moisture availability, emissivity, roughness and thermal inertia. The USGS splits land-use into 24 different categories; those that occur over the mainland UK these are shown in table 4.1.

The next sub-program that forms the MM5 model is the 'REGRID' module. This module creates the 'first-guess' lateral boundary conditions (LBC) and initial conditions. Since the MM5 is a regional mesoscale model there is a requirement for the lateral boundary conditions, these define the state around the borders of the domain throughout the simulation period. The initial conditions define the state at the start

---

[2]There is also the ability to change the map projection to allow for an even coverage of a simulation domain at various latitudes

of the simulation period.

Regrid is typically supplied with outputs from global NWP models; the data must contain several compulsory fields:

- air temperature (K),

- grid relative u-component of the horizontal wind $(\mathrm{m\,s^{-1}})$,

- grid relative v-component of the horizontal wind $(\mathrm{m\,s^{-1}})$,

- water vapour mix ratio,

- geopotential height (m),

- sea-level pressure (Pa),

- sea-surface temperature / skin temperature (K).

The regrid system can work with any input data in grib format although the system most commonly uses NCEP GDAS (Global Data Assimilation System) analysis. These analyses are archived every 3 hours, on a 2.5° by 2.5° grid from the mid 1970s. This data is freely available from NCAR (dataset ds082.0 and ds083.0).

Once the first-guess has been created the next task is to assimilate radiosonde measurements and surface observations in order to provide a more accurate estimation of the boundary conditions (lateral and initial). The subprogram 'little_r' performs this operation using a banana-Cressman weighting scheme. In the standard Cressman scheme each observation is assigned a radius of influence, each grid point in the first guess field is then adjusted by taking into account all observations which influence that point. The differences between the first-guess field and the observations are calculated and a distance weighted average of these difference values is added to the value of the first guess field. Once all the grid points in the first-guess field have been adjusted, the adjusted field becomes a new first-guess for the next iteration of the objective analysis; each iteration uses a smaller radius of influence. Using a 'banana' scheme implies that analyses of wind and humidity upon pressure levels have the circles of influence from the standard Cressman scheme elongated in the direction of the flow and curved around streamlines, this approximates the standard scheme in low-wind conditions.

The input format for the little_r sub-program is a simple format ASCII file, this makes it very easy to incorporate measurements from many different sources. However, there is a large global database of high resolution radiosonde and surface measurements available

at NCAR. The outputs from little_r are then converted from the vertical pressure levels grid to simulation grid levels.

This completes the MM5 pre-processing tasks, the fields that are then submitted to the MM5 simulation engine are the outputs from terrain and the LBC and initial conditions from little_r.

The MM5 model is a non-hydrostatic model; this means that it is possible to directly resolve vertical motion and hence convective cloud and rain (assuming that the grid scale and 'physics' options are suitable). It is possible for the model to be compiled and executed on Unix workstations and small clusters, this makes it very cost-effective to deploy. The system has been designed primarily for research purposes and hence there is no 'operational' model and there is a large degree of flexibility over the model, there are many different physics options each appropriate for different grid sizes.

The system can also be configured to use a four-dimensional data assimilation (4DDA) technique, this allows the assimilation of measurements into the model run. This significantly improves the accuracy of the model during analysis periods.

After the simulation has been completed the output files can be interpolated back from model based sigma levels to the pressure levels, and the output is normally converted to the flexible netcdf standard format.

Real time execution

The MM5 model can also be run as a 'real-time' automatic forecast model. Running a MM5 real-time forecast is a reality even on a low computation budget. At the University of Bath a 72 hour forecast system has been demonstrated. The system runs over two domains covering the south of the UK, one with a grid size of 18 km and a nested grid of 6 km grid size. The model was run on a 6 node cluster (each node is a 2.8 GHz Pentium 4 with 512 MB RAM) and the model execution time was of the order of 90 minutes. A brief review of the accuracy of this technique is included in appendix D.

4.1.2 UK Met Office - Unified Model

The UK Met Office Unified Model (UM) became operational in 1997 and the operational limited area mesoscale model has a grid size of 0.11° which is approximately 12 km. In

comparison to the MM5 model the UM is hydrostatic (the model assumes there is no vertical movement in the atmosphere, see appendix C) this means the system is unable to directly resolve convective events. However, in order to help model the transfer of heat, moisture and momentum in the atmosphere there is a parameterisation scheme. The parameterisation scheme can be broken down to four components:

1. Triggering, determining whether convection will occur in a grid area.

2. Cloud base closure, determining how much convection will occur. This is determined from the mass transported through the cloud base.

3. The transport model determines the changes to the model temperature, moisture and wind fields due to convection and the precipitation associated with the convection.

4. The convective cloud scheme calculates the amount of cloud associated with the convection and this is passed to the radiation schemes.

The UM historical data is archived at the British Atmospheric Data Centre (BADC), most of the information is archived at 6 hourly intervals. However there are several periods where the data is archived at hourly intervals.

4.1.3   Estimates of rainfall rate from NWP systems

As discussed in section 2, the attenuation due to rain creates the largest fades for EHF and SHF based systems. Hence, it is very important that the estimates of rainfall rate are accurate. The actual rainfall rate, in $\mathrm{mm\,hr^{-1}}$, can be thought of as a continuous function in x,y and t: $r(x,y,t)$. This is related to an accumulation over the simulation time thus:

$$R\left(x,y,t\right) = \int_0^t r(x,y,t)\,dt \qquad (4.1)$$

Within an NWP system the rainfall field, R', is modelled as an accumulation on a discrete grid of size δ, where time is discretized as $\tau \in \{0, \Delta t, 2\Delta t, \ldots, T\}$ thus:

$$R'\left(i,j,\tau\right) = A \int_{x_0+i\cdot\delta}^{x_0+(i+1)\cdot\delta} \int_{y_0+j\cdot\delta}^{y_0+(j+1)\cdot\delta} \int_0^\tau r(x,y,t)\,dx\,dy\,dt \qquad (4.2)$$

where A is a constant relating to the grid size. As δ becomes small it can be seen that $R' \to R$, however for large values of δ it is apparent that spatial averaging of the rain accumulation will occur.

In order to estimate the instantaneous rainfall rate from the NWP modelled accumulation a simple finite difference can be performed:

$$r'(i,j,\tau) = B\left(R'(i,j,\tau+1) - R'(i,j,\tau)\right) \qquad (4.3)$$

where B is a constant relating to Δ. It is apparent that as Δ becomes small and δ becomes small then $r' \rightarrow r$. It is also worth noting that since rainfall is represented within the simulation environment as R', a small grid size not only reduces the spatial averaging but also improves the modelling of the phenomena within the NWP system which in turn improves the 'quality' of the analysis and forecasts.

4.2   Weather radar systems

In addition to NWP techniques weather radar systems can be used to improve the estimates of rain rate. The previous section explained the effect of the spatial averaging that occurs with NWP system. This combined with the relative low temporal resolution of the archived UM data and the hydrostatic nature of the UM means that adding radar estimates of rainfall rate is expected to improve the characterisation of the meteorological environment.

It is worth noting that radar measurements are also subject to errors and rely on several assumptions and approximations, for example:

- Hydro-meteors are an incoherent target, and the backscattered signal therefore fluctuates. Spatial and temporal averaging is applied but there remains some level of statistical uncertainty (It should be noted that the level of averaging is not on the same scale as NWP techniques).

- Assumptions about particle composition have to be made which can lead to errors where non-rain precipitation (*e.g.*, sleet, hail etc.) is concerned, or in areas such as the melting layer where there is significant mixed phase liquid.

- The reflectivity of the rain is dependent on the rainfall drop size, and the wide range of rainfall drop size distributions can create further uncertainty.

- Radars measure the reflectivity upon a polar co-ordinate system around the antenna and hence there are smoothing and averaging effects from the conversion to a Cartesian plane.

It is also worth noting that radar systems are an analytical measurement tool and have

limited application within a forecast paradigm. The forecast abilities of radar systems are typically limited to nowcasting techniques for short duration predictions of rainfall rate in the near future, typically exploiting motion tracking techniques.

4.2.1 The UK 'NIMROD' weather radar system

The 'NIMROD' system is based around 15 C-band radars, located in order to cover most of the UK (only the Shetland Islands and part of Orkney remain uncovered). All the radars have a wavelength of approximately 5.3 cm with a beam-width of approximately 1° and a minimum beam elevation of 0.5°.

The process followed by the 'NIMROD' system is as follows:

1. Each radar site produces rainfall data in terms of reflectivity from the backscattered rainfall data. These are at 5 minute intervals at four different elevations typically (0.5, 1.5, 2.5 and 4 degrees).

2. Permanent echoes are removed from radar reflectivity images using ground clutter maps.

3. The measured reflectivity is converted to precipitation rate.

4. The images are converted from polar to planar coordinates incorporating spatial averaging.

5. The surface rainfall rate is evaluated from the radar data using a physically based correction scheme.

6. On a weekly time scale radar data is compared with hourly rain gauges reports and corrections are made.

7. The single site data is then sent to the Met Office where it is compiled to a Cartesian plane.

This results in composites of estimations of rainfall rate every 15 minutes on a 5 km regular spatial grid.

A comparison of the rainfall rate estimates from the UM NWP technique and the 'nimrod' technique.

When comparing measurements of precipitation from different measurement techniques it is important to appreciate what the rainfall rate measurement physically represents.

Radar measurements of rainfall rate are derived from the average reflectivity within a voxel, this reflectivity is a function of the instantaneous rainfall rate. Hence radar measurements provide a measurement of the average instantaneous rainfall rate within a voxel, in the nomenclature previously used:

$$r'_{rad}(i,j,t) = A_{rad} \int_{x_0+i\cdot\delta_{rad}}^{x_0+(i+1)\cdot\delta_{rad}} \int_{y_0+j\cdot\delta_{rad}}^{y_0(j+1)\cdot\delta_{rad}} r(x,y,t)\, dx\, dy \qquad (4.4)$$

As previously discussed the rainfall rate output from NWP systems is a finite difference of an spatially and temporally averaged accumulation, of the form:

$$r'_{nwp}(i,j,\tau) = \left[A_{nwp} \int_{x_0+i\cdot\delta_{nwp}}^{x_0+(i+1)\cdot\delta_{nwp}} \int_{y_0+j\cdot\delta_{nwp}}^{y_0+(j+1)\cdot\delta_{nwp}} \int_0^{\tau+1} r(x,y,t)\, dx\, dy\, dt \right] -$$
$$\left[A_{nwp} \int_{x_0+i\cdot\delta_{nwp}}^{x_0+(i+1)\cdot\delta_{nwp}} \int_{y_0+j\cdot\delta_{nwp}}^{y_0+(j+1)\cdot\delta_{nwp}} \int_0^{\tau} r(x,y,t)\, dx\, dy\, dt \right] \qquad (4.5)$$

where, δ refers to the relevant grid scale and $\tau \in \{0, \Delta t, 2\Delta t, \ldots, T\}$.

As an aside the reader is warned that these approximations assume that the 'real' rainfall rate $r(x,y,t)$ is perfectly modelled; hence there should strictly be an error term ϑ in equations 4.2 - 4.5.

It is important to understand that when these rainfall rates are compared they are not exactly identical, but different estimation of the same phenomena. However there is still merit in comparing the two data sources.

The archived radar data for the three month period from April 2003 to June 2003 has been used. This data is available at the British Atmospheric Data Center (BADC). The rainfall rate time series for the University of Bath (51.3758°N, 2.3599°W) has been extracted; to reiterate this data is available at 15 minute intervals with a spatial resolution of 5 km. The radar derived rainfall rate time series is compared to the rainfall rate time series from the Operational Unified Model which was calculated from the sum of the convective and large-scale accumulations and is shown in figure 4.4.
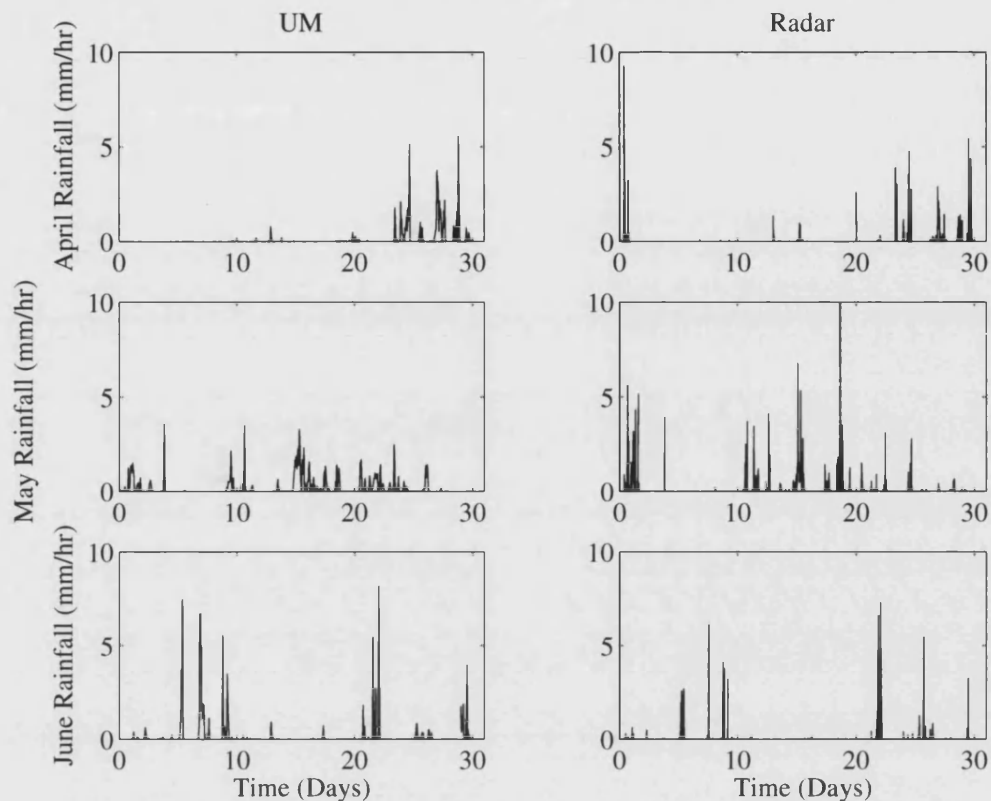
Figure 4.4: Rainfall rate time series for April, May and June 2004, from archived radar and UM measurements.

As can be seen there is some variation in the estimates from the different techniques. Both techniques concur broadly on a rain / no-rain basis however there is some variation in the magnitude; this is expected since the two time series have different time constants. It also worth noting that the radar time series are a time series of spatially averaged, but temporally instantaneous measures, whilst the time series from the UM measurements are temporally and spatially averaged and on a different grid size. Also of interest is the cumulative density function of the rainfall rate time series from the two techniques, this is shown in figure 4.5.

The cumulative distribution function of the rainfall rate time series also includes the prediction of the rainfall rate CDF from the ITU-R recommendation P.837-4. This has been derived from the statistics of the ECMWF forty year reanalysis. The two CDFs show good agreement with that provided by the recommendation. The radar CDF reaches a lower exceedance since the estimates are provided at 15 minute intervals whilst the NWP data is provided at hourly intervals. It is also important to remember that the smallest exceedances will be poorly sampled. Critically the ITU-R recommendation

64
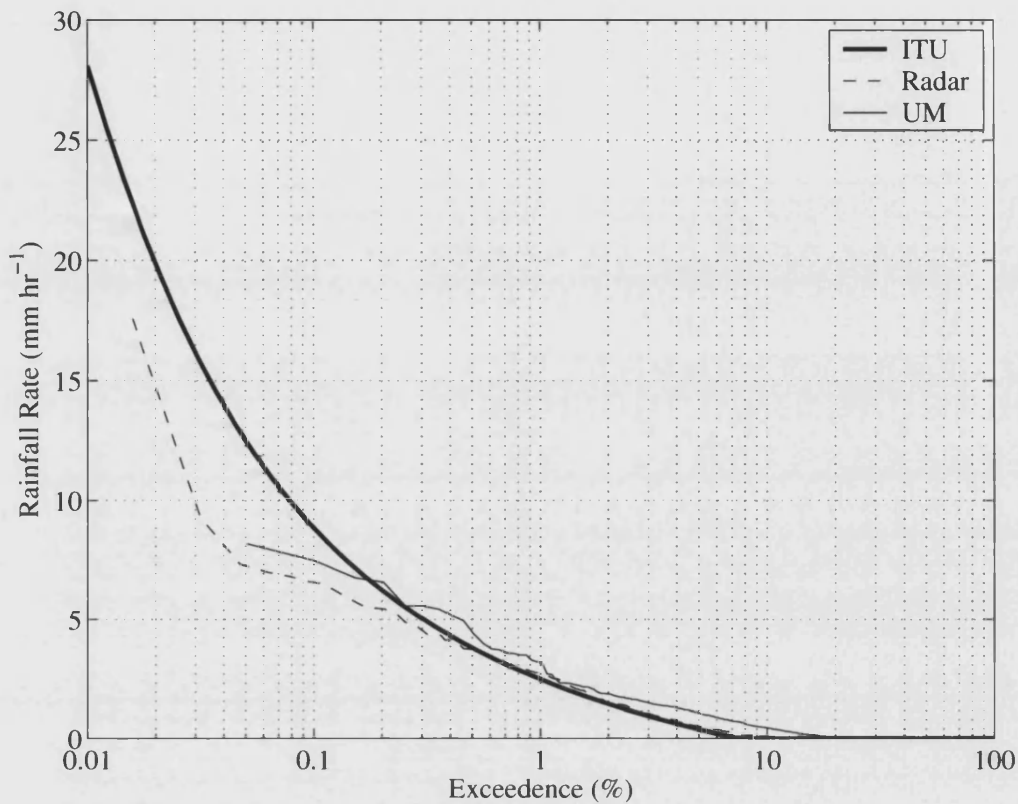
Figure 4.5: The CDF of the rainfall rate time series for April, May and June 2004, from archived radar and UM measurements.

provides an annual distribution whilst the time series from the radar and UM systems are for a three month period and hence will display seasonal variation.

The UM NWP system evaluates the occurrence of rain (i.e. a rainfall rate in excess of zero) more often than the radar estimates. This can be expected since the radar is an instantaneous measure of the rainfall rate provided every 15 minutes whilst the UM is a temporal average over an hour. The UM measure will include any rainfall during the integration period whilst the radar will only include rainfall occurring at the time of the scan.

Both techniques however demonstrate a similar distribution which implies that although some events are resolved differently the overall statistics of both data-sets are consistent.

4.3 The propagation forecast engine (PFE)

One of the goals in the design of the PFE is flexibility; this meant that the system was designed as a set of plug-ins which could be individually enabled and disabled easily allowing many different configurations. The ability to change the geometry of the problem is also very important, especially for creating time series for slant-path and terrestrial links.

Each of the plug-ins cover the modelling of one particular source of attenuation at EHF and SHF frequencies. The emphasis on the plug-ins has been to incorporate models with a physical background in preference to those empirically derived from beacon measurements or statistics of phenomena.

4.3.1 Cloud attenuation

As discussed previously the effects of clouds are rarely included in time series models since the attenuation due to rain causes the largest fades. However, since clouds typically occur more frequently than rain, for many climates the attenuation due to cloud can exceed the attenuation due to rain for much of the time [Konefal et al., 2000].

Cloud attenuation has been somewhat overlooked in the literature, however there has been some significant work on the modelling of cloud attenuation effects. For example, the work by Dissanayake et al. [2001] models the cloud attenuation probability density function using the average cloud occurrence rates with average cloud characteristics (such as average water contents, average height etc.). This approach works well for assessing the statistics of cloud attenuation, however it is not suitable for evaluating attenuation on an event-by-event basis.

An attempt to model cloud attenuation on an event-by-event basis was performed by Altshuler and Marr [1989], the cloud attenuation was estimated at two frequencies (15 GHz and 35 GHz) using a radiometer. The measurements were grouped into three different conditions, clear-sky, complete cloud cover (where the sun was not visible) and mixed clouds[3]. The two non-clear-sky events were then evaluated as a linear regression function of the surface absolute humidity, these results showed a considerable spread. This spread is expected since the zenith cloud liquid water can not be considered solely a function of the surface humidity. Similar work was documented by Dintelmann and Ortgies [1989] who used approximations relating the surface observa-

---

[3]In this case the term 'mixed' refers to partial cloud cover rather than mixed liquid phase

tions of relative humidity to vertical profiles of cloud liquid water which are valid for average conditions. The authors found that the model was representative over a year of data, however individual event-by-event cases were not modelled well.

The model that is used for the cloud attenuation modelling plug-in is that demonstrated by Salonen and Uppala [1991]. This model uses the vertical profiles of pressure, temperature and humidity in order to evaluate the cloud liquid water content and from this the attenuation due to cloud.

The first part of the process involves the detection of clouds. In order to investigate the presence of clouds, the critical humidity should be calculated thus:

$$U_c = 1 - \sigma (1 - \sigma) \left[1 + \sqrt{3} (\sigma - 0.5) \right]$$ (4.6)

where σ is the ratio of the pressure on the considered level and at the surface level. If the measured humidity is greater than U_c the level can be deemed to be a cloud layer.

Once a cloud layer has been detected it is possible to evaluate the liquid water content, w in $\mathrm{g\,m^{-3}}$, within that layer. The LWC, w, can be considered as a function of the vertical temperature profile and the height above the cloudbase, h_c:

$$w = w_0 (1 + ct) \left(\frac{h_c}{h_r} \right)^a p_w(t)$$ (4.7)

where $a = 0.14$ is height dependent coefficient, $c = 0.041\,^\circ\mathrm{C}^{-1}$ is a temperature dependent coefficient, h_r is the reference height typically taken as $1500\,\mathrm{m}$, $w_0 = 0.14\,\mathrm{g\,m^{-3}}$ is the liquid water content if $h_c = h_r$. The liquid water content fraction can be assumed to be approximated by a linear transition:

$$p_w(t) = \begin{cases} 1 & : & 0^\circ C < t \\ 1 + \frac{t}{20} & : & -20^\circ C < t < 0^\circ C \\ 0 & : & t < -20^\circ C \end{cases}$$ (4.8)

In liquid clouds attenuation due to absorption is the dominant loss. The Rayleigh approximation can be applied to compute the cloud specific attenuation. Whilst the Rayleigh approximation is valid the absorption coefficient is linearly proportional to the LWC and independent of the cloud drop size distribution [Li et al., 1997]. The Rayleigh approximation can be considered valid for:

$$|n| k_0 a \ll 1$$ (4.9)

where n is the refractive index of liquid water, k_0 is the free space wave number and a is the radius of the water droplet, at the upper limits of the EHF and SHF frequencies

the value $k_0 a$ only reaches values around 0.25.

The specific attenuation for cloud attenuation, γ_c, can be calculated from the MPM Liebe model [Liebe, 1989]:

$$\gamma_c = 0.182 f N_w''(f) \tag{4.10}$$

where f is the frequency and N_w'' is imaginary part of the complex refractivity. The complex refractivity can be calculated as:

$$N_w'' = \frac{9}{2} \frac{w}{\varepsilon'' (1 + \eta^2)} \tag{4.11}$$

$$\eta = \frac{2 + \varepsilon'}{\varepsilon''} \tag{4.12}$$

where w is the liquid water content and ε' and ε'' are the real and imaginary parts of the permittivity of water. The Debye relaxation model can be used to evaluate the permittivity of water for frequencies up to 1000 GHz and from -10°C to 30°C:

$$\varepsilon' = \varepsilon_2 + \frac{\varepsilon_0 - \varepsilon_1}{1 + \left(\frac{f}{f_p}\right)^2} + \frac{\varepsilon_1 - \varepsilon_2}{1 + \left(\frac{f}{f_s}\right)^2} \tag{4.13}$$

$$\varepsilon'' = \frac{f (\varepsilon_0 - \varepsilon_1)}{f_p \left[1 + \left(\frac{f}{f_p}\right)^2\right]} + \frac{f (\varepsilon_1 - \varepsilon_2)}{f_s \left[1 + \left(\frac{f}{f_s}\right)^2\right]} \tag{4.14}$$

where:

$$\varepsilon_0 = 77.66 + 103.3 \, (\theta - 1)$$
$$\varepsilon_1 = 5.48$$
$$\varepsilon_2 = 3.51$$
$$f_p = 20.09 - 142 \, (\theta - 1) + 294 \, (\theta - 1)^2$$
$$f_s = 590 - 1500 \, (\theta - 1)$$

and θ is the inverse temperature parameter $\theta = \frac{300}{T}$ where T is in Kelvin.

The attenuation from clouds can then be evaluated by integrating the specific attenuation along the path:

$$A_{\text{cloud}} = \int_0^L \gamma_c \, dP \tag{4.15}$$

where P represents the ray path of length, L through the modelled atmosphere[4].

As a brief example of this procedure, vertical meteorological profiles from the MM5

---

[4]NWP systems typically provide analysis to a level of approximately 100 mb. This is significantly far above the mixed phase transition which ends at -20°C

model were extracted for Bath on 23 October 2004, 22:00 (UT). The first task is to determine the critical humidity profile by implementing equation 4.6. Figure 4.6 shows the vertical profile of the critical humidity and the vertical profile of relative humidity.
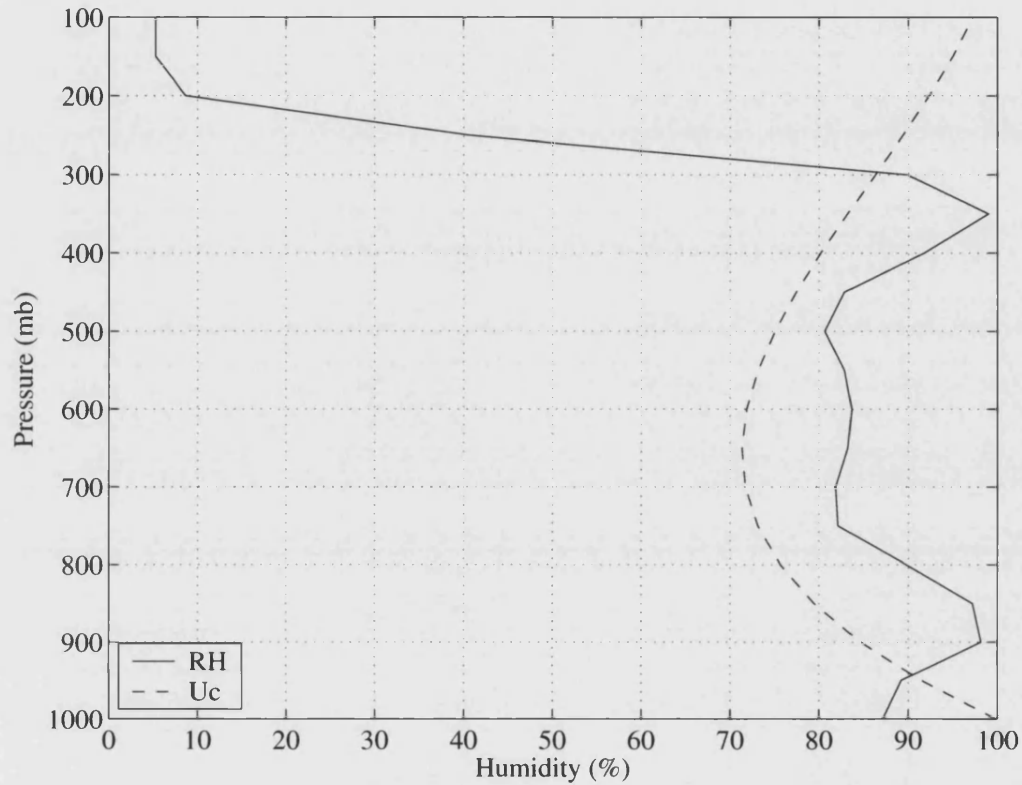
Figure 4.6: Vertical profiles of relative humidity and critical humidity, (23/10/2004 22:00UT).

From these profiles the presence of clouds can be detected when the relative humidity, RH, exceeds the critical humidity, U_c. The cloud liquid water fraction can then be determined using equation 4.8. Figure 4.7 shows the vertical profile of the cloud liquid fraction along with the presence of clouds.
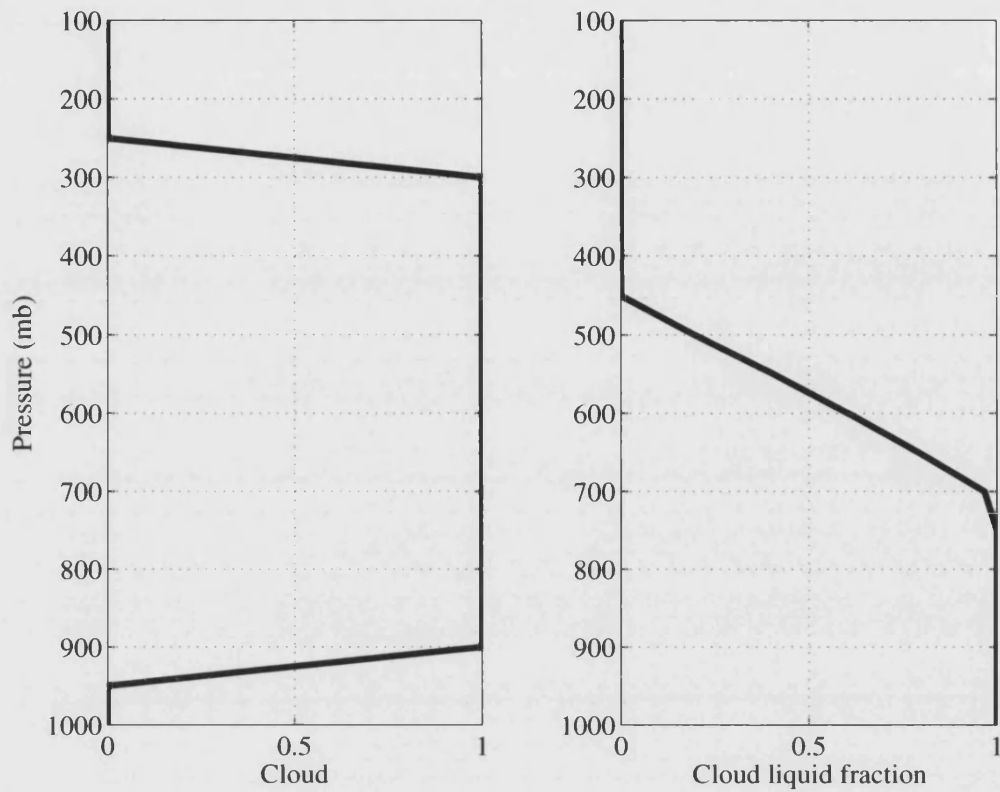
Figure 4.7: Cloud characteristic profiles, (23/10/2004 22:00UT).

Hence it can be seen that a liquid cloud layer exists approximately between the 900 mb and 450 mb layers, however a large portion of the upper layers of the cloud are cold cloud as evident by the decrease in cloud liquid fraction. Using equation 4.7 it is then possible to compute the vertical cloud liquid water content this is shown in figure 4.8.
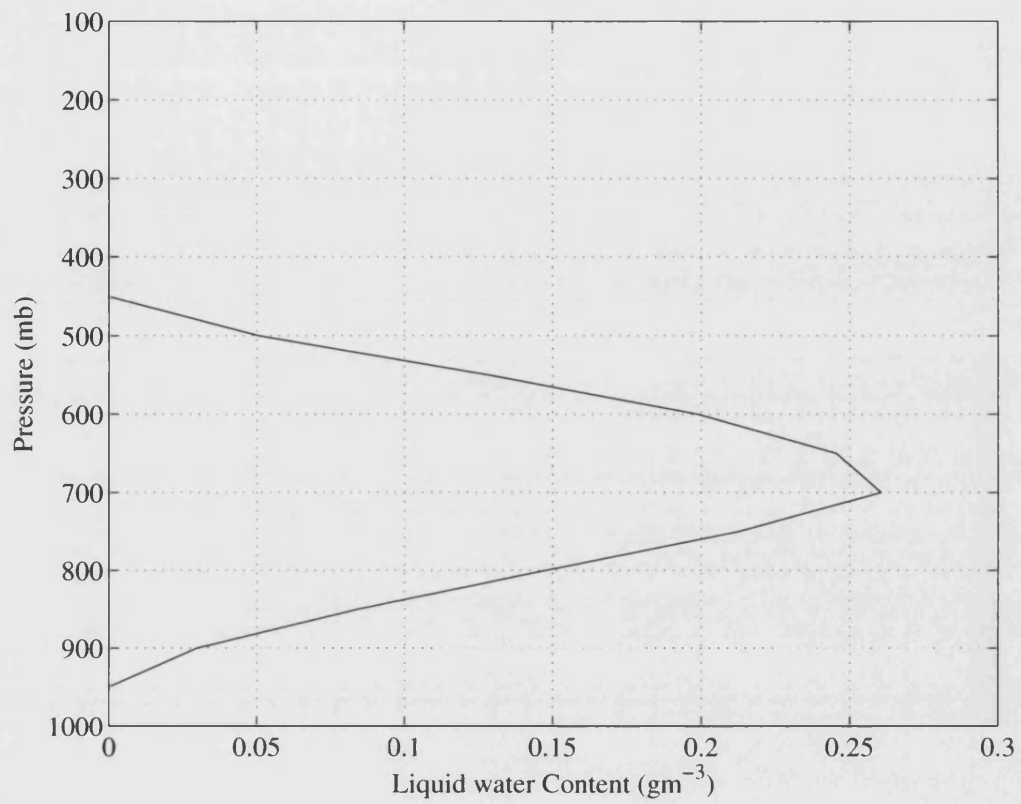
Figure 4.8: Vertical cloud LWC, (23/10/2004 22:00UT).

The specific attenuation can be evaluated from the cloud liquid water using equation 4.10, the vertical profile of specific attenuation is shown figure 4.9.
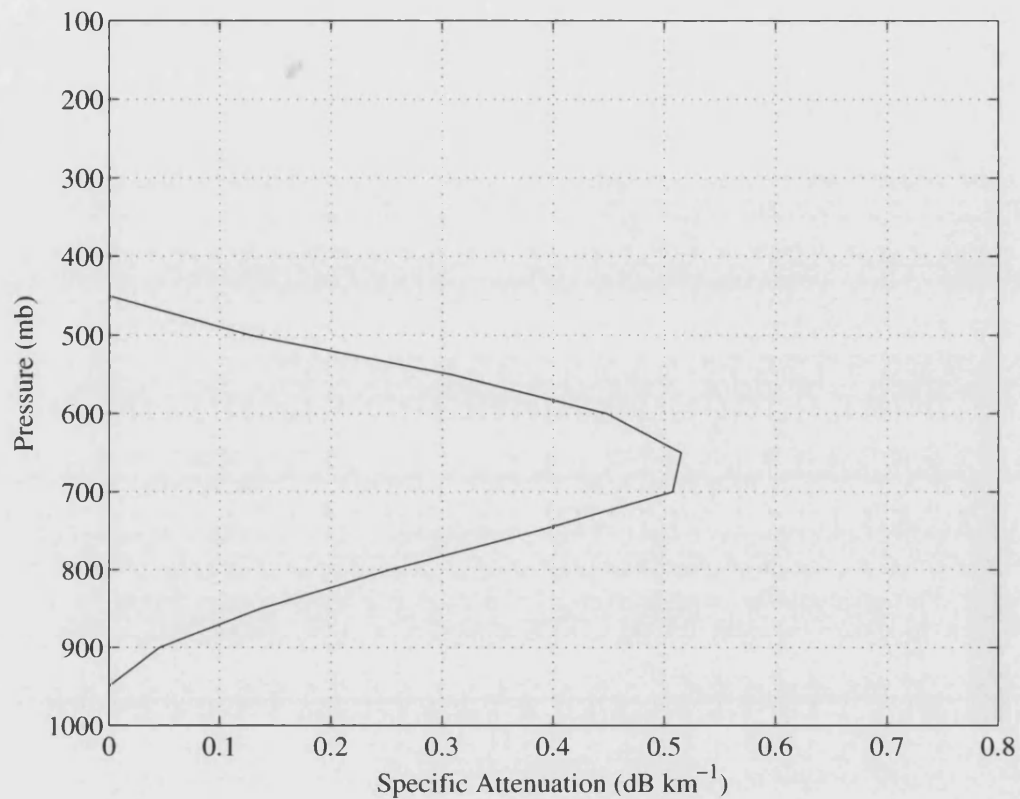
Figure 4.9: Zenith specific cloud attenuation at 50 GHz, (23/10/2004 22:00UT).

The vertical attenuation due to cloud is calculated by integrating the specific attenuation with respect to height. In order to evaluate a slant-path attenuation the same process is performed but along the slant path as opposed to the zenith path.

By following the same process as described the vertical cloud liquid water can be calculated over time, showing the movement of cloud layers over a given location. Using information from the MM5 NWP system the vertical cloud liquid water content was evaluated over a three day period, and is shown in figure 4.10.
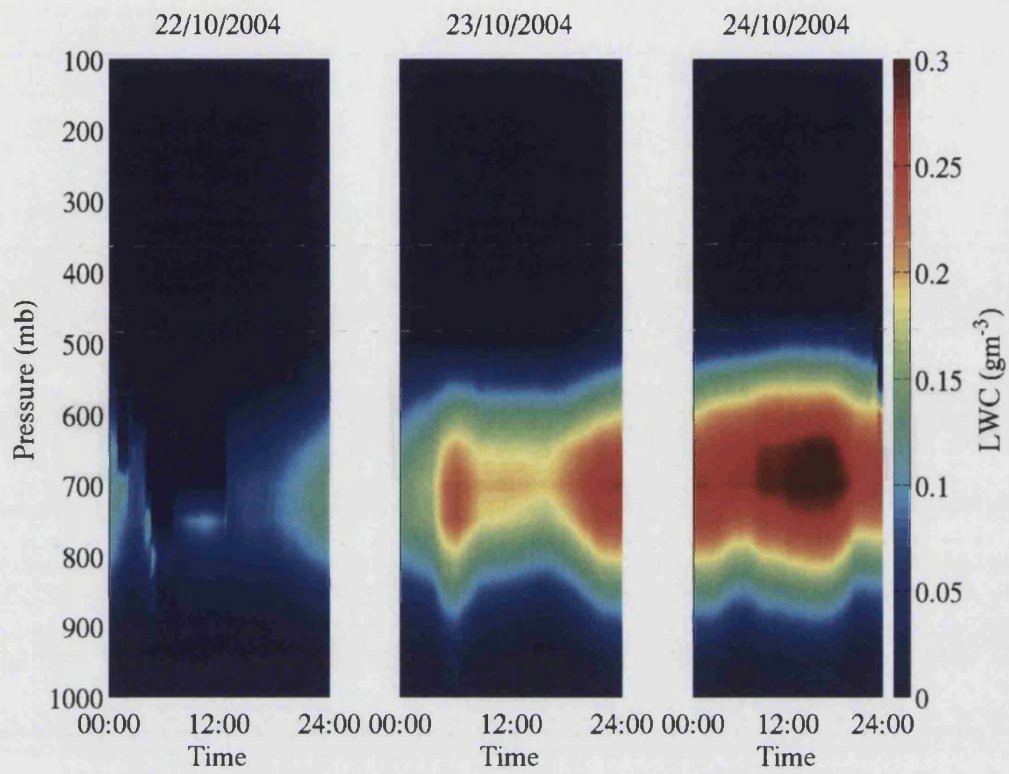
Figure 4.10: Vertical cloud LWC $(\mathrm{g\,m^{-3}})$ from the MM5 NWP for Bath.

For completeness, the resulting zenith cloud attenuation at $50\,\mathrm{GHz}$ for this period is shown in figure 4.11.
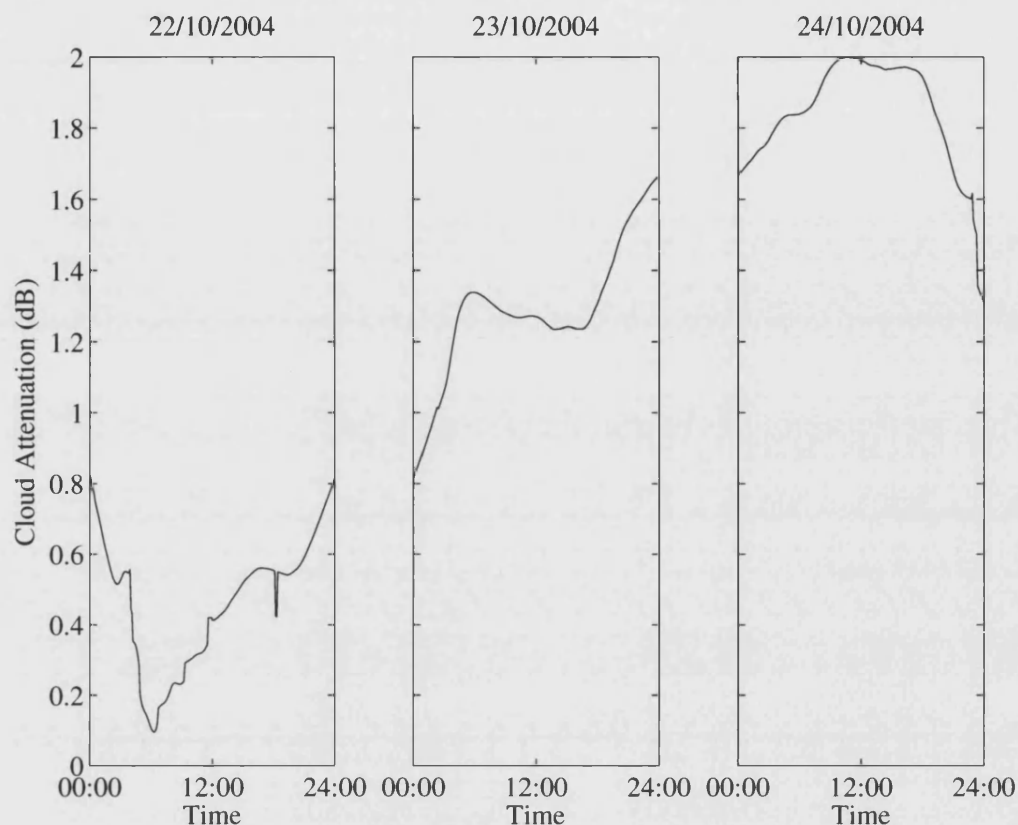
Figure 4.11: Vertical cloud attenuation from the MM5 NWP for Bath.

In summary this plug-in provides the cloud liquid water content along a slant path and the resulting path attenuation.

In theory since fog has a similar composition to warm clouds with a low LWC, it may be possible to use this module to evaluate the attenuation on terrestrial paths due to fog. However, this is beyond the scope of this thesis.

4.3.2 Gas attenuation

The attenuation at microwave frequencies from gaseous absorption is primarily due to absorption by water vapour molecules and oxygen molecules. The method defined in ITU-R recommendation P.676-5 provides a rigorous physical model of the attenuation from the absorption by atmospheric gases, this in turn is based upon the model of Liebe [1989]. The approximation of the algorithm which forms Annex 2 of P.676-5 can be

used to provide a rough approximation of the attenuation; however the approximation is only valid for the first 5 km of the atmosphere so it is not particularly suitable for earth-space slant paths.

The procedure used is a minor simplification of the 'line-by-line' calculation of the path attenuation. In order to simplify the plug-in, raybending is ignored and the path integral is evaluated from the simple geometry. This results in a slight underestimation of the attenuation since in each layer the refracted line integral will be slightly extended over the direct path. The variation in the media due to the slightly different path is unlikely to make a significant difference to the level of the attenuation. Due to the modular nature of the time series generator it is very simple to upgrade individual plug-ins, and the ability to integrate along the refracted path could be added in the future.

The gaseous attenuation can be evaluated by integrating the specific attenuation along the ray path:

$$A_{\text{gas}} = \int_0^L [\gamma_o + \gamma_w] \, dP \qquad (4.16)$$

where γ_o and γ_w represent the specific attenuation caused by dry-air and water vapour respectively. The specific attenuation can also be written thus:

$$\gamma_o + \gamma_w = 0.1820 f N'' (f) \qquad (4.17)$$

Where N'' is the imaginary part of the complex refractivity. The exact solution for this can be found from the line-by-line calculation:

$$N'' (f) = \left[\sum_i S_i F_i \right] + N_D'' (f) + N_W'' (f) \qquad (4.18)$$

where S_i and F_i are the strength and shape of the i-th line and $N_D'' (f)$ and $N_W'' (f)$ represent the wet and dry continuum spectra. The dry continuum spectra is required to help define the spectra below 10 GHz and pressure broadening above 100 GHz. The wet continuum is used to include effect of higher-frequency lines which are not included in the reduced calculation presented here.

Several meteorological parameters are required; the first is the water-vapour partial pressure, e (mb), which can be found from the relative humidity, RH (%), and the temperature T (Kelvin):

$$e = 0.061121 \cdot \text{RH} \cdot \exp \left[\frac{17.502 \, (T - 273.16)}{T - 32.19} \right] \qquad (4.19)$$

The techniques also require the dry air pressure, p (mb), if the total barometric, P, is used the dry air pressure can be found from $P = p + e$.

Once these parameters have been evaluated the wet-continuum can be found from:

$$N_W''(f) = f\left(3.57\left(\frac{300}{T}\right)^{7.5} e + 0.113p\right) 10^{-7} e\left(\frac{300}{T}\right)^3 \qquad (4.20)$$

and the dry continuum found from:

$$N_D''(f) = fp\left(\frac{300}{T}\right)^2 \cdot$$
$$\left[\frac{6.14 \times 10^{-5}}{d\left[1 + \left(\frac{f}{d}\right)^2\right]} + 1.4 \times 10^{-12}\left(1 - 1.2 \times 10^{-5} f^{1.5}\right) p\left(\frac{300}{T}\right)^{1.5}\right] \qquad (4.21)$$

where $d = 5.6 \times 10^{-4}\left(p + 1.1e\right)\left(\frac{300}{T}\right)$.

In order to calculate the specific attenuation there is still a requirement to calculate the line strength and shapes. The constants for these calculations are contained in tables B.2 and B.1.

The line strength of each 'dry-atmosphere' line can be calculated from:

$$S_i = a_1 \times 10^{-7} p\left(\frac{300}{T}\right)^3 \exp\left[a_2\left(1 - \left(\frac{300}{T}\right)\right)\right] \qquad (4.22)$$

and for each water vapour line:

$$S_i = b_1 \times 10^{-1} e\left(\frac{300}{T}\right)^{3.5} \exp\left[b_2\left(1 - \left(\frac{300}{T}\right)\right)\right] \qquad (4.23)$$

The remaining parameter in 4.18 is the line shape factor, F_i which is given by the following equation:

$$F_i = \frac{f}{f_i}\left[\frac{\Delta f - \delta\left(f_i - f\right)}{\left(f_i - f\right)^2 + \Delta f^2} + \frac{\Delta f - \delta\left(f_i - f\right)}{\left(f_i + f\right)^2 + \Delta f^2}\right] \qquad (4.24)$$

where f_i is the line frequency, Δf is the line width and δ accounts for the interference between the oxygen lines around 60 GHz. These can be calculated for the oxygen lines:

$$\Delta f = a_3 \times 10^{-4}\left(p\left(\frac{300}{T}\right)^{(0.8 - a_4)} + 1.1e\left(\frac{300}{T}\right)\right) \qquad (4.25)$$

$$\delta = \left(a_5 + a_6\left(\frac{300}{T}\right)\right) \times 10^{-4} p\left(\frac{300}{T}\right)^{0.8} \qquad (4.26)$$

and for the water vapour lines:

$$\Delta f = b_3 \times 10^{-4} \left(p \left(\frac{300}{T} \right)^{b_4} + b_5 e \left(\frac{300}{T} \right)^{b_6} \right) \tag{4.27}$$

$$\delta = 0 \tag{4.28}$$

In conclusion, the gaseous attenuation is calculated by evaluating the complex refractivity using a 'line-by-line' calculation which is then used by the MPM Liebe model [Liebe, 1989] to evaluate the specific attenuation along the approximate slant path. This is then path integrated in order to calculate the link attenuation.

4.3.3 Rain attenuation

The attenuation due to rain is a particularly difficult phenomenon to model since as described previously, there is great variability from rain event to rain event. This variability is typically in such characteristics as raindrop size distribution, canting angle, rain height, raindrop temperature and fall speed. Rainfall is typically measured as a surface measure and assumed to be vertically homogeneous to the rain height. This implies that two storms that produce an identical rainfall rate may have a very different structure, i.e. have very different rain heights, raindrop size distributions and vertical profiles of different water phases.

The specific attenuation due to rain ($dB\,km^{-1}$) is typically modelled using a power-law:

$$\gamma_{rain} = aR^b \tag{4.29}$$

Where R is the surface rainfall rate, and a and b are constants. The validity of this relationship has been considered by Olsen et al. [1978], where equation 4.29 was evaluated as an approximation of a more general relationship:

$$\gamma_{rain} = a'R^{b'} \left[1 + \sum_{n=2}^{\infty} c_n f^n R^{nd} \right] \tag{4.30}$$

In this equation a', b' and c_n are functions of frequency, rain temperature and the raindrop size distribution and d is a constant depending on the raindrop size distribution. This implies that the simplification in equation 4.29 is an exact solution as $f \rightarrow 0$ or R is small. This is consistent since log-log plots of γ_{rain} and R are not perfectly linear. Nevertheless over rainfall rates of practical interest such plots are almost linear, making equation 4.29 valid [Olsen et al., 1978].

