



**Impact of introducing Automatic Transmit
Power Control in P-P Fixed Service systems
operating in bands above 13 GHz**

**Extension: Soft Boundary Frequency
Assignment Techniques and Adaptive Coding
and Modulation**

Final Project Report

S.A. Callaghan, I. Inglis, P. Hansell.

On behalf of:

CCLRC Rutherford Appleton Laboratory and Aegis Systems Ltd.



28th November 2007

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and
Modulation: Final Report

Table of Contents

EXECUTIVE SUMMARY	3
INTRODUCTION	4
WORK PACKAGE 1: SOFT BOUNDARY FREQUENCY ASSIGNMENT	5
SERVICE EFFICIENCY CLASS	5
INITIAL TWO-COMMUNITY ASSIGNMENT	6
ASSIGNMENT EFFICIENCY IN REDUCED-BAND SIMULATIONS	8
ASSIGNMENT FAILURE RATES IN REDUCED-BAND SIMULATIONS	13
WORK PACKAGE 2: ADAPTIVE CODING & MODULATION	15
REVIEW OF TECHNIQUES	15
IMPLEMENTATION APPROACH.....	16
SIMULATION RESULTS	17
WORK PACKAGE 3: RAIN FIELD SIMULATOR	23
MODIFICATIONS MADE TO THE SIMULATOR.....	24
DISTRIBUTION OF STRATIFORM AND CONVECTIVE EVENTS THROUGH AN AVERAGE YEAR.....	24
CREATION OF A SIMULATED ANNUAL DATASET	26
CONCLUSIONS AND CAVEATS REGARDING THE RAIN FIELD SIMULATOR	33
REFERENCES	34
APPENDIX A: MATHEMATICAL BASIS FOR THE SIMULATOR	35
THE $K(q)$ FUNCTION FOR CHARACTERISATION OF MULTIFRACTALS	35
MULTIFRACTAL ANALYSIS OF RADAR DERIVED RAIN FIELDS.....	36
CONVERTING FROM A LINEAR ADDITIVE PROCESS TO A NON-LINEAR MULTIPLICATIVE PROCESS	38
APPENDIX B: METHOD USED TO DETERMINE THE SCALING PARAMETERS A AND B. 40	
APPENDIX C: IMPLEMENTATION TEST OF RAIN MODEL.	42
COMPARISON OF PROPAGATION MODELS.....	42
OUTAGE RATE IN A STANDARD PLAN EXPOSED TO ANNUALISED RAIN	44
RAIN RATE DISTRIBUTIONS APPLICABLE TO FIXED LINK PLANNING IN THE UNITED KINGDOM	45

Executive Summary

Previous work has established a capability to simulate the Ofcom fixed link planning process and to generate rain fields with which to confront those plans. This study exploits that development by applying the same approach to a new assignment method (soft boundary assignment) and a new technology (adaptive coding and modulation); in addition, the rain field simulator has been extended to create a rain field database that reproduces annual rain statistics, thereby allowing planned annual link availabilities to be tested. This study uses data from the 38 GHz fixed link band.

Soft boundary frequency assignment is a link assignment method in which each link belongs to a community, which is then subject to community-specific frequency assignment rules. For example, one community might be assigned from the bottom of a band and another from the top. The boundaries of the communities are said to be 'soft' because the communities can overlap in frequency. It is expected that the soft boundary method would excel when a band is congested and simulations have therefore been conducted in which congestion is induced by artificially reducing the size of the band. Using the number of links assigned to the starting channel of each community as a measure of spectral efficiency, simulations show that there is no crossover point because the soft boundary method is always more efficient than conventional single-community assignment. Other measures based on average frequency or channel offset give the same result. However, examining the number of links for which no compatible channel can be found shows that the soft boundary method is nearly always worse than the conventional approach, and if each community is homogeneous then the assignment failure rate is dramatically higher than for mixed or single-community assignment.