In order to evaluate the attenuation due to rain the model described by Leitao and Watson [1986] is used. This model is particularly useful since no assumptions are made about the structure of the rain fields. The model was derived from dual-polarisation radar measurements of the micro-structure of rain. The model has also two different formulations, one for widespread rain and one for convective rain. Having two formulations helps to reduce the effects of the variable raindrop size distribution. This model also performed well in the assessments that were performed as part of the COST 205 programme.

The Leitao-Watson model uses two power-law relationships, one to model the specific attenuation up to rain rates of $20\,\mathrm{mm\,hr^{-1}}$ and another for rain rates in excess of $20\,\mathrm{mm\,hr^{-1}}$, this can be justified by the piecewise linear approximation to the summation in equation 4.30. The model can be summarised as:

$$\gamma_{\mathrm{rain}} = \begin{cases} a\left(\frac{R}{20}\right)^x & : \quad R < 20\,\mathrm{mm\,hr^{-1}} \\ a\left(\frac{R}{20}\right)^y & : \quad R \geq 20\,\mathrm{mm\,hr^{-1}} \end{cases} \tag{4.31}$$

The values of the constants a, x and y can be calculated for any polarisation and elevation angle using:

$$a = \frac{\left|a_H + a_V + (a_H - a_V)\cos^2\theta\cos 2\tau\right|}{2} \tag{4.32}$$

$$x = \frac{\left|a_H x_H + a_V x_V + (a_H x_H - a_V x_V)\cos^2\theta\cos 2\tau\right|}{2a} \tag{4.33}$$

$$y = \frac{\left|a_H y_H + a_V y_V + (a_H y_H - a_V y_V)\cos^2\theta\cos 2\tau\right|}{2a} \tag{4.34}$$

where θ is the elevation angle and τ is the polarisation angle with respect to the horizontal (circular polarisations can be considered as $\tau = 45°$). The values for the vertical and horizontal constants can be calculated :

$$a_{H,V} = \frac{s_{H,V}}{1 + 0.01L} \tag{4.35}$$

$$x_{H,V} = t_{H,V}\log_{10}(L+1) + u_{H,V} \tag{4.36}$$

$$y_{H,V} = v_{H,V}\log_{10}(L+1) + w_{H,V} \tag{4.37}$$

where L is the horizontal projection of the slant path. The constants s, t, u, v and w have been evaluated for both widespread rain (i.e. stratiform rain which demonstrates a bright-band on radar measurements) and showery rain (i.e. convective rain which, because of the turbulent structure, typically has no bright-band) and are shown in table B.3[5]. The ability to evaluate the attenuation from convective and large-scale rain independently is a very powerful tool which helps to reduce the errors caused by the variable structure of rainfall and allows the use of the separate convective and large

---

[5]The authors recommend that a logarithmic interpolation scheme is used to find the value of these

scale NWP fields.

L is the horizontal projection of the slant path, and can simply be approximated as $L = \frac{H_r}{\tan \theta}$ where H_r is the effective rain height. Once the specific attenuation has been calculated, the slant-path attenuation caused by rain can be approximated as:

$$A_{\text{rain}} = \frac{H_r}{\sin \theta} \cdot \gamma_{\text{rain}} \tag{4.41}$$

Rain Height

An important parameter within the modelling of the attenuation due to rain is the effective rain height. The attenuation due to to rain is typically derived from surface rain measurements by assuming the rain is homogeneous from the surface to the effective rain height. Hence the $\frac{H_r}{\sin \theta}$ term in equation 4.41, which defines the length of the slant path through the homogeneous media.

The rain height is typically considered to extend from the top of the melting layer to the surface. The ITU-R recommendation P.839.3, 'Rain height model for prediction methods', suggests that the rain height can be considered as being 360 m above the 0°C isotherm, whist Leitao and Watson [1986] recommend 300 m above the 0°. To be consistent with the Leitao-Watson model we assume that the rain height is 300 m above the freezing height. This is true for the on-average condition, however it varies from event-to-event.

Among the future improvements to NWP techniques are the inclusion of the vertical structure of both convective and large-scale rain i.e. rain as a three dimensional phenomena. This will greatly improve the modelling of the attenuation due to rain. Indeed, the newest versions of the MM5 and WRF models are able to estimate the mixing ratio of mixed phase phenomena (i.e. cloud ice, cloud water, rain water, etc.) on model levels.

If the rain height is considered to be 300 m above the zero-degree isotherm it is still

---

coefficients for other frequencies thus:

$$\log_{10} a_{H,V} = \left(\log_{10} a''_{H,V} - \log_{10} a'_{H,V} \right) \frac{\log_{10} f - \log_{10} f'}{\log_{10} f'' - \log_{10} f'} + log_{10} a'_{H,V} \tag{4.38}$$

$$x_{H,V} = \left(x''_{H,V} - x'_{H,V} \right) \frac{\log_{10} f - \log_{10} f'}{\log_{10} f'' - \log_{10} f'} + x'_{H,V} \tag{4.39}$$

$$y_{H,V} = \left(y''_{H,V} - y'_{H,V} \right) \frac{\log_{10} f - \log_{10} f'}{\log_{10} f'' - \log_{10} f'} + y'_{H,V} \tag{4.40}$$

where f is the required frequency between f' and f''

important to calculate the zero degree isotherm as accurately as possible. The ITU-R recommendation P.839-3 provides a yearly average height for the zero degree isotherm, shown in figure 4.12.

Annual average zero degree isotherm height
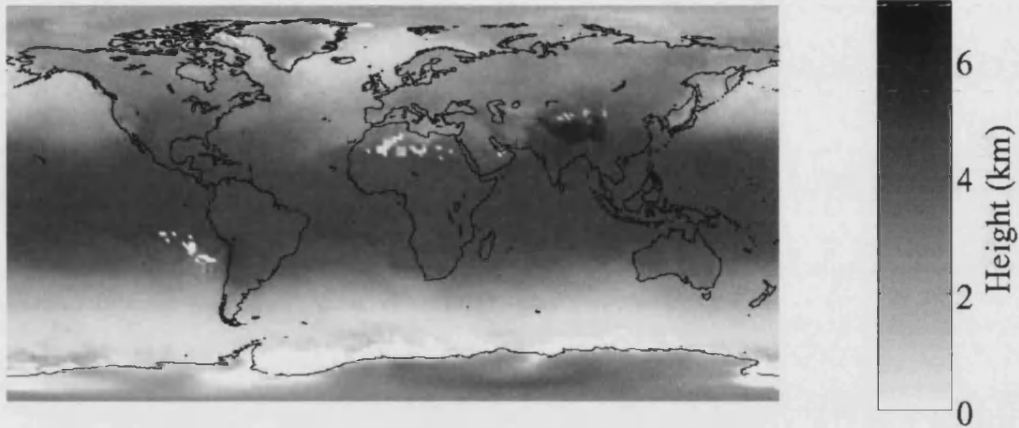
Figure 4.12: The annual average zero degree isotherm height (km) from ITU-R recommendation P.839-3.

However, this is not suitable for instantaneous evaluation of the rain height since the isotherm will display significant seasonal variation. In order to calculate the instantaneous zero-degree isotherm height it is possible to use the vertical temperature profile from NWP techniques. Temperature is a prognostic variable and is well modelled by NWP techniques.

Using the UM NWP model, an example time series of the height of the zero-degree isotherm over three months in 2003, has been evaluated for the location of Bath, UK (2.3599°W, 51.3758°N). The convective and large scale rain rate time series have also been determined and are shown in figure 4.13.
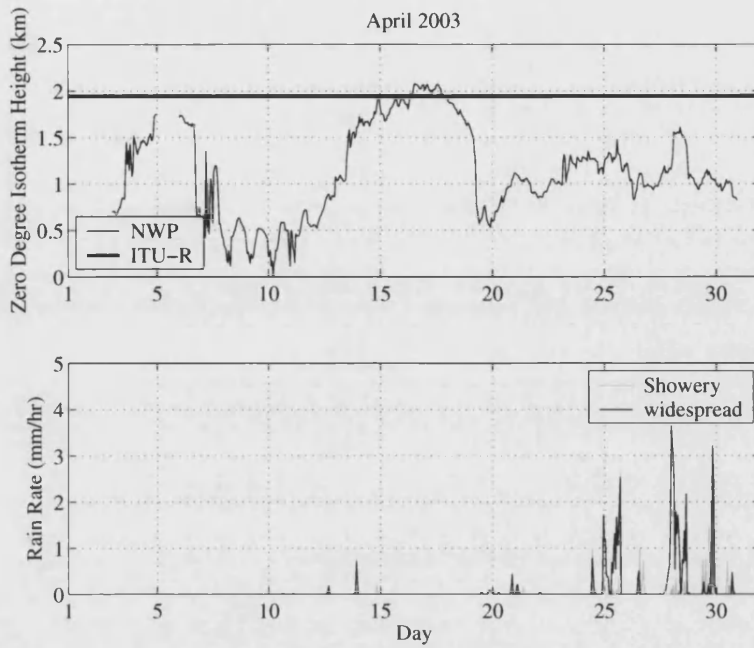
(a) April 2003
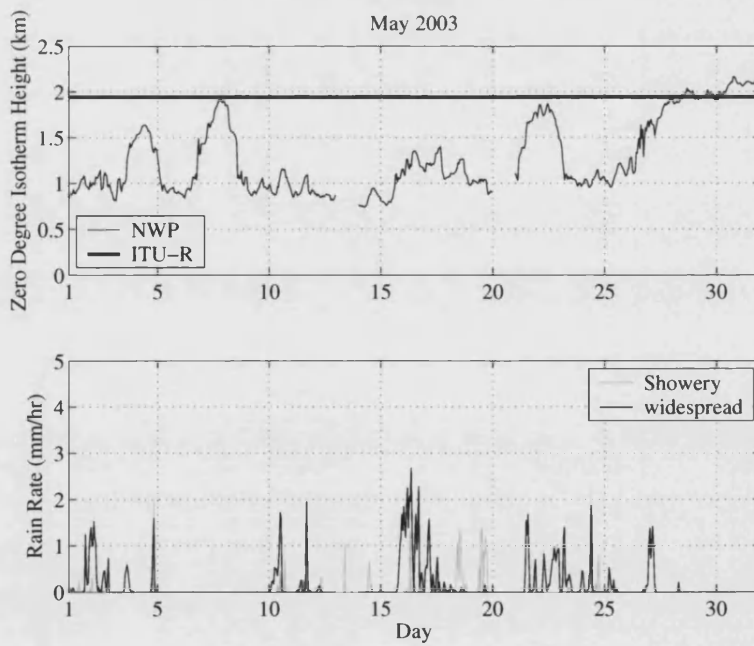
(b) May 2003

Figure 4.13: Time series of the zero degree isotherm height as measured by the UM NWP system.
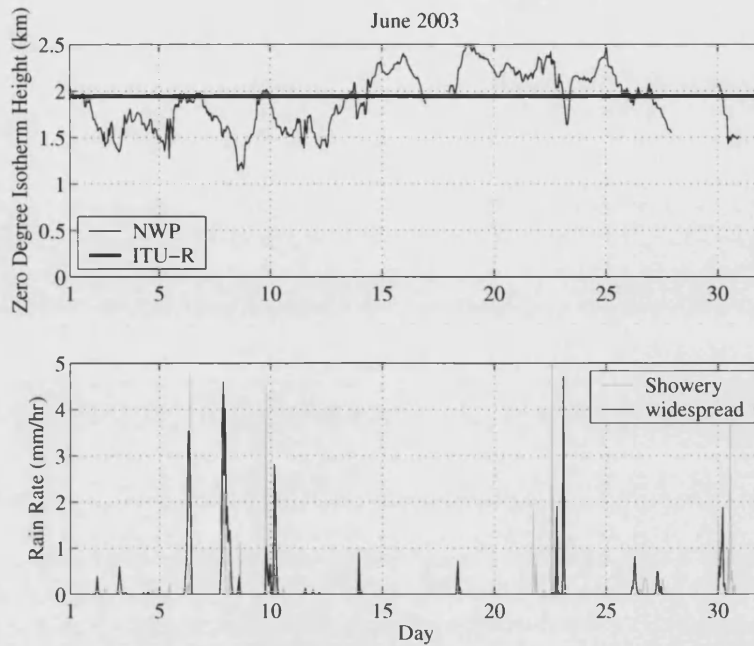
(c) June 2003

Figure 4.13: Time series of the zero degree isotherm height as measured by the UM NWP system.

The seasonal, daily, and indeed hourly, variation shown in figure 4.13 is not modelled by the ITU-R annual average. This is a significant variation about the value suggested by the ITU-R recommendation. It is very important to realise that from equation 4.41 a 50 % error in the rain height will result in a 50 % error in the attenuation in dB.

4.3.4   Scintillation effects

The effects of scintillation are modelled in two stages. The first stage is to model the scintillation intensity. The scintillation intensity is the short-term variance of the scintillation component. This value can also be of use within a propagation forecasting paradigm since it can be used to provide a warning of possible large scintillation intensities.

Since it is impossible to forecast the instantaneous value of the scintillation component practical systems will typically have a small scintillation fade-margin. By exploiting a forecast of the scintillation intensity it may be possible to modulate this fade margin, however the gains from this are likely to be small.

Within a simulation paradigm the time series of the scintillation intensity can be used to create a stochastic synthetic scintillation time series which can be combined with the 'slow' fading time series from the other plug-ins.

Ideally, to estimate the scintillation intensity the turbulent structure constant C_n^2 would be evaluated and following Tatarski's seminal theory [Tatarski, 1961] the scintillation intensity at millimeter frequencies could then be evaluated thus:

$$\sigma^2 = 42.48 \frac{k^{\frac{7}{6}}}{(\sin\theta)^{\frac{11}{6}}} \int_0^L C_n^2(z) z^{\frac{5}{6}} dz \tag{4.42}$$

where k is the wave number, θ is the elevation angle. However, in order to evaluate C_n^2 accurately enough to resolve all the significant turbulent structures would require vertical data at a resolution of a few meters [Vasseur, 1999], this is not practical for NWP techniques and even 'high-resolution' radiosonde ascents can only achieve resolutions of tens of meters.

At present there are no reliable models to generate the short term (10 minutes) scintillation variance. However, a procedure attempting to model the scintillation intensity on time series of the order of 3 hours has been demonstrated by Tervonen et al. [1988]. This used the average cumulus cloud coverage which is well correlated with the diurnal variation in scintillation intensity seen in the spring and summer. Cloud coverage can be measured by an integer parameter between 0 and 8 indicating how many eighth parts of the sky are covered. This parameter is not well suited for a physical model, such as the one presented here.

The cumulus cloud cover parameter can be considered an indicator of turbulence. The model of Tervonen et al. [1988] was later refined by van de Kamp [1999] to use a more rigorous parameter: the water content of heavy clouds, W_{hc}[6]. The model is recommended to be used with annual averages for N_{wet} and W_{hc}, however it has been used as an instantaneous model as the most appropriate plug-in.

The model is formulated thus [van de Kamp, 1999]:

$$\sigma_n = 0.98 \times 10^{-4} \left(N_{wet} - 39.2 + 56 W_{hc} \right) \tag{4.43}$$

$$\sigma = \sigma_n g\left(D_e\right) f^{0.45} \sin^{1.3}\theta \tag{4.44}$$

where N_{wet} is the wet term of the refractive index, W_{hc} is the water content of 'heavy clouds', f is the frequency (GHz), θ is the elevation angle and $g(D_e)$ is the antenna averaging function.

---

[6]As discussed previously these are clouds with an integrated water content greater than $0.7\,\mathrm{kg\,m^{-2}}$.

The antenna averaging function is given by:

$$g(De) = \begin{cases} \sqrt{1.0 - 0.7\left(\frac{D_e}{\sqrt{\lambda L}}\right)} & : \quad 0 \leq \frac{D_e}{\sqrt{\lambda L}} \leq 1 \\ \sqrt{0.5 - 0.2\left(\frac{D_e}{\sqrt{\lambda L}}\right)} & : \quad 1 < \frac{D_e}{\sqrt{\lambda L}} \leq 2 \\ \sqrt{0.1} & : \quad 2 < \frac{D_e}{\sqrt{\lambda L}} \end{cases} \qquad (4.45)$$

The effective turbulent path length can be approximated using [Karasawa et al., 1988]:

$$L = \frac{2h}{\sqrt{\sin^2 \theta + \frac{2h}{R_e}} + \sin \theta} \qquad (4.46)$$

where h is the height of the turbulent layer (generally taken as $h = 2000\,m$), θ is the elevation angle and R_e is the radius of the earth ($8.5 \times 10^6\,m$).

This only leaves the two meteorological parameters in the model, N_{wet} and W_{hc}, to be evaluated. N_{wet} is the wet term of the refractive index and can simply be evaluated using:

$$N_{wet} = \frac{22790 \text{RH} \exp\left(\frac{19.7T}{T+273}\right)}{(T+273)^2} \qquad (4.47)$$

where RH is the relative humidity (%) and T is the temperature (°C).

The water content of 'heavy-clouds' can be evaluated using the Salanon and Uppalla cloud model described previously. The critical humidity can be extracted using equation 4.6, from this the water content can be calculated using:

$$W_{hc} = w_0 \left(1 + ct\right) \left(\frac{h_c}{h_r}\right)^a \qquad (4.48)$$

Note the similarity between this equation and equation 4.7 which was used for the evaluation of the cloud *liquid* water content. This equation is identical except since W_{hc} includes the ice content the parameter defining the liquid fraction is removed.

Once the scintillation variance has been evaluated the time series can be synthesised. A vector the length of the time series at a sample time of 1 seconds is created from a Gaussian random variate. The use of the Gaussian random variable is valid, since for most levels of scintillation intensity the scintillation component follows a Gaussian distribution. It should be noted that at very large levels of scintillation intensity the distribution begins to exhibit some asymmetry [van de Kamp, 1998].

This scintillation time series is then filtered in order to provide the correct spectral characteristics, the filter is defined by wave propagation theory in turbulent media,

[Ishimaru, 1978].

The average spectrum of scintillation can be considered to be [Ishimaru, 1978]:

$$
\begin{aligned}
W_x^0(\omega) &= 2.765\frac{\sigma^2}{\omega_t} &&: \quad \omega \to 0 \\
W_x^\infty(\omega) &= 7.13\frac{\sigma^2}{\omega_t}\left(\frac{\omega}{\omega_t}\right)^{-\frac{8}{3}} &&: \quad \omega \to \infty
\end{aligned}
\tag{4.49}
$$

The value of σ^2 represents the scintillation intensity. These two functions define the two theoretical average asymptotes, $\omega \to 0$ and $\omega \to \infty$, which meet at the frequency $\omega = 1.43\omega_t$.

The variable ω_t is defined as:

$$
\omega_t = W_t\sqrt{\frac{2\pi}{\lambda L}}
\tag{4.50}
$$

where W_t is the transverse wind speed, λ is the wavelength and L represents the effective turbulent path length (see equation 4.46).

This behaviour has been seen to agree with beacon measurements from the Olympus measurement campaign [Savvaris et al., 2004] (the observed spread around the average was quite large, this is likely due to the local measurements of surface winds, which were actually taken from about 20 km away).

In conclusion the low pass filter cut-off can be described as:

$$
\omega_c = 1.43W_t\sqrt{\frac{2\pi\left(\sqrt{\sin^2\theta + \frac{2h}{R_e}} + \sin\theta\right)}{\lambda 2h}}
\tag{4.51}
$$

It is worth noting that NWP techniques typically provide u and v winds, which are the winds with respect to the model grid. These grids are typically rotated pole grids hence u and v winds do not represent true westerly and southerly winds, this is done in order to maintain an even grid size across the domains. These u and v winds can be converted to true westerly and southerly winds using simple spherical geometry:

$$
\begin{aligned}
u_t &= c_1 u + c_2 v \tag{4.52} \\
v_t &= c_1 u + c_2 v \tag{4.53}
\end{aligned}
$$

where u_t and v_t represent the true westerly (from the west) and true southerly (from the south) winds respectively. The constants c_1 and c_2 are functions of the rotated grid parameters: λ and ϕ , the true latitude and longitude; λ' and ϕ', the rotated latitude

and longitude and λ_0 and ϕ_0 the latitude and longitude of the rotated pole.

$$
\begin{aligned}
c_1 &= -\sin(\lambda - \lambda_0)\sin\lambda'\sin\phi_0 + \cos(\lambda - \lambda_0)\cos\lambda' & (4.54)\\
c_2 &= \sin\phi\sin\lambda'\cos(\lambda - \lambda_0)\sin\phi_0 - \sin\phi\cos\lambda'\sin(\lambda - \lambda_0) - \\
&\quad \sin\lambda'\cos\phi\cos\phi_0 & (4.55)
\end{aligned}
$$

Once the time series has been filtered it can then be modulated by the scintillation intensity to provide the final scintillation time series.

4.3.5 Example time series

In order to better illustrate the procedure an example time series has been created using MM5 NWP data. This provided data at 6 minute intervals on two grid scales, the top domain uses a grid scale of 18 km and a nested grid with a grid size of 6 km. The simulated link is a 50 GHz circularly polarized link from Bath ($2.3599°$W, $51.3758°$N) to a geostationary payload at $2°$West on 23/10/2004.

The attenuation due to gases is shown in figure 4.14, adding the attenuation due to clouds results in the time series shown in figure 4.15. The final time series shown in figure 4.16 includes the attenuation due to rain. In figure 4.17 the scintillation variance has been plotted about the attenuation level from the smallest grid size.
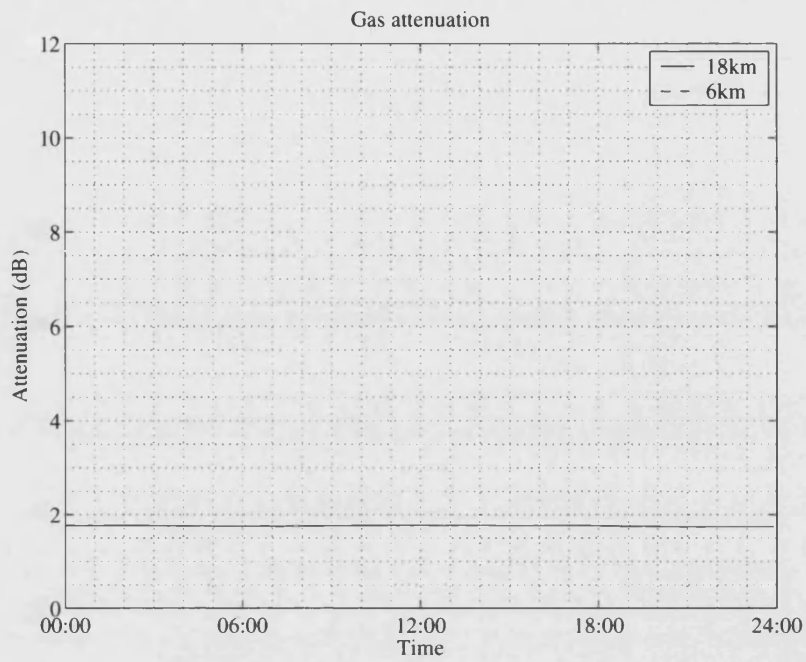
Figure 4.14: The time series for the attenuation due to gases.
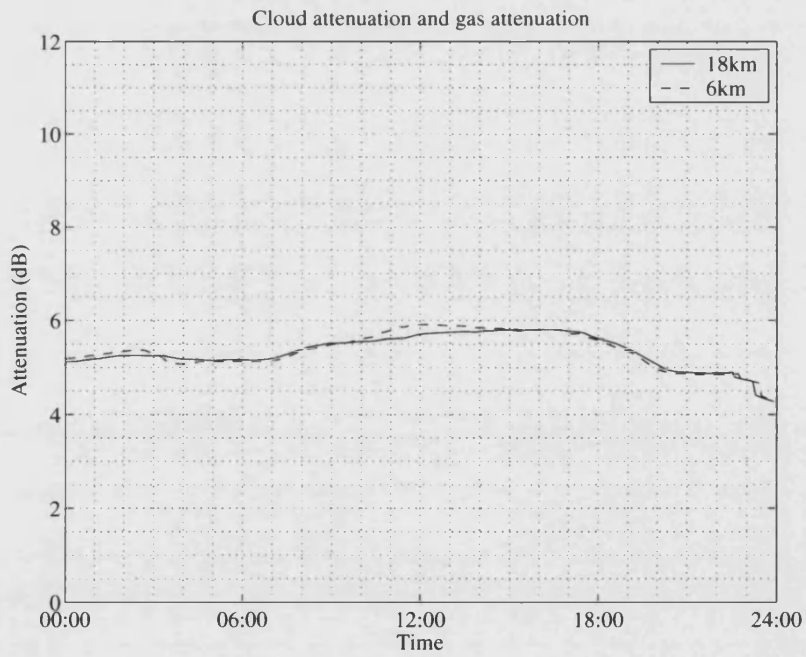
Figure 4.15: The time series for the attenuation due to gases and clouds.
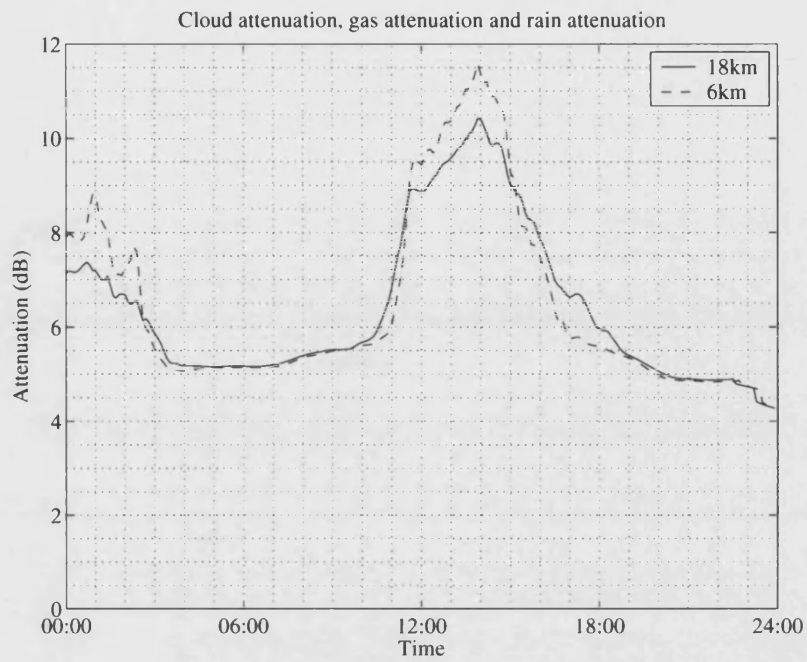
Figure 4.16: The time series for the attenuation due to gases, clouds and rain.
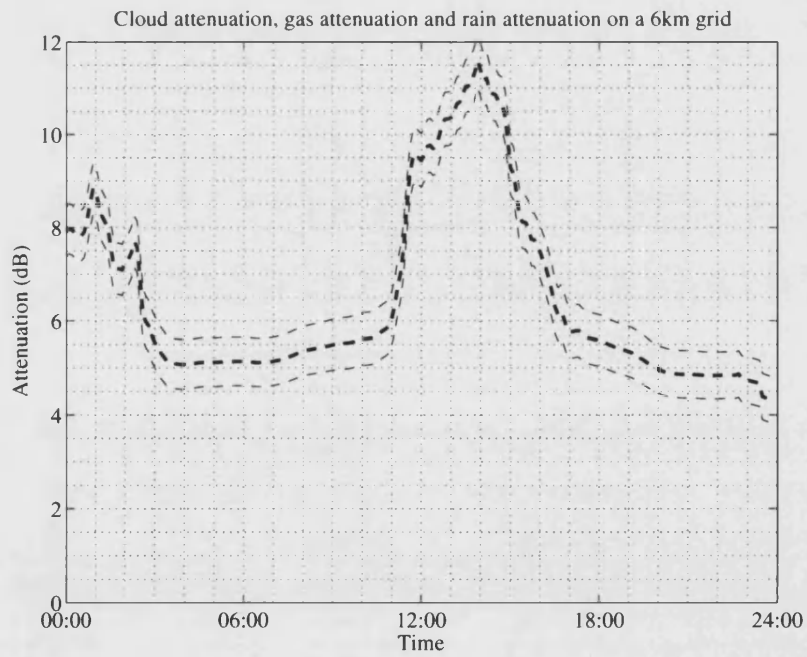
Figure 4.17: The time series for the attenuation due to gases, clouds and rain for the smallest grid scale.

Table 4.1: Example of the USGS land use characteristics.

Vegetation Integer Identification	Vegetation Description	Albedo (%)		Moisture Avail. (%)		Emissivity (% at $9\,\mu$ m)		Roughness Length (cm)		Thermal Inertia (cal cm^{-2} k^{-1} s$^{-1/2}$)	
		Sum	Win	Sum	Win	Sum	Win	Sum	Win	Sum	Win
1	Urban	18	18	10	10	88	88	50	50	0.03	0.03
2	DryInd Crop. Past.	17	23	30	60	92	92	15	5	0.04	0.04
5	Crop. / Grassland	18	23	25	40	92	92	14	5	0.04	0.04
6	Crop. / Wood	16	20	35	60	93	93	20	20	0.04	0.04
7	Grassland	19	23	15	30	92	92	12	10	0.03	0.04
11	Decids. Broadleaf	16	17	30	60	93	93	50	50	0.04	0.05
16	Water Bodies	8	8	100	100	98	98	0.01	0.01	0.06	0.06
18	Wooded wetland	14	14	35	70	95	95	40	40	0.05	0.06

Chapter 5

Downscaling Techniques

One problem with evaluating the time series from NWP techniques is the temporal and spatial resolution of the data sets. Typically the rain estimates are provided with a spatial resolution that is far worse than that of the phenomena. This can result in time series that do not display the effects of the microstructure of rain, since the phenomena at this resolution is not resolved.

One way to alleviate the problems of finite resolution in the 'fields' of rainfall rate, or indeed in rainfall rate time series, is to introduce a synthetic small scale structure. The introduced structure is representative of the small scale structure seen in the relevant phenomena. It is very important to understand that the structure is synthetic and hence can be considered one possible realisation of the small-scale structure. This process is very similar to the process used to create the scintillation time series; since it is not possible to resolve the small scale variations that lead to scintillation a synthetic scintillation time series is created that exhibits the correct statistics (in essence the time series is one possible realisation of the scintillation time series).

It is also worth noting that since the small scale structure is only one possible solution (of, in theory, an infinite set) these techniques should not be used when the time series generator is working within a forecasting paradigm. Since the small-scale structure is stochastic the likelihood of it being 'correct' is very small. However for simulation purposes the ability to create representative small-scale structures is nonetheless very useful.

In order to create the small scale structure it is possible to exploit the fractal nature of rainfall.

5.1 Introduction to rainfall and fractal behaviour

The term 'fractal' technically refers to an object that displays self-similarity on all scales. The object need not exhibit exactly the same structure at all scales, but the same 'type' of structures must appear on all scales. A plot of this quantity on a log-log graph versus scale then gives a straight line, whose slope is said to be the fractal dimension.

Fractals are characterised by their fractal dimension, if the linear dimension is changed by a factor λ then the values of the fractal quantity change by a factor λ^D. For example, a cube exhibits a dimension of 3, since if the edge length is doubled the number of instances of the original cube increases to 8.

Lines or contours exhibit a fractal dimension between 1 and 2, whilst surfaces exhibit a fractal dimension between 2 and 3. Generally the higher the fractal dimension the rougher the surface. It is important to understand that the fractal surface of a two-dimensional discrete field is very difficult to evaluate, there are several ways of estimating the fractal dimension of surface fields and each process demonstrates some bias and errors.

One of the simplest and yet most rigorous ways of evaluating the fractal dimension of fractal two-dimensional surfaces is using a Fourier technique, such as the one described by Russ [1994]. When a fractal surface undergoes a Fourier transform it is expected that the log magnitude will decay linearly with log frequency. Rain fields on-average tend to exhibit a Kolmogorov-like spectrum, implying a linear fall off between log-magnitude and log-spatial-frequency[1].

The fractal dimension can be extracted from the Fourier transform as:

$$D = \frac{6 + \beta}{2} \tag{5.1}$$

where β is the slope of the log-magnitude squared of the Fourier transform plotted against log-frequency.

In order to determine the fractal dimension accurately the two-dimensional FFT is calculated over the entire field, and subsequently shifted such that the central point of the image represents the zero frequency component. To evaluate the magnitude of the Fourier transform the transform is extracted from 30 radial distances around 24

---

[1]One interesting phenomena of the Fourier transform of a fractal surface is that the phase component is random, this property can be used as another verification of fractal behaviour.

equally spaced radials. These are then averaged across each of the radials to provide the final log-magnitude squared vs. log-frequency plot. The β variable can then be evaluated using a simple least mean squares algorithm. It has been found that this process of evaluating the fractal dimension provides a value that is less than or equal to the Hausdorf Dimension [Russ, 1994].

An example taken from the MM5 field shown in figure 5.1(a). From this example field the Fourier estimation of the fractal dimension can be determined and is shown in figure 5.1(b).

(a) Example field

Fourier estimation of fractal dimension: β=−1.2259, D=2.387
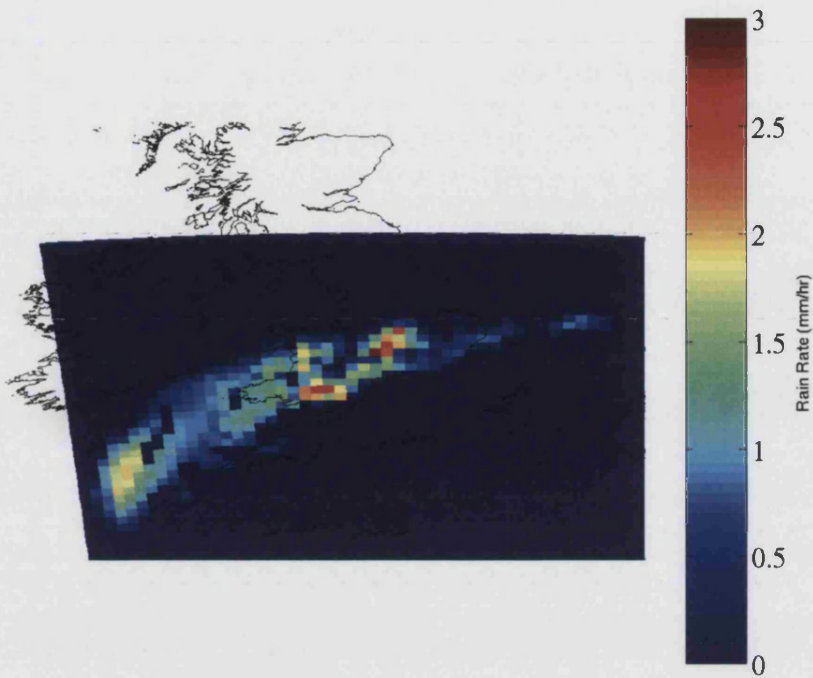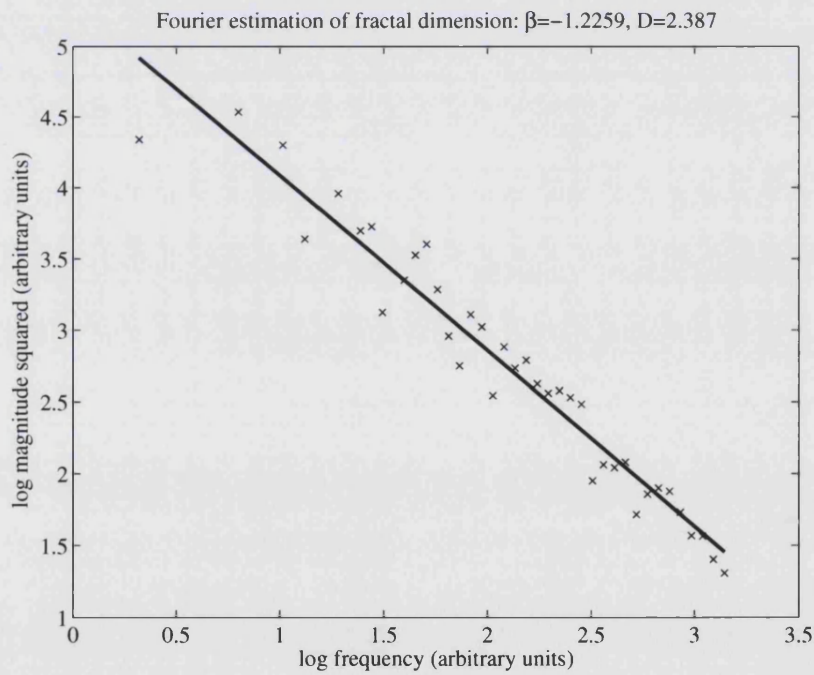
(b) Fractal dimension of the example field

Figure 5.1: An example MM5 rain field and the fractal dimension calculation, 22/10-/2004 12:00UT, grid scale 18 km.

The fractal dimensions were evaluated from three different data sources, NWP rain fields from two different MM5 grids, the CAMRa radar and NWP rain fields from the operational UM. The two MM5 data sets have grid sizes of 18 km and 6 km and are over a period of 5 days at an interval of 6 minutes. The CAMRa radar data has a grid scale of approximately 330 m with each scan separated by about 2 minutes. The data from the operational UM have a grid size of 12 km with each estimate separated by an hour. The data sets have been restricted to fields that demonstrate a mean rainfall in excess of $0.02 \, \mathrm{mm \, hr^{-1}}$ across the rain field. This spread of fractal dimensions are shown in figure 5.2.

(a) NWP - MM5 system   (b) Radar - CAMRa

(c) NWP - UM system
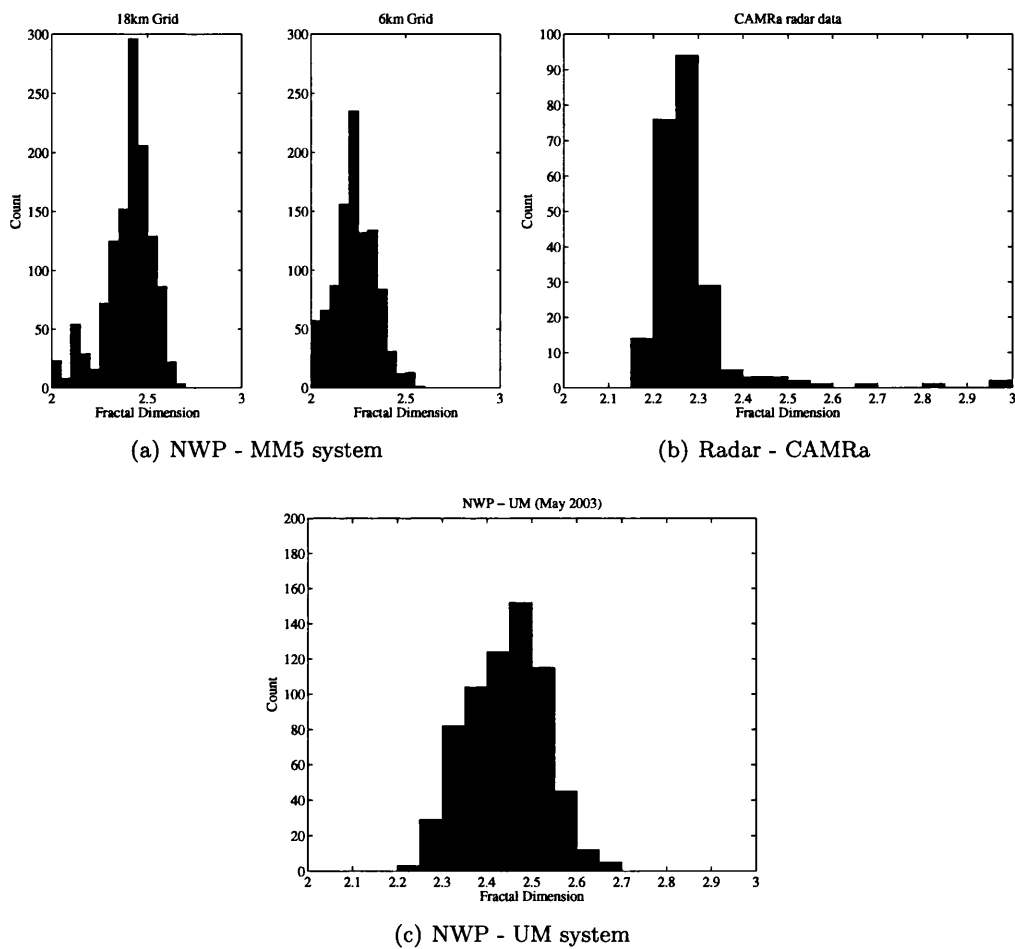
Figure 5.2: Histograms of fractal dimension from three different data sets.

It can be seen there is a large amount of variation in the distributions of fractal dimension between the different data types. This variation is due to several phenomena. Each of the data types is upon a different grid scale and is resolving slightly different phenomena. For example, the UM and MM5 models model rain with slightly different

algorithms whilst the radar provides a measure of instantaneous rainfall rate.

It should also be noted that many quantities can appear to exhibit fractal scaling. However, caution should always be used since the techniques mentioned generally use linear regression, and hence will always produce a solution. If we assume that rain fields exhibit fractal scaling the level of the scaling needs to be evaluated. The most common way of doing this is to determine the generalised structure function.

The use of the structure function as a tool to investigate small-scale turbulence was introduced by Kolmogrov in the early 1940s, and then extended by Benzi et al. [1993] to become the generalised structure function. The generalised structure function in the x and y directions of a discretized rain field R on a regular grid (x, y) are defined thus:

$$S_x(\lambda, q) = \langle |R(x + \lambda, y) - R(x, y)|^q \rangle \qquad (5.2)$$

$$S_y(\lambda, q) = \langle |R(x, y + \lambda) - R(x, y)|^q \rangle \qquad (5.3)$$

where q represents the order of the structure function and angle brackets represent the ensemble average. The structure functions were evaluated for the three previous data sources. For isotropic fields, S_x and S_y should be identical, however rain fields typically demonstrate a small level of anisotropy which results in slight differences between the two measures. To be representative an average of the two directions is typically taken:

$$S(\lambda, q) = \frac{S_x(\lambda, q) + S_y(\lambda, q)}{2} \qquad (5.4)$$

The first-order structure functions were evaluated for a set of rainfall fields; these are shown in figure 5.3. The MM5 NWP results were taken from an example day (2004/10/24) from a group of domains with a varying grid size at a temporal time step of 6 minutes. The CAMRa radar data set is at approximately a 2 minute resolution upon a grid size of approximately 330 m. The NWP UM data set was taken from the first 10 days of May 2003 at 1 hour time step on a 12 km grid.

(a) NWP - MM5 system

(b) Radar - CAMRa

(c) NWP - UM system
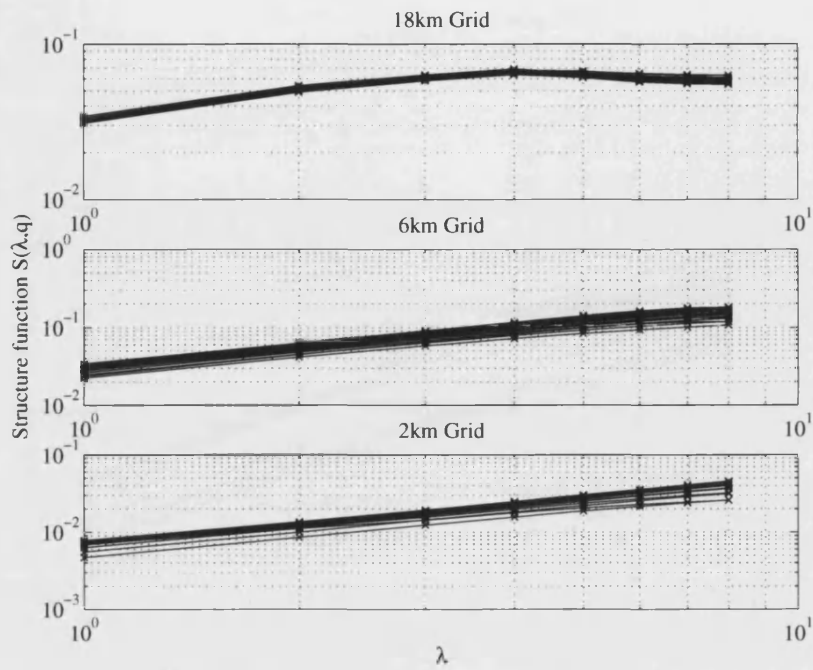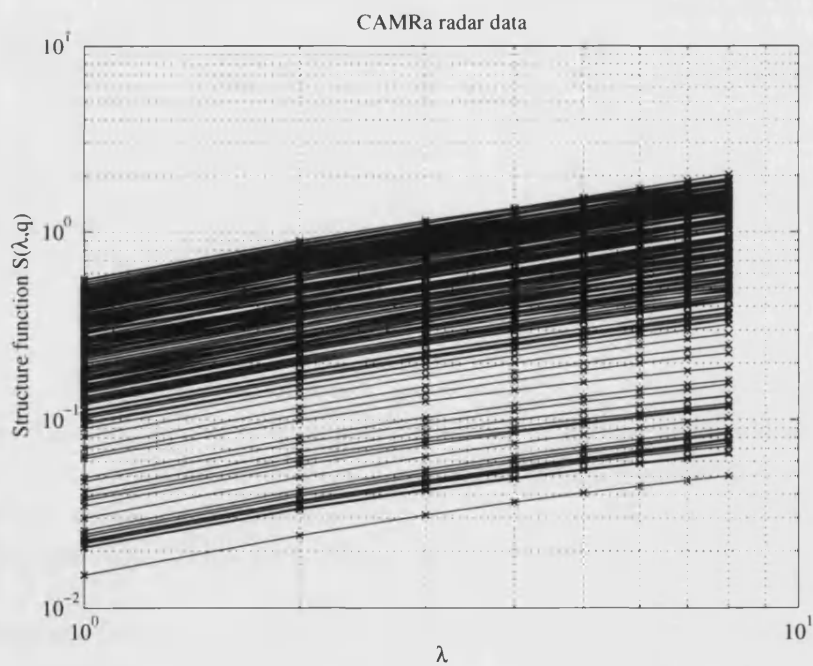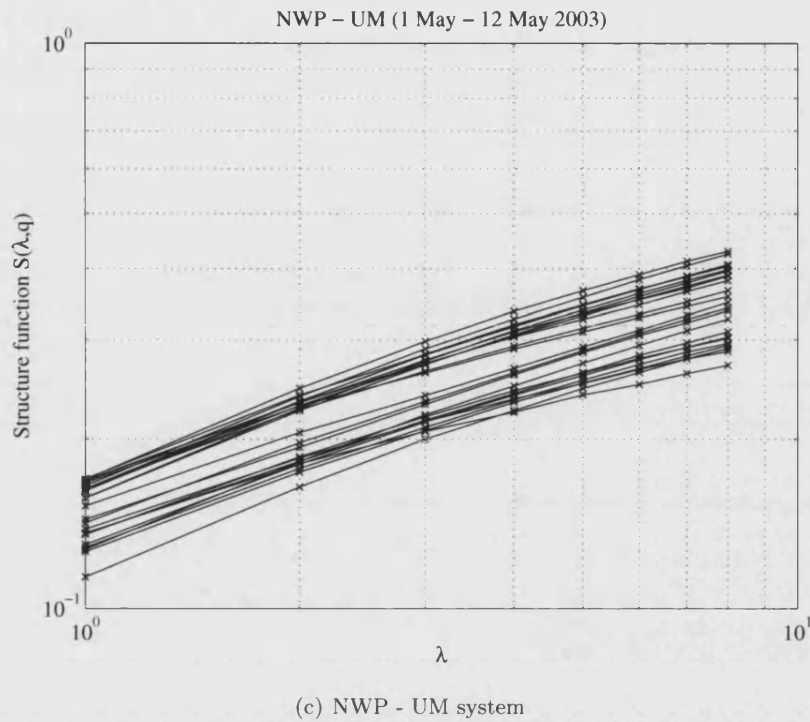
Figure 5.3: The average structure functions, from three different data sets.

If a field demonstrates scaling behaviour then as $q \to 0$ [Ferraris et al., 2003]:

$$S \propto \lambda^{\zeta(q)} \tag{5.5}$$

where ζ is the scaling exponent and $\zeta(q) = qH$, where H is typically called the Hurst coefficient.

By reference to figure 5.3, it can be seen that the first-order generalised structure function of the MM5 2 km run exhibits fractal scaling over at least eight grid scales, whilst the 6 km grid begins to show some deviation at around eight grid scales. The 18 km grid scale demonstrates fractal scaling only over the first three grid scales. This indicates that the rain phenomena in these two data sets exhibit fractal scaling up to about $6 \times 8 = 48$ km and $3 \times 18 = 54$ km respectively.

These MM5 runs are from the same time period, however since the smaller domains are nested within the larger domains the coverage areas are slightly reduced; if these areas were identical it might be expected that the structure functions would appear to be overlaid. However this is not entirely true. When using NWP modelling the decrease in grid size does not only allow more features to be visible it also improves the internal

97

modelling of the rain phenomena. These effects in combination improve the 'quality' of the rain modelling. Conversely, if the radar could be increased in resolution it can be expected that the structure functions would be overlaid. In this case the increased resolution of the measurement technique allows more small-scale structure to become visible, as opposed to significantly influencing the modelling of the rain phenomena.