Adaptive coding and modulation is a link technology that can be used to enhance the performance or efficiency of a link. A variety of scenarios was examined to determine whether band efficiency gains could be achieved and whether ACM links might be susceptible to interference in the presence of rain. On current planning assumptions, ACM appears to be a benign technology: substantial increases in data rates are possible, with modest efficiency gains if fade margins are reduced. The fixed link planning simulator detected no interference-induced outages when confronting high-penetration ACM plans with annualised rain.

Previous work attempting to assess the impact of new assignment methods or technologies on planned link availabilities was restricted to relative comparisons because annual outage rates could not be determined. This study has successfully developed a method for generating sequences of simulated rain fields that, together, represent annual rain statistics. This allows simulated link availabilities to be compared directly with planned availabilities. The method relies on the generation of simulated stratiform and convective rain fields, which are then scaled to fit the tail of the ITU-R P.837-4 rain rate distribution and combined in proportions appropriate to annual rain statistics. The method has been proved by generating a plan with known link availabilities and then testing to see whether the links respond in the appropriate way to the scaled, mixed, annualised collection of rain fields.

Introduction

This study is an extension of the SES project “Impact of introducing Automatic Transmit Power Control in P-P Fixed Service systems operating in bands above 13 GHz”. Two software tools are outputs of the project:

1. The existing planning and analysis tool has been extended to model the soft boundary assignment method. It has also been extended to provide support for ACM as an additional technology.
2. The rain field simulator has been extended to produce 2-D rain fields representing annual statistics.

It is envisaged that the use of ACM may be appropriate on access or business-to-business links, deployed in traditional fixed links bands, where some sacrifice in data throughput is acceptable in faded conditions. This is more appropriate to access links, which use IP, and hence can tolerate variable data rates. Infrastructure links, which require a fixed data rate, might not benefit from the use of ACM.

The software developed for the ATPC project has been successfully adapted for the new approaches considered here. Results are compared, in terms of spectral utilisation, using the same methods as have been used previously in assessing ATPC as a technology. A detailed analysis of the methods used in simulating ‘annualised’ rain fields is also presented in this report.

Work Package 1: Soft boundary frequency assignment

The soft boundary frequency assignment method is based on the idea that multiple communities of links each sharing a characteristic—such as a common W/U ratio—might be more efficiently planned than a single, varied community of links.

Service efficiency class

In an earlier investigation [Bacon, 2003] it was found that planning efficiency in the 38 GHz band could be improved by taking into account the efficiency class of the link when assigning; simulated plans were generated in which the actual service code of the link was promoted to a compatible higher efficiency service code or demoted to a compatible lower efficiency service code. The possible transitions for those service codes present in the 38 GHz band are shown in Figure 1. The investigation found that efficiency was improved in both cases (i.e. promotion and demotion), which supported the idea that exploiting similarities in link communities—or reducing variety—leads to improved efficiency¹.

Because of the importance of this early work in establishing the basis for soft boundary assignment, an equivalent simulation was performed using the planning simulator developed for this study. The conclusions do not completely reproduce the earlier results (see Table 1)—in particular, whereas the efficiency is improved by promoting link service codes, the efficiency is reduced when service codes were swapped with compatible lower efficiency codes.

		Approx capacity, Mbits/s					
		2	4	8	17	34	155
Channel width MHz	3.5	DH ↕	DI ↕				
	7	DJ	DK EK	DL DM ↕			
	14			DN	DP DQ		
	28					DR DS	DT ↕
	56						DU

Figure 1: Possible service code swaps.

	Conventional	Promotion	Demotion
Assignments in channel 1	50.3%	50.9%	48.3%
Maximum bandwidth (MHz)	252	224	350

Table 1: Efficiency from code swaps.

¹ Note, however, that all links were assigned as a single community subject to the same assignment rule.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

There are two differences between the simulations that may explain the variation:

- the fixed link database used here is almost double the size of the earlier database; it is likely therefore to lead to a more constrained plan than the earlier database
- the W/U ratios used here have identical or very similar values for the allowed swaps, resulting in plans that simply reflect the lower bandwidth consumption of the more efficient services rather than any reduction in link variety.