The CAMRa data set is provided on a significantly smaller grid scale than any of the other data sets. As the domain is smaller, for the events selected there is almost complete rain coverage. This is unique in the datasets (and indeed very few measurement techniques could produce such a complete rain field at this spatio-temporal resolution). However, the data still exhibit fractal scaling over eight grid scales.

Structure functions for the UM NWP model is shown in figure 5.3. It can be seen that almost all fields exhibit fractal scaling extending over eight grid scales.

The aim of fractal downscaling is to exploit the linearity of the structure function by adding a stochastic small-scale structure which continues the linear relationship between the log-structure function and the log-scale as in figure 5.4.
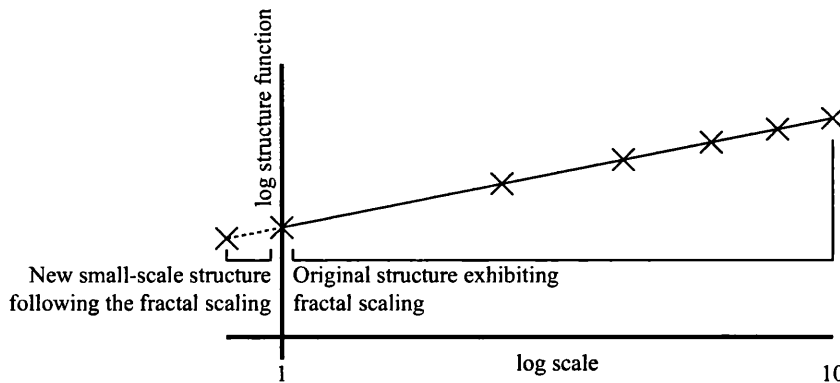
Figure 5.4: The use of the structure function in fractal scaling.

Hence the new small-scale structure is a fractal since it exhibits a scaled version of the large structure. In general it is also desirable to ensure that the new rain field has the same probability density as the original rain field. The fractal dimension should also be conserved since the straight line shown in figure 5.1(b) should be continued over the new scales.

98

5.2 Multifractal bounded cascades

The downscaling technique presented here can be explained by taking the rainfall rate field on a discrete (x, y) grid as $R_0(x_0, y_0)$. At the first stage of the downscale each pixel is split into four pixels: $R_1(x_{11}, y_{11})$, $R_1(x_{12}, y_{11})$, $R_1(x_{11}, y_{12})$ and $R_1(x_{12}, y_{12})$, as shown in figure 5.5. Each of these new pixels is then given the value:

$$R_1(x_{1(1,2)}, y_{1(1,2)}) = wR_0(x_0, y_0) \tag{5.6}$$

where each value of w is the independent realisation of a random variate defined by the generator function, W. This results in the four pixels each perturbed by a small amount from their original value thereby adding small-scale structure. The process is relatively simple however the derivation of the generator function, W, is not. The
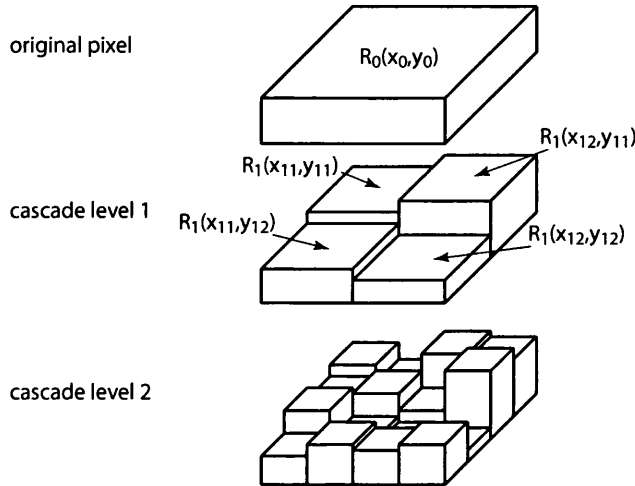
Figure 5.5: The downscaling procedure.

generator function should depend explicitly upon the cascade level and should decrease with increasing spatial resolution. This has the effect of reducing the power in the higher frequencies. For this procedure the log-normal generator has been chosen as it is easy to evaluate the log-normal parameters. The log-Poisson distribution has also been used but this requires the solution of a minimisation problem further complicating the cascade process Deidda [1999].

Following the procedure used by Menabde et al. [1999] and Ferraris et al. [2003] the generator function for cascade level n can be defined as:

$$W_n = C_n \exp(\sigma_n X) \tag{5.7}$$

where, C_n is a normalisation constant, σ_n is the scale parameter that decreases with

the cascade level and $X = \text{Norm}(0,1)$ is a standard normal random variate with zero mean and unit variance.

The normalisation constant can be eliminated by requiring that the mean of the downscaled field be the same as the mean of the original field. The moments of the lognormal distribution are given by $\langle \exp(tX) \rangle = \exp(t^2)$. With this simplification equation 5.7 becomes:

$$W_n = \exp(-\sigma_n^2 + \sigma_n X) \tag{5.8}$$

As discussed previously, the scale parameter σ_n should decrease as the cascade increases. Menabde et al. [1999] suggest the use of an exponentially decreasing scale parameter:

$$\sigma_n = \sigma_0 2^{-nH} \tag{5.9}$$

These simplifications mean that the generator function can now be written in terms of two parameters, σ_0 and H. The parameter H is the 'Hurst coefficient' (see equation 5.5).

Normal random distributions demonstrate stability under addition, i.e.:

$$\sum_{i=1}^{n} [\sigma_i X_i] = \sqrt{\sum_{i=1}^{n} [\sigma_i^2]} \cdot X \tag{5.10}$$

Hence, it can be seen that after N cascades the downscaled rainfield will be a lognormally distributed random variable $\exp(\sigma^2 - \sigma X)$. The final scale parameter can be evaluated using the stability property in equation 5.10:

$$\sigma^2 = \sum_{n=1}^{N} \sigma_n^2$$
$$= \sum_{n=1}^{N} \sigma_0^2 4^{-nH} \tag{5.11}$$

This is a finite geometric sum of the form:

$$S_k = \sum_{k=0}^{K} ar^k = \frac{a(1 - r^K)}{1 - r} \tag{5.12}$$

and hence:

$$\sigma^2 = \sigma_0^2 \frac{1 - 4^{-(N+1)H}}{1 - 4^{-H}} \tag{5.13}$$

By requiring that the scale parameter of the final downscaled field, σ, is the same as the initial field and rearranging equation 5.13 it is possible to extract the final parameter, σ_0.

The example field from figure 5.1(b) is downscaled in order to provide an example field. The original field is shown in Figure 5.6(a) whilst the downscaled field is shown in figure 5.6(b). This field has been downscaled by a factor of four (i.e. two cascades) from a grid size of 18 km to one of 4.5 km.
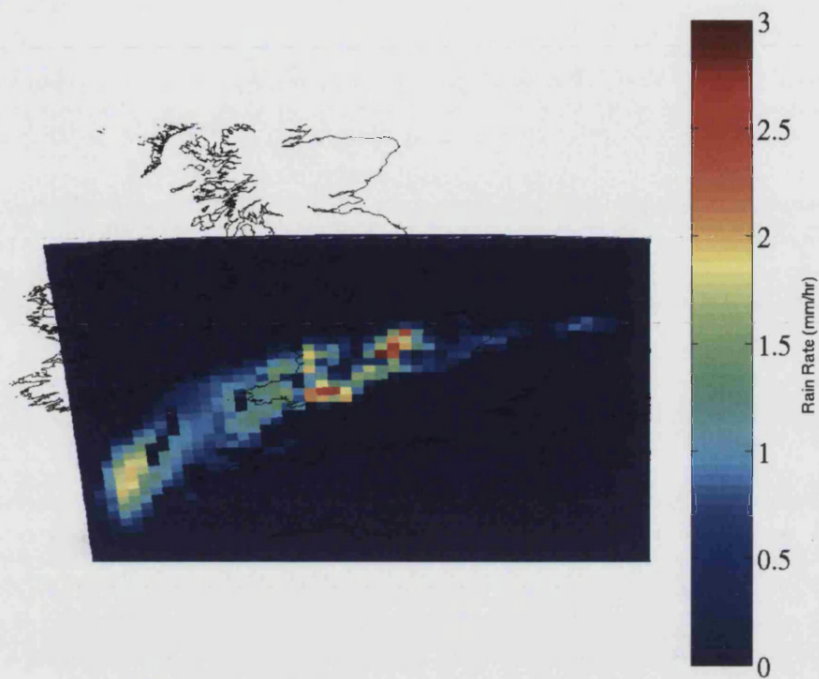
(a) Original field
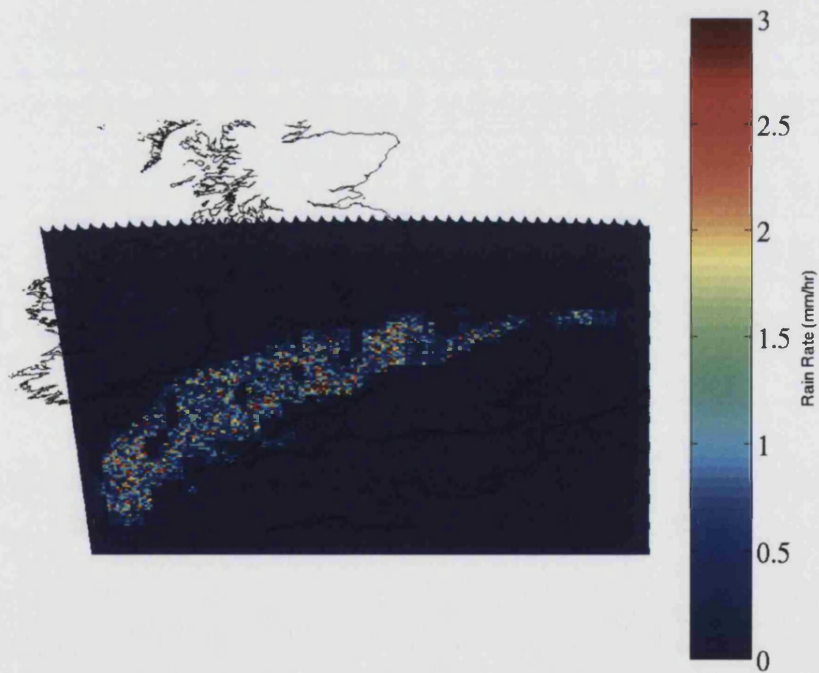
(b) Downscaled field

Figure 5.6: An example MM5 rain field downscaled (22/10/2004 12:00UT) on a grid scale of 18 km.

5.2.1 Filtering

Whilst the multifractal cascade approach in section 5.2 has the correct statistical distribution, the data will not necessarily have the expected frequency response. The easiest way to understand this problem is to consider the autocorrelation function. The autocorrelation can be considered to be directly related to the Fourier transform by the Wiener-Khinchin theorem, which states that the auto-correlation, C, is the Fourier transform of the absolute square of a time-series, E:

$$C = F\left\{|E|^2\right\} \tag{5.14}$$

When the autocorrelation is considered on scales in excess of the original grid scale the autocorrelation is defined by the initial rainfall field. However, when the field is examined as a whole there is also a contribution to the correlation at that grid scale between the offspring of adjacent pixels. In essence as each pixel is downscaled individually with no regard for the surrounding original pixels this results in the fields displaying the correct power spectral density / autocorrelation within a grid square, but the incorrect power spectral density / autocorrelation when the image is considered as a whole, see figure 5.7.
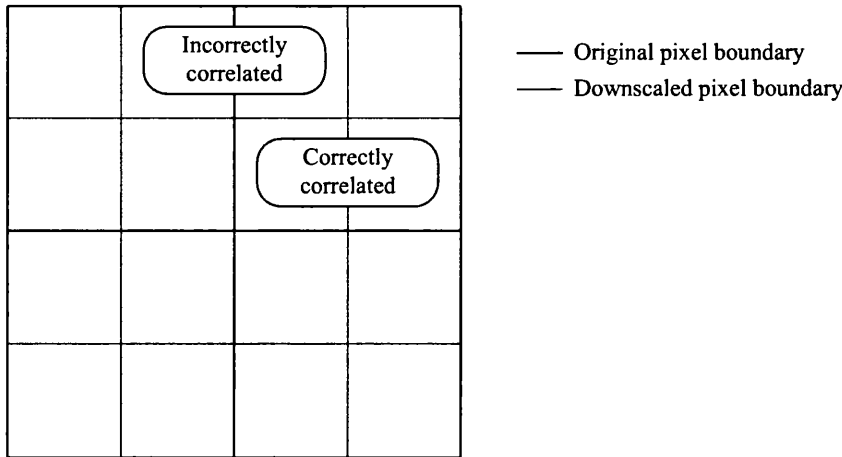
Figure 5.7: The correlation errors from a discrete cascade.

This is visible in the fields in figure 5.6(b) as a 'blocking' effect. The ensemble power spectral density of figure 5.6(b) is shown in figure 5.8. With reference to the downscaled field, it demonstrates a flatter power spectrum at frequencies around 1.5. This indicates the incorrect power spectra and poor correlation at these frequencies.
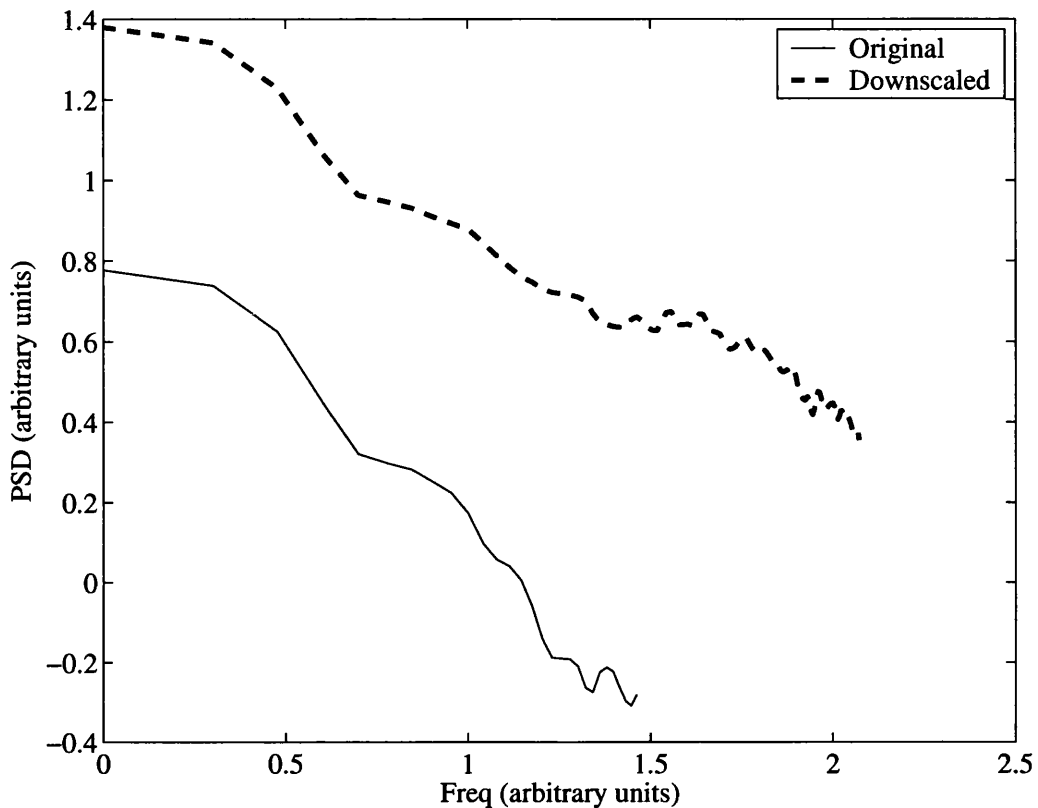
Figure 5.8: The ensemble power spectral density of the downscaled and original field (22/10/2004 12:00UT) on a grid scale 18 km.

One way to improve the autocorrelation / power spectral density for the rain field is to use a spatial filter. The aim is to extend the straight line of the ensemble power spectra as seen for example in figure 5.1(b). The slope of the ensemble power spectra of the final field should be the same as that of original field (and hence a function of the fractal dimension). The cut-off frequency of the low-pass filter can be assumed to be the scale at which the stochastic structure has been added. The cut-off frequency is, in essence, a function of the final cascade level, N. The power spectral density from this process is shown in Figure 5.9, along with a thick line that represents the gradient of the filter roll-off.
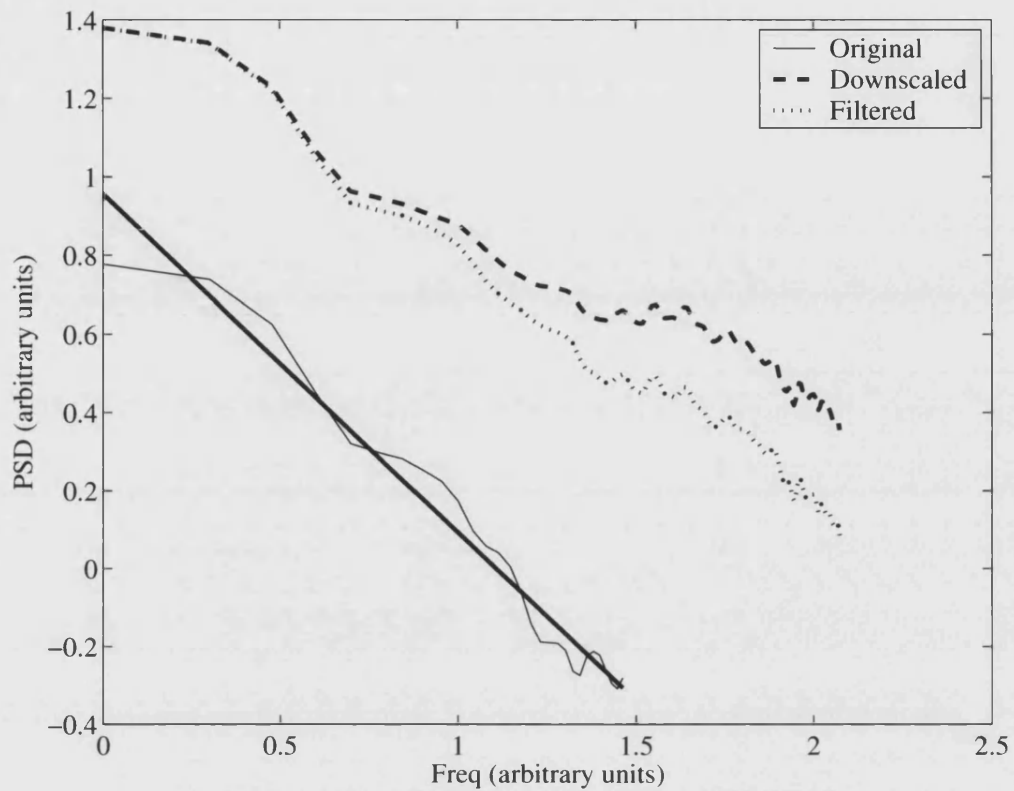
Figure 5.9: The ensemble power spectral density of the filtered, downscaled and original field (22/10/2004 12:00UT) on a grid scale 18 km.

This results in a final downscaled field that looks like figure 5.10(b). It can be seen that the small-scale structure has been added yet the downscaled field does not exhibit excessive 'blocking' effects and displays characteristics similar to those seen in higher resolution data sets.
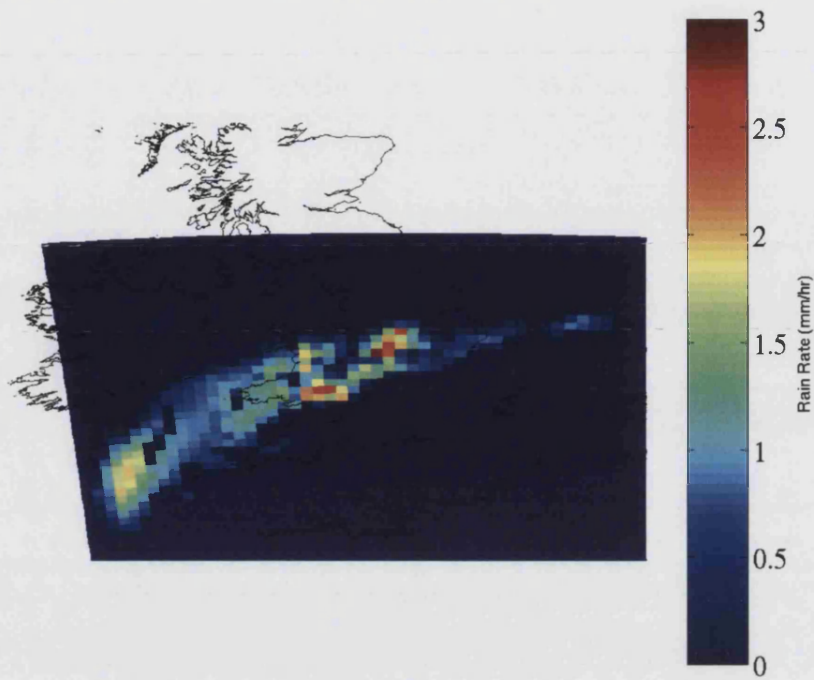
(a) Original field
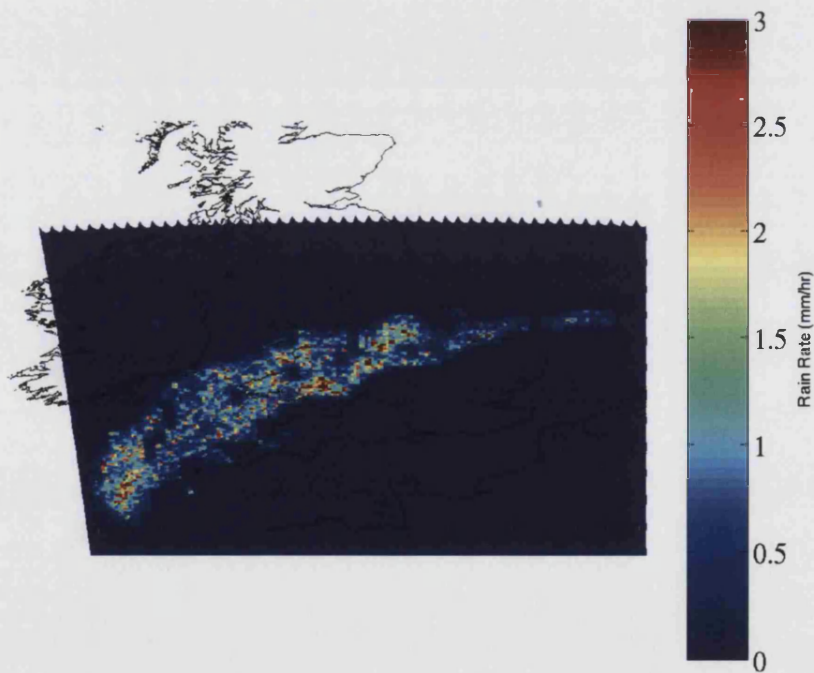
(b) Downscaled and filtered field

Figure 5.10: An example MM5 rain field downscaled and filtered (22/10/2004 12:00 UT) with an original grid scale of 18 km and a downscaled grid scale of 4 km.

The procedure described here is a spatial downscaling procedure, since the stochastic small scale is distributed spatially. The stochastic structure is not temporally correlated, so this procedure can be used for simulations that require instantaneous maps (i.e. terrestrial link load balancing). However, caution should be used if these maps are used for simulations that require temporal evolution of the rain field. This caution should be extended to the case where an interpolation scheme is used to 'join' two downscaled rain fields, since the structure will not evolve correctly.

The most appropriate way to use these downscaled rainfall fields is to use these as instantaneous rain fields at the sample period. That is to say, the rainfall field is undefined between the samples. In addition the temporal scale should be such that the small scale structure can be considered temporally uncorrelated.

This problem can be solved by downscaling within a spatio-temporal framework, i.e. the downscaling procedure provides stochastic small-scale structure which shows the correct evolution in time. This is a very complex procedure since spatio-temporal rainfall fields (i.e. three-dimensional fields) typically exhibit multi-affine scaling [Deidda, 1999]. Multi-affine scaling describes an object that exhibits fractal scaling in all dissensions but by different scaling exponents. In essence the temporal dimension follows a different scaling exponent to the spatial dimensions.

5.2.2   Conservation of statistics

In order for the downscaled rainfall field to be used in place of the original field it is important that certain characteristics of the original rainfall field are conserved. The probability density function should also be conserved. After the filtering, the rainfall field should exhibit the same β component and hence the same fractal dimension as the original field.

By comparing the means and standard deviations of the rainfall fields before and after downscaling (by a factor of 4) it can be confirmed that the downscaling procedure conserves the mean (figure 5.11) and standard deviation (figure 5.12) of the example rainfields[2] (the thin line represents a guide line for a perfect 1:1 fit).

---

[2]Only the two smaller scale MM5 rainfall fields were used since they best demonstrate scaling behaviour
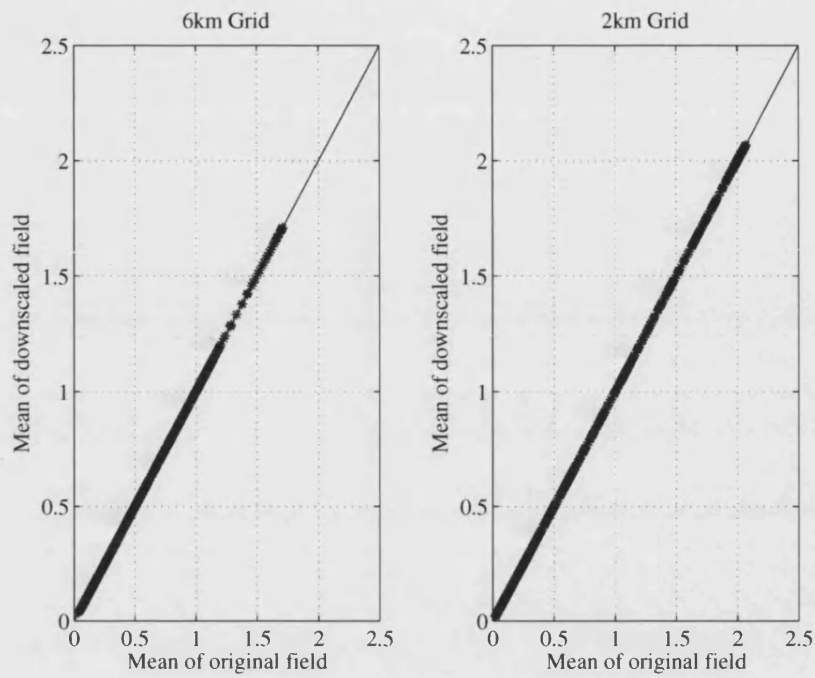
(a) NWP - MM5 system $(\mathrm{mm\,hr^{-1}})$
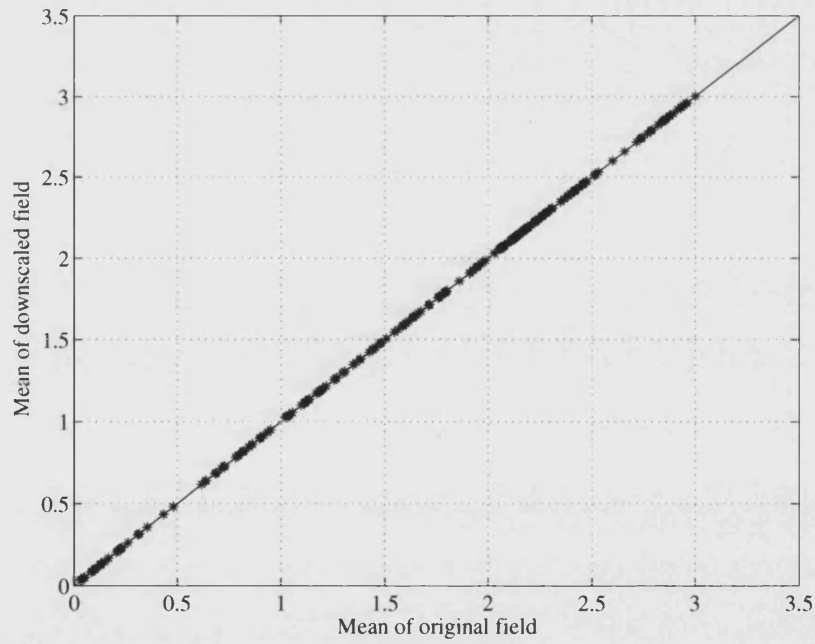
(b) Radar - CAMRa $(\mathrm{mm\,hr^{-1}})$

Figure 5.11: The conservation of the mean, from the three different data sets.

As can be seen there is a slight tendency for the standard deviation of the new field to

108
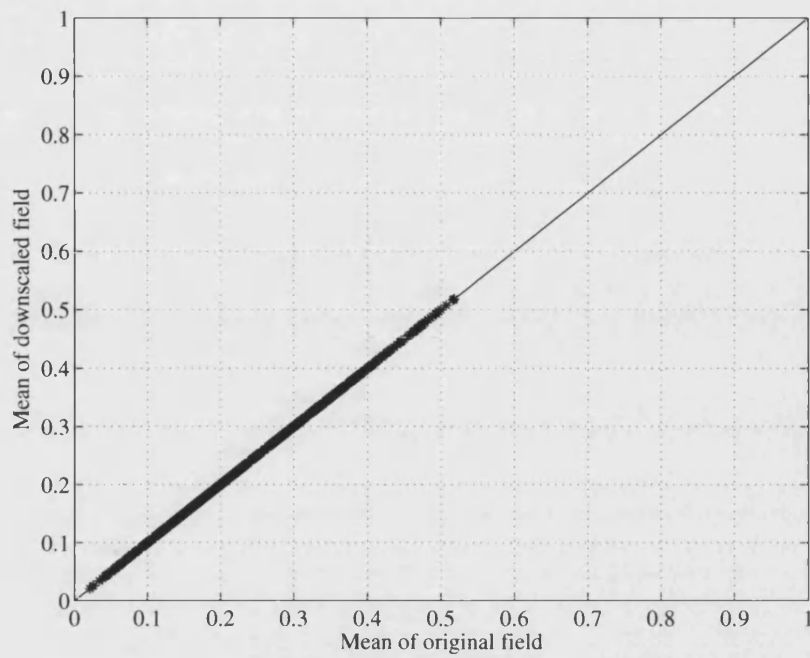
(c) NWP - UM system $(\mathrm{mm\,hr^{-1}})$

Figure 5.11: The conservation of the mean, from the three different data sets.

be slightly higher than the original field. This is especially true for fields that display a higher initial standard deviation. Even with this caveat it can be seen the statistics of the rain field are conserved well.
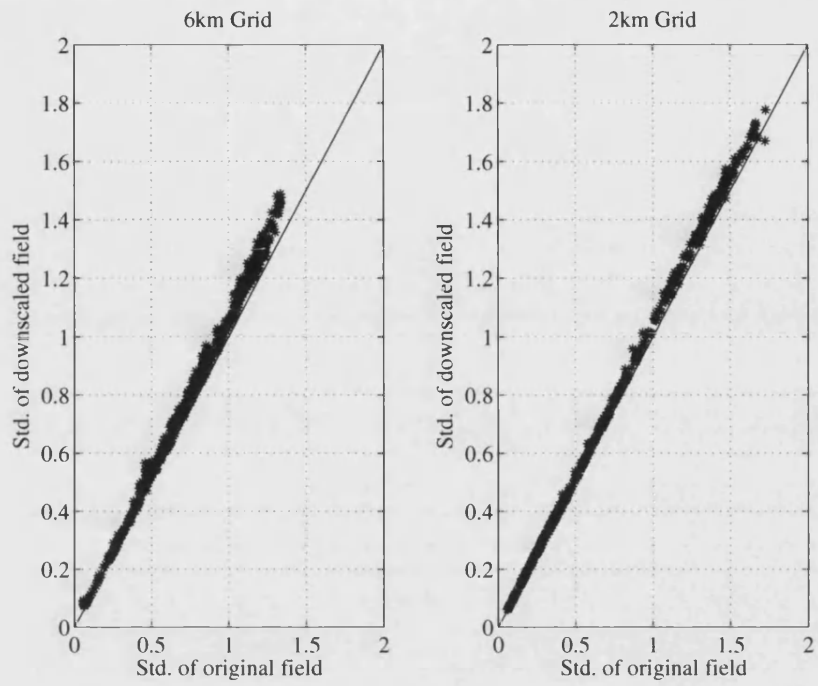
(a) NWP - MM5 system $(\mathrm{mm\,hr^{-1}})$
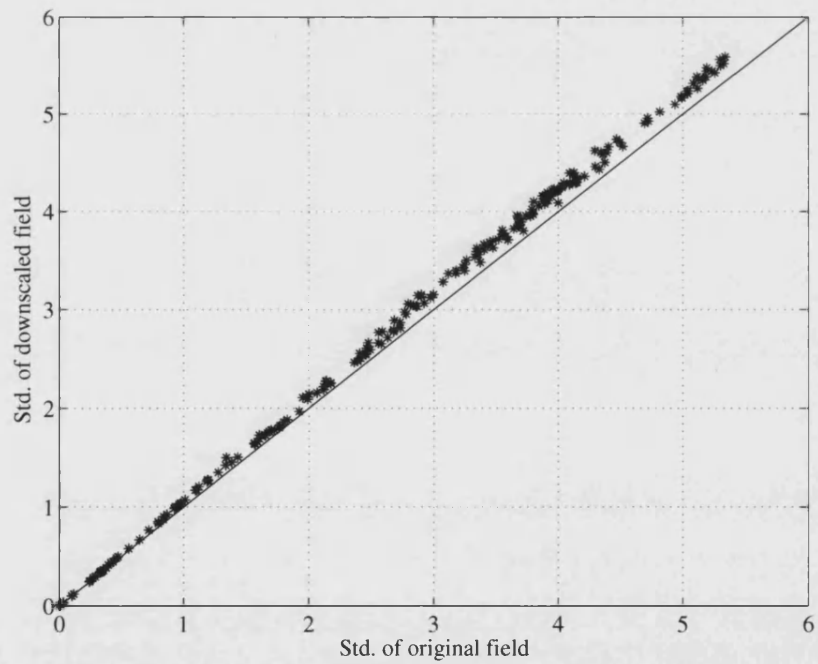
(b) Radar - CAMRa $(\mathrm{mm\,hr^{-1}})$

Figure 5.12: The conservation of the standard deviation, from the three different data sets.

(c) NWP - UM system $(\mathrm{mm\,hr^{-1}})$
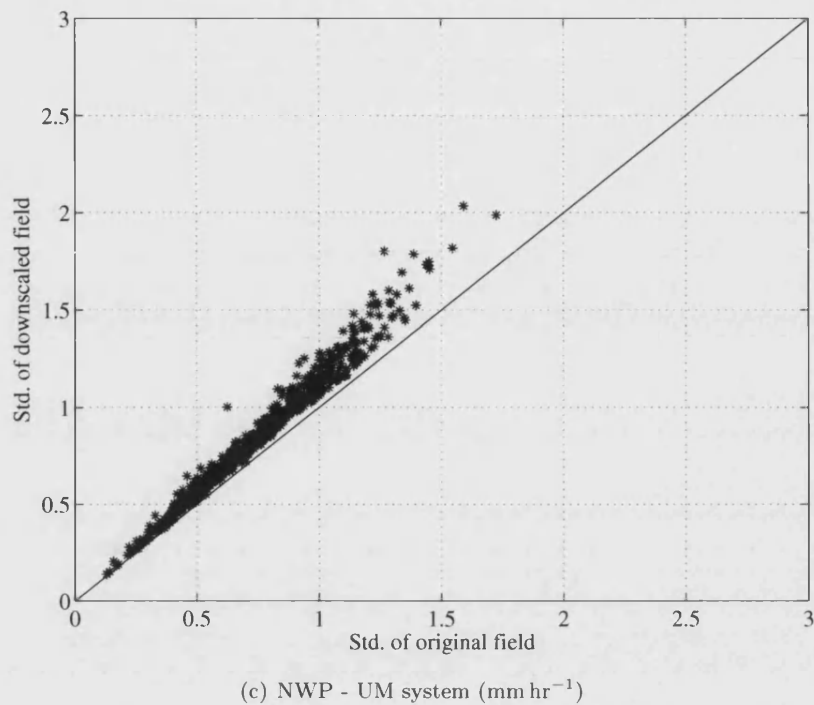
Figure 5.12: The conservation of the standard deviation, from the three different data sets.

This means that this multifractal downscaling scheme is suitable for use in order to produce higher resolution rain fields for the propagation forecast engine.

5.3 Monofractal interpolation

For some applications it is more convenient to extract the rainfall rate time series at a larger grid scale and then perform a temporal interpolation scheme which adds the small scale structure. Caution should now be employed in the spatial domain since each time series will exhibit a stochastic small structure that will not be properly correlated in space.

The procedure presented here is a mono-fractal interpolation scheme. There is significant evidence to suggest that time series of log-rainfall-rate exhibit mono-fractal scaling properties [Paulson, 2004]. Mono-fractal scaling allows us to use a much simpler interpolation scheme.

The scheme can be conceptually explained by considering a series of equally spaced samples. If new samples are introduced halfway between two original samples using a smooth interpolation scheme there is now a frequency component at the higher frequency however it contains little or no power. If we consider the same situation but this time the new sample consists of a sample from a normally distributed random variate summed with the contribution from the linear interpolation technique. If the random variate has a variance of σ^2 at the first downscale level and the variance is reduced geometrically, it is apparent the power spectral density of the time series will exhibit a power law.

Expressed formally, it can be considered that the log-rainfall-rate time series, X,

$$X = \{X_i; i = 0, \ldots, N\} \tag{5.15}$$

is interpolated with the small scale log rainfall rate time series, Y,

$$Y = \{Y_i; i = 1, \ldots, N\} \tag{5.16}$$

to produce the time series, Z, such that:

$$Z = \{X_0, Y_1, X_1, Y_2, \ldots, Y_N, X_N\} \tag{5.17}$$

As discussed previously the value of Y should be a sample from a zero-mean normally distributed random variate summed with the linear interpolation between the two relevant points such that:

$$Y_i = \frac{1}{2}(X_{i-1} + X_i) + \zeta \tag{5.18}$$

where $\zeta = \text{Norm}(0, \sigma_0^2)$

This procedure will conserve the mean, i.e. $E(X) = E(Y) = E(Z)$, and the variance is conserved if [Paulson, 2004]:

$$\sigma_0^2 = E(X^2)\frac{1}{2}[1 - \rho(1)] \tag{5.19}$$

Where $\rho(1)$ is the autocorrelation at unit lag of the time series, X. The procedure can then be iterated with the ζ term evaluated from a normally distributed random variate with a variance that geometrically decreases:

$$\sigma_{i+1}^2 = \frac{\sigma_i^2}{2} \tag{5.20}$$

It is important to also appreciate that since this is an additive (not multiplicative) cascade it should only be applied to individual rain events on an event-by-event basis, not to a rainfall rate time series as a whole. An example rainfall rate time series is

shown in figure 5.13. The time series has been linearly interpolated and interpolated using the monofractal interpolation technique, the new stochastic structure is clearly evident in the time series.
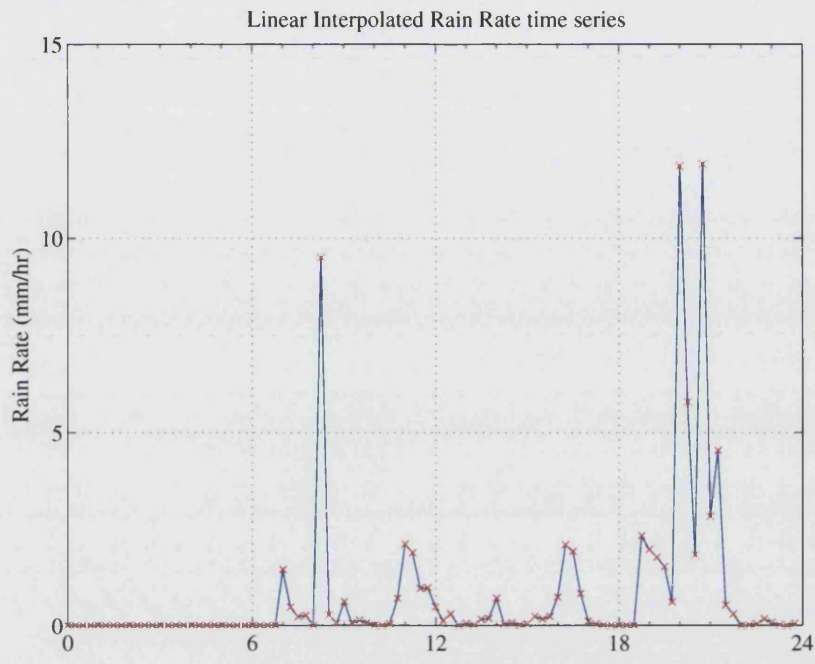
(a) Linear interpolation
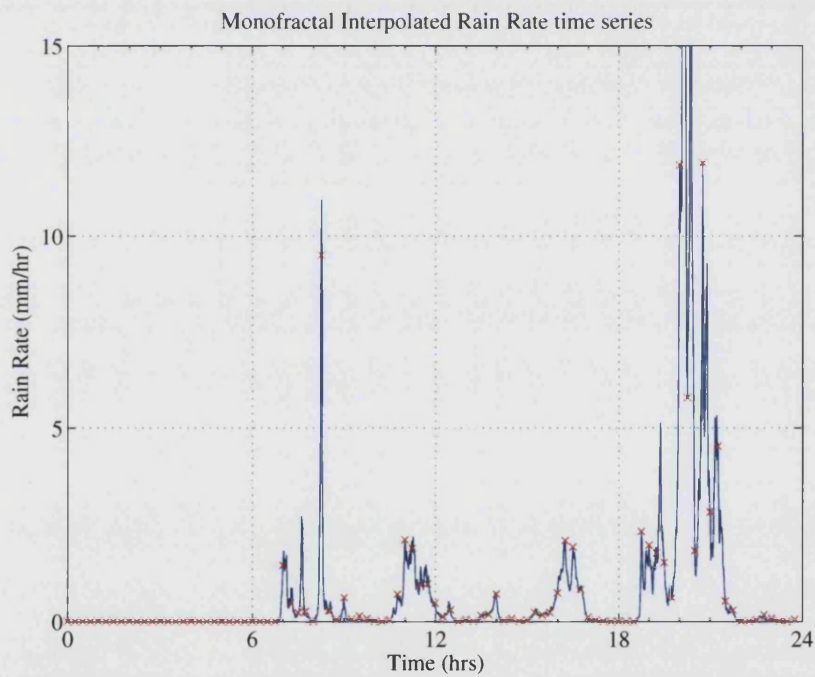
(b) Monofractal interpolation

Figure 5.13: Monofractal interpolation on an radar derived rainfall rate time series 22/06/2003 - red crosses denote original radar scans.

114

Linear Interpolated Rain Rate time series
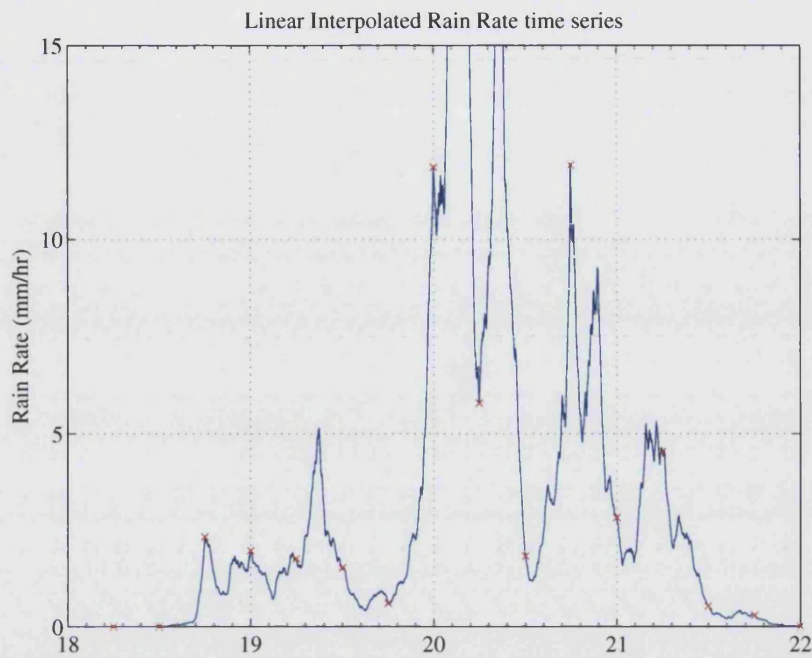
(c) Detail of the monofractal approach

Figure 5.13: Monofractal interpolation on an radar derived rainfall rate time series 22/06/2003 - red crosses denote original radar scans.

The power spectral density of the rainy part of the time series from figure 5.13 is shown in figure 5.14. The linearly interpolated time series displays significantly reduced power in the high frequency components as well as the characteristic harmonic pattern from the interpolation scheme[3].

---

[3]It should be noted that since the two time series are not of the same length the PSD of the initial time series (with a temporal sampling time of 900 seconds) has been scaled by adding $20 \log_{10}(900)$, in order to overlay the curves.
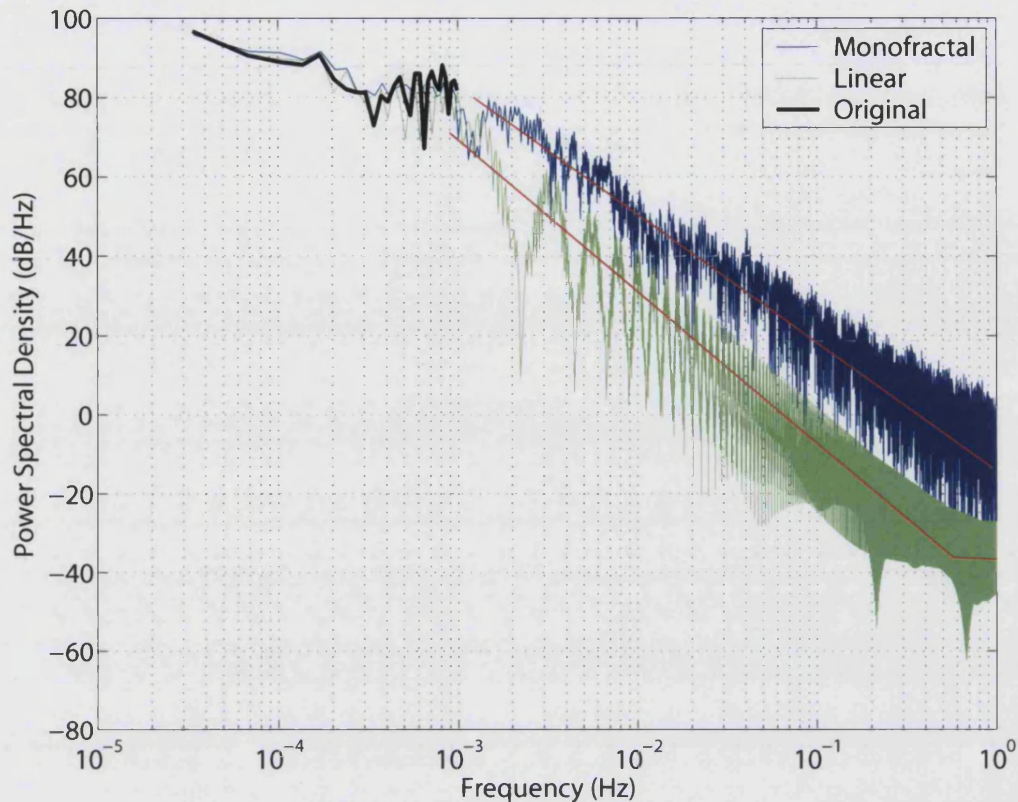
Figure 5.14: The effect of interpolation upon the PSD of rainfall rate time series.

As can be seen this monofractal interpolation scheme produces a stochastic small scale structure with a realistic power spectral density whilst conserving the statistics of the original time series. Using these techniques it is possible to produce higher resolution rainfall rate time series which are suitable for use with the propagation forecast engine.

In this chapter techniques for improving the spatial or temporal resolution of rainfall rate estimates have been presented. A multi-fractal spatial downscaling procedure has been demonstrated which conserves the probability density of the rainfall field. In addition the procedure has been shown to conserve the ensemble power spectral density. A mono-fractal temporal interpolation scheme has also been demonstrated, this interpolation scheme allows individual rain-events to be downscaled.

These downscaling techniques ensure that the rainfall rate estimates exhibit the correct power spectral density. A correct power spectral density ensures that the time series of attenuation exhibit the correct dynamic characteristics such as power spectral density and fade slope. It should be noted that since the small-scale structure that is introduced

116

is stochastic these time series should not be used for forecast systems.

Chapter 6

Verification of PFE generated time series

The Propagation Forecast Engine (PFE) is capable of being reconfigured to produce time series for many different link geometries. The aim of this Chapter is to demonstrate some example link attenuation time series and attenuation maps. In this Chapter we also consider verification of the generated time series.

6.1 Example PFE outputs

The simplest outputs from the time series generator are multiple-site time series. These time series are of interest for network simulation since an unlimited number of spatially consistent, time-coincident time series can be created. A set of eight example time-coincident time series are presented in figure 6.1. The meteorological information was taken from the real time MM5 forecasts from the model run at the University of Bath[1]. The time series are created at second intervals using monofractal interpolation to downscale the rainfall rate time series. Each time series is for a 50 GHz circularly polarised downlink from a geostationary payload at 2°West.

---

[1]To reiterate the data is provided at a 15 minute resolution and the 18 km grid was used for this example in order to provide a greater coverage area, which is shown as the dark thick line in figure 6.1(a)
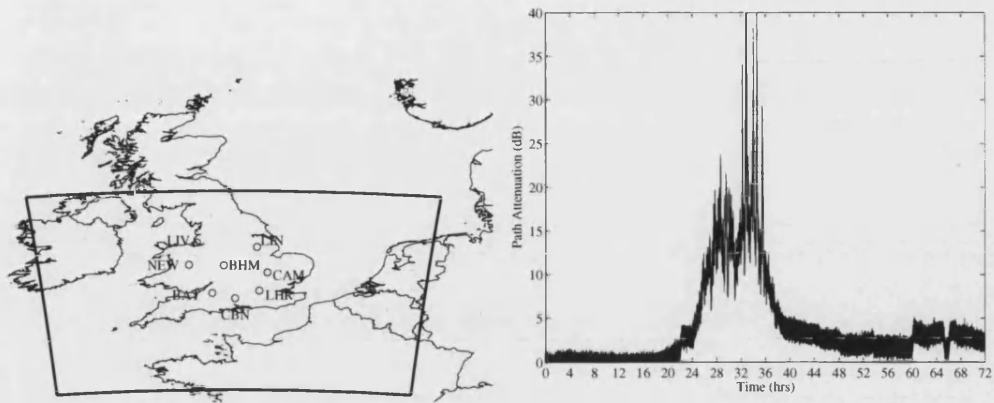
(a) Geographical locations of the generated time series
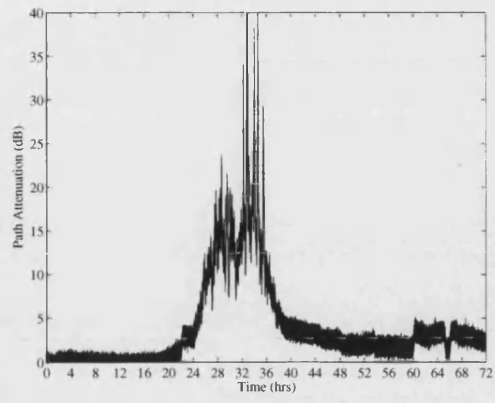
(b) Liverpool (LIV)
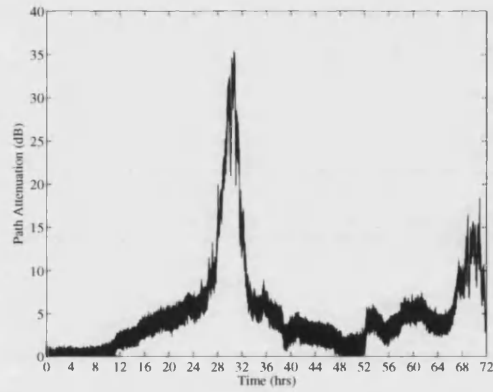
(c) Lincoln (LIN)
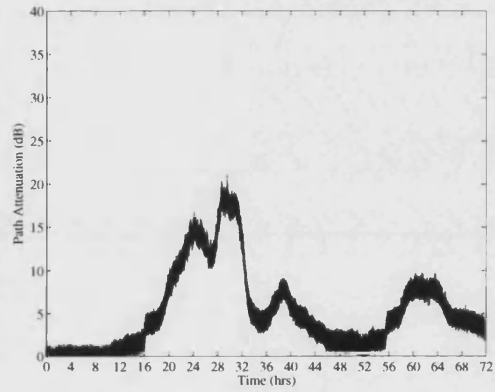
(d) Newtown (NEW)
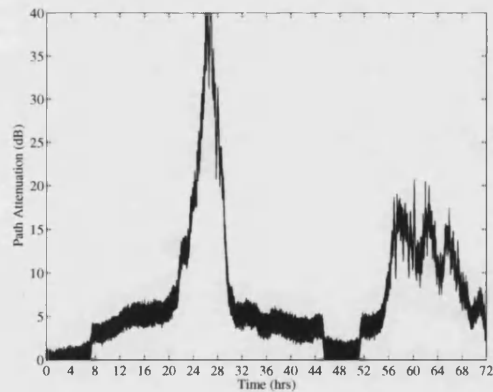
(e) Birmingham (BHM)
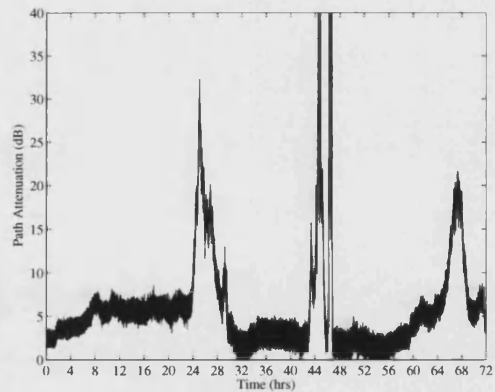
(f) Cambridge (CAM)

Figure 6.1: Example multiple site time-coincident time series, 27/07/2005 00:00UT to 30/07/2005 00:00UT.

(g) Bath (BAT)

(h) Chilbolton (CBN)

(i) Heathrow (LHR)
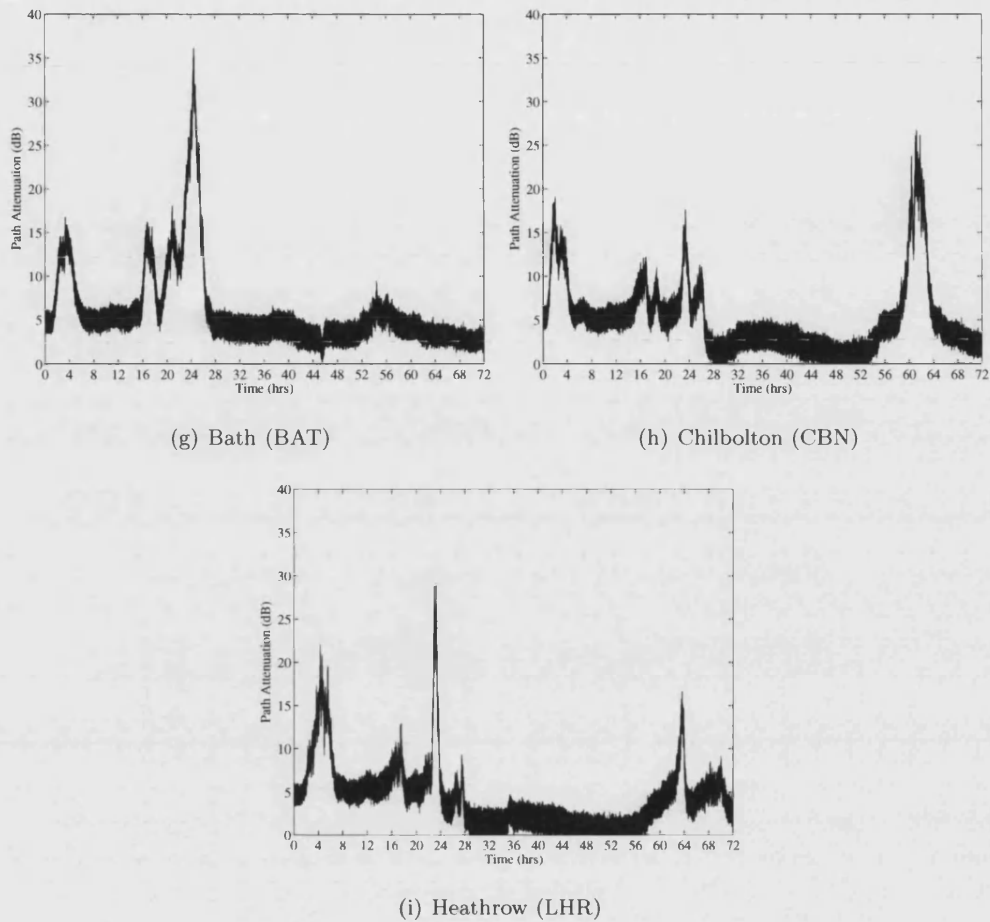
Figure 6.1: Example multiple site time-coincident time series, 27/07/2005 00:00UT to 30/07/2005 00:00UT.

These time series demonstrate the 'slow-fading' effects from rain, clouds and gas. The main rain event enters the simulation area after the end of the first day of simulations, the three sites along the southern edge (Bath, Chilbolton and Heathrow) show significant cloud fading.

These time series can be extended to create 'maps' of attenuation that provide an image of the instantaneous link attenuation across a coverage area, for example a satellite spot beam footprint. The examples in figure 6.2 show an attenuation map for a 50 GHz link to a transponder at 2°West. The first subfigure (a) represents the largest scale domain and the black square represents the coverage of the nested domain subfigure (b), in this case the first sub-domain. This sub-domain has a smaller coverage area at 6 km grid scale. The black square on this image represents the coverage area of the nested domain of this subfigure (c). This smallest domain has a grid scale of 2 km.

120
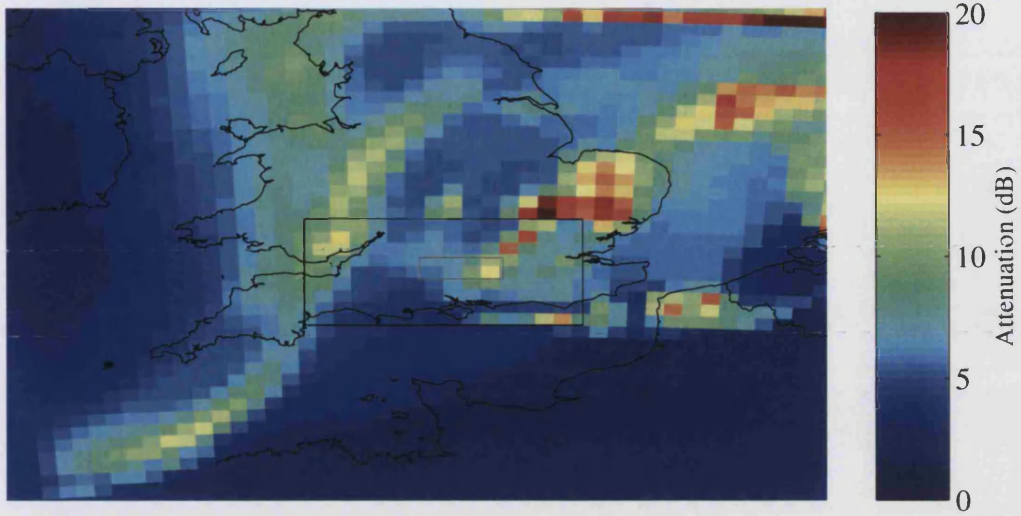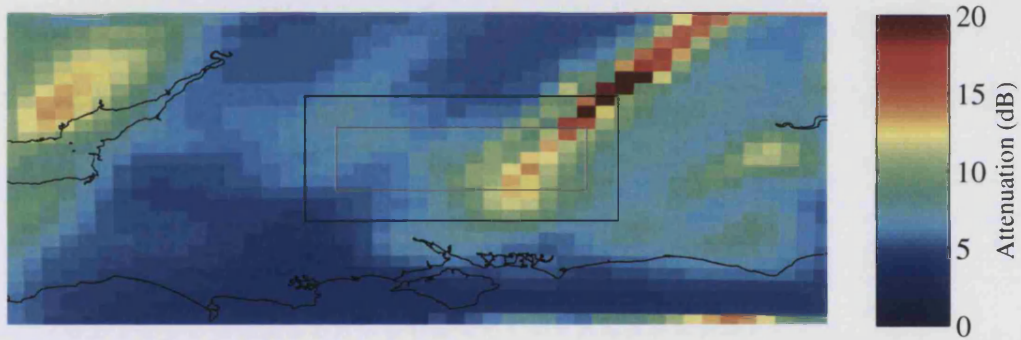
(a) 18 km grid
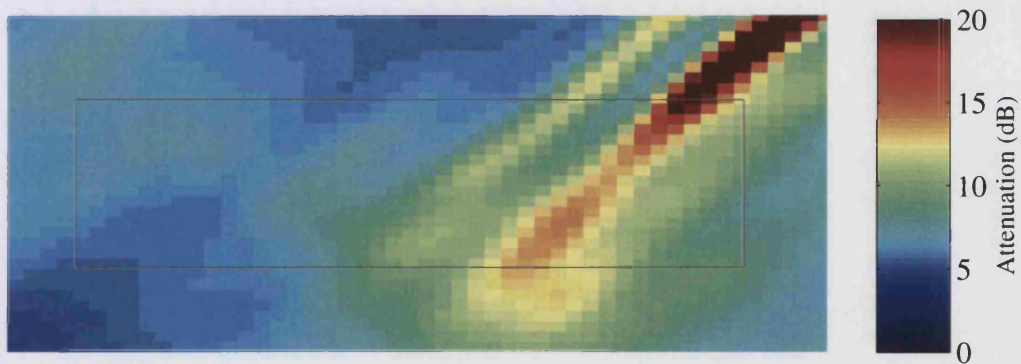
(b) 6 km grid

(c) 2 km grid

Figure 6.2: Example attenuation map on an 18 km grid, 6 km grid and 2 km grid, 50 GHz link to 2°West, 20/10/2004 - 15:50 UT.

The northern edge of the largest coverage domain, shown in subfigure (a), demonstrates some edge effects from the NWP system. The boundary conditions for this large domain only include the mix-ratio for water vapour, yet within the domain the physics derives the mix-ratios for many other phases (e.g. cloud-water, cloud-ice, rain-water, graupel etc.). This causes a discontinuity around the edges of the outer domain and the only action the model can take at these points is to drop the moisture as rain. Caution should always be exhibited around the edges of domains and using data from the eight grid points of the boundary edges is generally not advised.

In order to examine the effect of the grid scale a close up of the area shown in grey in figure 6.2 is shown in figure 6.3 for the three respective grid scales.

(a) 18 km grid
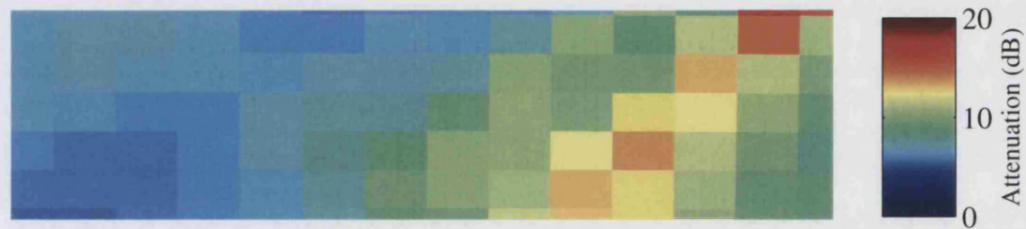
(b) 6 km grid
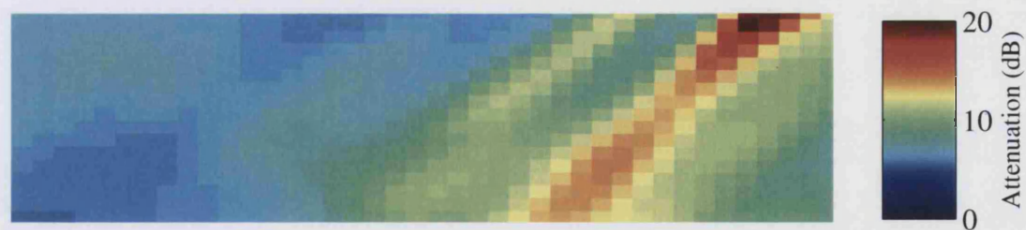
(c) 2 km grid

Figure 6.3: Close-up of example attenuation maps on the 18 km, 6 km and 2 km grids, 50 GHz link to 2°West, 20/10/2004 - 15:50 UT.

The variation in grid size demonstrated in figure 6.2 show the significant effects of reducing grid scales. As the grid scale decreases the line of attenuation becomes more defined. On the smallest scale it even becomes visible as two distinct bands such is the improvement in the quality of the forecast. This improvement in forecast 'quality' and the reduction in spatial averaging results in an increase in modelled attenuation of up to 10 dB.

In addition, it is possible to reconfigure the link geometry to provide time series for terrestrial point-to-point links such as those used in terrestrial overlay or backhaul networks. The main focus of this work has been slant-path links. The application in terrestrial links has been a secondary consideration. Many problems still remain for propagation forecasting applied to terrestrial links. However the propagation forecast engine in the current state can provide a first-order estimate of the attenuation on such links.

To demonstrate the ability of the forecast engine to forecast link attenuation on terrestrial link we consider an experimental link that has been operating at the University of Bath. The TSaR group at the University of Bath operates a 28 km long 32 GHz link for interferometric ray-bending experiments. The transmitter is situated at Blagdon (51.3275, -2.7192) and the receiver is situated at the University of Bath. An example is taken from the MM5 model on a 6 km grid, the meteorological conditions along the slant path are shown in figure 6.4.
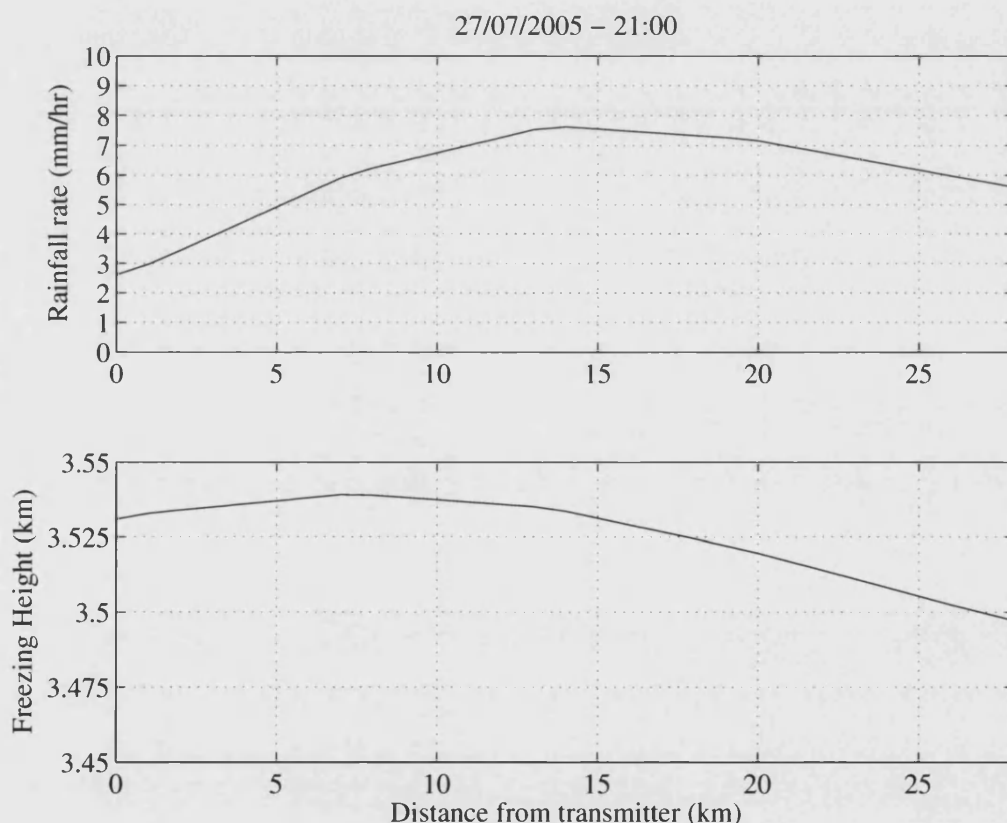
Figure 6.4: Meteorological parameters along the example terrestrial path (Blagdon to Bath).

124

The meteorological conditions are then used to estimate a specific attenuation using the propagation forecasting algorithm (figure 6.5).
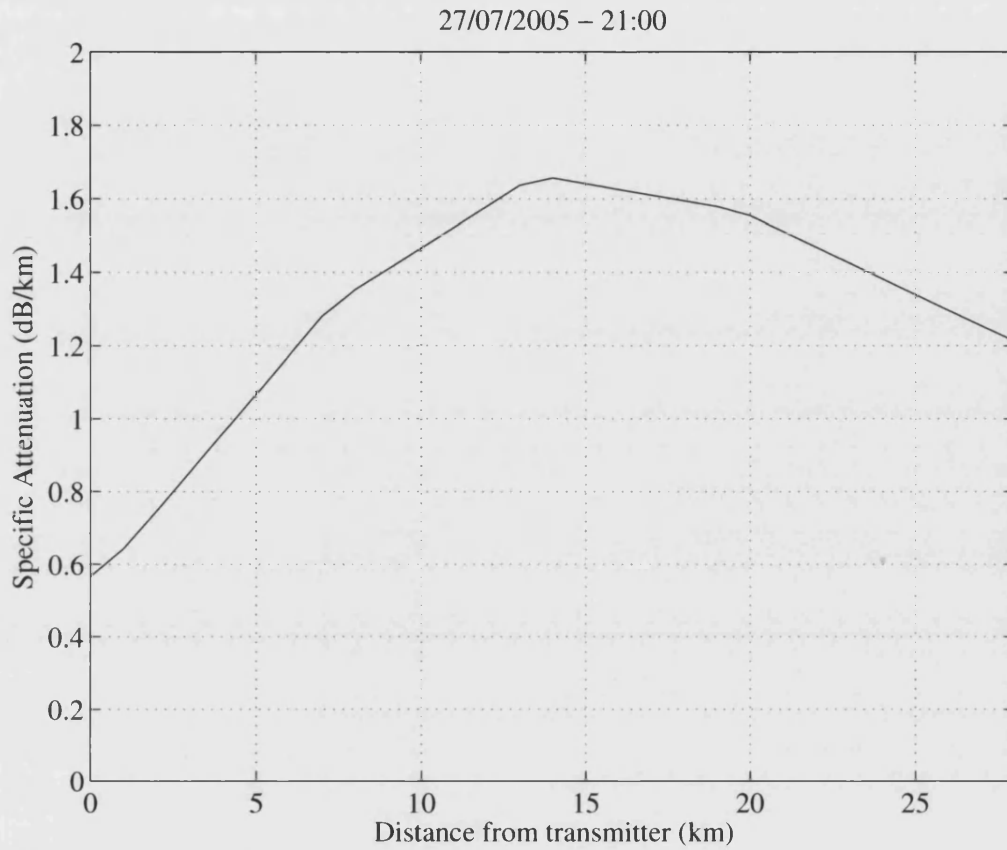
Figure 6.5: Specific attenuation along the example terrestrial path (Blagdon to Bath).

This specific attenuation is then integrated to provide the total link attenuation as demonstrated in figure 6.6.
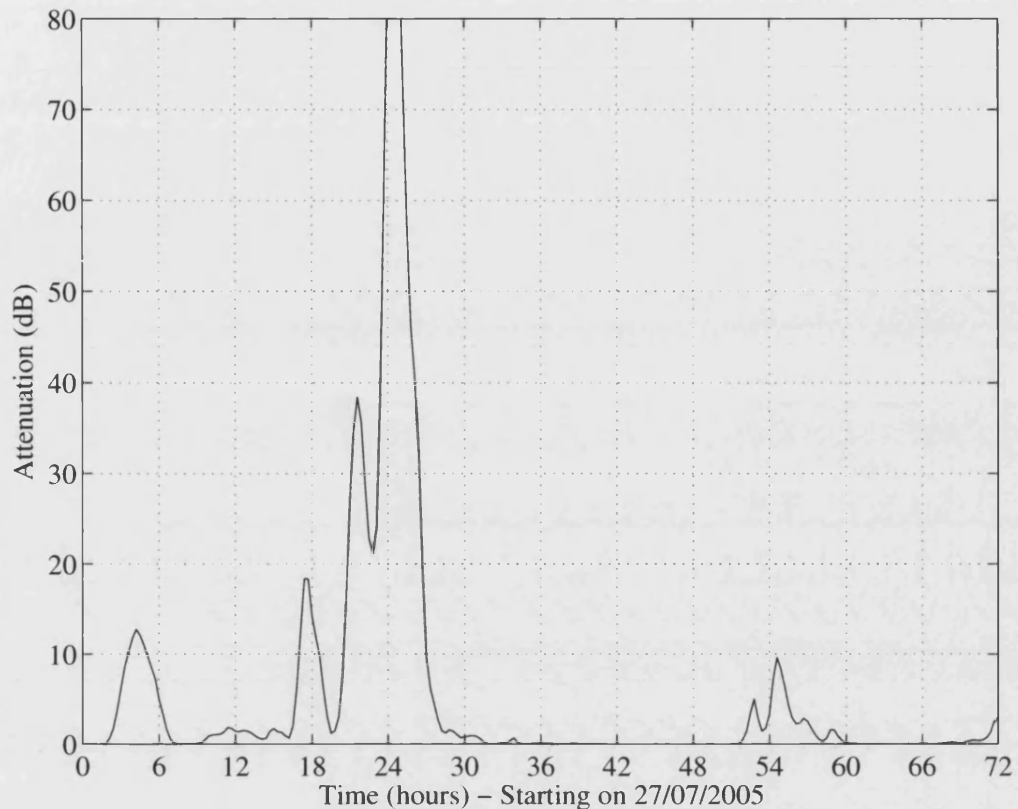
Figure 6.6: The attenuation time series for the example terrestrial link.

6.2 Time series verification against ITU-R long term statistics

The time series from the propagation forecast engine should exhibit the correct long term statistics. The ITU-R recommendations (as described in section 3.1) can be used as an estimate of the 'correct' annual CDF. It is worth noting these are average annual CDFs and real time series may exhibit a considerable amount of seasonal variation. However, the ITU-R recommendations can still provide a good estimate of the approximate attenuation levels to be expected for given link parameters and exceedance levels.

Currently the longest continuous set of meteorological data available is a three month period of time-coincident UM and NIMROD radar data. This three month period covers April, May and June 2004. Time series were calculated for a 50 GHz link to a transponder at 2°West. Three locations were investigated at varying latitudes and the

126

CDFs are shown in figure 6.7.
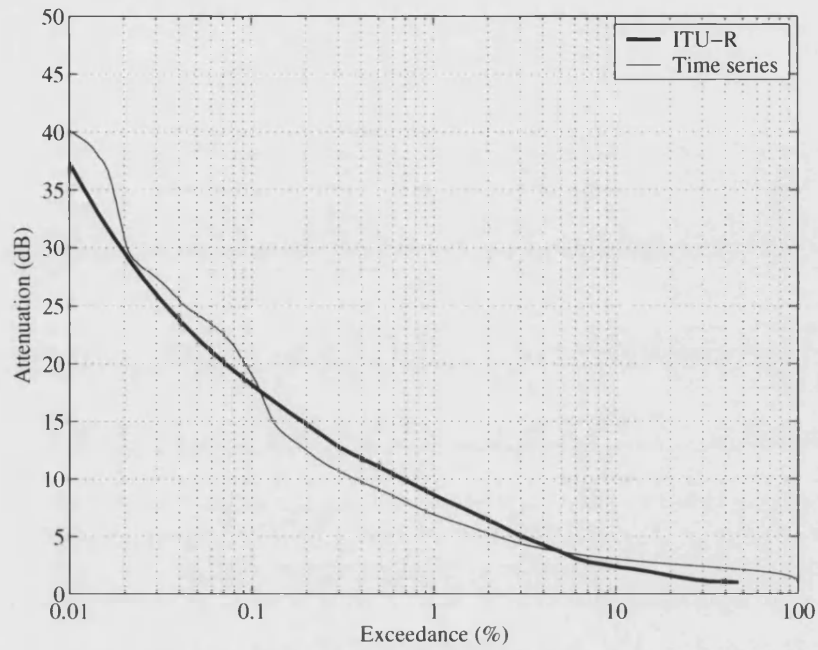
(a) Aberdeen
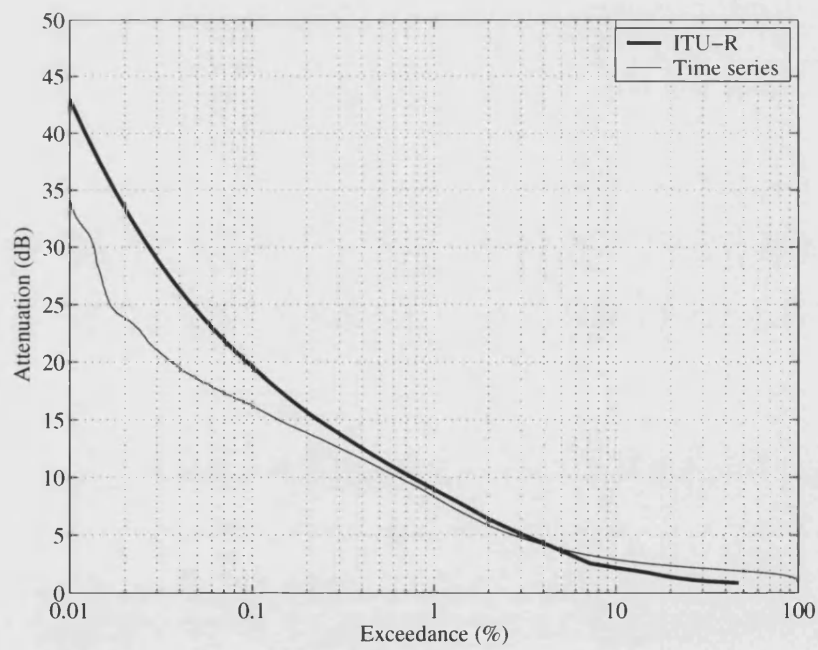
(b) Birmingham

Figure 6.7: Long term CDF for three locations compared to the ITU-R recommendations.

127
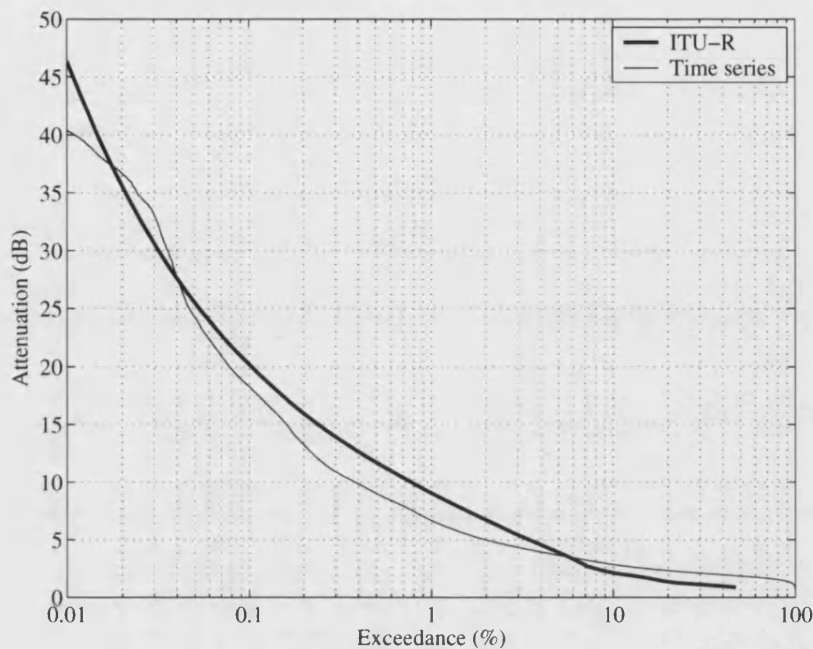
(c) Bath

Figure 6.7: Long term CDF for three locations compared to the ITU-R recommendations.

It can be seen from figure 6.7 that while there is some variation away from the annual ITU-R recommendations there is broad agreement. The ITU-R recommendations used do not model the effects of tropospheric scintillation, this results in the ITU-R underestimating the attenuation at high exceedances. The time series from the PFE covers a three-month period, this is not ideal since the very low exceedances will not be statistically valid.

Despite this the forecast engine produces time series that display long term statistics that match well with the expected annual statistics. This is the case down to at least 0.01 % exceedances (about 52 minutes a year).

6.3   Time series verification against fade slope models

As discussed in Section 2.7 the second order behaviour of the attenuation time series are very important for the design of FMT control systems. Hence, it is very important that the time series from the propagation forecast engine display the correct second

128

order behaviour.

The method for extracting fade slope characteristics from attenuation time series is described by van de Kamp [2003], this process was performed on time series from many different propagation measurement campaigns and resulted in the fade slope model described in Section 2.7.

The time series from April, May and June 2004 that was used to generate the long-term statistics for Section 6.2 were processed following the method of van de Kamp [2003, 1999]. These were then compared to the conditional fade slope probability predicted by the van de Kamp model.

The time series is first quantized into 1 dB bins. The standard deviation of the fade slope in each attenuation bin should be approximately linear with attenuation level [van de Kamp, 2003]. It should be noted that the results presented here are not normalised to a clear-sky, while the beacon measurements from van de Kamp [2003] were normalised to clear sky therefore demonstrating significant fade slope below the 2 dB fade level.

Figure 6.8 shows the standard deviation of the fade slope in each attenuation bin, (*), and a line-of best-fit derived from a least-mean squares algorithm. As can be seen the data shows good agreement with the conclusions of van de Kamp [2003].
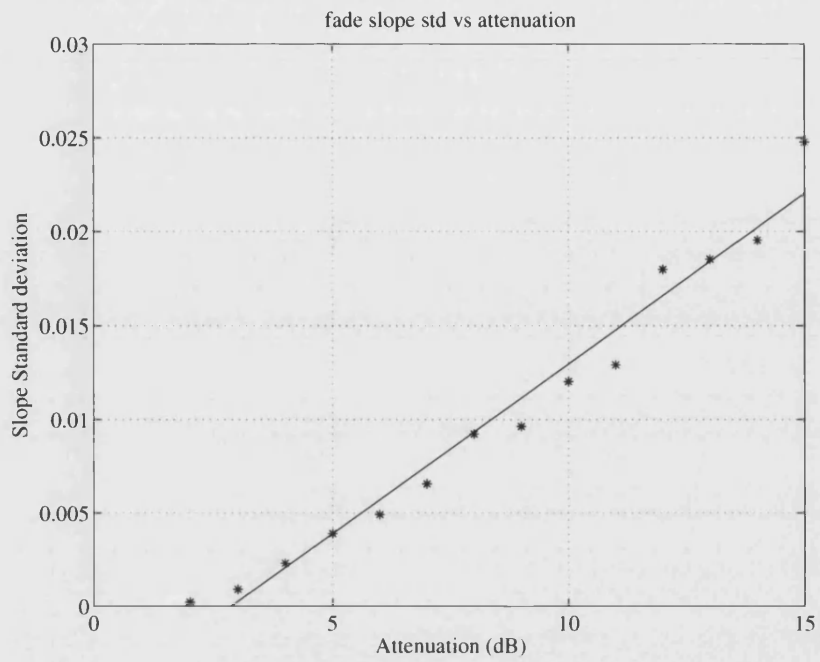
(a) Aberdeen
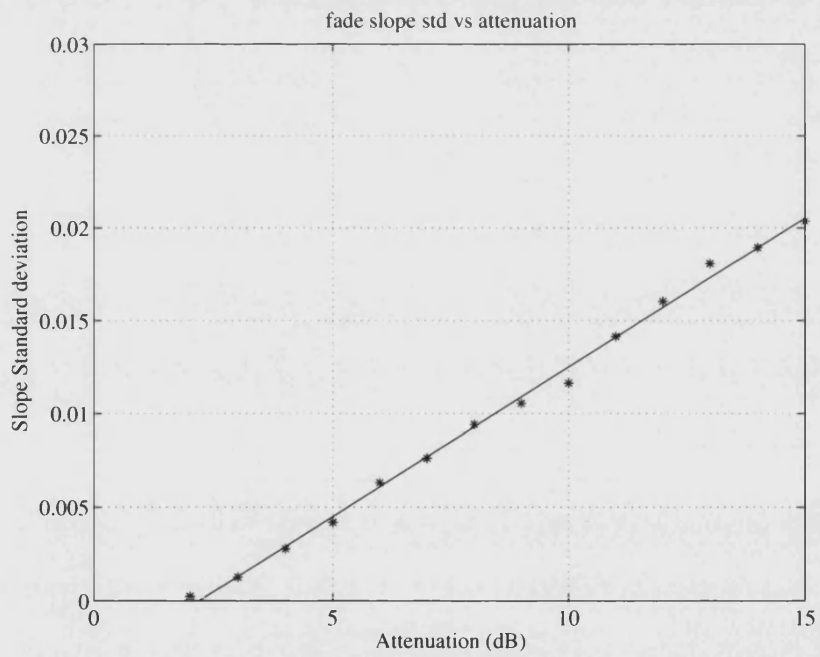
(b) Birmingham

Figure 6.8: The standard deviation of the fade slope in each attenuation bin.

fade slope std vs attenuation

(c) Bath
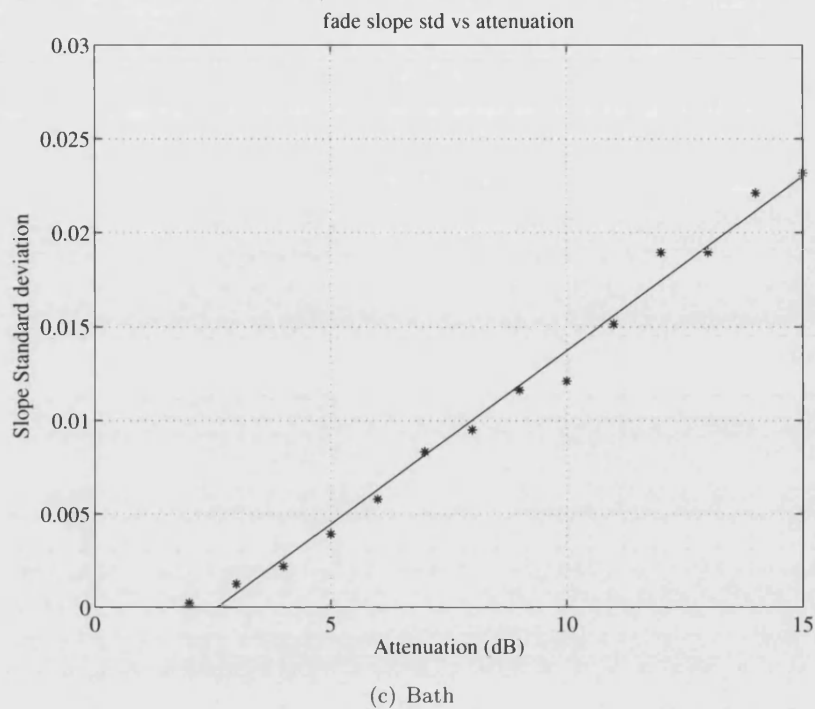
Figure 6.8: The standard deviation of the fade slope in each attenuation bin.

The probability distribution of the fade slope $(\mathrm{dB\,s}^{-1})$ conditioned by the fade level is calculated by binning the fade slope in each attenuation bin into fade slope bins of $0.001\,\mathrm{dB\,s}^{-1}$. The result is shown as the thick lines in figure 6.9. The corresponding model for the fade slope conditional probability [van de Kamp, 2003] is shown in the same figure with a thin line. The attenuation bins have been shown up to $8\,\mathrm{dB}$. For the available datasets the attenuation bins above $8\,\mathrm{dB}$ do not contain a statistically significant amount of data.
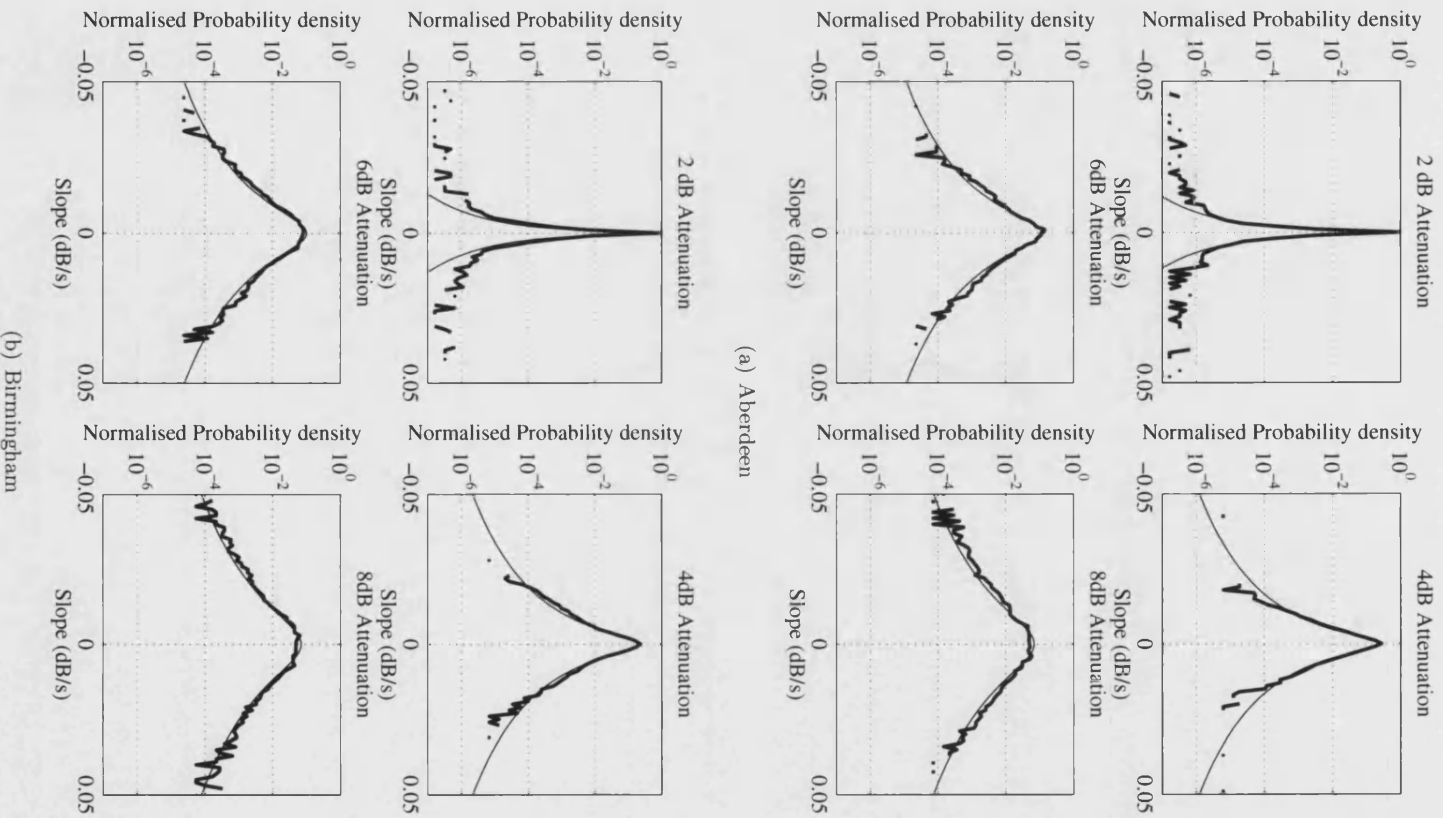
Figure 6.9: The conditional fade slope probability density in the attenuation bins.

132
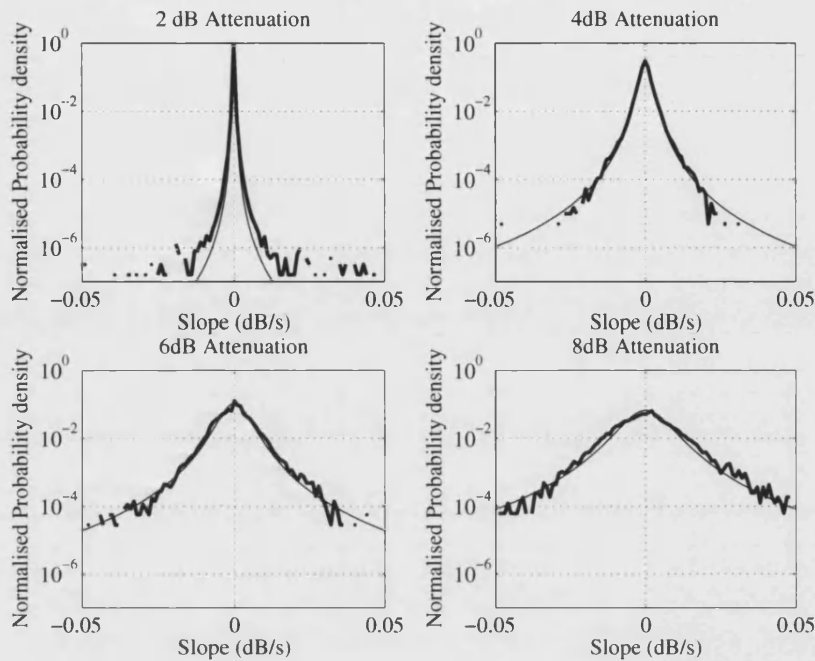
(c) Bath

Figure 6.9: The conditional fade slope probability density in the attenuation bins.

From figure 6.9 it is apparent that the conditional fade slope probability compares well to the van de Kamp model. Since the van de Kamp model was derived from beacon measurements it is not unreasonable to conclude that the time series demonstrate the same fade slope characteristics as beacon measurements.

6.4 Time series verification against GBS beacon measurements

Within a propagation forecasting paradigm the comparison between the outputs of the time series generator and measurements of EHF beacons is critical. This comparison highlights the ability to not only produce time series that demonstrate the correct statistics but the ability to correctly resolve individual fade events.

Currently, the only active beacon measurements in the UK that are suitable for propagation are being made by the Rutherford Appleton Laboratory (RAL). These measurements are of the 20.7 GHz beacon on the Global Broadcasting System (GBS) carried

133

aboard the US Department of Defence UFO-9 satellite. The geosynchronous satellite is at approximately 23°West, however it is on a slightly inclined orbit of approximately 3° although the ephemeris is not precisely known.

The slightly inclined orbit causes some problems with the beacon measurements. The data is post processed at RAL although the time series do still exhibit some residual effects. The time series is then filtered to remove the scintillation component using the procedure defined by van de Kamp [1999] with a cut-off frequency of 20 mHz. This procedure involves performing a FFT and then abruptly removing all the power in the components above the cut-off frequency. An inverse FFT is then performed to return the data to the time domain.

It is worth noting that these beacons measurements are made at 20.7 GHz, this is at the very bottom of the Ka- band. At this frequency the model is unable to demonstrate its ability to model deep cloud induced fading and very large rain fades.

6.4.1 Historical comparisons

The model was run with the UM and NIMROD data and the output was compared to the receivers at Sparsholt and Chilbolton in southern England. The comparison for a twelve-day period in October 2004 is shown in figure 6.10. It can be seen that there is broad agreement between the two time series.

The UM and radar meteorological data is provided at a 15 minute resolution. This time resolution is too coarse to resolve many events and shows evidence of spatial and temporal averaging. To minimise the effect of the temporal averaging, the beacon data has been filtered with a 15 minute block-average filter. This is not mathematically rigorous (since the attenuation is not linearly related to the rainfall rate) and does not completely mitigate the effects. However, this process is sufficient to allow a coarse comparison. A close-up of one day is included in figure 6.11.

There are several effects that may introduce errors in the output of the time series model that lead to differences between the beacons measurements and the output of the model. The meteorological estimates from NWP systems suffer from spatial and temporal averaging (in this case the highest resolution information is at a temporal resolution of 15 minutes and spatial grid squares of 5 km). This averaging results in the possibility of temporal shifting or fade depth errors when compared to the beacon measurements. The structure and organisation of rainfall and raindrops varies greatly from event to event which results in variations in the vertical profile of the attenuation

134
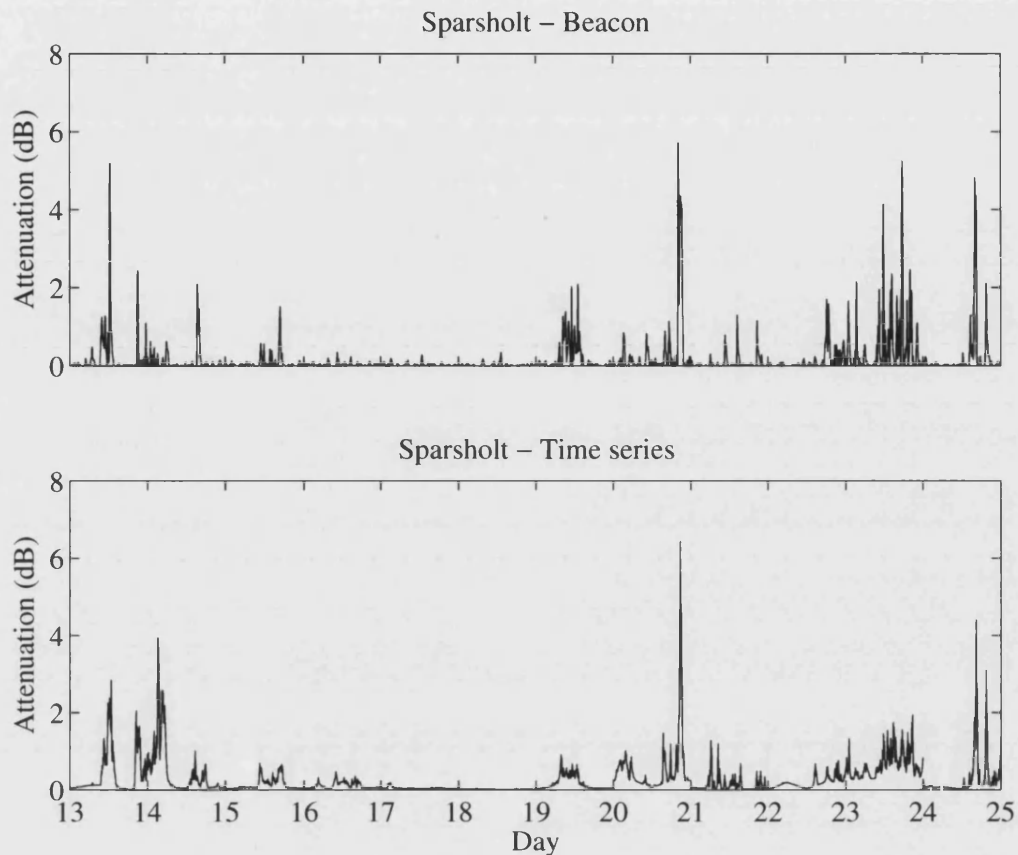
Figure 6.10: Comparisons of 20.7 GHz beacon measurements and the time series generator (15 minute intervals).

and variation of the drop size distribution. Whilst precautions have been taken to reduce these effects (for example the use of the showery and widespread versions of the Leitao-Watson rain attenuation model) they cannot be completely eliminated and will manifest themselves as fade depth errors.

From the close-up shown in figure 6.11 several effects are evident. The main event at 21:00 is modelled well although the small-scale structure is lost due to the inability to resolve very small-scale rain structure and drop size variations. In the early afternoon, between 15:00 and 18:00, we identify two rain events which are also evident in the beacon measurements. However, these are shifted temporally most likely due to the spatio-temporal resolution of the datasets.

The attenuation event between 09:00 and 12:00 is very poorly modelled in both depth and time. This rain event has been poorly modelled by the NWP system and poorly resolved by the radar system. This is also most likely a resolution error or an abnormal weather event, e.g. heavily convective, with supercooled liquid above the melting layer,
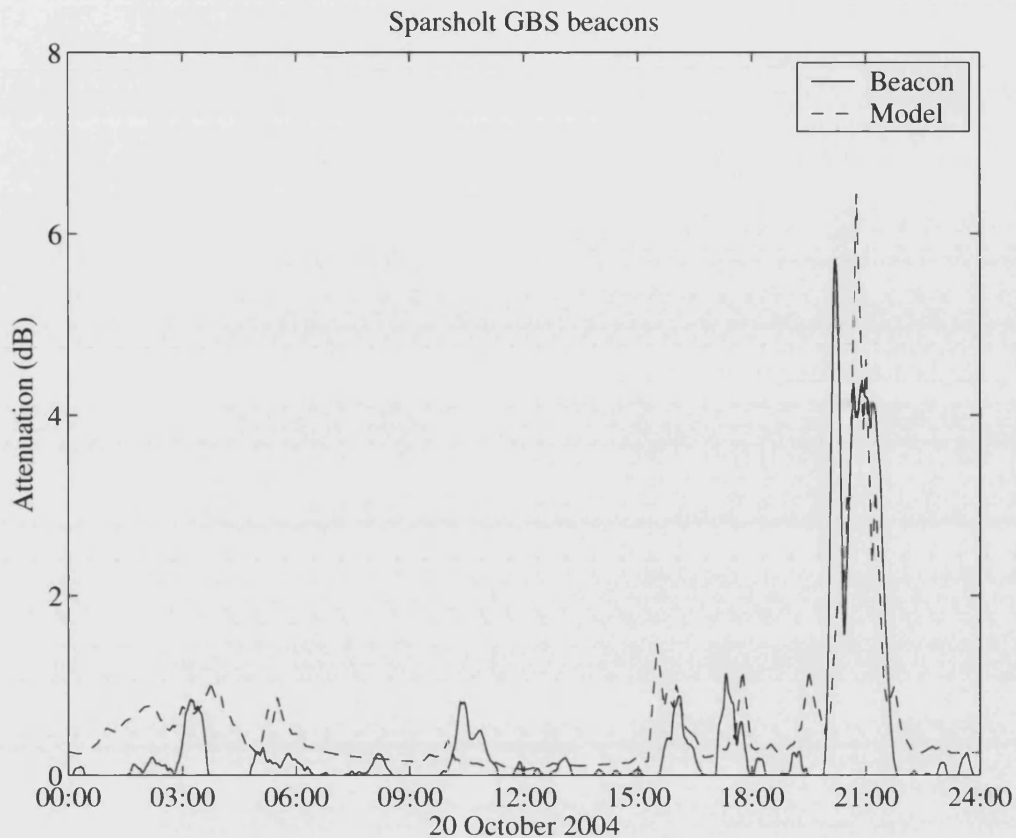
Figure 6.11: Close-up of the comparisons of 20.7 GHz beacon measurements and the time series generator (15 minute intervals).

etc.

These comparisons have been performed with historical NWP data, of more interest within a forecasting environment is the performance of the propagation forecast engine with forecast NWP data. Using the 'real-time' MM5 NWP system at the University of Bath (for more information on this model see Appendix D) it is possible to get an appreciation of the forecast accuracy. It is also worth noting that the real time system has no observational nudging, therefore it is possible to significantly improve the forecast accuracy of the NWP technique.

6.4.2 Forecast comparisons

The data from April 2005 was used to validate the forecast performance. This NWP data is produced at hourly intervals and time series were compared with the receivers at both Sparsholt and Chilbolton. The forecasts from the NWP system are provided

for the day of the NWP execution ('today'), the day after the forecast run ('+24H') and the day after that ('+48H').

Figures 6.12 and 6.13 show the performance of the forecast system. Note that the beacon measurements are shown with a different vertical scale. Since the forecasts are made on a day-by-day basis the forecast attenuation is not joined to form a continuous time series but separated into discrete daily predictions.
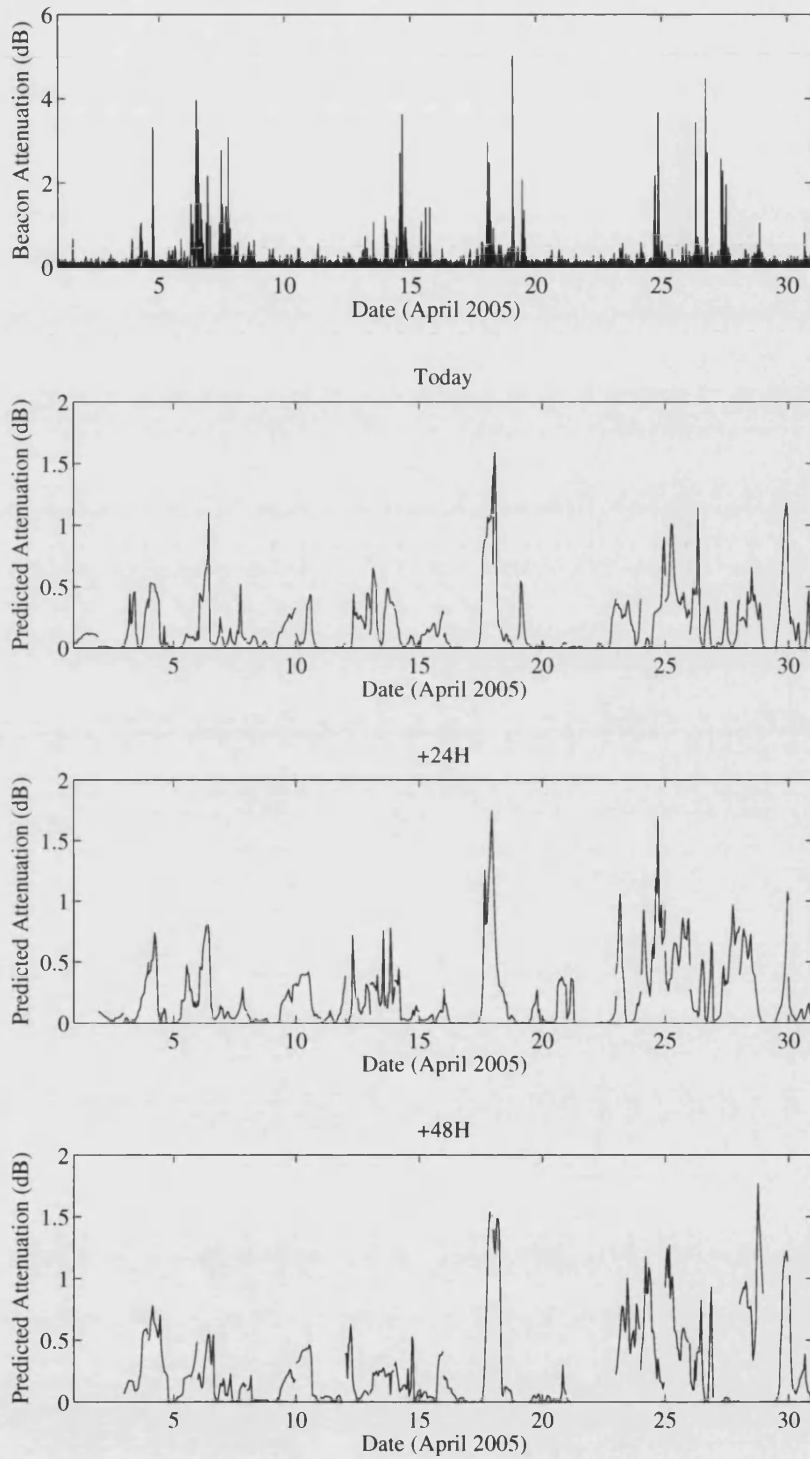
Figure 6.12: The ability of the propagation forecast engine to predict the attenuation of a 20.7 GHz beacon signal - Sparsholt receiver April 2005.
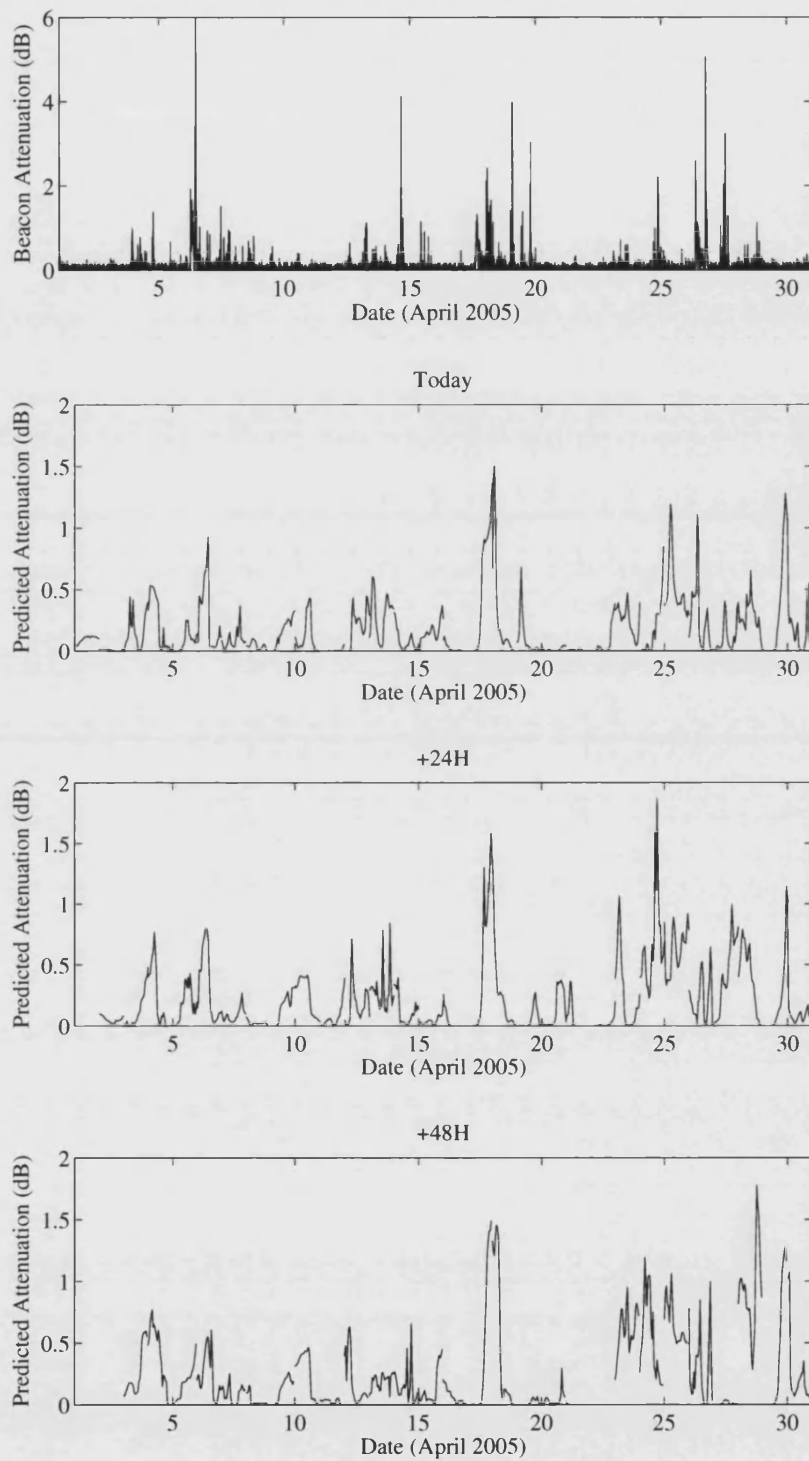
Figure 6.13: The ability of the propagation forecast engine to predict the attenuation of a 20.7 GHz beacon signal - Chilbolton receiver April 2005.

The forecast system typically underestimates the attenuation that is observed. The most likely cause of this is temporal averaging (since the attenuation is predicted at an hourly resolution), despite this it can be seen that the ability to resolve fade events is still good. In order to provide a more meaningful comparison the beacon data is block averaged over an hour. Figures 6.14 and 6.15 repeat the forecast data with the averaged beacon measurements. It should be noted that all the figures are now on the same vertical scale.
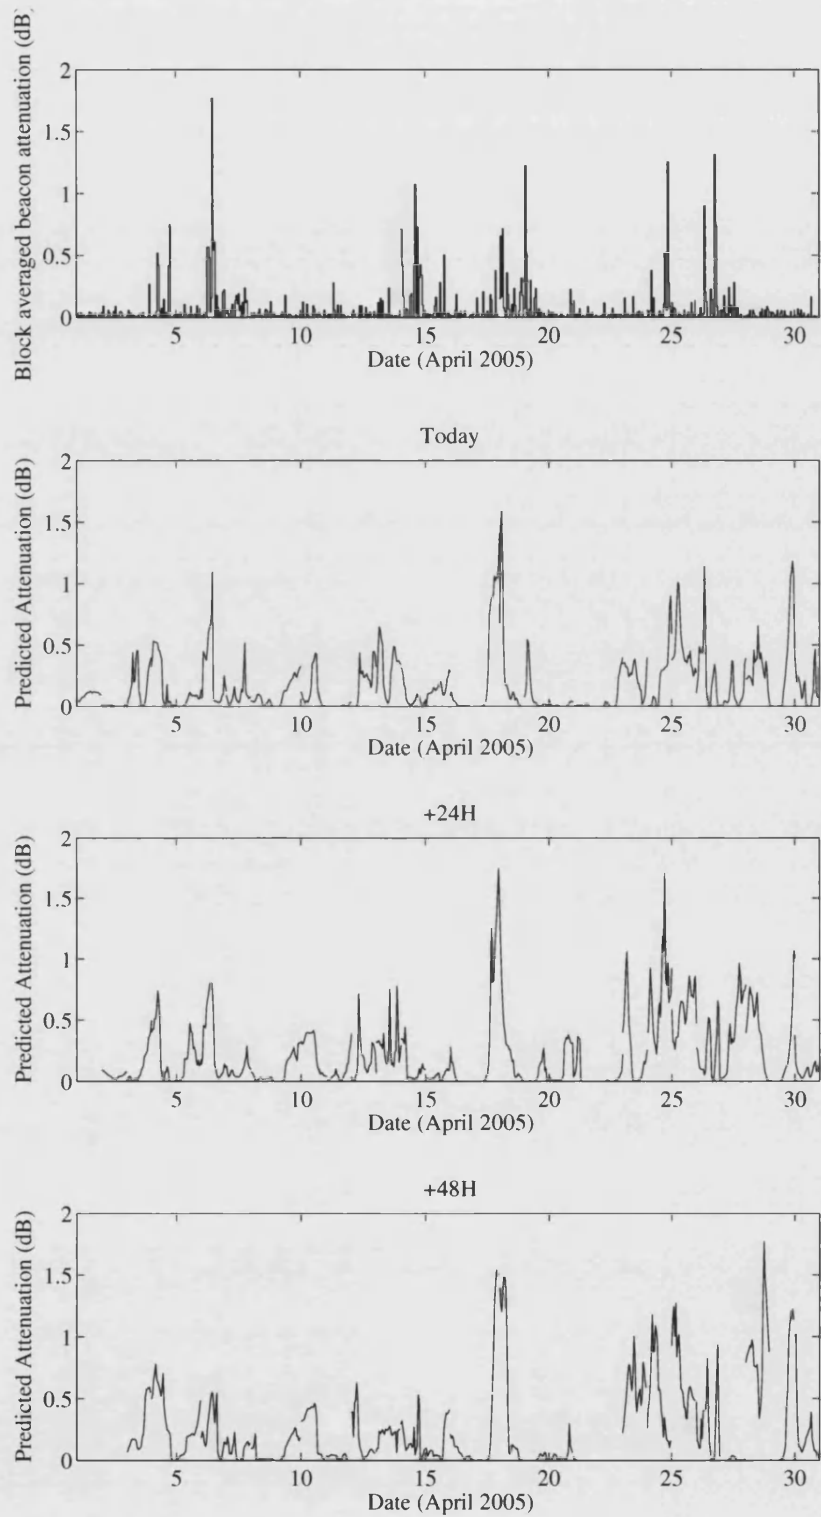
Figure 6.14: The ability of the propagation forecast engine to predict the block averaged attenuation of a 20.7 GHz beacon signal - Sparsholt receiver April 2005.
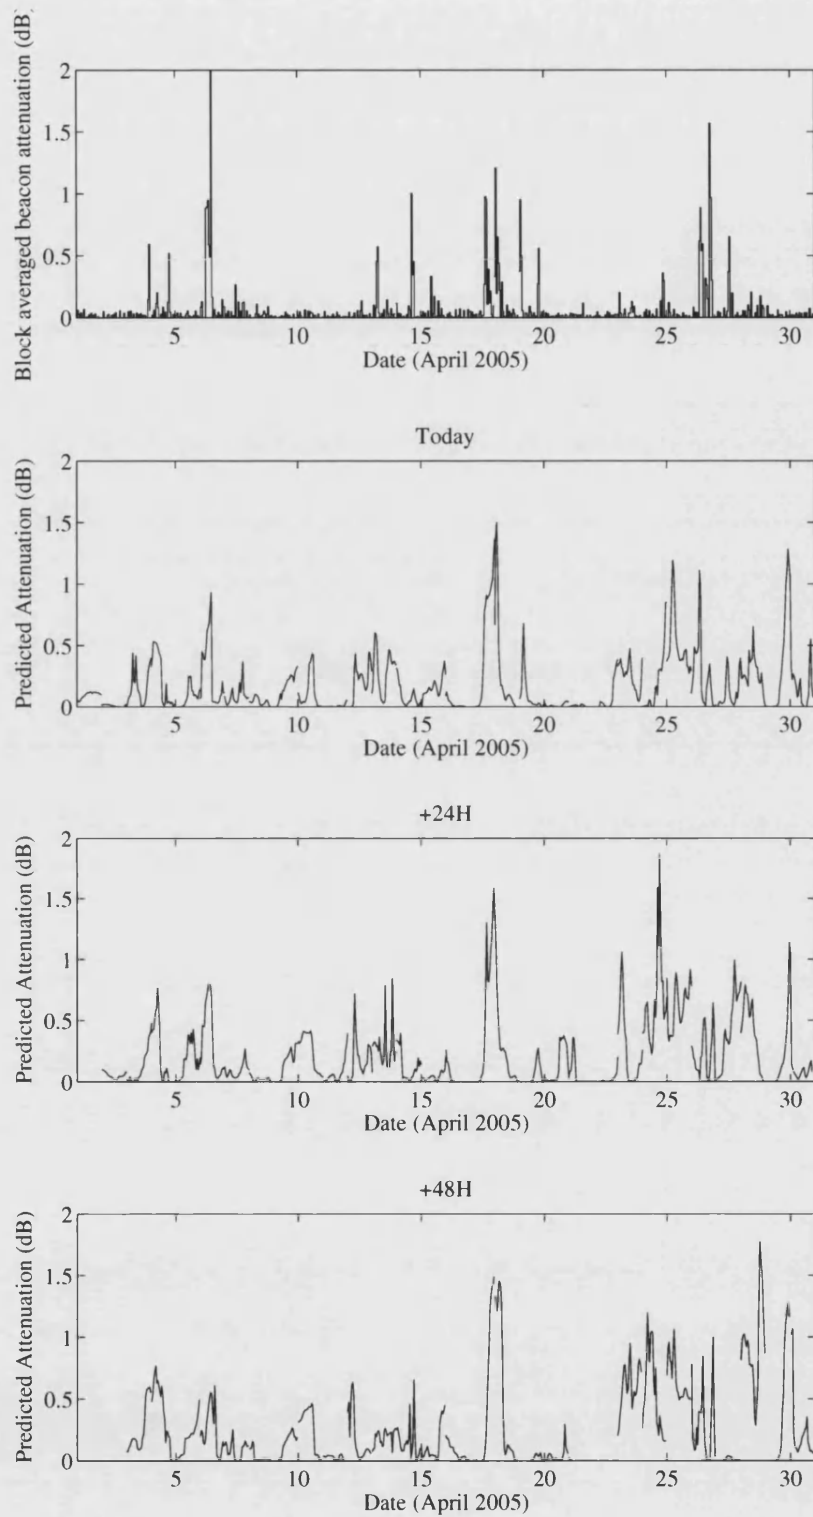
Figure 6.15: The ability of the propagation forecast engine to predict the block averaged attenuation of a 20.7 GHz beacon signal - Chilbolton receiver April 2005.

In order to provide an estimation of the accuracy for attenuation levels at the higher frequency within the Ka- band, the slant path attenuation for the same period was calculated with the frequency increased to 40.7 GHz. This is shown in figure 6.16.
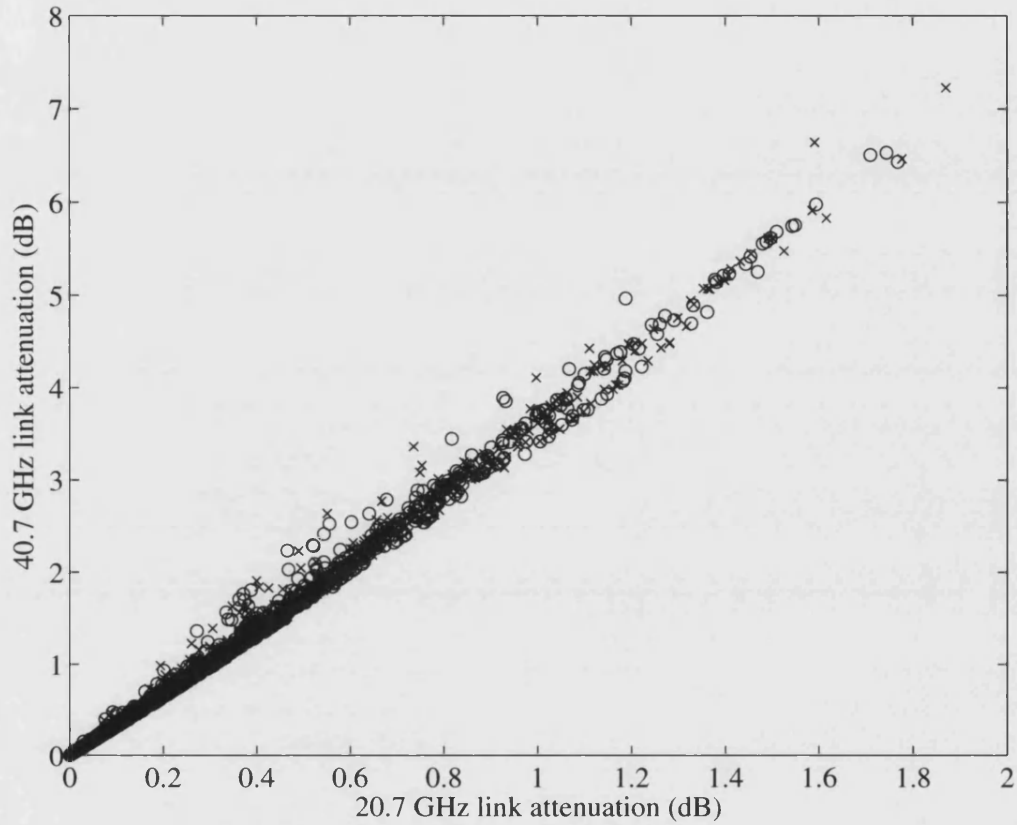
Figure 6.16: Coarse attenuation scaling for the GBS forecasted data, Chilbolton, 'x', and Sparsholt, 'o', receivers - April 2005.

Estimates of the attenuation scaling are shown in in table 6.1, assuming a least mean squares linear regression. This shows that a correct prediction of an attenuation level of 2 dB at 20.7 GHz is roughly comparable to a correct prediction of an attenuation level of 7 dB at 40.7 GHz.

Table 6.1: Coarse attenuation scaling for the GBS forecasted data - April 2005.

Link Attenuation (dB)	
20.7 GHz	40.7 GHz
0.25	0.873
0.50	1.769
0.75	2.664
1.00	3.559
1.25	4.454
1.50	5.349
1.75	6.244
2.00	7.139

6.4.3 Forecast skill

Within an operational 'tactical' forecasting system the ability to predict the exceedance of a particular attenuation level is of critical importance. To evaluate the ability of the propagation forecast model to perform this task it is possible to examine the forecast as a problem with a binary outcome, e.g. does the link attenuation exceed 2 dB? From this point a contingency table can be constructed as in table 6.2.

Table 6.2: Example contingency table for a binary categorical forecast.

Event forecast	Event observed	
	Yes	No
Yes	Hit (a)	False alarm (b)
No	Miss (c)	Correct rejection (d)

In order to evaluate some statistics and draw conclusions of the forecast ability some parameters must be defined in terms of the contingency table variables (a, b, c and d).

The first two metrics are the hit ratio, H, and the false-alarm ratio, F, these are simple metrics to understand and need little explanation:

$$H = \frac{a}{a+c} \tag{6.1}$$

$$F = \frac{b}{b+d} \tag{6.2}$$

The other simple metric is the base rate, s, this is the probability of the phenomena

occurring i.e.,

$$s = \frac{a + c}{a + b + c + d} \qquad (6.3)$$

From this point we can now define some forecast metrics. Caution should always be exhibited with the use of individual 'skill-scores' in isolation and several different metrics should always be used to try provide a thorough examination of the forecast characteristics [Jolliffe and Stephenson, 2003]. It is also important to recognise that these binary skill scores punish temporal errors twice since the error causes a miss followed by a false alarm (or vice versa).

The most commonly used skill score is the equitable threat score (ETS) (also known as the Gilbert skill score) and this is defined as:

$$\text{ETS} = \frac{a - a_r}{a - a_r + b + c} \qquad (6.4)$$

where,

$$a_r = \frac{(a + b) \cdot (a + c)}{a + b + c + d} \qquad (6.5)$$

represents the number of hits expected from pure chance. This score highlights the ability for the forecast system to provide forecasts that are better than a random chance. The score ranges from $-\frac{1}{3}$ to 1 where 0 represents 'no-skill'.

The ETS at various attenuation exceedances is shown in figure 6.17 for both the Sparsholt and Chilbolton receivers. As can be seen the forecast engine scores higher for the Chilbolton site than the Sparsholt site. Since the sites are relatively close, this is likely due to one of three factors:

1. The analysis was performed upon a relatively short time series. This could be proved by taking longer time series in excess of a year to provide a more rigorous example.

2. The attenuation has been block averaged in order to simulate the effect of a large temporal step size. It is possible that short differences in the rain arrival time at the two sites caused significant differences in the block averaged time series.

3. Differences in the beacon receivers (or indeed the post processing to remove the orbit inclination) which would affect the receivers independently.

Of these factors the most likely reason will be the first.

There are several other effects that are shown in figure 6.17 worthy of note. It can be

seen that as the forecast length increases, the ETS reduces quite dramatically. This can be explained since the MM5 model with no observational nudging is very good at resolving prognostic variables out to 72 hours. However, beyond about 26 hours it is poor at resolving rain at the accuracy required for it not be significantly penalised by the ETS[2].
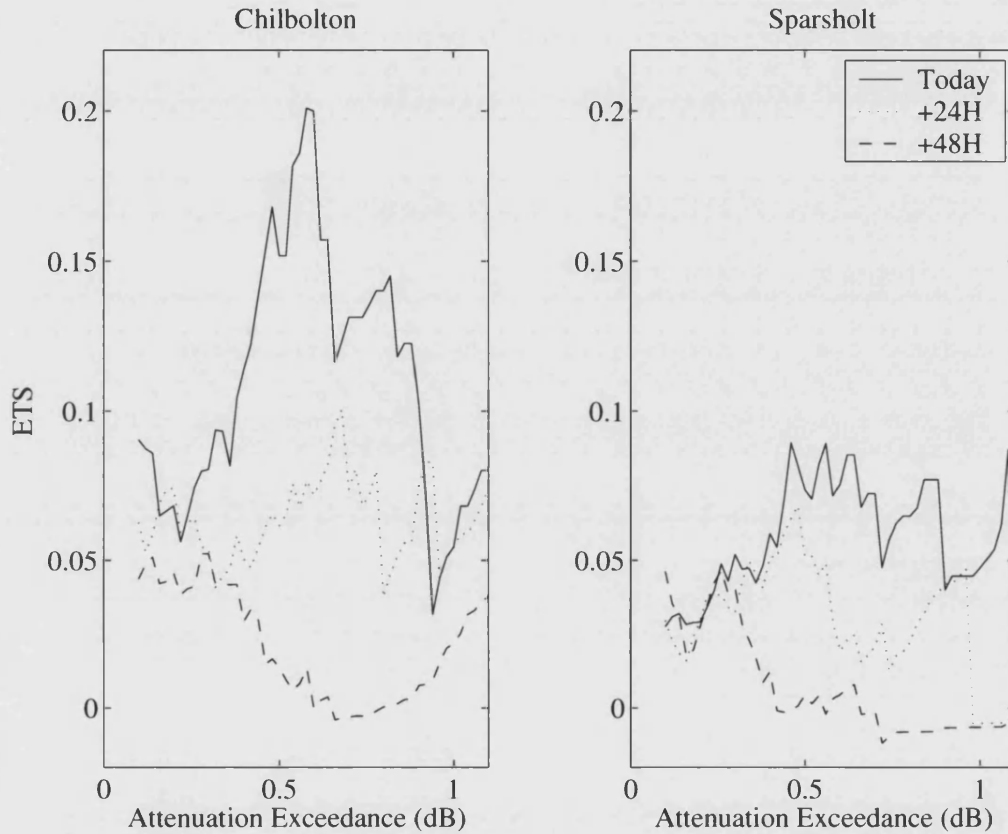
Figure 6.17: The equitable threat score (ETS) of the GBS beacon predictions.

The ETS, as defined in equation 6.4, can be rearranged in terms of the base rate thus:

$$\text{ETS} = \frac{H - F}{\frac{1 - sH}{1 - s} + \frac{F(1 - s)}{s}} \tag{6.6}$$

As the base rate, s, becomes small the first term in the denominator tends to one whilst the second term becomes very large. This results in the ETS becoming small independently of the hit rate. In essence, the ETS measure is biased for statistically rare events.

---

[2]This is true of all systems that use no observational nudging, for more information see Appendix D

In order to provide a skill score that does not become small for small base rates the extreme dependency score (EDS) can be used. The EDS does not tend to zero for rare events where the base rate is small. In addition it is not explicitly dependent on the bias and is a very reliable measure of the skill for 'extreme forecasts' [Stephenson, 2005, Coles et al., 1999]. The extreme dependency score is defined as:

$$\text{EDS} = \frac{2 \log_{10}\left(\frac{a+c}{a+b+c+d}\right)}{\log_{10}\left(\frac{a}{a+b+c+d}\right)} - 1 \tag{6.7}$$

The forecast engine performance is evaluated as an EDS and is shown in figure 6.18. This provides a much more constant skill score, with the Chilbolton predictions still performing slightly better than those for Sparsholt. However, the performance for the prediction for the +24H is much more comparable to that for the day of the forecast whilst the +48H forecast is still poor. To put this skill score in context, table 6.3 contains the extreme dependency score for a number of NWP systems.

Table 6.3: Example EDS scores for several different NWP techniques.

Model Name	Organisation	Rainfall rate threshold	
		5 mm / day	10 mm / day
Hirlam	Finnish Meteorological Institute	0.55	0.47
Lokal Modell LM	Deutsche Wetterdienst	0.61	0.59
Unified Model	UK Met Office	0.72	0.66
Aladin	Météo France	0.61	0.56

As can be seen the skill score of the forecasting technique is typically of an order comparable to that of NWP techniques rain prediction capability. This is an important conclusion since this infers that the propagation forecasting technique does not add significantly to the errors beyond those typically associated with NWP systems.

The final skill metrics that are used are the false alarm and hit ratios. These are shown in figure 6.19. It can be seen that there is a significant decay in skill over the duration of the +48H forecast. The relative operating characteristic (ROC) plot is also shown in figure 6.19. This is a plot of false alarm ratio versus the hit ratio. As can be seen from these plots at relatively high attenuations (for this example) the false alarm ratio is very low whilst the hit ratio is considerable higher. This indicates that the main causes of errors are missed events rather than false alarms. This is to be expected with a temporal resolution of 1 hour and no observational nudging for the NWP system. In this configuration the forecast will be unable to resolve small scale rain events.

At low attenuations the separation between the hit rate and the false alarm ratio are
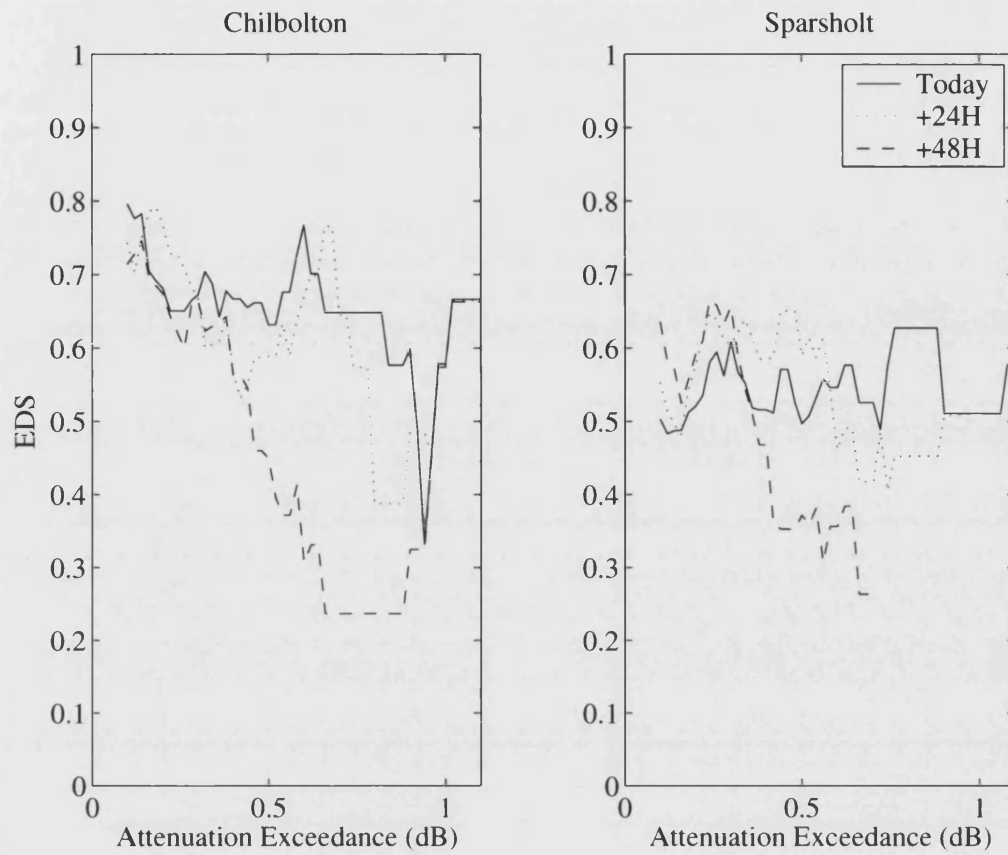
Figure 6.18: The extreme dependency score (EDS) of the GBS beacon predictions.

relatively constant (although as previously discussed they are still implicitly tied to the base rate). This relatively constant separation enhances the conclusion that the skill score should be relatively constant.
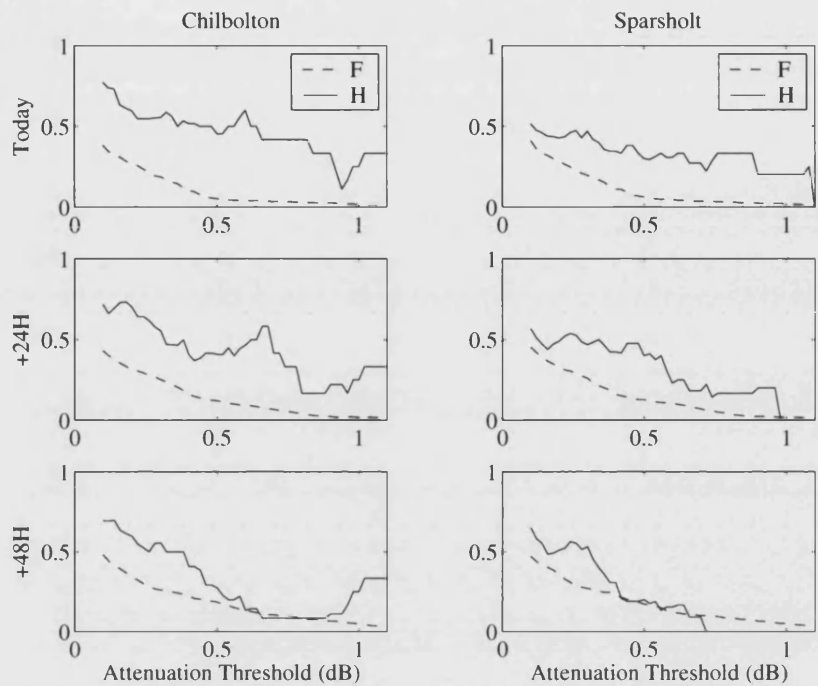
(a) False-alarm / Hit ratio plot
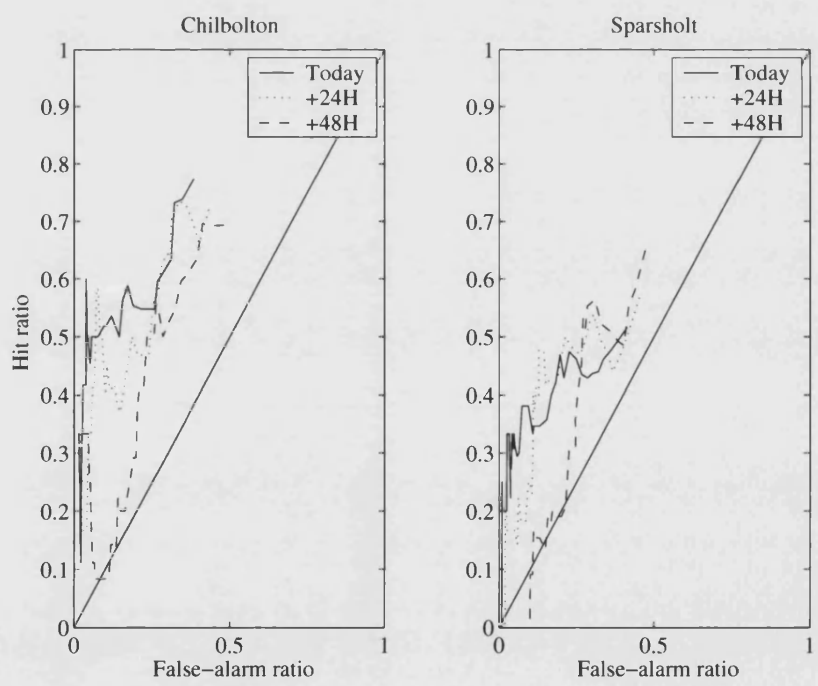
(b) ROC plot

Figure 6.19: The hit rate and false alarm characteristics of the propagation forecast engine.

Chapter 7

Network-centric simulations

This chapter is dedicated to the design of a network simulator for VSAT networks. This network simulator can be used to demonstrate the use of time-coincident time series for multiple sites. It is important to understand that it is impossible to perform simulations for all possible network configurations. The problem space is massive in terms of number of receivers, carrier frequency and other link characteristics, transponder bandwidth, link protocols, MAC schemes etc.

The network simulations performed for this chapter consider what could be described as a typical future network design. The network configuration is discussed in section 7.3.

These simulations consider the availability of a network affected by rain, clouds and atmospheric gases since the simulation only considers the network performance to be linked purely to signal power. In order to fully assess the performance of the network it should be linked to the ratio of energy per bit to the spectral noise density $\left(\frac{E_b}{N_0}\right)$ and encapsulate the increase in the noise power cause by a warm attenuating media, such as clouds and rain.

7.1 Opnet simulator

Initial simulations were performed using the OPNET simulation tool (version 9.1). This tool is primarily for IP based network configuration and originally for wired communications. The system has been improved with the addition of a wireless radio module. The wireless module creates a 'radio-pipeline' that joins a radio transmitter that has a

packet to send, with **all** radio receivers within the simulation environment.

The first simulations that were performed were based upon a full DVB-S protocol stack. However, this quickly becomes very unsuitable for the sort of 'long term' (of the order of days rather than minutes) simulation that is required. A more suitable simulation type is a high level simulation with a reduced protocol stack. For example, the contents of a data bit stream is not (generally) important. At the transport layers the data is repackaged, scrambled and interleaved to a degree that it can be assumed to be a constant rate 'random' data stream.

The initial high-level simulations attempted to characterise a simple satellite broadcast environment. The overview of the transmitter is shown in figure 7.1. The transmitter was designed as a single DVB multiplex made of eight packetised elementary streams (PES). The PES can consist of any data packaged as MPEG-2 transport packets. There are well defined protocols for such packing, one of the most common protocol is the Multi-Protocol Encapsulation (MPE) scheme. Several of these PES are then combined to form a Transport Stream Multiplex (TS-MUX) which is then placed into the radio pipeline.
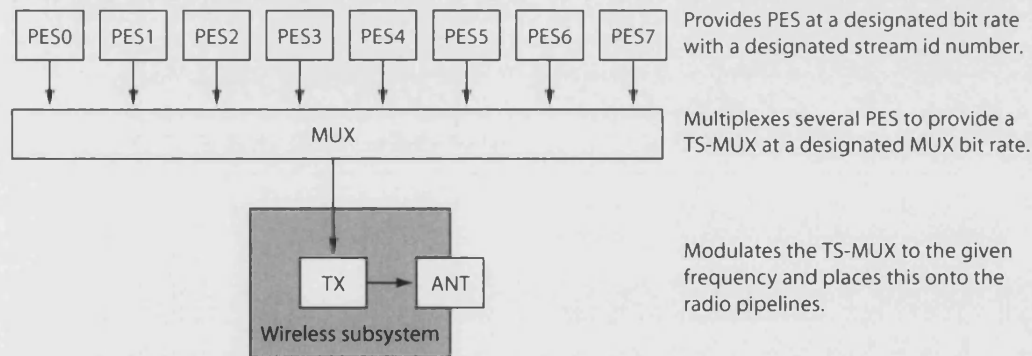
Figure 7.1: The initial high-level design for a transmitter for an Opnet simulation of a simple broadcast DVB network.

The simplified radio pipeline as used for this project is shown in figure 7.2. This pipeline is executed for every packet that is sent; each receiver has it's own individual pipeline joining every transmitter. Several of the pipeline stages were assumed to be constant over the coverage area which reduced the calculation time in the pipeline, for example, antenna gains, propagation delay and interference noise were assumed to be constant. The attenuation time series for each receiver is loaded at the initialisation of the network simulation, and the fade level from the time series is added to the free space path loss calculated in the received power stage.

151
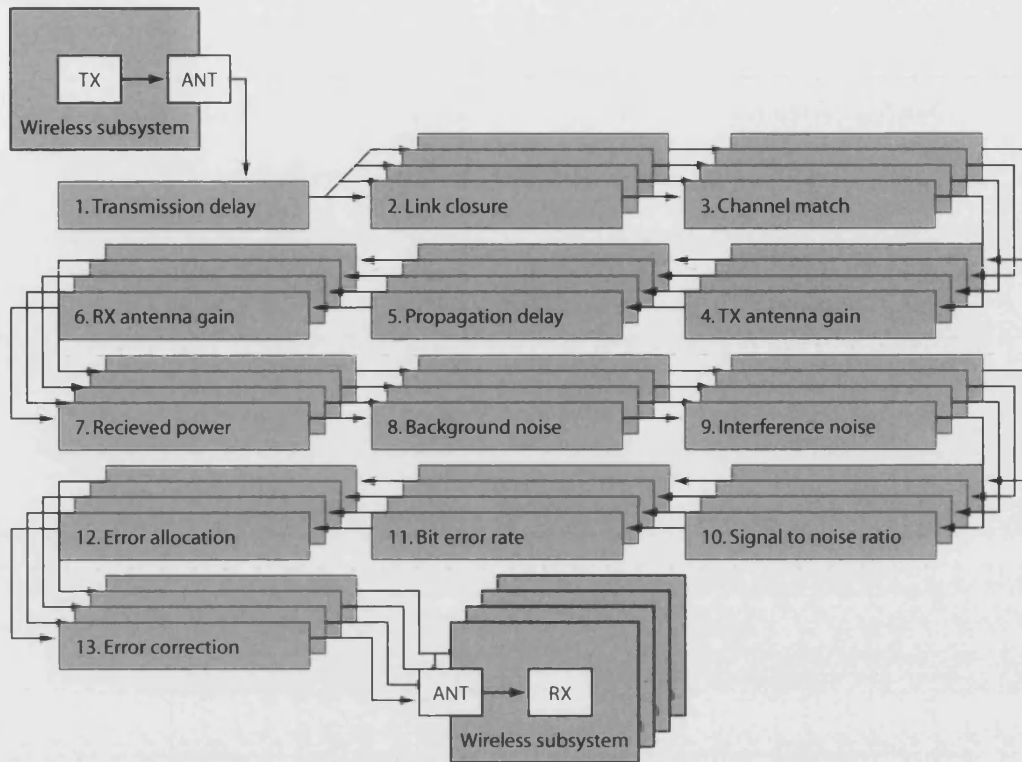
Figure 7.2: A simplified radio pipeline for four receivers.

The space segment can be modelled using a simple bent pipe architecture[1], as seen in figure 7.3. This was placed in a geostationary orbit and simply takes the uplink off a radio pipeline and places it back into a downlink pipeline.
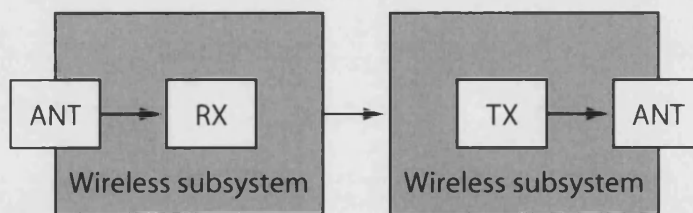
Figure 7.3: The bent pipe satellite architecture.

Each receiver is an instance of the very simple receiver shown in figure 7.4. This simply takes the downlink off the radio pipeline, retrieves the requested PES in the TS-MUX and logs and then sinks the packet.

---

[1]Typically systems working at Ka-band and above require some on-board processing (OBP) in order to provide quality of service however as a simple first approximation this was not considered

152
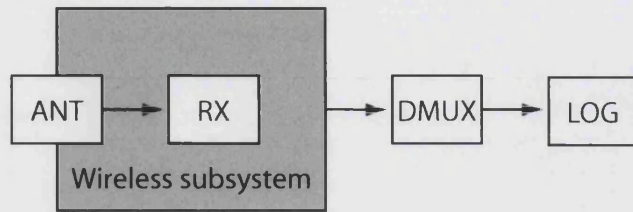
Figure 7.4: The receiver architecture used in the initial Opnet simulations.

These models were configured with each of the eight PES configured with a 2 Mbps channel rate, with a carrier frequency 48.2 GHz for the uplink and 39.5 GHz for the downlink. The system was configured using a static link budget in order to provide a clear sky fade margin of 8 dB. An example time series for one receiver is shown in figure 7.5. This figure shows the received SNR and the throughput of both the entire TS-MUX and the PES to which the receiver was tuned.
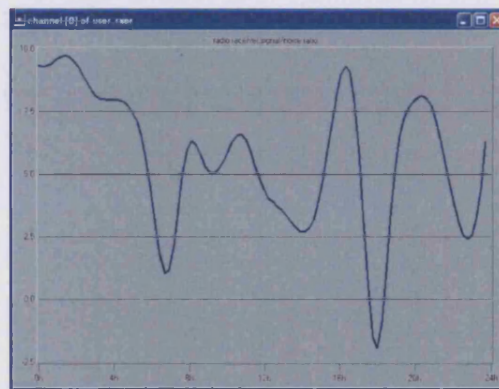
(a) Received SNR
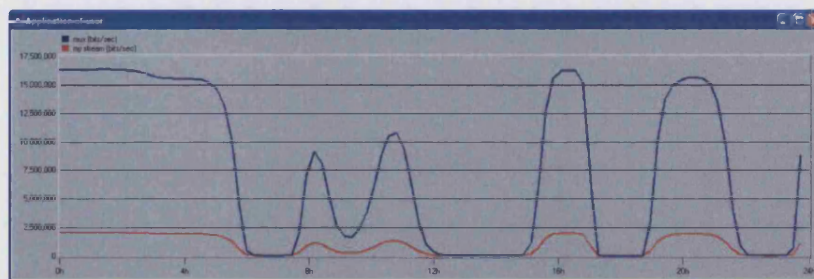
(b) Throughput

Figure 7.5: The results of the simple Opnet simulation.

With just two receivers this simulation takes in excess of 2 days of computing time[2]. Using the Opnet performance meter it is possible to estimate the length of time that higher density receiver simulations would take; these estimates are shown in table 7.1. It can be seen that for a month of computing time it is only possible to simulate an 11 receiver network for 1 day with a very simple broadcast network with no return channel, fade mitigation technique or resource management scheme.

Table 7.1: Estimated simulation times for the simple Opnet simulation

Number of receivers	Simulation time (days)
1	1.31
2	2.59
3	4.03
4	5.91
8	15.86
11	29.97
20	72.45
25	93.88

Given this poor scaling we conclude the Opnet simulation package is not suitable for high level network simulations such as those required for this project.

7.2 Bespoke network simulator

In order to provide a flexible efficient simulation environment a high level network simulator was built using Matlab. The system was designed to be able to simulate many FMT or resource allocation schemes at once; this allows the performance of various different FMT techniques to be compared over one link attenuation time series. The system is designed to be able to simulate proactive FMT / resource management strategies. In order to create an efficient simulation environment no high-level network layers are included, since at the transport layer each packet appears to be full of a binary symmetric random data stream. The system does also not simulate a full MAC layer, some of the tasks of the MAC layer are still required and these are included.

The propagation forecast engine creates two attenuation time series, one for the purposes of fade timetabling which has no scintillation component (denoted the 'forecast time series'); the other time series is used to simulate the channel conditions (denoted the 'attenuation time series') which includes the scintillation component. It is possible to introduce errors into the forecast time series in order to investigate the effects of an

---

[2]Using a Windows workstation (2.8 GHz Pentium 4 with 512 MB of RAM)

incorrect propagation forecast.
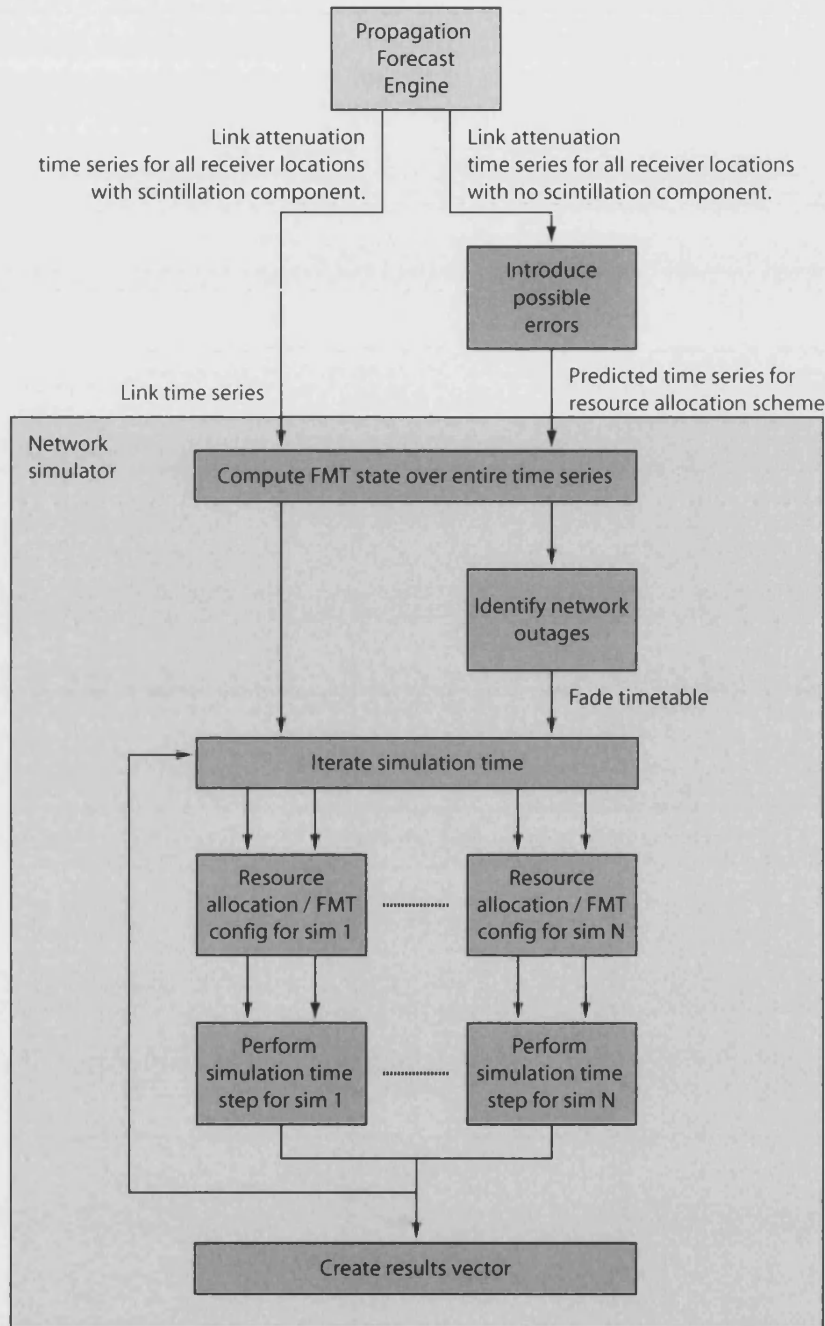
Figure 7.6: The overview of the bespoke network simulation.

Once the time series have been passed to the network simulator the time series are converted to a 'FMT level'. This level is a classification of the fade level, and is

discussed in later sections.

The next process is to analyse the forecast time series and construct a fade time table. This defines the time and length of forecasted fade events that can not be combated using the FMT techniques. These forecasted fade events indicate times where the link will be dropped.

After this preparation the system can then begin to step through the simulation time. At each time step the FMT and the resource management is configured and then the simulation step is performed. The results and statistics for that step are then saved and the process is iterated for the length of the simulation.

7.3 Configuration

This section is dedicated to the network configuration for the simulation system. As discussed previously there are so many different network configurations it is impossible to simulate all possible networks. A generic network has been designed around a typical DVB-RCS network (ETSI standard 102 402). The system uses two 36 MHz transponders, one for uplink and another downlink. A QPSK modulation scheme with a root-raised-cosine roll-off filter $\alpha = 0.35$ results in a transponder bit rate of:

$$
\begin{aligned}
R_{SB} = \frac{B}{1+\alpha} &= \frac{36 \times 10^6}{1 + 0.35} \approx 26.5\,\text{MBaud} \\
&= 26.5 \times 10^6 \times 2 = 53\,\text{M bps}
\end{aligned}
\tag{7.1}
$$

The system is simplified by assuming that all data is sent in MPEG-2 transport packets[3]. These consist of 204 bytes, of which 187 bytes are data with the remaining 17 bytes used for header and sync bytes. This results in 32,475 packets second^{-1} which are available for the link.

The carrier frequencies are 48.2 GHz for the uplink and 39.5 GHz for the downlink. The link budget is considered purely by the fade margin at the terminal. Each terminal has a minimum bandwidth of 1 Mbps which is required for the service to function. Each terminal also has a total desired bandwidth of 2 Mbps which can be used to provide either a better QoS (e.g., higher bit rate video codecs) or value added services.

Three different types of network are now considered: a traditional network, a simple network and a complex network.

---

[3]MPEG-2 transport packets are typically packaged into a packetised elementary stream (PES) which are then multiplexed together to construct the transport stream

7.3.1 Traditional system

In this work we have defined a traditional system as one that has no form of FMT. There is no power control or changes of code rate. The power is defined with a clear-sky fade margin of 7 dB. Using the ITU-R recommendations (see section 3.1) as described previously this fade margin results in a 99 % available uplink and a 99.9 % available downlink.

The coding for DVB-S and DVB-RCS systems is a variant of the JPL deep space code and is essentially two concatenated codes. The outer code is a shortened version of the systematic Reed-Soloman (255,239) code whilst the inner code is a punctured rate $\frac{1}{2}$ constraint length seven code with possible code rates of $\frac{1}{2}$, $\frac{2}{3}$, $\frac{3}{4}$, $\frac{5}{6}$ and $\frac{7}{8}$. The performance of the DVB decoder is defined in ETSI 300 421 standard. Annex D of this standard describes an example signal-to-noise ratio required for each code rate to get 'quasi-error-free' QEF performance[4]. The example SNR is shown in table 7.2 for an example system which is close to the one described in this chapter.

Table 7.2: Required $\frac{E_b}{N_0}$ at the input to the RS decoder for QEF after RS decoder.

Code rate	SNR required	Packet Efficiency
1/2	4.1 dB	61.0 %
2/3	5.8 dB	68.7 %
3/4	6.8 dB	73.4 %
5/6	7.8 dB	78.7 %
7/8	8.4 dB	81.6 %

This system does not have the ability to change code rates, however when the system is deployed the code rate can be chosen from the set in table 7.2. Since there is no significant on-board-processing the terminals can only be allocated the minimum data rate.

7.3.2 Simple instantaneous reactive FMT systems

The simple FMT system uses the traditional system explained in the previous section but with the added protection of two FMT techniques. The first is 7 dB of power control which is deployed to achieve the highest data rate possible at all times.

The system also has the ability to change the code rate in order to provide the most appropriate code rate for the propagation conditions. The data rate is chosen from

---

[4]This is defined as an error ratio of 10^{-11}

157

those in table 7.2 depending upon the fade level. The number of time slots a terminal is allocated is no longer constant since the efficiency varies with the code rate. Hence, the system requires an adaptive TDMA and associated resource allocation schemes. This is also the case since the system will attempt to allocate any 'spare' bandwidth to terminals for the extra services.

The high level bandwidth allocation is performed using the procedure shown in figure 7.7.
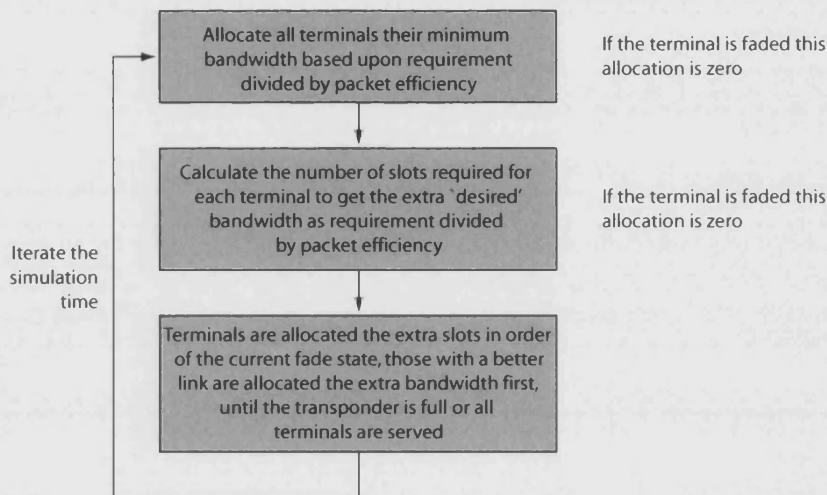
Figure 7.7: Overview of the bandwidth allocation procedure for the simple FMT network.

This procedure attempts to ensure that all terminals are allocated their minimum allocations and the extra bandwidth is allocated in preference to those with a good link. In essence this ensures that the maximum number of terminals can be allocated the extra 'desirable' bandwidth [Kreuer and Schmidt, 1994].

Essentially this system aims to provide the best possible service to as many terminals as possible whilst maintaining a lower quality of service for those terminals that are undergoing a fade.

7.3.3 Simple FMT systems with an A-TDMA time diversity system

This more complex system allows the network control centre (NCC) to use the forecasted fade timetable to identify incoming fades and provide content earlier than it is required in order to 'buffer-up' terminals which have a fade approaching. This proce-

dure follows the previous system but allocates any remaining bandwidth for terminals to exploit time diversity techniques. This extra bandwidth is allocated by exploiting concepts from game theory.

The principles from game theory allow a framework where terminals can *'bargain'* for bandwidth, in this case the term bandwidth is reflected as the bartering for TDMA time slots. The bargaining occurs within a framework that ensures that the results are fair and optimal.

The initial work on game theory was performed by Von Neumann and Morgenstern [1944] and Nash [1950]. This work was to provide insight into economic problems such as state trading between two nations or negotiations between employer and labour union. The game theoretic framework can be broken down using some simple definitions.

The Nash bargaining solution will be, by definition, a Pareto-optimal point. Formally a Pareto-optimal point is a point where, if all possible bargaining points form a set U then the point $u \in U$ is Pareto-optimal if for each $v \in U, v \geq u$, then $v = u$ [Yaïche et al., 2000a]. More practically a Pareto-optimal point can be regarded as a point where it is impossible to find another point which offers better performance for all players (in this case terminals). Any deviation from this point means that at least one terminal will suffer a reduction in performance. In general, in a game with N players the Pareto-optimal points form an $N - 1$ dimension hypersurface.

For non-trivial problems (i.e. games with more than two players) there are a theoretically infinite number of Pareto-optimal points, however not all of these Pareto-optimal points correspond to Nash bargaining solutions. A point in the set of possible solutions can be regarded as a Nash bargaining solution if it follows several constraints:

- The solution is Pareto-optimal

- The solution satisfies the linearity axiom

- The solution satisfies the irrelevant alternatives axiom

- The solution satisfies the symmetry axiom.

The last three requirements are the axioms of fairness. These provide a rigorous mathematical framework for what can be regarded as 'fair'. This fairness can be described as a situation in which no individual class or terminal is unfairly denied access to the network or overly penalised [Mazumdar et al., 1991].

The linearity axiom property of the solution ensures that the solution does not change if the objectives are scaled linearly. This means if the amount the terminals are willing to 'pay' for bandwidth are all scaled linearly the solution point will not change.

The irrelevant alternatives axiom property states the Nash bargaining point for a larger set of strategies is the same as that of the smaller game if the Nash bargaining point is a valid point for the smaller game. In our application this is the effect of adding new terminals to the problem. However, for the Nash bargaining point to be a valid point in the new larger problem the added terminals must have zero bandwidth requirement and zero value of that bandwidth. This axiom represents a slightly trivial mathematical legacy from the game theory applications.

The symmetry property states that the bargaining point does not depend on the specific labels. This infers that terminals with the same initial points and same objectives will achieve the same performance at the Nash bargaining solution.

In trivial examples terminals can each value their bandwidth allocation by the same amount. In effect the bandwidth demand is the allocation value to the terminal. In more complicated systems the value that each terminal gives their bandwidth allocation can be a function of the type of traffic, current propagation environment and possible future fade timetable.

In these more complicated systems the aim of the network control centre is to allocate the bandwidth in such a way that the sum of the amount the serviced terminals are prepared to 'pay' is maximised and the axioms of fairness are satisfied.

In the simulated network system an approach to the allocation of the bandwidth follows a procedure which allocates bandwidth at a Nash bargaining solution [Yaïche et al., 2000a,b]. This procedure has been shown to achieve a Pareto-optimal point and satisfy the axioms of fairness.

The procedure is defined such that each terminal, i, calculates the value of the extra bandwidth they require, v_{bi}, as a function of the length of the forecasted fade currently unbuffered, f_d, and the time until the forecasted fade starts, f_s, such:

$$v_{bi}(t) = \sum_{l=1}^{N} \frac{f_{dl}}{f_{sl} - t} \qquad (7.2)$$

Essentially, the total value is a sum of the value for individual fades, l, for a time in the future such that there are N fades. This is demonstrated graphically in figure 7.8 for three terminals. It is typical that the buffer action of the network is constrained

to some time ahead, denoted the 'lookahead-time'. This has the effect of limiting the buffer size requirement. All of the simulations have been performed with a lookahead time of 1 hour which results in a maximum buffer requirement of just around 430 MB[5]. However fades in excess of 1 hour can not be completely buffered. The trade-off between buffer size and the length of the maximum fade must be considered a network specific parameter. In essence further increasing an already very large problem space.

The resulting bandwidth values for the three terminals shown in figure 7.8 can be seen:

$$v_{b1}(t) = \frac{f_{d11}}{f_{s11}} + \frac{f_{d12}}{f_{s12}} \tag{7.3}$$

$$v_{b2}(t) = \frac{f_{d21}}{f_{s21}} + \frac{f_{d22}}{f_{s22}} \tag{7.4}$$

$$v_{b3}(t) = \frac{f_{d31}}{f_{s31}} \tag{7.5}$$
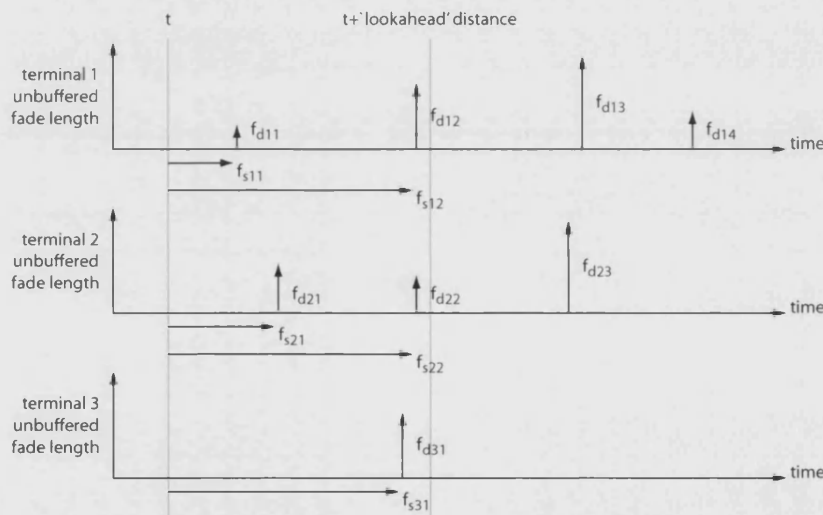
Figure 7.8: A diagram demonstrating the approach used to calculate the value of extra bandwidth for three terminals.

It can be seen that as an unbuffered fade approaches the current time, t, the value of the bandwidth increases very quickly, since $v_{bi}(t)$ is effectively the tangent of the geometric angle between the time (distance) to the fade and the unbuffered depth.

A number of spare time slots, B, are then allocated to the terminals following the simple pro-rata form:

$$B_i(t) = \frac{v_{bi}(t)}{\sum_{i=1}^{N} v_{bi}(t)} \cdot B(t) \tag{7.6}$$

---

[5]1 M bps × 3600 secs in an hour ÷ 8 bits per byte ÷ 1024^2 bytes per MB

161

where $B_i(t)$ is the allocation to the i-th terminal.

The aim of this system is to use the maximum bandwidth that the system has available in order to provide as high a service availability to as many terminals as possible.

7.4   Results

It should be noted the results presented in this chapter are only valid for the 'high-level' networks described in the previous sections. Changes to the network parameters will produce different results. The solution space for the infinite number of different networks and configuration is so massive that it is impossible to simulate a 'generic' network that is representative of all network configurations. This section describes the results for the networks defined previously and demonstrates the use of both historical multiple-site time-coincident time series and an application of forecast time series.

The network consists of 19 terminals deployed over the UK and Eire as in figure 7.9. Each terminal has a minimum requirement of 1 Mbps and a desired requirement of 2 M bps.

Figure 7.9: The terminal locations used for the simulations.

The simulations were performed over a selection of nine days from June 2003. Since the simulation procedure is still a lengthy procedure a selection of days which were considered to be a good sample of the month were used as a proxy for the entire month.

7.4.1 Traditional system

The 'network availability' (which is the average of the terminal availabilities) of the traditional system over individual simulation days is shown in the bar charts in figure 7.10. The crosses represent the individual terminal availability over that day.
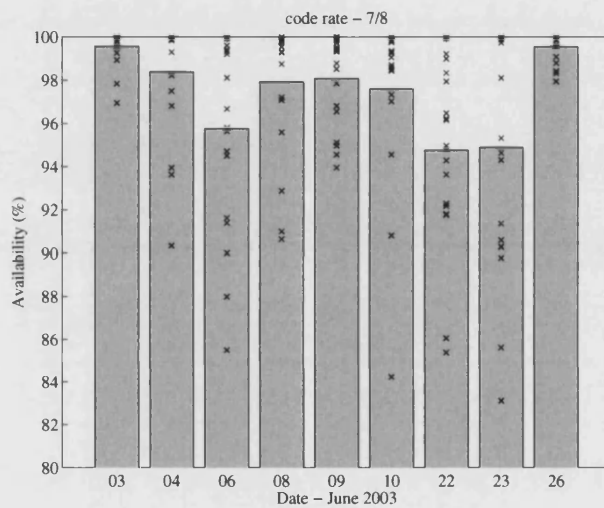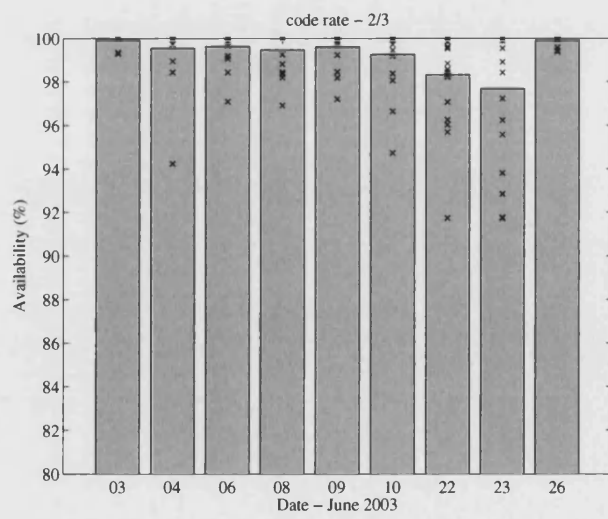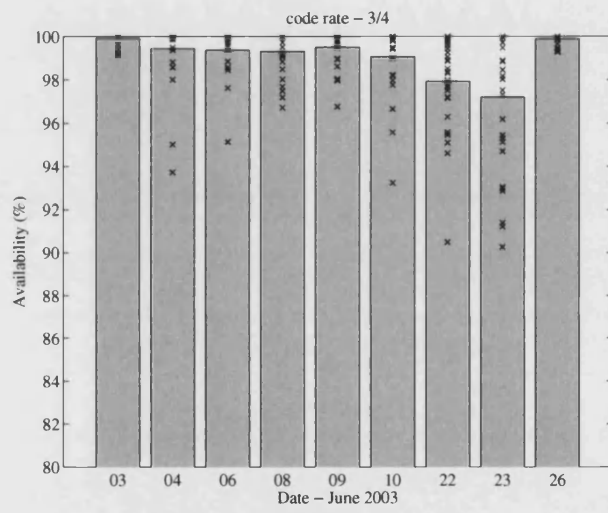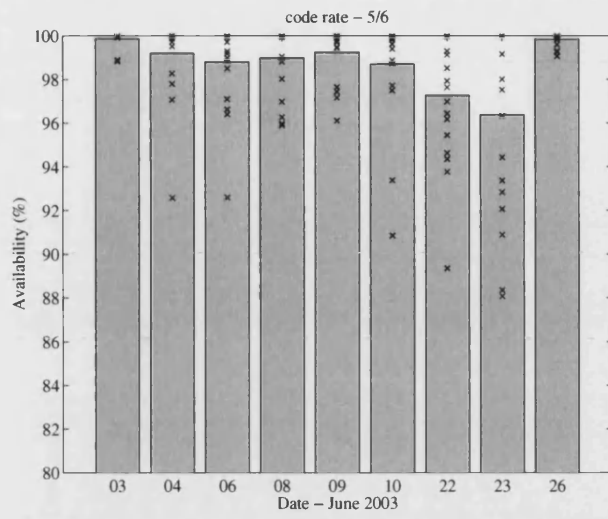
Figure 7.10: The performance of the 'traditional' network for the five different code rate configurations.

Figure 7.10: The performance of the 'traditional' network for the five different code rate configurations.

164
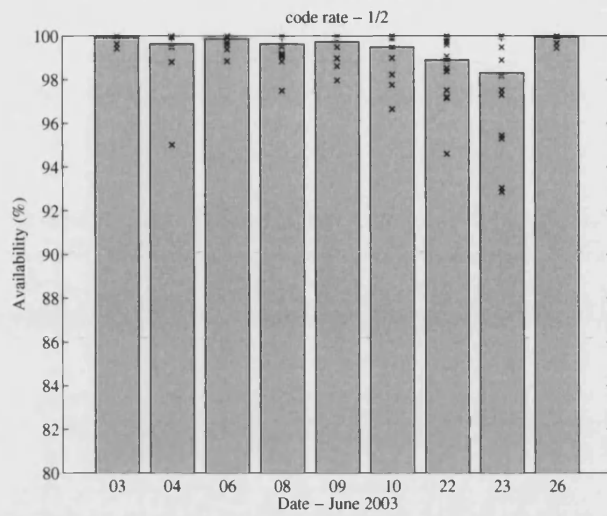
Figure 7.10: The performance of the 'traditional' network for the five different code rate configurations.

Figure 7.11 shows the information from figure 7.10 redrawn in terms of packet efficiency (as a metric for the code rate). The solid lines represent the 'network availability' and the stars represent the individual terminals.
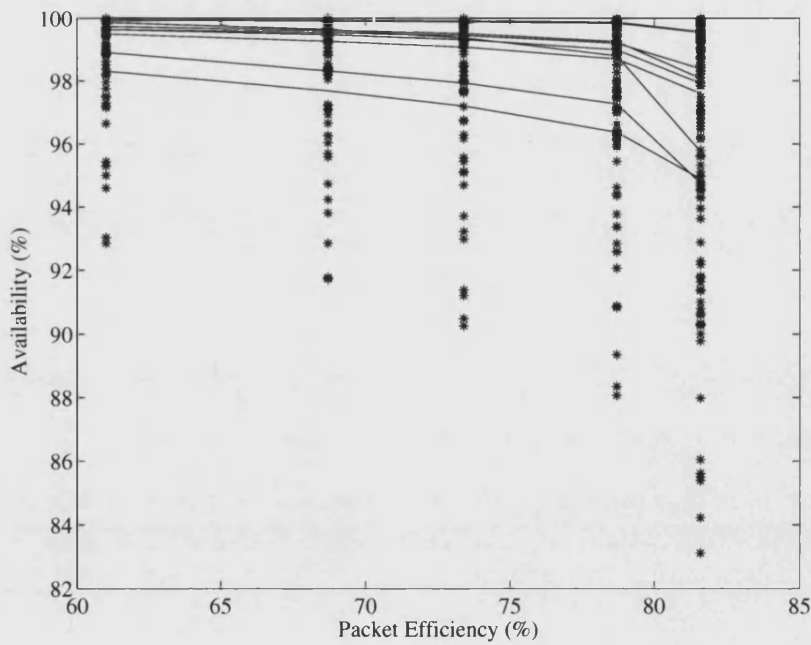
Figure 7.11: The performance of the 'traditional' network with varying packet efficiency.

This provides the expected result; the network availability increases with decreasing packet efficiency but there are some individual terminals that do not achieve significantly high availabilities. Indeed, by helping these terminals achieve the slightly higher (but ultimately still poor) availabilities the efficiency of the entire network is dramatically decreased. The terminals which are not significantly faded are forced to operate with poor packet efficiency. This results in poor network efficiency and 'wasted' network resource (in this case bandwidth).

7.4.2 Simple FMT systems

This system introduces the adaptive code rate as described previously. It is important to understand that this system will not improve the total availability when compared to the $\frac{1}{2}$ code rate from the traditional system. However, using this procedure it is possible to achieve these availabilities with a much lower bandwidth requirement.

Figure 7.12 shows the number of time slots per second required to supply the 1 Mbps minimum bandwidth requirement to all terminals. This is for no fade mitigation techniques applied (dashed line) when configured with the $\frac{1}{2}$ rate code. The thin sold line is the number of slots per second required to provide a minimum requirement of 1 Mbps but a desirable requirement of 2 Mbps (as described in section 7.3.2). The thick solid line represents the total number of slots available on the transponder per second. The histogram on each plot shows the distribution of the bit rate provided to each terminal over the day.
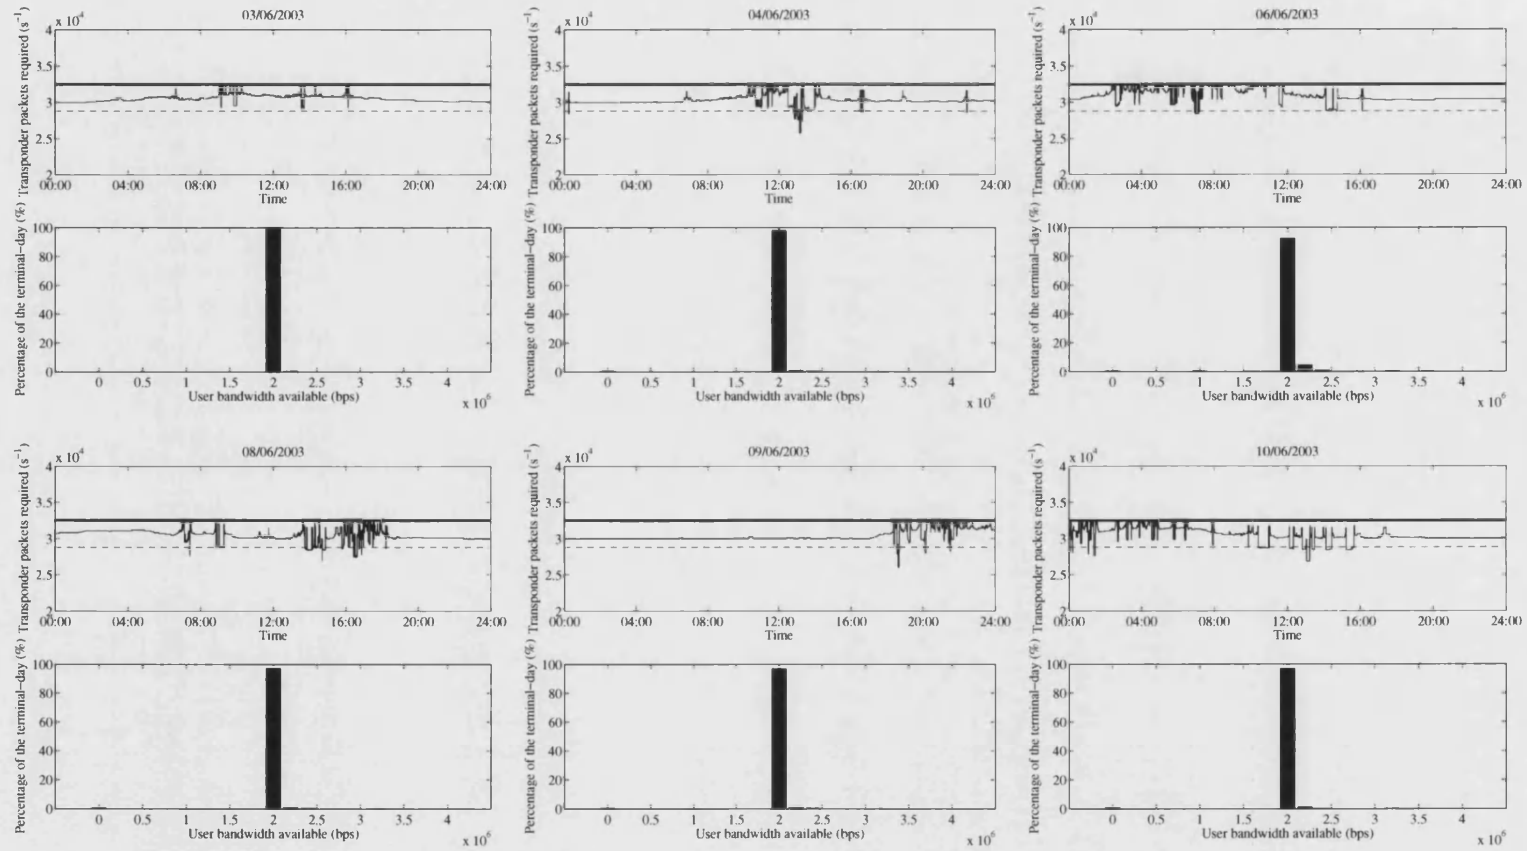
Figure 7.12: The improvement in the bandwidth requirement for the simple FMT system.
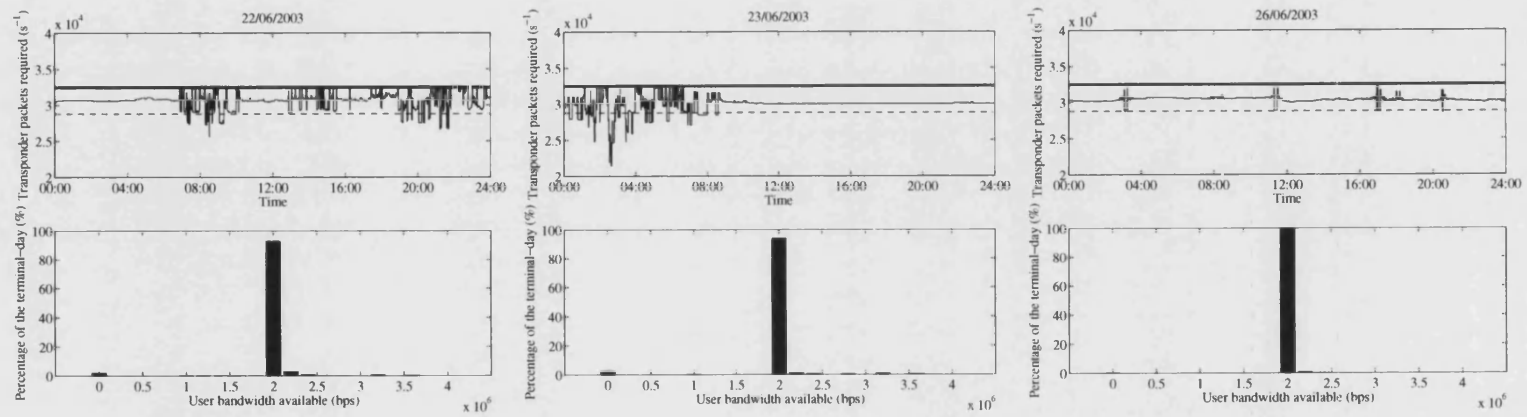
Figure 7.12: The improvement in the bandwidth requirement for the simple FMT system.

As can be seen, by varying the code-rate to the most appropriate configuration at any instant in time, the desirable bandwidth (2 Mbps) can be provided to most terminals. In providing this a small increase in network throughput is required over that to provide a minimum service (1 Mbps) to the same number of terminals in the traditional system. In essence the network provision has been doubled for a small increase in the amount of data packets required.

The histograms also show that few terminals are being provided with their minimum capacity and most terminals are being provided with the highest level of service. However, the terminal availabilities are still fairly low. There are still fades that are so deep they can not be combated using 'instantaneous' FMTs, this results in the availabilities shown in figure 7.13.
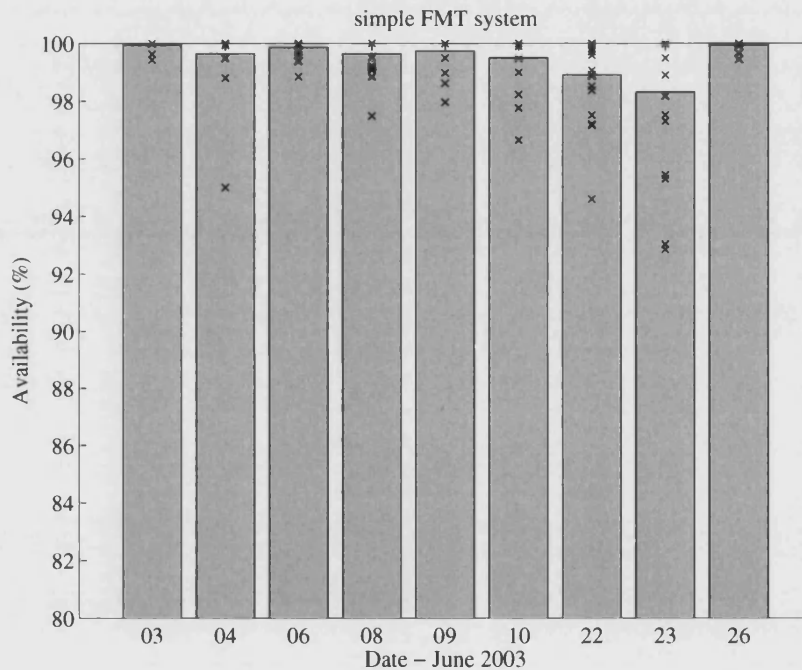
Figure 7.13: The availability of the network configured with simple FMTs.

7.4.3 Simple FMT systems with an A-TDMA time diversity system

In order to improve the terminal availabilities a time diversity system is implemented allowing terminals to receive content before it is required. In essence this can be used to provide a service availability in excess of the link availability. Typically, there will always be fades that are too deep to correct with an instantaneous FMT technique.

169

In this configuration the terminals bargain for the spare capacity using the game theoretic framework previously described. In this instance it is assumed that the forecast is perfect i.e. the forecasted time series is simply the link time series with no scintillation component.

However the forecast time series includes a 2 dB scintillation fade margin. Since the scintillation component is random it is impossible to predict instantaneously and hence the only way to combat scintillation fading is to introduce a small fade margin.

In order to illustrate the procedure, figure 7.14 shows the results for one terminal on one day. Figure 7.14(a) shows the link attenuation and the forecasted time series, since the forecast is assumed to be perfect these time series correlate well. Figure 7.14(b) shows the code rate that the link is configured to use, both as a forecasted level and as an actual level. A state of 0 represents the $\frac{7}{8}$ code through to state 4 representing the $\frac{1}{2}$ code, and a state of 5 represents a fade that can not be combated instantaneously.

Figure 7.14(c) shows the number of slots per second that are allocated to the terminal, both for 'data' which provides the direct content, and the allocation for 'buffering' which provides the content during link outages. One particular interesting point in this figure is the buffer allocation at 13:00. The system is configured to run from the buffer as there is a forecasted fade, however the terminal is able to detect it has the ability to connect to the network allowing further buffer data to be streamed to the terminal, whilst still providing content from the buffer. The terminal will not switch back to direct content provision since the link can be considered to be poor because of the forecasted fade.

The terminal buffer content level is shown in figure 7.14(d). This shows the provision from both the direct content and the buffered content. This also demonstrates one problem with this technique. The system requires the forecasted time series includes a scintillation fade margin. In many cases the terminal is assumed to be in a 'worst' code rate state than it is, so is allocated more slots than it technically requires. This results in each packet having a greater data capacity than expected which results in a greater provision than is strictly required.

(a) Link attenuation performance

(b) The channel code rate state

(c) The number of slots allocated per second

(d) Bandwidth provided to the terminal
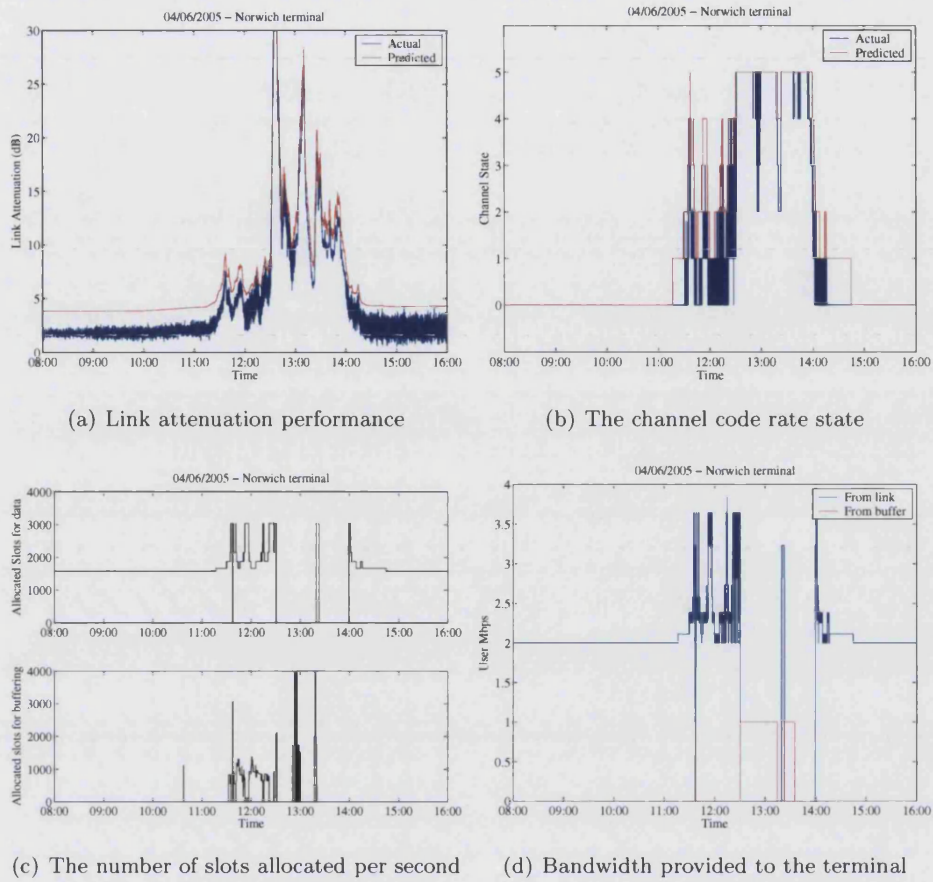
Figure 7.14: An example of the time diversity system.

The goal of this system is to improve upon the system availability and to demonstrate the improvement in the availability of individual terminals. Figure 7.15 is a scatter plot of the daily outage with and without time diversity for each of the terminals.

171

Figure 7.15: The improvement in the availability for a system with time diversity over a system with no time diversity.

As can be seen there is a drastic reduction in the terminal outages, for example the terminal which had an outage of just over 10% of the day ($\sim$145 minutes) has had that outage reduced to around 5% ($\sim$65 minutes). Also of interest is the transponder throughput in packets per second. Figure 7.16 shows the number of packets per second required to supply the 1 Mbps minimum bandwidth requirement to all terminals with no fade mitigation techniques (dashed line) configured with the $\frac{1}{2}$ rate code. The thin solid line is the number of packets per second required for the simple FMT to provide the 2 Mbps desired bandwidth (with no time diversity). The thick solid line represents the total number of packets available on the transponder per second. The red line shows the number of packets that are required to provide the time diversity technique.

172

Figure 7.16: The transponder throughput required for the time diversity technique.

Figure 7.16: The transponder throughput required for the time diversity technique.

As can be seen the increase in the number of packets required is relatively small, whilst the increase in the availability (decrease in outage) is significant. This demonstrates the level of improvement that is possible from this sort of time diversity technique. This simulation assumes that it is possible to perfectly predict the link attenuation with one second accuracy, this is evidently not the case. The impact of this is discussed in the following sections.

7.4.4 Simple FMT systems with an A-TDMA time diversity system - forecast errors

Attenuation forecasts will suffer from a number of different possible errors, these have been split into three different types [Watson and Hodges, 2004]:

- Temporal errors occur where the forecast is correct in terms of spatial organisation but is advected at the incorrect velocity. Time series will show fades arriving slightly before or after expected. This is the most likely type of error when frontal rain is considered, since most NWP techniques predict the presence of frontal rain well.

- Spatial errors occur when the forecast is initially correct but the birth and death of rain-cells is not correctly modelled. This results in forecasts that 'miss' fade events. These are most likely to occur with small-scale convection, which current NWP systems model relatively poorly.

- Fade depth errors occur if the temporal and spatial arrangement of the forecast is correct but the depth of the fade is incorrectly modelled. This may be the result of assuming an average raindrop size distribution or the inherent errors in the conversion from rainfall rate to attenuation level, or it can simply originate from an incorrect estimate of the rainfall rate.

It is very difficult to isolate the error type from the 'forecasted' and 'actual' time series. However, it is possible to artificially introduce the errors into the forecasted time series in order to provide a very simplified understanding of the error resilience of time diversity techniques.

Temporal forecast errors

In order to simulate the temporal forecast errors the forecast time series is delayed or advanced with respect to the actual time series. This provides a simple approximation of a forecast with a temporal error. This approach is simplified considerably since it assumes that the entire rainfield is incorrectly advected by the same amount yet it can be considered to provides an estimation of the level of error resilience.

An example day was chosen (22 June 2003) and a number of simulations were performed. In each simulation a different level of forecast temporal fade error was assumed ranging from 30 minutes late to 30 minutes early. The daily outage for each terminal is plotted as a function of the temporal fade error in figure 7.17, where each line represents a different terminal.
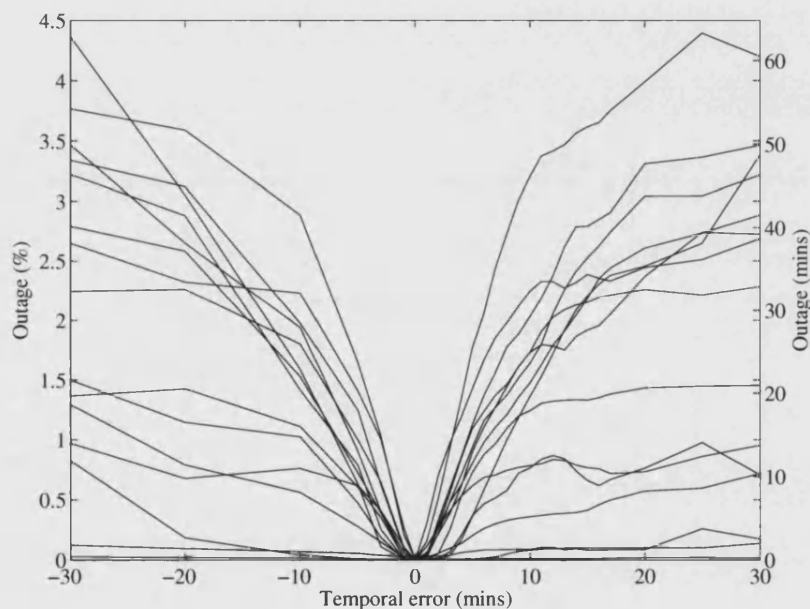
Figure 7.17: The effect of temporal errors upon the terminal outages (22 June 2003).

Unsurprisingly as the temporal error increases so do the network outages. However, the most important part is the performance of a time diversity system with temporal errors when compared to a system with no time diversity. The outage improvement (in daily %) is shown in figure 7.18. This is the algebraic difference between the outage with the time diversity system and with no time diversity system.
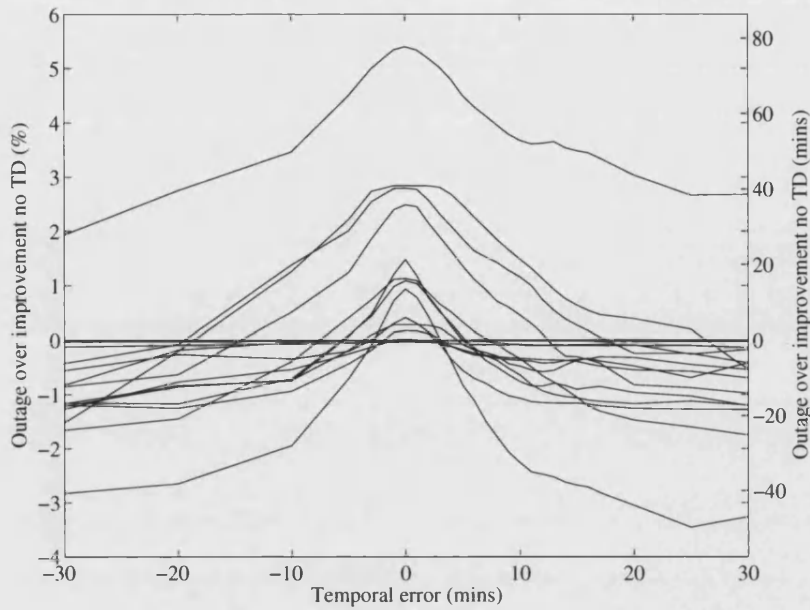
Figure 7.18: The improvement in terminal outages using a time diversity system (22 June 2003).

As can be seen there is a significant improvement available even with a relatively large forecast temporal error. The bound of the improvement for each of the 18 terminals is shown in table 7.3.

The shape of the error resilience curves for each of the receivers shown in figure 7.18 is a characteristic shape which can be explained using several simplifications. The outage can be considered to be caused by a number of fades which are shorter than the temporal error, N_s, and a number of fades which are longer than the temporal error, N_L. If it is also assumed that the resource allocation system is good enough to be able to buffer all of the fades then each individual fade results in the outage as described in figure 7.19, where t denotes the temporal fade error and f_L denotes the length of the fade.
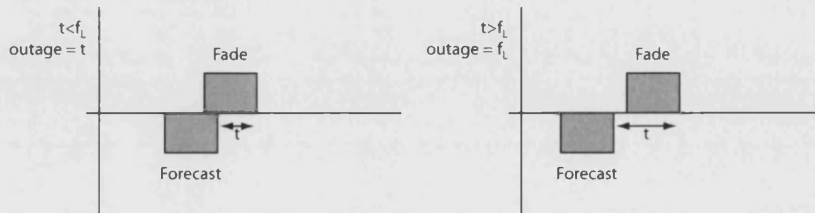
Figure 7.19: A simplified model of the effect of the temporal errors in the fade forecasts.

Table 7.3: The error bound required to provide an improvement over not using time diversity

Terminal Number	Maximum outage improvement (%)	Temporal error bound (mins)	
		Late forecast	Early forecast
1	0.000	N/A	N/A
2	0.000	N/A	N/A
3	0.282	3.72	4.51
4	0.000	N/A	N/A
5	0.000	N/A	N/A
6	0.172	1.91	2.63
7	0.414	3.06	3.07
8	1.131	7.57	7.29
9	0.000	N/A	N/A
10	1.089	5.59	4.98
11	2.794	19.4	17.57
12	0.936	2.78	3.05
13	0.016	1.07	1.03
14	0.000	N/A	N/A
15	1.481	5.46	5.23
16	2.487	14.4	12.47
17	2.840	18.6	26.3
18	5.400	> 30	> 30

Therefore the daily outage, E, for a given temporal error, t, can be considered to be the sum of the outages from each individual fade, i.e. :

$$E(t) = N_L(t) \cdot t + \sum_{i=1}^{N_s(t)} f_{Li} \qquad (7.7)$$

This explains several phenomena observed in figure 7.18. For large temporal errors the outage improvement tends to a constant value, since at this point $N_L = 0$. All fades are considered 'short' and the outage is simply the sum of the length of these fades. The decrease in the outage improvement is not linear since as the value of t increases the ratio $N_L : N_s$ varies.

If the performance is considered from a network-centric viewpoint the network outage can be calculated in terms of the average outage across the network over the day. It can be seen that there is a network improvement over no time diversity for temporal forecast errors between $\sim$10 minutes late and $\sim$10 minutes early. The network performance as a function of the forecast error is shown in figure 7.20.
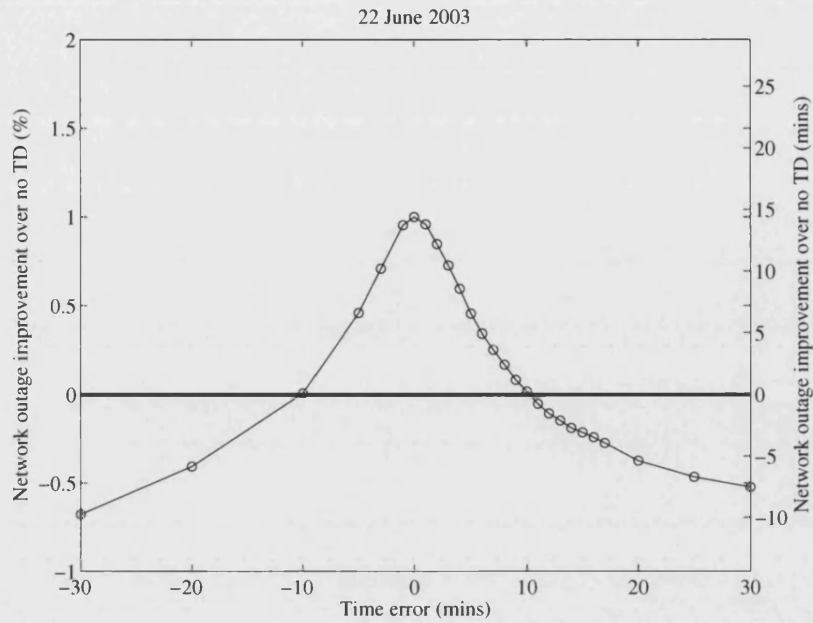
Figure 7.20: The improvement in average network terminal outage using a time diversity system (22 June 2003).

To summarise, for a network-centric outage improvement upon the 22 June 2003 it is necessary to predict fades with an accuracy of approximately 10 minutes.

Figure 7.21 shows the same analysis for the other simulation periods. The results are summarised in table 7.4.
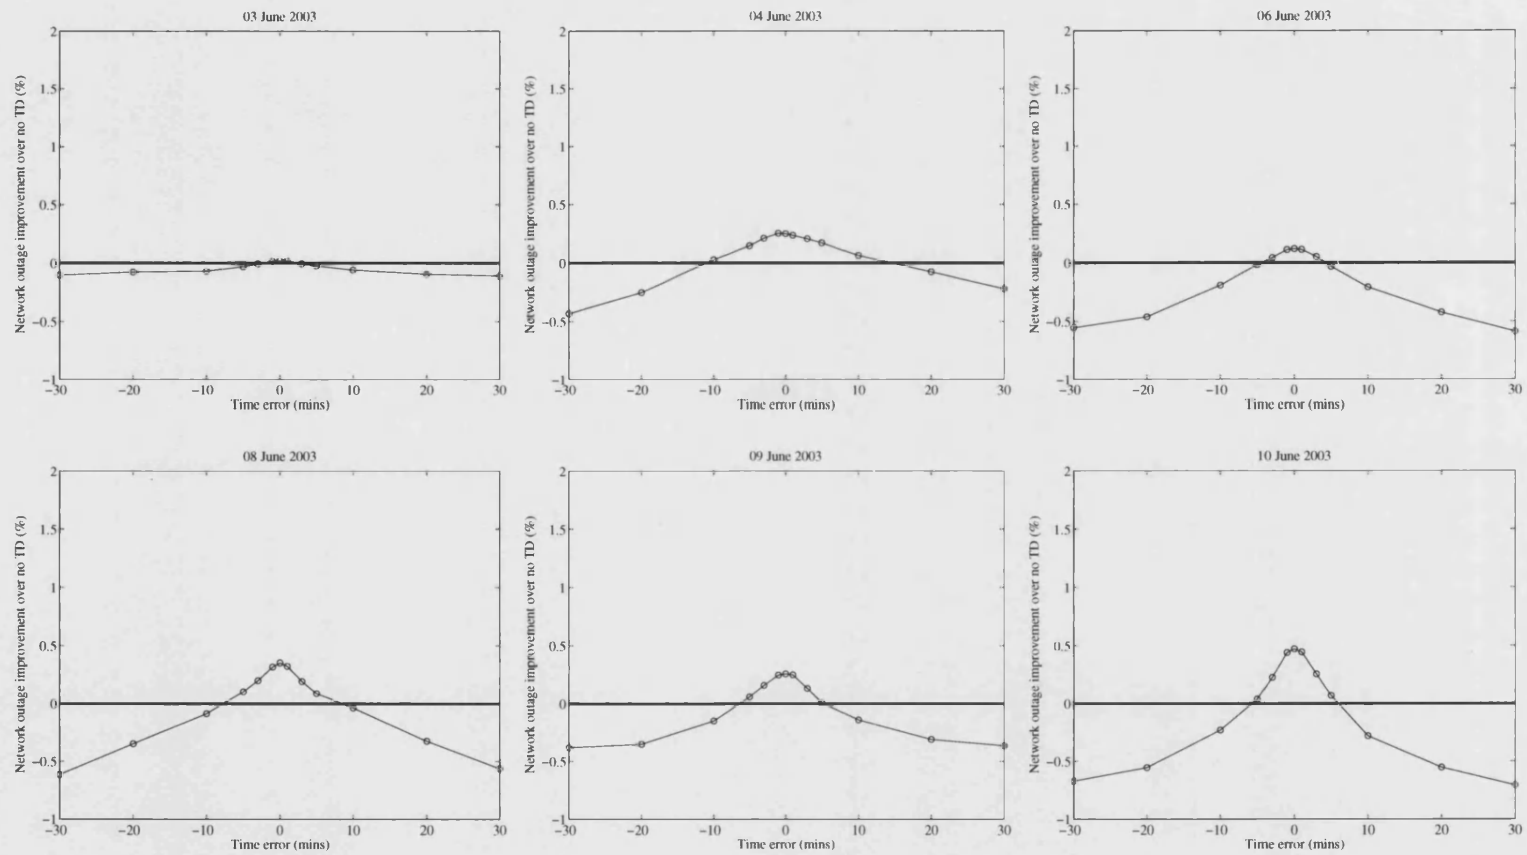
Figure 7.21: The improvement in average network terminal outage using a time diversity system.
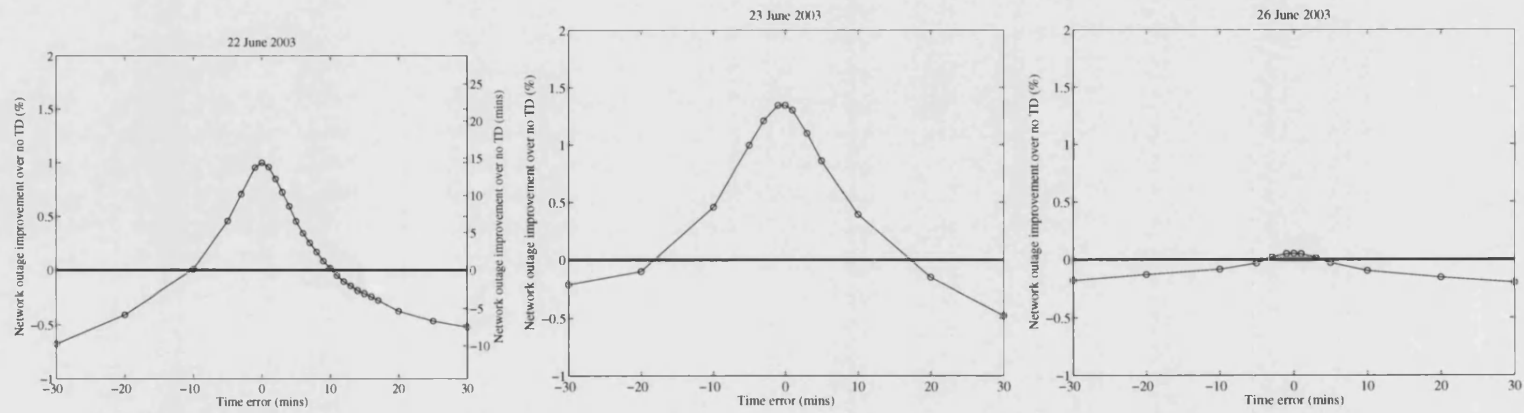
Figure 7.21: The improvement in average network terminal outage using a time diversity system.

Table 7.4: The error bound required to provide a network improvement over not using time diversity

Date (June 2003)	Maximum network outage improvement (%)	Temporal error bound (mins)	
		Late forecast	Early forecast
3	0.018	3.12	3.58
4	0.253	10.90	13.98
6	0.119	4.35	4.14
8	0.355	7.74	8.53
9	0.270	5.97	5.13
10	0.471	5.66	5.98
22	1.002	10.17	10.25
23	1.343	17.27	16.50
26	0.051	3.76	3.70

The overall network performance does not provide a good indication of the improvement for individual terminal outage improvements since the metric is averaged across all terminals. The network performance provides a good summary of the fade timetabling accuracy required to demonstrate an improvement for the entire network over not using time diversity. However, the performance of individual terminals as the temporal fade error increases provides a better indication of the resilience to temporal fade errors.

A scatter plot of the outage improvement for individual terminals over the nine days is shown in figure 7.22. Analysis of the performance of individual terminals provides several interesting insights. By considering the small availability improvements ($< 1.5\%$) it can be seen that the performance is close to a linear relationship. This linear relationship can be explained since to create relatively small outages there are likely to be single small fades and hence the error bound is the length of the fade. However, longer outages are likely to be comprised of multiple fades which complicates the analysis and also increases the probability of the resource allocation scheme failing to buffer the fade. More simulations are required to provide a detailed insight into this behaviour.
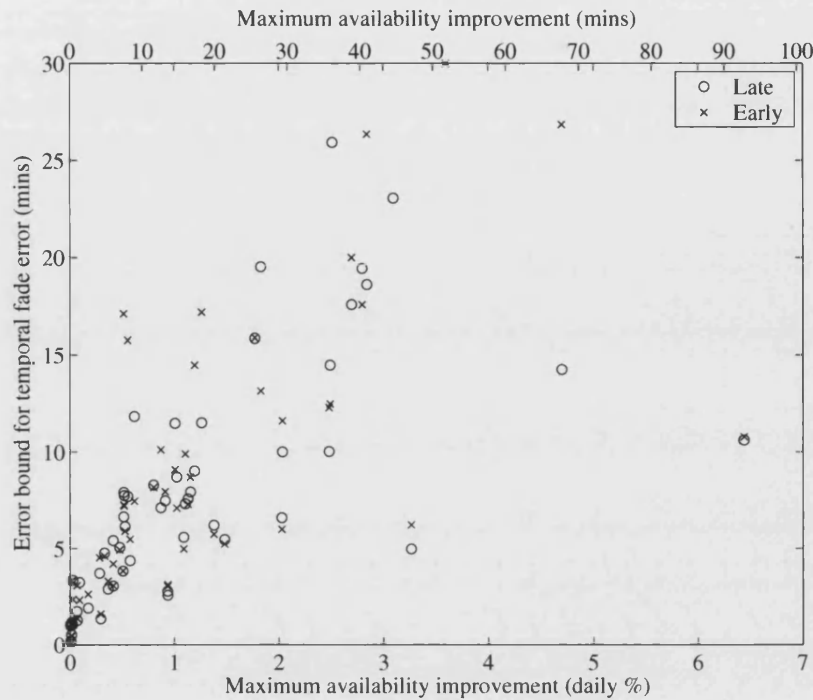
Figure 7.22: The error resilience of individual terminals to temporal fade errors.

Fade depth errors

To perform a high level simulation fade depth errors it is possible to introduce fade level errors by altering the coefficients in the relationship between the rain attenuation, A_r and the rainfall rate, R. Typically the equation is a power-law relationship of the form $A_r = aR^b$. In order to introduce a fade depth error this is changed to $A_r = a'R^{b'}$, where $a' = a_{scale} \cdot a$ and likewise $b' = b_{scale} \cdot b$. The values a_{scale} and b_{scale} represent the level of error, where 1 represents no error and 0.8 and 1.2 represent an error of -20% and +20% respectively.

These errors were applied to the 'a' and 'b' coefficients individually. Initially the error was simulated only on the 'a' coefficient. This simulates either an incorrect estimate of the rainfall rate or the variation from the average raindrop size distribution[6]. The results for the 22/06/2003 are shown in figure 7.23.

Interestingly, by overestimating the rainfall rate (i.e. $a_{scale} > 1$) there is only one terminal whose performance degrades significantly. However, the network efficiency

---

[6]This could also be caused by the incorrect rain height
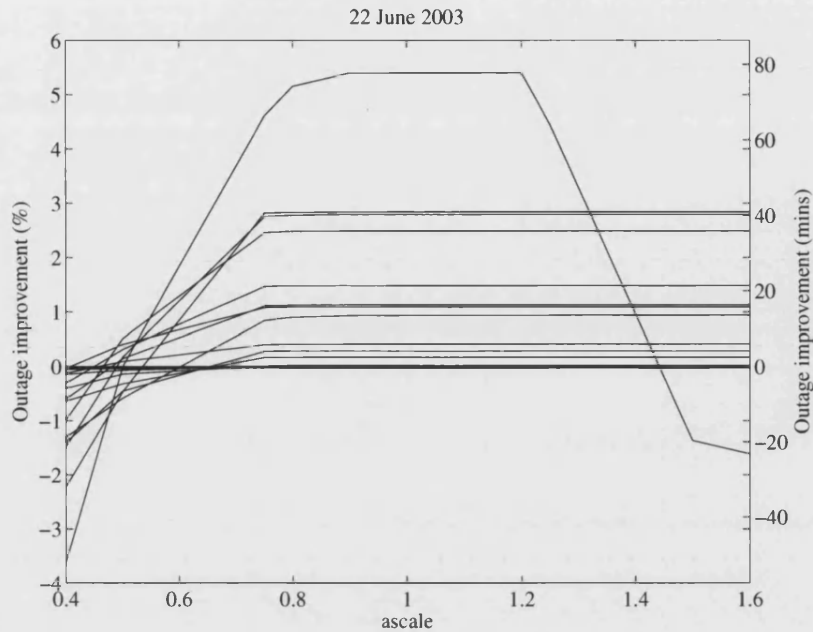
Figure 7.23: The improvement in terminal outage using a time diversity system with fade errors.

decreases since the system is buffering information which in reality the network does not need to. This results in content being provided from the buffer as opposed to the direct stream, hence the link which the forecast says is faded is in reality available but not used.

If the rainfall rate is underestimated (i.e. $a_{scale} < 1$) there is a distinct fall off in network performance. However, this does not occur until an underestimate in excess of 20 %. This implies that we can underestimate the rainfall rate by 20 % with **no** decrease from the maximum attainable performance. However, it can be seen from figure 7.23 we can have much poorer performance ($\sim$40%) before the majority of terminals are disadvantaged. This will be partly due to the damping provided by the scintillation fade margin.

This example demonstrates tremendous error resilience to poorly estimating the rainfall rate. This is to be expected at 30/40 GHz, since the rain attenuation causes such severe fading. In order to cause an outage a low rainfall rate is required, hence there must be a significant underestimation of the rainfall rate in order to not cause a predicted outage. However, even systems with large fade margins and hybrid FMTs are unlikely to require discrimination between very large rain fades.

In order to further simulate fade errors the value of the b_{scale} error value was linearly

scaled. It is worth noting that since this is a power term the variation of the attenuation with this term can be significant. The variation in the performance is shown in figure 7.24.
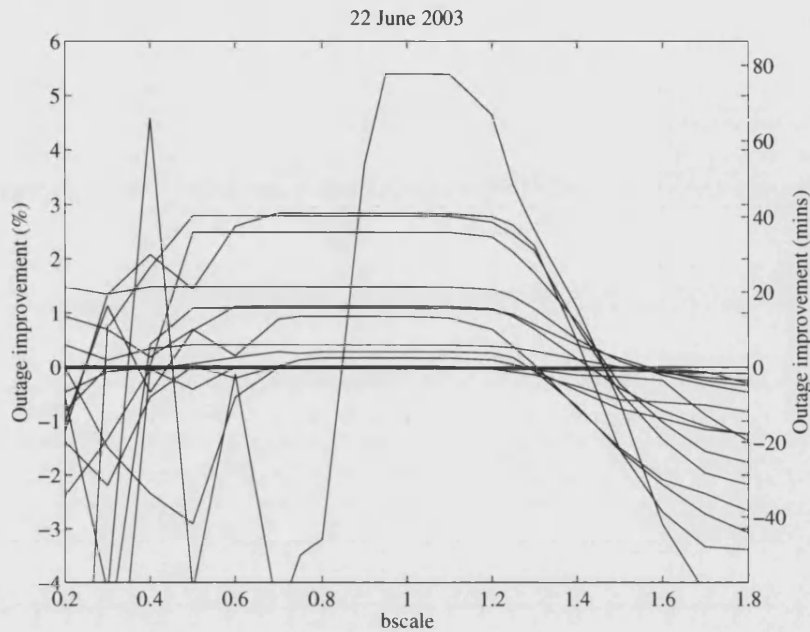
Figure 7.24: The improvement in terminal outage using a time diversity system with fade errors.

From this figure it can be seen that it is possible to overestimate the b coefficient by 20 % with no deviation from the maximum possible improvement. Since there is a power law relationship between the attenuation and the b coefficient, significantly decreasing the b coefficient creates a large amount of variation in the attenuation and hence the network performance. Even with the highly variable results below a b_{scale} value of 0.6 it can still be seen that there is no reduction in performance for an underestimation of the b coefficient of around 20 to 25 %.

In conclusion it can be seen that the system is particular resilient to fade depth errors, to a far greater degree than the temporal fade errors.

Spatial errors

It is very difficult to create realistic synthetic spatial errors and these have not been considered for this project. It may be possible to simulate spatial errors using image processing techniques such as image warping. These techniques could be used to

identify rain events which could then be artificially advected or removed as required.

Chapter 8

Conclusions

In the preceding chapters a novel approach to radiowave propagation modelling has been demonstrated. This approach differs from other techniques since the radiowave propagation modelling is performed by making better use of knowledge of the meteorological environment. In particular there is a demonstration of the application of this technique with meteorological information extracted from numerical weather prediction (NWP) techniques and weather radars.

This technique has concentrated upon the application of radiowave propagation modelling at EHF and SHF frequencies, primarily for slant-path applications, although a brief exploration of the application of the technique for point-to-point links has been performed. At these frequencies the attenuation from atmospheric gases, clouds and rain are very significant and hence the meteorological conditions are a good indicator of the radiowave propagation environment.

The current focus of radiowave propagation modelling for Ka-band applications relies upon a limited resource of beacon measurements from a small selection of areas. This reduces the ability of such generators to provide time series for many different meteorological climates. These 'synthetic time series generators' typically provide time series for individual sites and do not have the ability to provide time-coincident multiple-site time series with the correct spatial correlation. This restriction means that they are not suitable for the modelling of dense networks since each link must be placed far enough away to be statistically independent.

This is further complicated since, for example, a satellite link in London and a satellite link in Bath may be instantaneously statistically independent. Equally on occasions the performance may be correlated for example with a cold front coming in from the south

west across the UK, where the performance in London may be a time-shifted version of the performance in Bath. At other times the performance may be identical for example during a stable summer high pressure. It is this variability in instantaneous attenuation that is very hard to predict without understanding the meteorological conditions. It is even more difficult to suitably characterise this behaviour with an 'on-average' statistical approach.

The technique presented in this thesis also has the ability to work in a forecast paradigm. The NWP forecasts can be used to characterise the radiowave propagation conditions in order to provide a radiowave propagation forecast. This is a unique ability. The MM5 NWP system was used to provide the inputs to the PFE, this was installed and adapted to work on a linux cluster, work was also done to evaluate the optimum configuration in terms of coverage area and grid-scale. Since the MM5 is a research model there was a significant amount of work to be done in order to select the correct model options (in-terms of cloud micro-physics models, boundary layer models, convective parameterisations etc.).

Examples of the time-coincident multiple site time series have been shown (figure 6.1). The ability to provide 'maps' of attenuation has also been demonstrated (figure 6.2).

Link attenuation time series from this approach have been shown to demonstrate the correct first order characteristics, as measured by the cumulative density function. Comparison of three-month time series for three locations with the annual time series distribution from the ITU-R recommendations have been presented (figure 6.7). The time series from the generator show some variation from the ITU-R prediction but this can be considered typical of the seasonal variation expected from a short time series.

The time series have been shown to demonstrate the correct second-order characteristics, as measured by the conditional fade slope probability (figure 6.9). Good agreement with the conditional fade slope probability from ACTS, Italsat and OPEX propagation experiments is also shown.

To verify the ability of the system to produce accurate time series a comparison of generated time series with those from beacon measurements is important. Initially, historical time series were retrospectively created using radar and NWP data. Good agreement with the 20.7 GHz GBS beacon was seen (figures 6.10 and 6.11). To more accurately simulate the system in a 'tactical' forecast environment the outputs from a real-time MM5 system were used with the propagation forecast engine and compared to the beacon measurements over a one-month period (figures 6.12 to 6.15). This demonstrated a skill score (EDS) of around 0.65 for the Chilbolton receiver for a forecast up to 48 hours in the future.

An analysis of the importance of the grid scale upon which the meteorological information is provided has been performed. The improvement in forecast 'quality' as the grid scale decreases from 18 km to 6 km and then to 2 km has been demonstrated (figure 6.3). When creating historical time series of attenuation maps for the purposes of network simulation it is possible to introduce a representative stochastic small-scale structure. This small scale structure is one possible solution of the 'actual' small scale structure and therefore should not be used within a forecast system.

A multi-fractal downscaling technique has been demonstrated which exploits the linear variation of the log-structure-functions with log-scale (figure 5.3). The procedure can be improved by applying a spatial filter in order to reduce the blocking effects. The technique has demonstrated the conservation of the statistical moments (figures 5.11 and 5.12). The downscaled fields also demonstrate a realistic PSD (figure 5.9).

A mono-fractal temporal interpolation scheme has also been demonstrated. This technique introduces a stochastic small scale structure which provides a more realistic PSD for higher frequency components than linear interpolation techniques.

An example network simulation has also been described. Using this technique the application of time coincident multiple site time series have been demonstrated within a simulation framework.

A time diversity technique using propagation forecasts has been demonstrated. A brief high-level assessment of the error-bounds has also been performed. It should be noted that is heavily dependent upon the network configuration. The network configuration simulated was found to be more tolerant of fade depth errors than temporal errors.

Chapter 9

Further work

This chapter provides a brief overview of several areas where the technique can be improved, as well as other areas of radio science where this technique can be exploited.

9.1 Improvements in NWP modelling

In essence the core of the technique, particularly within a forecast paradigm, is the ability of NWP techniques to accurately model the meteorological environment. Of particular importance is the modelling of convection. At present this is not resolved as accurately as the technique ideally would require, particularly at forecast lengths in excessive of 24 hours. The propagation model also has the ability to exploit more information about rain events (e.g. raindrop size distributions, melting layer information and the three dimensional structure of rain) which some NWP techniques are able to provide. The new versions of WRF and MM5 are able to provide vertical profiles of the mixing ratio of mixed-phase liquids, this should be of considerable interest to enhance the model. Figure 9.1 shows vertical cross-sections along a latitude of 51°N, these sections show the mixing ratio of vapour, cloud-ice, cloud-water and rain-water from an example WRF model run.
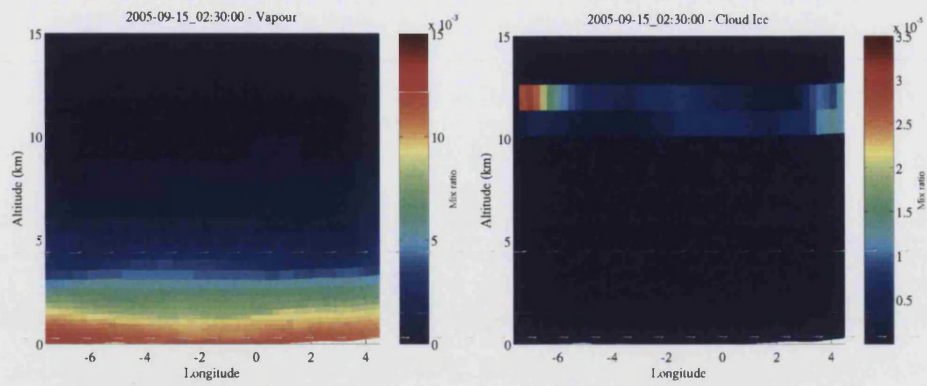
(a) The vapour mix ratio over all heights.

(b) The ice mix ratio over all heights.
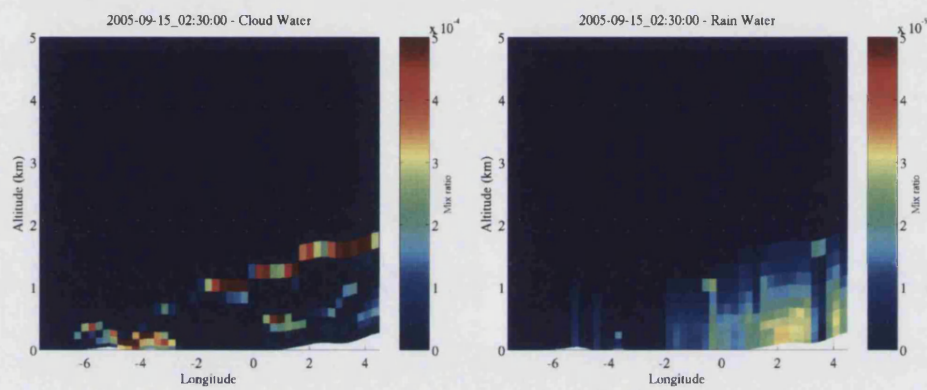
(c) The cloud water mix ratio over heights of interest.

(d) The rain water mix ratio over heights of interest.
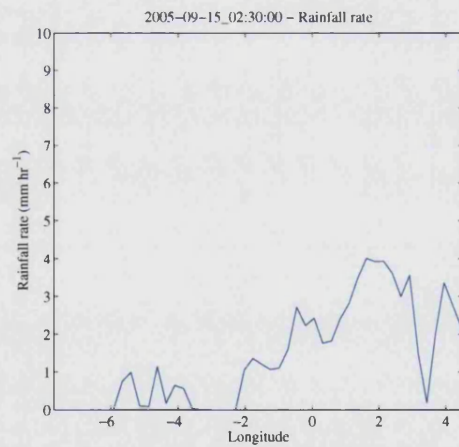
(e) The rainfall rate across the cross section.

Figure 9.1: The vertical cross section of the mixing ratio of mixed-phase liquids - WRF model 2005/09/15 (02:30UT) at a latitude of 51°N.

New large-coverage mesoscale models are also becoming available. For example, the North Atlantic and European (NAE) model run by the UKMO provides a large enough coverage area to encompass European antenna footprints, far in excess of the size of the current Eutelsat HotBird satellites. The NAE model is currently run upon a 12 km grid, however a 2 km version is expected in the near future. An example NAE temperature field is shown in figure 9.2
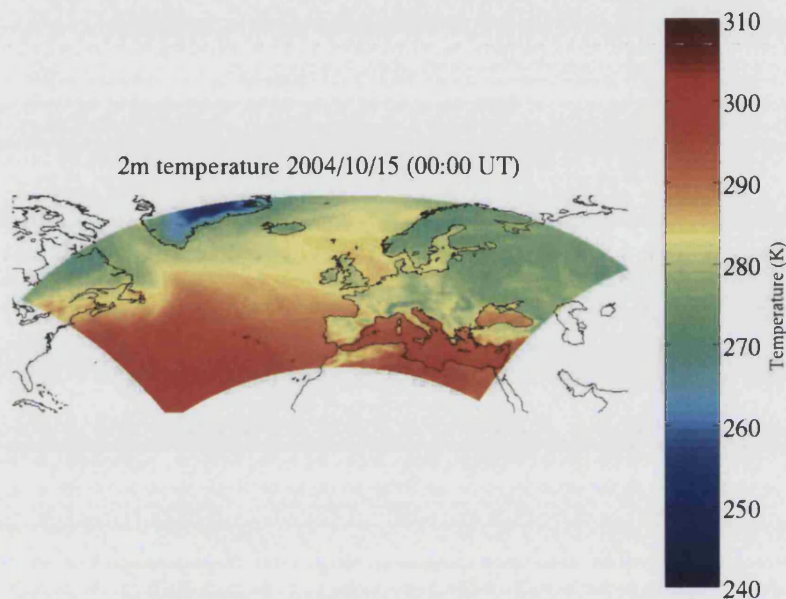
Figure 9.2: The 2 m temperature field from an example NAE field.

It is important that the propagation modelling techniques are able to improve as NWP techniques improve, to exploit all available data.

9.2 Improvement in the scintillation variance modelling

The scintillation variance is currently not modelled using a physical model. This is an area where the technique could be improved; typically NWP measurements are not at a vertical resolution to be able to resolve C_n^2 to a satisfactory instantaneous value [Vasseur, 1999]. In essence, the small-scale turbulent inhomogeneities that cause scintillation are far too small to be resolved by NWP techniques. This is unlikely to change in the near-future since resolutions on the scale of 10 m are required.

However, further work directed at using indicators of convection as indicators of high

scintillation variance should prove fruitful. This has already been demonstrated with the application of heavy clouds [Tervonen et al., 1988, van de Kamp, 1999]. However convective-available-potential-energy (CAPE) and the lifting index provide a very good indication of the level of convection, particularly from the boundary layer to a level of 800 mb. The presence of convective inhibitions (CINs) are less important since up to the CIN there will still be considerable turbulent mixing and hence scintillation.

9.3   Further terrestrial point-to-point investigations

The application of the work to terrestrial point-to-point applications has only been briefly investigated. This technique has a clear application within the modelling of very dense overlay / backhaul networks where traffic load balancing and the ability to increase power and bandwidth efficiencies (and in addition frequency re-use and spectral efficiency) are very important.

Comparison with data from mobile telephony overlay networks would provide an interesting insight into this technique. From an intuitive perspective the issues with spatial averaging are less of a concern for terrestrial links since the radiowave path itself will average the propagation conditions.

In addition the technique could be applied for long low-elevation slant-paths (with elevation angles around 4°), these links are unique in that they can not be considered 'terrestrial' paths since the link exists above the melting layer and they can not be considered typical slant-paths since the elevation angle is so low. This unique problem is a perfect example of the reconfigurability of this technique; in essence with the technique described in this thesis there are very few EHF / SHF propagation configurations that can not be evaluated.

9.4   Clear air fading and fog

In order to improve the modelling of terrestrial links several propagation conditions should be taken into account. These are conditions that are relatively insignificant for typical earth-space links. The presence of fog is a good example of this. For millimetre-wave terrestrial links, fog should be considered since although the water-content is small, the path length could be very large. It is expected that the attenuation due to fog may be accurately represented by the cloud attenuation models.

A particular requirement for terrestrial link modelling is the ability to model clear-air fading, typically from large scale refractive effects. One area for investigation is the ability of this technique to model the vertical refractive index profiles accurately enough and at a high enough vertical resolution to be able to predict refractive duct fading.

9.5 Radiometric modelling

For many applications it is more appropriate to measure the link performance radiometricaly (i.e. as a noise temperature). One particular application is with deep-space links, such as those used by NASAs Jet Propulsion Labs (JPL). With recent exploration missions, in particular the Mars exploration missions (Mars Reconnaissance Orbiter, Spirit and Opportunity rovers and the Mars Telecommunications Orbiter), there is a considerable demand for high bandwidth links. This demand is well served using links at Ka- band and above. Indeed, the MRO has the Ka- band demonstrator transponder.

Since Mars is a considerable distance from Earth the link time-of-flight is very high, a one-way time-of-flight which at times can exceed 15 minutes. This time-of-flight is too long to use reactive FMT and will typically require some degree of fade forecasting. However, for this particular application the system is significantly noise limited and hence a measure of the noise temperature is more appropriate than a link attenuation [Davarian et al., 2004].

9.6 Modelling of water vapour delay for Global Navigation Satellite Systems (GNSS)

The modelling of GNSS requires an approximation for the delay caused in the troposphere by the presence of water vapour. Initial work was performed by Bevis et al. [1996], which showed some interesting results, however the modelling was compromised by the NWP performance. Since this work has been performed the standard of NWP systems has improved dramatically and heuristically it can anticipated that it should be possible to produce more accurate forecasts of water vapour delay.

9.7 Further verification

Further verification of the forecasting system could also be performed. The new satellite platforms with EHF transponders including the A-EHF satellite, Syracuse 3A and the NASA Ka- band vehicles all provide an opportunity to further verify and improve the forecast performance.

It is also possible to retroactively investigate the performance of the system by creating historical time series (the NWP techniques could be run as a forecast system using archived boundary conditions) and retroactively comparing them to beacon measurements from the Italsat campaign. Since the boundary conditions have been archived by NCAR the data sets would have to be purchased. To consider the three year period of the Italsat campaign (January 1994 to March 1997) the 15.7 GB of boundary conditions are required. This is certainly a verification that should be considered. The meteorological modelling would also provide a valuable input into the scintillation modelling previously described.

9.8 Improved network modelling

Simple high-level network simulations have been performed as part of this project. However there is a considerable scope for improvements in network modelling. The techniques used to create realistic forecast errors need to be improved. Specifically, the techniques used to simulate realistic spatial errors must be improved.

The current simulation technique could also be expanded by performing simulations over longer time-scales and further exploring the solution space; e.g., the ratio of terminals to bandwidth etc.

Appendix A

Kohler curves and cloud droplet condensation

Droplets formed of a pure water droplet by condensation from a supersaturated vapour, without the aid of an aerosol[1], are said to have been created by a process called spontaneous nucleation. In this process a number of water molecules collide and form water droplets that are large enough to remain intact. In a vapour that is unsaturated very small droplets continually form and evaporate and do not grow large enough to become visible as a cloud of droplets. However, in supersaturated vapour droplets above a certain size will be supported and continue to grow. The size at which this behaviour is supported is given by Kelvin's formula. Kelvin's formula defines the relative humidity with which a droplet is in equilibrium, Kelvin's equation is:

$$\frac{e}{e_s} = \exp\left[\frac{2\sigma}{\rho R_v T r}\right] \tag{A.1}$$

where σ is the surface tension of water ($0.073\,\mathrm{N\,m^{-1}}$), ρ is the density of water, R_v is the gas constant of water vapour ($482\,\mathrm{J\,kg^{-1}\,K^{-1}}$), T is the temperature (K) and r is the drop radius.

The variation of the equilibrium drop size with supersaturation is shown in figure A.1. The largest size of droplets likely to created by the collision of water molecules is around $0.01\,\mu\mathrm{m}$, in order for these drops to not evaporate requires a supersaturation in excess of 12.5% [Wallace and Hobbs, 1977]. Supersaturations caused by adiabatic ascents rarely exceed 1%. This implies that droplets do not form in clouds by the spontaneous nucleation of pure water.

---

[1]An aerosol is a suspension of a liquid within a gas.
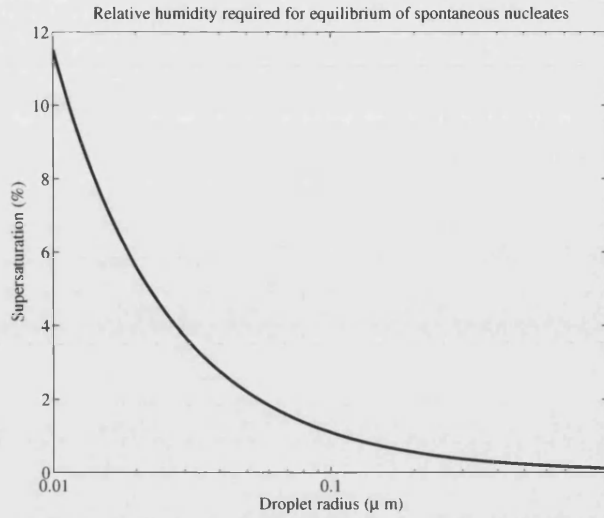
Figure A.1: Kelvin's equation at a temperature of 278 K.

Most cloud droplets are formed by the condensation of water vapour on aerosols. This process allows droplets to form and grow at a much lower supersaturation than required for spontaneous nucleation.

In this case the Kelvin equation becomes:

$$\frac{e}{e_s} = \exp\left[\frac{2\sigma}{\rho R_v T r}\right]\left[1 + \frac{imM_w}{M_s\left(\frac{4}{3}\pi r^3\rho - m\right)}\right]^{-1} \tag{A.2}$$

where i is the number of ions each molecule of the solute disassociates to in water, m is the mass of the solute, M_w is the molecular mass of water, M_s is the molecular mass of the solute.

By plotting equation A.2 as a function of droplet radius a family of curves called the 'Kohler curves' is produced. A set of these curves for three different dry mass of NaCl are shown in figure A.2.

Consider the curve in figure A.2 that refers to a 10^{-19} kg of NaCl. If the air is at a supersaturation of 0.2 % a solution droplet would form with a radius of approximately 0.1 μm. The size of this droplet is in equilibrium since if it grows the vapour pressure around the droplet would rise above the ambient air and the droplet would evaporate back to it's equilibrium size. However, if the equilibrium droplet evaporates a little the water vapour pressure would drop below that of the ambient air and the drop would grow by condensation to it's equilibrium size.
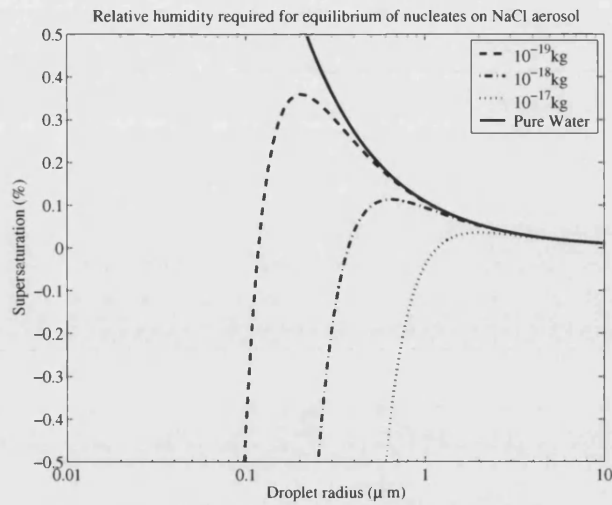
Figure A.2: Kohler curves for NaCl at a temperature of 278 K.

Consider now a droplet at the peak of the curve, a supersaturation of approximately 0.35 % and a droplet of approximately $0.2\,\mu$m. If this droplet was to evaporate slightly the supersaturation of the air would drop slightly below that of the ambient air and the droplet would condensate back to the original size. However, if the droplet should grow slightly the supersaturation of the air close to the drop would fall below that of the ambient air and the droplet would grow. As the droplet grows further the supersaturation continues to drop and the droplet grows rapidly.

A droplet which has passed the peak of the Kohler curve is said to be activated.

Appendix B

Data required for the PFE

Table B.1: Spectroscopic data for water vapour.

f_0	b_1	b_2	b_3	b_4	b_5	b_6
22.235080	0.1090	2.143	28.11	0.69	4.80	1.00
67.813960	0.0011	8.735	28.58	0.69	4.93	0.82
119.995941	0.0007	8.356	29.48	0.70	4.78	0.79
183.310074	2.3000	0.668	28.13	0.64	5.30	0.85
321.225644	0.0464	6.181	23.03	0.67	4.69	0.54
325.152919	1.5400	1.540	27.83	0.68	4.85	0.74
336.187000	0.0010	9.829	26.93	0.69	4.74	0.61
380.197372	11.9000	1.048	28.73	0.69	5.38	0.84
390.134508	0.0044	7.350	21.52	0.63	4.81	0.55
437.346667	0.0637	5.050	18.45	0.60	4.23	0.48
439.150812	0.9210	3.596	21.00	0.63	4.29	0.52
443.018295	0.1940	5.050	18.60	0.60	4.23	0.50
448.001075	10.6000	1.405	26.32	0.66	4.84	0.67
470.888947	0.3300	3.599	21.52	0.66	4.57	0.65
474.689127	1.2800	2.381	23.55	0.65	4.65	0.64
488.491133	0.2530	2.853	26.02	0.69	5.04	0.72
503.568532	0.0374	6.733	16.12	0.61	3.98	0.43
504.482692	0.0125	6.733	16.12	0.61	4.01	0.45
556.936002	510.0000	0.159	32.10	0.69	4.11	1.00
620.700807	5.0900	2.200	24.38	0.71	4.68	0.68
658.006500	0.2740	7.820	32.10	0.69	4.14	1.00
752.033227	250.0000	0.396	30.60	0.68	4.09	0.84
841.073593	0.0130	8.180	15.90	0.33	5.76	0.45
859.865000	0.1330	7.989	30.60	0.68	4.09	0.84
899.407000	0.0550	7.917	29.85	0.68	4.53	0.90
902.555000	0.0380	8.432	28.65	0.70	5.10	0.95
906.205524	0.1830	5.111	24.08	0.70	4.70	0.53
916.171582	8.5600	1.442	26.70	0.70	4.78	0.78
970.315022	9.1600	1.920	2.50	0.64	4.94	0.67
987.926764	138.0000	0.258	29.85	0.68	4.55	0.90

Table B.2: Spectroscopic data for the dry atmosphere.

f_0	a_1	a_2	a_3	a_4	a_5	a_6
50.474238	0.94	9.694	8.60	0	1.600	5.520
50.987749	2.46	8.694	8.70	0	1.400	5.520
51.503350	6.08	7.744	8.90	0	1.165	5.520
52.021410	14.14	6.844	9.20	0	0.883	5.520
52.542394	31.02	6.004	9.40	0	0.579	5.520
53.066907	64.10	5.224	9.70	0	0.252	5.520
53.595749	124.70	4.484	10.00	0	-0.066	5.520
54.130000	228.00	3.814	10.20	0	-0.314	5.520
54.671159	391.80	3.194	10.50	0	-0.706	5.520
55.221367	631.60	2.624	10.79	0	-1.151	5.514
55.783802	953.50	2.119	11.10	0	-0.920	5.025
56.264775	548.90	0.015	16.46	0	2.881	-0.069
56.363389	1344.00	1.660	11.44	0	-0.596	4.750
56.968206	1763.00	1.260	11.81	0	-0.556	4.104
57.612484	2141.00	0.915	12.21	0	-2.414	3.536
58.323877	2386.00	0.626	12.66	0	-2.635	2.686
58.446590	1457.00	0.084	14.49	0	6.848	-0.647
59.164207	2404.00	0.391	13.19	0	-6.032	1.858
59.590983	2112.00	0.212	13.60	0	8.266	-1.413
60.306061	2124.00	0.212	13.82	0	-7.170	0.916
60.434776	2461.00	0.391	12.97	0	5.664	-2.323
61.150560	2504.00	0.626	12.48	0	1.731	-3.039
61.800154	2298.00	0.915	12.07	0	1.738	-3.797
62.411215	1933.00	1.260	11.71	0	-0.048	-4.277
62.486260	1517.00	0.083	14.68	0	-4.290	0.238
62.997977	1503.00	1.665	11.39	0	0.134	-4.860
63.568515	1087.00	2.115	11.08	0	0.541	-5.079
64.127767	733.50	2.620	10.78	0	0.814	-5.525
64.678903	463.50	3.195	10.50	0	0.415	-5.520
65.224071	274.80	3.815	10.20	0	0.069	-5.520
65.764772	153.00	4.485	10.00	0	-0.413	-5.520
66.302091	80.09	5.225	9.70	0	-0.428	-5.520
66.836830	39.46	6.005	9.40	0	-0.726	-5.520
67.369598	18.32	6.845	9.20	0	-1.002	-5.520
67.900867	8.01	7.745	8.90	0	-1.255	-5.520
68.431005	3.30	8.695	8.70	0	-1.500	-5.520
68.960311	1.28	9.695	8.60	0	-1.700	-5.520
118.750343	945.00	0.009	16.30	0	-0.247	0.003
368.498350	67.90	0.049	19.20	0.6	0	0
424.763124	638.00	0.044	19.16	0.6	0	0
487.249370	235.00	0.049	19.20	0.6	0	0
715.393150	99.60	0.145	18.10	0.6	0	0
773.839675	671.00	0.130	18.10	0.6	0	0
834.145330	180.00	0.147	18.10	0.6	0	0

Table B.3: Model constants required for the Leitao-Watson rain attenuation model.

Frequency	Polarisation	Model Constants				
		s	t	u	v	w
Widespread Rain						
10 GHz	H	0.277	-0.295	1.307	-0.475	1.298
	V	0.248	-0.248	1.239	-0.456	1.258
11 GHz	H	0.366	-0.238	1.263	-0.455	1.267
	V	0.329	-0.267	1.255	-0.440	1.234
12 GHz	H	0.464	-0.238	1.259	-0.449	1.249
	V	0.420	-0.217	1.215	-0.433	1.214
17 GHz	H	1.070	-0.191	1.180	-0.456	1.193
	V	0.978	-0.202	1.162	-0.420	1.144
20 GHz	H	1.530	-0.212	1.171	-0.509	1.200
	V	1.390	-0.198	1.135	-0.470	1.147
24 GHz	H	2.280	-0.196	1.151	-0.497	1.173
	V	2.060	-0.212	1.131	-0.492	1.135
30 GHz	H	3.640	-0.166	1.118	-0.413	1.113
	V	3.280	-0.182	1.094	-0.430	1.088
40 GHz	H	6.168	-0.151	1.107	-0.456	1.008
	V	5.444	-0.170	1.070	-0.459	1.052
50 GHz	H	8.997	-0.138	1.102	-0.454	1.059
	V	7.785	-0.160	1.053	-0.461	1.023
Showery Rain						
10 GHz	H	0.449	-0.634	1.587	-0.563	1.318
	V	0.375	-0.567	1.490	-0.563	1.287
11 GHz	H	0.562	-0.624	1.553	-0.533	1.289
	V	0.474	-0.559	1.466	-0.538	1.245
12 GHz	H	0.684	-0.597	1.510	-0.561	1.277
	V	0.578	-0.572	1.454	-0.53	1.221
17 GHz	H	1.390	-0.472	1.341	-0.553	1.218
	V	1.170	-0.447	1.277	-0.515	1.148
20 GHz	H	1.900	-0.441	1.290	-0.468	1.162
	V	1.580	-0.417	1.222	-0.444	1.094
24 GHz	H	2.660	-0.414	1.237	-0.445	1.114
	V	2.240	-0.394	1.181	-0.452	1.075
30 GHz	H	3.920	-0.389	1.180	-0.410	1.051
	V	3.380	-0.375	1.134	-0.428	1.028
40 GHz	H	7.250	-0.361	1.123	-0.365	1.006
	V	6.180	-0.343	1.080	-0.397	0.966
50 GHz	H	11.253	-0.350	1.091	-0.329	0.963
	V	9.629	-0.329	1.052	-0.373	0.924

Appendix C

Hydrostatic approximations

In a hydrostatic state the upward force at a given point caused by the pressure gradient is balanced by the weight of the air above that point. This is true for large-scale situations where 99 % of the kinetic energy is contained in the horizontal components [Wallace and Hobbs, 1977].

The situation in figure C.1 shows a vertical column of air with a unit cross-sectional area.
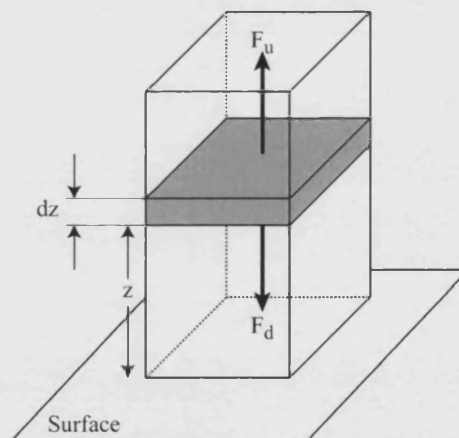
Figure C.1: Hydrostatic balance forces.

The mass of the slab of air between the heights z and $z + dz$ is $\rho(z)dz$, where $\rho(z)$ is the air density. The force downwards from this mass of air, F_d, is therefore $g(z)\rho(z)dz$, where $g(z)$ is the acceleration due to gravity. The pressure change from height z to height $z + dz$ is dp and since pressure decreases with height dp must be a negative

quantity. This also infers that the upward pressure on the bottom of the slab must be greater than the downward pressure on the upward face of the slab. Thus the net force on the slab due to the pressure gradient is upward with the value $-dp$.

If the slab is in hydrostatic equilibrium it can be seen that:

$$- dp = g(z)\rho(z)dz \tag{C.1}$$

Equation C.1 is known as the hydrostatic equation. By defining the pressure at height z as $p(z)$ it can be shown that:

$$- \int_{p(z)}^{p(\infty)} dp = \int_{z}^{\infty} g(z)\rho(z)dz \tag{C.2}$$

and since $P(\infty) = 0$

$$p(z) = \int_{z}^{\infty} g(z)\rho(z)dz \tag{C.3}$$

Hence the pressure at a level z is equal to the weight of the air in a vertical column of a unit cross sectional area lying above that area. If the mass of the atmosphere was uniformly distributed over the globe the surface pressure would be 1013 hPa.

These hydrostatic assumptions are true on a large scale, however on smaller scales convective effects become evident. Hydrostatic simulations are, by definition, unable to directly resolve these effects.

Appendix D

Accuracy of a real-time MM5 model

The real time forecasting system is an implementation of the NCAR MM5 model. The system has been compiled using Intel Fortran and C++ compilers, whilst the entire system runs on a ROCKS cluster (Makula version). The system runs forecasts out to 72 hours, upon two grids one at an 18 km grid size whilst a nested grid runs at a 6 km grid size. The time step is 45 seconds, whilst outputs are saved at hourly intervals[1]. The automated scripts are run via `crontab`. The first script is run at 06:00 UT which downloads the boundary conditions from NCEP (ftpprd.NCEP.noaa.gov), these are at 3-hourly intervals and from the GFS system at 1 degree resolution.

Once the boundary conditions have been downloaded `pregrid`, `regrid` and `interpf` are executed. At this point all the information that is required for the simulation is acquired and in the correct format. The next stage is to execute the MM5 model. The simulation run takes around 90 minutes depending upon the meteorological conditions that modelling requires (periods with significant mixed-phase rain or convective activity typically lengthen the simulation time). Once the simulation is completed the results are archived and `interpb` is executed. The results are then converted to a `netcdf` format before images and animations of the forecasts are uploaded to the internet. The final step is to execute the propagation forecast model and upload the information to the relevant locations.

It is worth noting that there is no observational analysis, i.e. `little_r` or `rawins` have not been used, the output from `regrid` goes straight to `interpf`. This creates a simple system with the minimum of downloaded material. Observational analysis will

---

[1] This has been recently been reduced to 15 minutes

significantly improve the accuracy of the system particularly for diagnostic variables such as rain and forecasts a significant time out from the simulation run.

The performance of the MM5 system to accurately model the prognostic variables (such as temperature, pressure and humidity) is best examined by comparing the outputs from the MM5 system with data from radiosonde ascents. Data from radiosonde ascents are available from the Met Office sites. However, the free data has a very low vertical resolution.

In July and August the TSaR Group at the University of Bath had been launching radiosondes for a NERC funded research project, the convective storm initiation project (CSIP). CSIP was run by a joint group including the institute for Atmospheric Science at Leeds University, the University Facility for Atmospheric Measurements (UFAM) and the Rutherford Appleton Laboratory (RAL) and the project was headed by Professor Alan Blythe (Leeds and UFAM).

Each radiosonde lunch directly measures temperature, pressure and humidity through the ascent and using a differential GPS system the wind direction and velocity are estimated. The radiosonde unit is shown in figure D.1(a) which shows the main measurement sensors and the transmit antenna, this is then attached to the parachute unit and a helium filled balloon D.1(b). When released these ascend at approximately $300 \, \text{m} \, \text{min}^{-1}$ through the atmosphere providing 'in-situ' measurements until the balloon bursts[2] - typically at around 60 mb / 20 km.

---

[2]During intensive CSIP operations launches were required on an hourly basis, this required terminating the ascent after approximately 40 minutes in order to use the receive equipment to prepare for the next launch, during these intensive operation periods ascents were only made to approximately 200 mb / 12 km

(a) Radiosonde measurement unit.

(b) Radiosonde launch.

Figure D.1: The radiosondes used at the University of Bath as part of CSIP.

These radiosonde ascents provide a good source of data for comparison with the MM5 model. Since this data is attached to the CSIP project all operational periods are likely

to exhibit convection which provide the greatest challenge to NWP systems hence this can be considered to be the most challenging environment that an NWP system in the UK is likely to consider.

The reader should also be made aware of the weaknesses of radiosondes measurements. There are three main issues that arise in this particular application of radiosonde information.

1. The vertical profile is not instantaneous. Since the ascent occurs at around $300\,\text{m}\,\text{minute}^{-1}$, the ascent takes a long period of time. Hence, the measurements at the top of the profile are not time-coincident with those at the base.

2. The radiosonde 'drifts' providing a non-vertical section. During the ascent the balloon will be affected by the horizontal winds this will cause the radiosonde to drift and not provide a perfectly vertical profile. Indeed, this process is used to derive the horizontal wind speed and direction by the variation in position as measured by DGPS. With ascents to 12 km the typical horizontal drift was around 30 km.

3. In high humidity parts of the ascents (primarily in clouds) the balloon can disturb the airflow and cause an under-estimation of humidity. The balloon effectively 'pushes' the cloud out of the path of the humidity sensor.

Whilst it is important to appreciate the first two weaknesses, they do not have a significant effect upon the resulting comparison. However, the final issue is important to note.

Typically radiosonde launches (and vertical profiles in general) are viewed as 'skew-T plots'. Skew-T diagrams can be difficult to read and understand, so in order to make these comparisons as accessible to all readers the data is presented as simple plots of temperature versus pressure and dewpoint temperature versus pressure.

The dewpoint temperature is the temperature at which air can no longer hold all of the water vapour. At this point some of the water vapour must condense to liquid water. At 100% humidity the dewpoint-temperature and temperature are identical. The dewpoint temperature can be considered to be a measure of how much water there is in the air and can be calculated using equation D.1; where T_d is the dewpoint temperature (°C), where T is the temperature (°C), RH is the relative humidity and

e_s is the saturated water vapour pressure calculated using equation D.2.

$$T_d = \frac{237.7 \log_{10}\left(\frac{e_s \cdot RH}{611}\right)}{7.5 - \log_{10}\left(\frac{e_s \cdot RH}{611}\right)} \tag{D.1}$$

$$e_s = 6.11 \times 10^{\left(\frac{7.5 \cdot T}{237.7 + T}\right)} \tag{D.2}$$

The radiosonde launches from the CSIP intensive operation periods in July and August 2005 from the University of Bath have been compared to the outputs from the MM5 NWP model.

In order to provide an assessment of the rainfall rate forecasts, the UK NIMROD weather radar system has also been compared to the University of Bath MM5 model. Caution should be exercised whilst performing these comparisons since the radar provides estimates of instantaneous rainfall rate at 15 minute intervals whilst the UM provides an estimation of the accumulation over the same period. In addition to this, the systems are on different grid scales. However, they can be used to provide an appreciation of the accuracy.

Several case studies follow the along with comments adapted from the CSIP summary for that day.

D.1 Case study 1 - 18/07/2005

This observation period saw a narrow line of convective showers associated with the passage of a cold front. The main band of frontal cloud was situated ahead of the convective showers. The narrow line of convection was visible for many hours as it advected towards southern central UK but part of it intensified as it moved across the area and produced heavy rain as it passed close to London around 16:45. A convective lid at around 3 km (700 mb) was present over the Bath site, this later was raised by the front to around 4 km (620 mb) which instigated the showers.

The system does a very good job of forecasting the prognostic variables, as seen in the vertical profiles in figure D.2. The model also resolves the intensive rain band across the Midlands and up to Hull, figure D.3. However the squall line across the south coast is resolved far weaker and advecting much slower.

Figure D.2: A comparison between the University of Bath NWP model and the high resolution radiosondes launched by the University of Bath as part of CSIP - 18/07/2005.
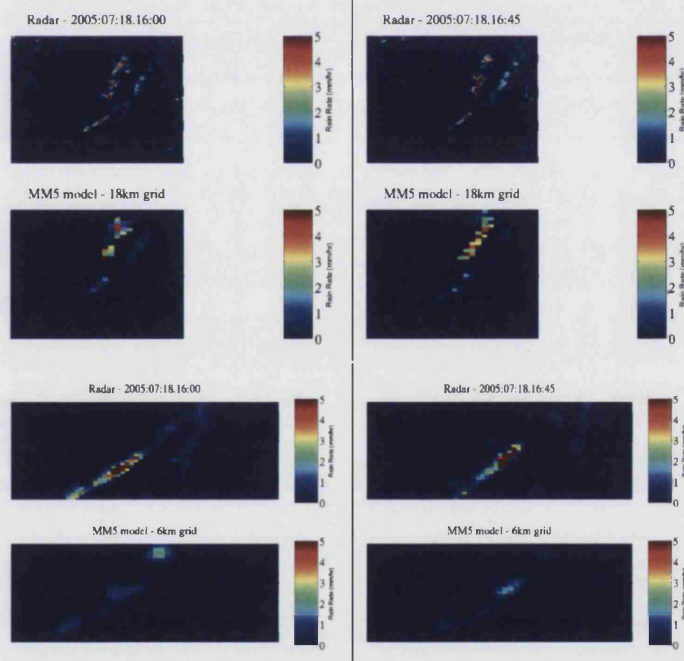
Figure D.3: A comparison between the University of Bath NWP model and the NIM-ROD radar system- 18/07/2005.

211

D.2 Case study 2 - 23/07/2005

This intensive all-day observation was performed during which a lid at 2 km suppressed convection until late afternoon. This is evident as a temperature inversion at around 850 mb in the radiosonde ascents shown in figure D.4. More showers had been predicted and a contributory factor to the convective inhibition may have been the extensive cover of low cloud which did not show any breaks until the late afternoon. On-shore winds also suppressed convective activity close to the south coast.

For this case study the system is able to predict the water vapour at low altitudes well although the radiosonde underestimation of the humidity is evident.
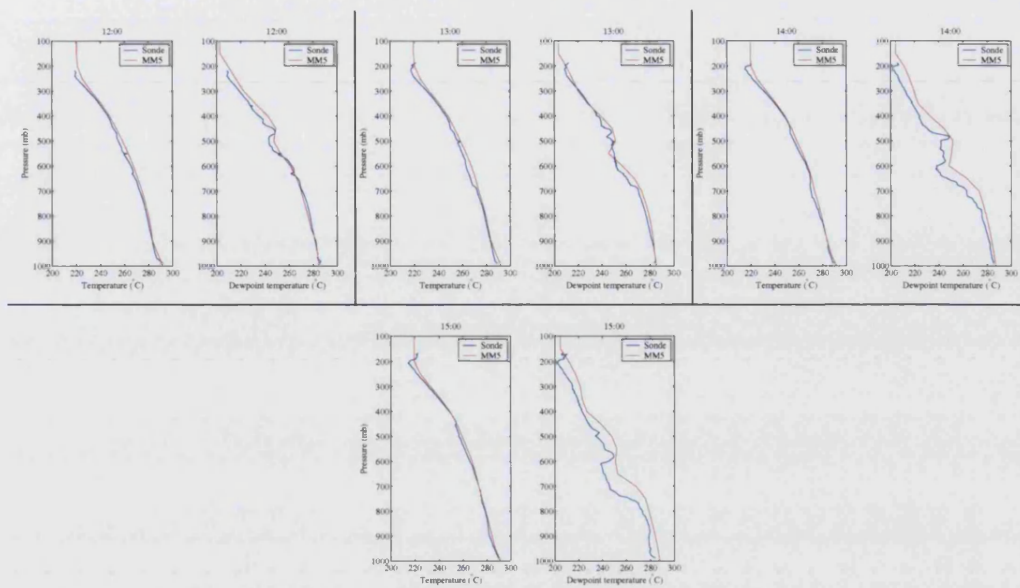
Figure D.4: A comparison between the University of Bath NWP model and the high resolution radiosondes launched by the University of Bath as part of CSIP - 23/07/2005.

D.3 Case study 3 - 24/07/2005

A low pressure area just south of Ireland was associated with a broad belt of heavy rain associated northwards and cleared Southern England by early afternoon. The main operation period commenced at this point. Shallow convective cloud lines, transverse to

the flow, formed beneath a lid at 2 to 3 km, visible as a temperature inversion at around 750 mb until 13:00 (figure D.5). Later, well defined convective rainbands approached the area from SW England and broke through the lid and produced thunderstorms in south Wales. Near Bath, there were moderate showers with tops to 5 km (550 mb) in the late afternoon.
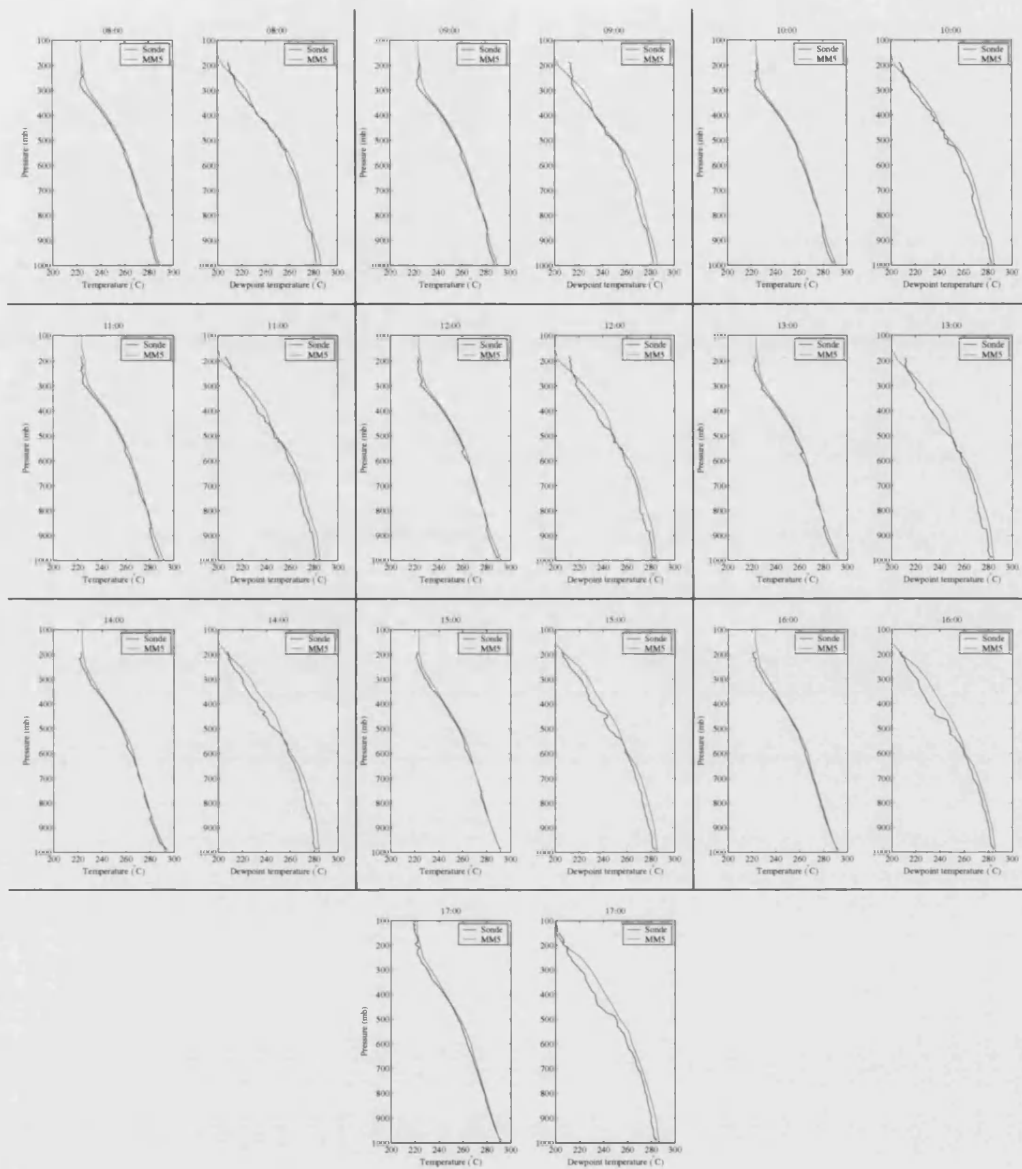

Figure D.5: A comparison between the University of Bath NWP model and the high resolution radiosondes launched by the University of Bath as part of CSIP - 24/07/2005.

At 11:15 the frontal belt can clearly be seen in both the MM5 and radar scans in figure D.6, whilst the MM5 model shows evidence of convective events following the front. At 18:30 the modelling of the convective rain is very unclear, the MM5 model evidently resolves some convective activity but the actual individual rain events are not isolated on the 18 km grid.
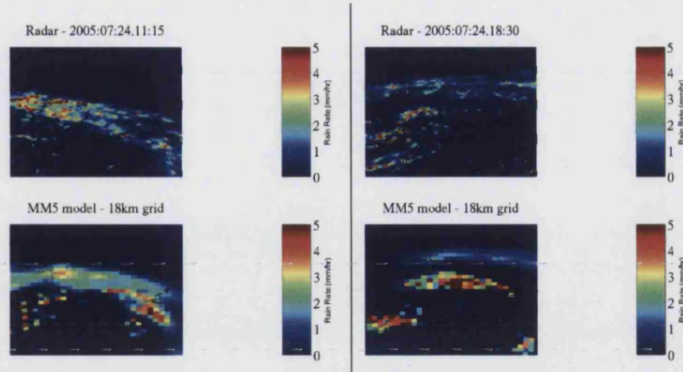
Figure D.6: A comparison between the University of Bath NWP model and the NIM-ROD radar system- 24/07/2005.

D.4 Case study 4 - 28/07/2005

This intensive observation period was characterised by a dry slot ahead of a low that was centred over SW England. Bands of showery rain at the leading edge of the dry slot cleared away during the early morning and for a while it was dry in Southern England with only shallow cumulus clouds, note the temperature inversion at 800 mb at 09:00. Because of the strong wind, the cumulus were aligned in cloud streets with a narrow spacing of about 6 km at this time.

By 10:30 the cloud streets were being modulated so as to give transverse bands of brighter, taller clouds. Some of the first rain echoes formed within three of these bands between 10:00 and 11:00. By 11:30 there were bands of showers associated with two of these bands.

The southernmost of these shower bands soon decayed but the showers in the other band intensified and developed in lines along each individual longitudinal cloud street. The spacing of which had increased to about 12 km by 12:15.

One of the showers went on to give a tornado in Birmingham, around 13:30. A preliminary assessment suggests that the storm began as a developing cumulus north of Dorchester at 10:00 which the cloud top then proceeded to rise rapidly from roughly 3 km (700 mb) at 10:45 to nearer 9 km (300 mb) at 11:15 as it passed over Bath. This is evident in the poor modelling of the dewpoint temperature demonstrated in figure D.7.
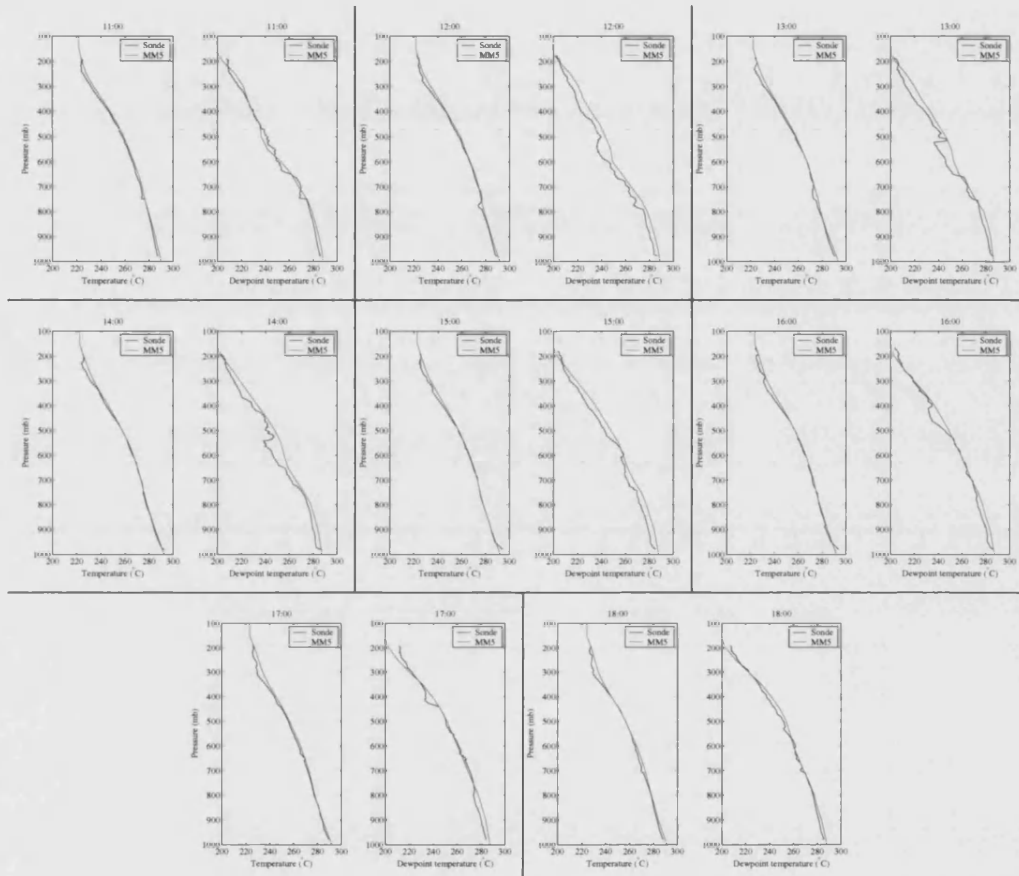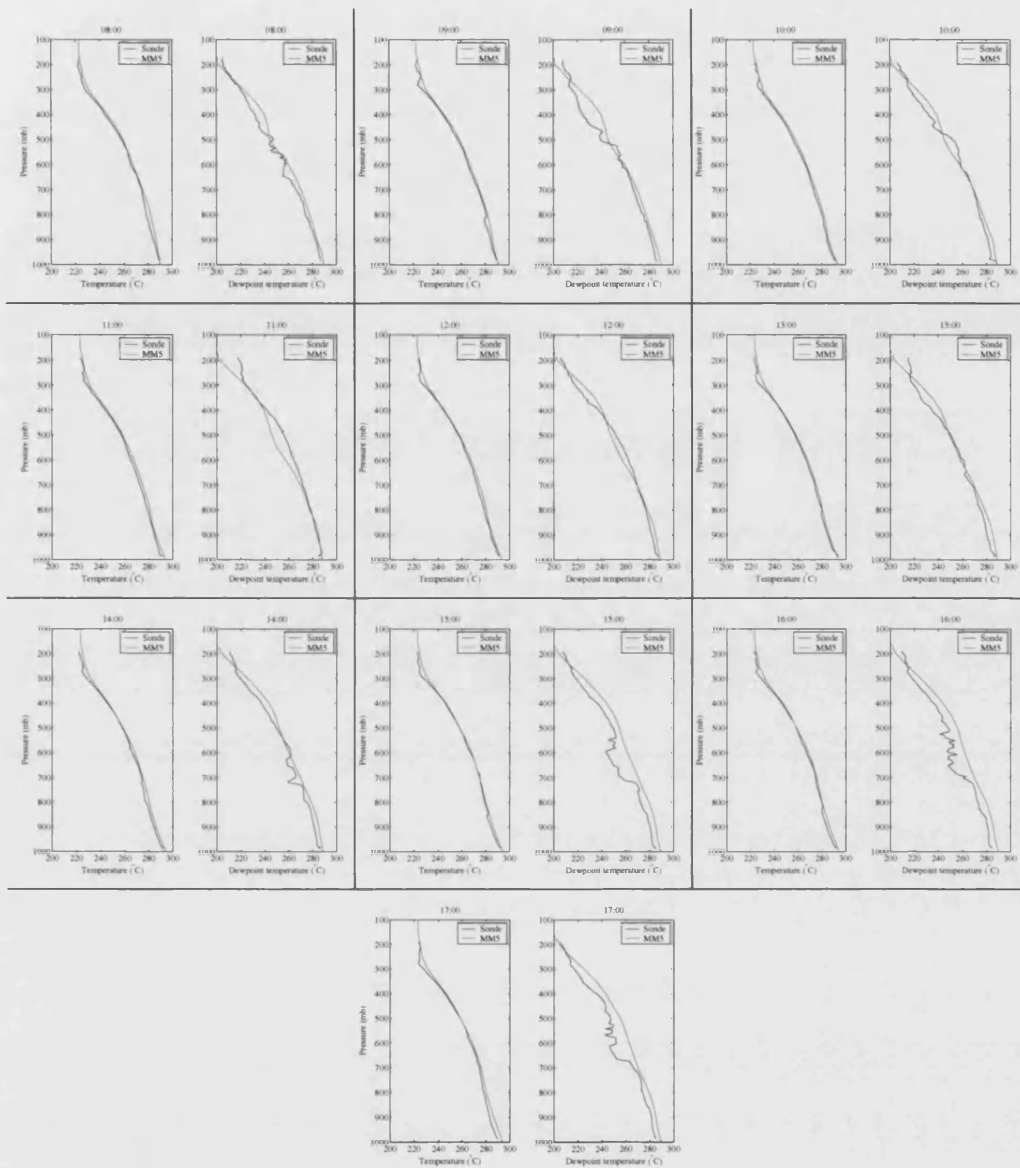
215

Figure D.7: A comparison between the University of Bath NWP model and the high resolution radiosondes launched by the University of Bath as part of CSIP - 28/07/2005.
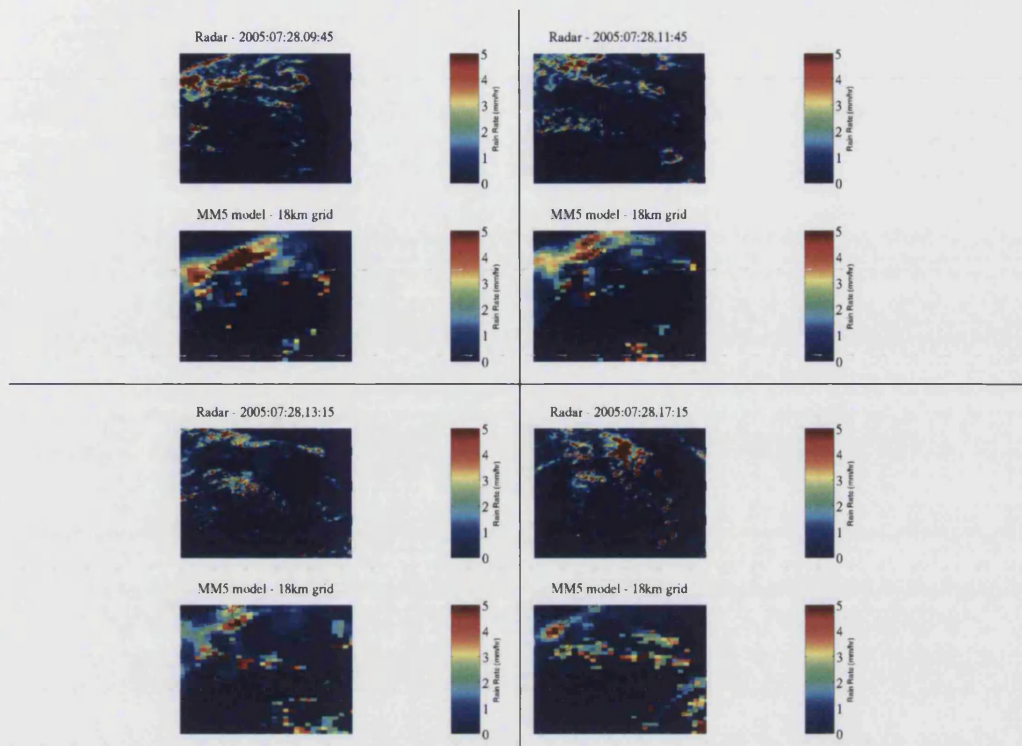
Figure D.8: A comparison between the University of Bath NWP model and the NIM-ROD radar system - 28/07/2005.

The poor modelling of the dewpoint has a direct effect on the quality of the rainfall modelling. The estimate for 09:45 shown in figure D.8 demonstrates identifiable features that are present on both the radar scan and the MM5 analysis. This had started to break down by the 11:45 and the 13:15 scan although areas of intense rainfall were still in the same areas. However by 17:15 the modelling is so poor that the radar and MM5 model show little or no consistency. In conclusion this day provided a serious challenge to the MM5 model (and unsurprisingly most NWP techniques).

D.5   Case study 5 - 18/08/2005

Southern England was covered by a weak, broadly southerly flow ahead of a band of cold-frontal rain advancing slowly from the west. Much of the region experienced a warm mainly sunny day with only shallow convection, note the temperature inversion at 800 mb in figure D.9. The first cumulus were tied very closely to topography which then became more widespread but their development was for a long time restricted

217

by the strong lid. By late afternoon, however, the lid had weakened and a small area of showers with tops to 6 km (450 mb) developed, note the temperature inversion at 450 mb as the lid at 800 mb is penetrated.
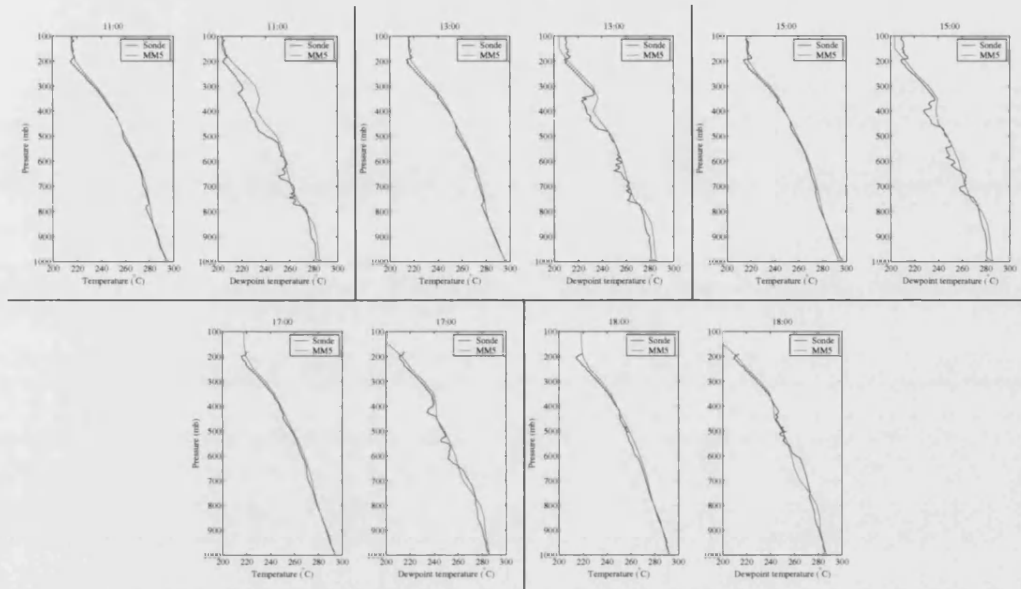
Figure D.9: A comparison between the University of Bath NWP model and the high resolution radiosondes launched by the University of Bath as part of CSIP - 18/08/2005.

Once again the frontal rain is well-modelled (figure D.10). The band is well defined at both 12:00 and 16:00. As with previous results presented here, the system goes on to identify that convection is likely to occur however is unable to resolve the convective activity in the correct areas.
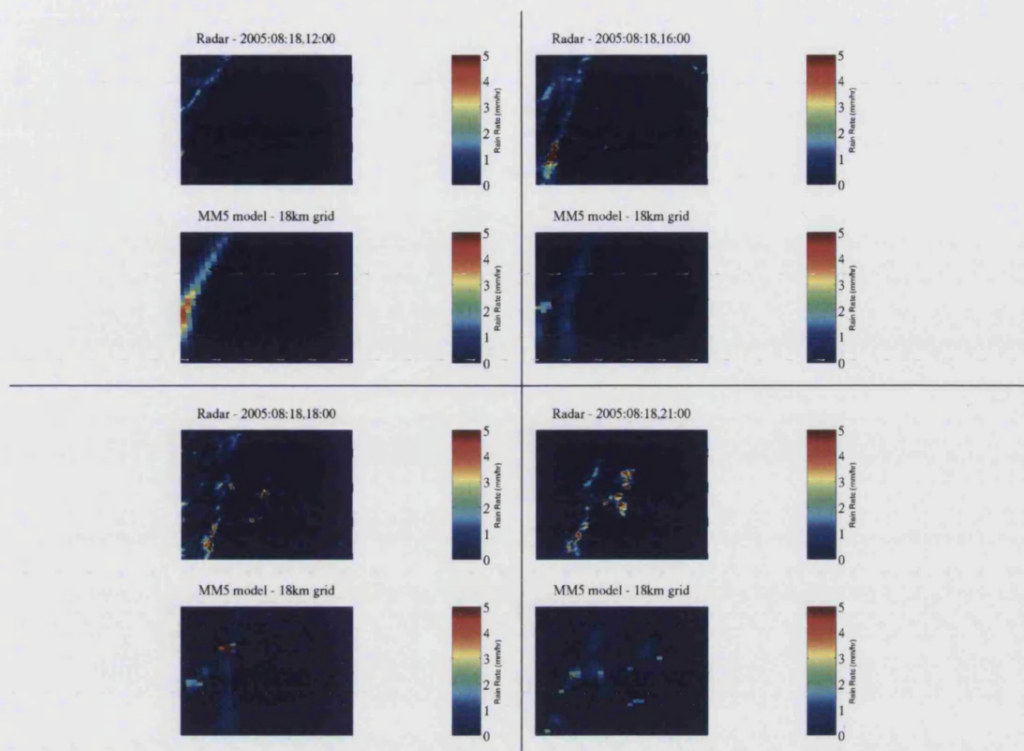
Figure D.10: A comparison between the University of Bath NWP model and the NIM-ROD radar system - 18/08/2005.

D.6 Case study - conclusion

In general the MM5 and the radiosondes compare well, even during this period of convective activity. Since the vertical resolution of the MM5 model is of the order of 50 mb it is unable to resolve a lot of the effects that are visible in the radiosonde data. This is one of the contributing to factors to the poor ability to model convection since it is unable to resolve small temperature inversions and other convective inhibitors. In addition the continual underestimation of humidity by the radiosondes is evident in the plots of the dew-point temperature. In conclusion the MM5 model is very competent at estimating the prognostic variables even in these challenging convective situations.

The rainfall predictions for the frontal rain are typically very good, however the modelling of convection is generally relatively poor. Heuristically it can be seen that in general the system is able to predict convective activity, but is unable to place areas where convective inhibition (CIN) are broken and rain occurs. This performance can be dramatically improved by using observational nudging. Other improvements could

219

include an implementation of the ensemble MM5 system (with the associated processing overhead) or moving to the new WRF model which claims to improve the modelling of small-scale effects.

Appendix E

Geopotential height

Geopotential height is used as a vertical co-ordinate system in most atmospheric modelling. The effect of the Earth's gravity is reduced as the distance from the surface is increased. In order to mitigate this effect the geopotential height is used. This ensures that the work required to move a given mass a given geopotential distance in the lower atmosphere is identical to the work done moving the same mass the same distance in the upper atmosphere.

Geopotential height, Z, is defined as:

$$Z = \frac{1}{g_0} \int_0^z g(z) \, dz \qquad \text{(E.1)}$$

where z is the height above sea-level, g_0 is the globally averaged acceleration due to gravity whilst $g(z)$ is the acceleration due to gravity at a height z.

Some example values of the geopotential height are shown in table E.1

As can be seen, within the troposphere there is little difference between the geopotential and geometric height.

Table E.1: Comparison between the geopotential height, Z, and the geometric height, z, at a latitude of 40°, [Wallace and Hobbs, 1977].

Geometric Height (km)	Geopotential Height (km)	acceleration due gravity $(\mathrm{m\,s^{-1}})$
0	0	9.802
1	1.000	9.798
10	9.986	9.771
20	19.941	9.741
30	29.864	9.710
60	59.449	9.620
90	88.758	9.531
120	117.795	9.443
300	286.520	8.940

References

E E Altshuler and R A Marr. Cloud attenuation at millimeter wavelengths. *IEEE Transactions on Antennas and Propagation*, 37(11):1473–1479, November 1989.

D Atlas, R C Srivasta, and R S Sekhon. Doppler radar characteristics of precipitation at vertial incidence. *Rev. Geophys. Space Phys.*, 2:1–35, 1973.

B Audoire. *Modélisation fine du canal de propagation Terre-Espace dans les bandes Ka- et V-*. PhD thesis, University of Toulouse, 2001.

P D Baxter and P Garcia. The extraction of scintillation statistics from Italsat 50 GHz beacon data using wavelets. In *COST 280: MCM4 meeting*, November 2004.

R Benzi, S Ciliberto, R Tripiccione, C Baudet, F Massaioli, and S Succi. Extended self-similarity in turbulent flows. *Physics review E*, 48(1):347–350, 1993.

A C Best. Empirical formulae for the terminal velocity of water drops falling through the atmosphere. *Quart. J. R. Met. Soc.*, 76:302–311, July 1950.

M Bevis, S Chiswell, S Businger, T A Herring, and Y Bock. Estimating wet delays using numerical weather analyses and predictions. *Radio Science*, 31(3):477–487, May 1996.

G Brussaard. A meteorological model for rain-induced cross polarization. *IEEE Transactions on Antennas and Propagation*, 24(1):5–11, January 1976.

L Castanet. *Techniques adaptives de lutte contre les affaiblissements de propagation pour les systémes de télécommunications par satellite en EHF*. PhD thesis, University of Toulouse, 2001.

L Castanet, T Deloues, and J Lemorton. Channel modelling based on n-state Markov chains for satcom systems simulation. *Proc. International Conference on Ant. and Prop.*, pages 119 – 122, 2003.

S Coles, J Heffernan, and J Tawn. Dependence measures for extreme value analyses. *Extremes*, 2(4):339–365, December 1999.

F Davarian, S Shambayati, and S Slobin. Deep space Ka-band link management and Mars Reconnaissance Orbiter: Long-term weather statistics versus forecasting. *Proceedings of the IEEE*, 92(12):1879–1894, December 2004.

R Deidda. Rainfall downscaling in a space-time multifractal framework. *Water Resour. Res.*, 36(7), 1999.

F Dintelmann and G Ortgies. Semi-empirical model for cloud attenuation prediction. *Electronics Letters*, 25(22):1487–1488, October 1989.

A Dissanayake, J Allnutt, and F Haidara. Cloud attenuation modelling for SHF and EHF applications. *International Journal of Sat Comms*, 19:335–345, 2001.

L Dossi. Real-time prediction of attenuation for applications to fade countermeasures in satellite communications. *Electronics Letters*, 26(4):250–251, 1990.

G Drufuca. Rain attenuation statistics for frequencies above 10 GHz from rain guage observations. *J. Rech. Atmos.*, 1(2):399–411, 1974.

J V Evans and A Dissanayake. Prospects for commercial services at Q- and V-bands. *MilCom 1998*, 1:18–21, October 1998.

J Feil, L J Ippolito, H Helmken, C E Mayer, S Horan, and R E Henning. Fade slope analysis for Alaska, Florida and New Mexico ACTS propagation data at 20 and 27.5 GHz. *Proceedings of the IEEE*, 85:926–935, June 1997.

L Ferraris, S Gabellani, and N Rebora. A comparison of stochastic models for spatial rainfall downscaling. *Water Resour. Res.*, 39(12), 2003.

U C Fiebig. Modelling rain fading in satellite communication links. *Proc. Vehicular Technology Conference*, pages 1422–1426, 1999.

B C Grémont. *Fade Countermeasure Modelling for Ka- band digital satellite links.* PhD thesis, Coventry University, 1997.

B C Grémont and M Filip. Spatio-temporal rain attenuation model for application to fade mitigation techniques. *IEEE Transactions on Antennas and Propagation*, 52 (5):1245–1256, May 2004.

B C Grémont, A P Gallois, and S Bate. Efficient fade compensation for Ka-band VSAT systems. *Proc. 2nd Int. Ka-band Utilisation Conference*, pages 439–443, 1996.

B C Grémont, M Filip, A P Gallois, and S Bate. Comparative analysis and performance of two predictive fade detection schemes for ka-band fade countermeasures. *IEEE Journal on Selected Areas in Communications*, 17(2):180–192, February 1999.

R Gunn and G D Kinzer. The terminal veolcity of fall for water droplets in stagnant air. *J. Meteorol.*, 6:243–248, August 1949.

A Ishimaru. *Wave Propagation and Scattering in Random Media*. Oxford University Press, 1978.

I T Jolliffe and D B Stephenson. *Forecast Verification: A practitioner's guide in atmospheric science*. Wiley and Sons Ltd, 2003.

Y Karasawa, M Yamada, and J E Allnutt. A new prediction method for tropospheric scintillation on earth-space paths. *IEEE Transactions on Antennas and Propagation*, 36:1608–1614, November 1988.

T Konefal, P A Watson, A K Shukla, and A Akram. Prediction of monthly and annual availabilities on 10-50 ghz satellite-earth and aircraft-to-aircraft links. *IEE Proc. Microw. Antennas Propag.*, 147(2):122–127, April 2000.

D Kreuer and A Schmidt. Adaptive resource sharing strategies for TDMA satellite networks. In *PIMRC 1994 : Satellite systems session*, pages 1192–1196, 1994.

J O Laws and D A Parsons. The relation of raindrop size to intensity. *Trans. Amer. Geophys. Union*, 24:452–460, 1943.

M J Leitao and P A Watson. Method for prediction of attenuation on earth-space links based on radar measurements of the physical structure of rainfall. *Proceedings of the IEEE*, 133(4):429–441, July 1986.

Li Li, J Vivekanandan, C H Chan, and L Tsang. Microwave radiometric technique to retrieve vapour, liquid and ice: Part 1 - development of a neural network based inversion method. *IEEE Journal on Geoscience and Remote Sensing*, 35(2):224–236, March 1997.

H J Liebe. MPM - an atmospheric millimeter-wave propagation model. *International Journal of Infrared and Millimeter Waves*, 10(6):631–650, 1989.

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

T Maseng and P M Bakken. A stochastic dynamic model of rain attenuation. *IEEE Transactions on Antennas and Propagation*, 29:660–669, May 1981.

E Matricciani. Are concurrent rain attenation and rain rate measurements really concurrent ? The impact on prediction methods used to design satellite communication systems with high availability. *Proc. 10th Int. Ka-band Utilisation Conference*, pages 119–126, September 2004.

E Matricciani. Physical-mathematical model of the dynamics of rain attenuation based on rain rate time series and a two-layer vertical structure of precipitation. *Radio Science*, 31(2):281–295, 1996.

R Mazumdar, L G Mason, and C Douligeris. Fairness in network optimal flow control : Optimality of product forms. *IEEE Transactions of Communications*, 39(5):775–782, May 1991.

M Menabde, A Seed, D Harris, and G Austin. Multiaffine random field model of rainfall. *Water Resour. Res.*, 35(2), 1999.

J F Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

T Oguchi. Electromagnetic wave propagation and scattering in rain and other hydrometeors. *Proceedings of the IEEE*, 71(9):1029–1078, September 1983.

R L Olsen, D V Rogers, and D B Hodge. The aR^b relation in the calculation of rain attenuation. *IEEE Transactions on Antennas and Propagation*, 26(2):318–329, March 1978.

A J Page, R J Watson, and P A Watson. Time-series of attenuation on ehf and shf fixed radio links derived from meteorological forecast and radar data. *IEE Proceedings on Microwaves, Antennas and Propagation*, 152(2):124–128, April 2005.

A Paraboni and C Riva. A new method for the prediction of fade duration statistics in satellite links above 10 GHz. *International Journal of Sat Comms*, 12:387–394, 1994.

K S Paulson. Fractal interpolation of rain rate time series. *J. Geophys. Res.*, 109, 2004.

H R Pruppacher and K V Beard. A wind tunnel investigation of the internal circulation and shape of water drops falling at terminal velocity in air. *Quart. J. R. Met. Soc*, 96:247–256, 1970.

H R Pruppacher and R L Pitter. A semi-empirical determination of the shape of cloud and rain drops. *J. Atmos. Sci.*, 28:86–94, 1971.

J C Russ. *Fractal Surfaces*. Plenum Press, 1994.

E Salonen and S Uppala. New prediction method of cloud attenuation. *Electronics Letters*, 27(12):1106–1108, June 1991.

A Savvaris, C N Kassianides, and I E Otung. Observed effects of cloud and wind on the intensity and spectrum of scintillation. *IEEE Transactions on Antennas and Propagation*, 52(6):1492–1498, June 2004.

SpaceWay. Mesh point-to-point via satellite broadband - the best of both worlds. www.hns.com/default.asp?CurrentPath=spaceway/benefits.htm.

D B Stephenson. Extreme rainfall: from weather forecasts to future climate. In *Rainmap 2005: Kick off conference*, January 2005.

V I Tatarski. *Wave propagation in turbulent medium*. McGraw Hill, 1961.

G I Taylor. The spectrum of turbulence. *Proc. R. Soc.*, A164:476–490, 1938.

J K Tervonen, M M J L van de Kamp, and E T Salonen. Prediction model for the diurnal behaviour of the tropospheric scintillation variance. *IEEE Transactions on Antennas and Propagation*, 46:1372–1378, September 1988.

F T Ulaby, R K Moore, and A K Fung. *Microwave remote sensing, active and passive, Volume 1:Microwave remote sensing and fundamentals and radiometry*. Artech House Books, 1986.

C W Ulbrich. The natural variations in the analytical form of the raindrop size distribution. *Journal of climate and applied meteorology*, 22(10):1764–1775, November 1983.

M M J L van de Kamp. Efficient fade compensation for Ka-band VSAT systems. *Proc. 8th Int. Ka-band Utilisation Conference*, pages 257–264, September 2002a.

M M J L van de Kamp. Short-term prediction of rain attenuation using two samples. *Electronics Letters*, pages 1476–1477, November 2002b.

M M J L van de Kamp. Asymmetric signal level distribution due to tropospheric scintillation. *Electronics Letters*, 34(11):1145–1146, May 1998.

M M J L van de Kamp. Statistical analysis of rain fade slope. *IEEE Transactions on Antennas and Propagation*, 51(8):1750–1759, August 2003.

M M J L van de Kamp. *Climatic radiowave propagation models for the design of satellite communications systems*. PhD thesis, University of Eindhoven, 1999.

M M J L van de Kamp, J K Tervonen, E T Salonen, and J P V P Baptista. Improved models for long-term prediction of tropospheric scintillation on slant paths. *IEEE Transactions on Antennas and Propagation*, 47(2):249–260, February 1999.

H Vasseur. Prediction of tropospheric scintillation on satellite links from radiosonde data. *IEEE Transactions on Antennas and Propagation*, 47(2):293–301, February 1999.

J Von Neumann and O Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press, 1944.

J M Wallace and P V Hobbs. *Atmospheric Science : An introductory survey*. Academic Press, 1977.

J M Warnock, T E Vanzandt, and J L Green. A statistical model to estimate mean values of parameters of turbulence in the free atmosphere. *Seventh Symposium of turbulence diffusion*, 1985.

R J Watson and D D Hodges. The use of short term attenuation forecasts for resource management. *10th Ka and Broadband communications conference*, pages 295–302, October 2004.

C J Wiesner. *Hydrometeorology*. Chapman and Hall Ltd., 1970.

M J Willis. Fade counter-measures applied to transmissions at 20/30 GHz. *Electronics and Communication Engineering Journal*, pages 88–96, April 1991.

H Yaïche, R R Mazumdar, and C Rosenberg. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking*, 8(5):667–678, October 2000a.

H Yaïche, R R Mazumdar, and C Rosenberg. Distributed algorithms for fair bandwidth allocation to elastic services in broadband networks. In *INFOCOMM*, March 2000b.

W Zhang. Scattering of radiowaves by a melting layer precipitation in backward and forward directions. *IEEE Transactions on Antennas and Propagation*, 42(3):347–356, September 1994.

W Zhang, J K Tervonen, and E T Salonen. Backward and forward scattering by the melting layer composed of spheroidal hydrometeors at 5-100 GHz. *IEEE Transactions on Antennas and Propagation*, 40(9):1208–1219, September 1996.