Initial two-community assignment

The first soft boundary assignment simulations were based on the identification of lower- and higher-order modulation classes; lower-order is taken to be 16-states and below; higher-order is then 32 states and above. The process of determining the modulation order defined in this way is indirect, since service codes do not convert directly to modulation type (see following tables).

Class	States	Modulation
1	2	2-FSK, 2-PSK
2	4	4-FSK, 4-QAM
3	8	8-PSK
4	16/32	16-QAM, 32-QAM
5	64/128	64-QAM, 128-QAM

Table 2: Spectrum efficiency class.

Service Code	Description	bps/Hz	W/U (co-channel)	Class	Number in Database
DH	2 Mbps in 3.5 MHz	0.6	27	1	41
DJ	2 Mbps in 7 MHz	0.3	27	1	232
DI	4 or 2x2 Mbps in 3.5 MHz	1.1	27	2	984
DK	4 or 2x2 Mbps in 7 MHz	0.6	27	n/a	1134
EK	4 or 2x2 Mbps in 7 MHz	0.6	27	n/a	12
DL	8 or 4x2 Mbps in 7 MHz	1.1	27	2	2374
DM	8 or 4x2 Mbps in 7 MHz	1.1	27	2	3547
DN	8 or 4x2 Mbps in 14 MHz	0.6	27	1	886
DP	17 or 2x8 Mbps in 14 MHz	1.2	27	2	1334
DQ	17 or 2x8 Mbps in 14 MHz	1.2	27	2	60
DR	34 or 2x17 Mbps in 28 MHz	1.2	27	2	1053
DS	34 or 2x17 Mbps in 28 MHz	1.2	27	2	2312
DT	155 Mbps in 28 MHz	5.5	38	5a	101
DU	155 Mbps in 56 MHz	2.8	36	4	70

Table 3: Service code characteristics.

On this basis service codes DH, DJ, DI, DK, EK, DL, DM, DN, DP, DQ, DR, DS were assigned to the lower-order group and service codes DT and DU were assigned to the higher-order group.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

The result of re-planning the band is shown below.

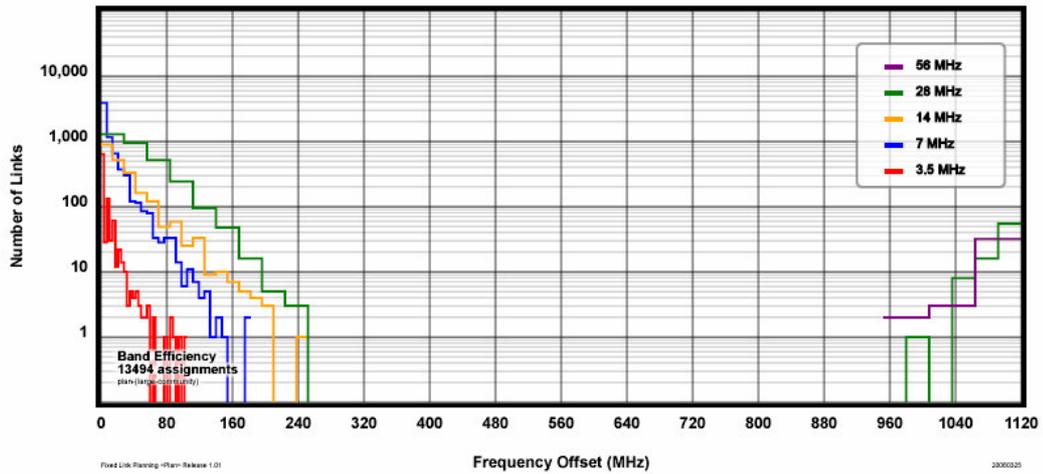


Figure 2: Two-community soft boundary assignment.

The effect on planning efficiency, defined as the sum of the channels assigned to the starting channel for each community, is shown below. The number of links in the higher-order community is relatively small, so the difference in plan efficiency between the conventional approach (i.e. single community) and the soft boundary assignment is not large. However, the trend is clear: as measured by this statistic, the soft boundary approach is always more efficient than the conventional approach. This simply reflects the reduced number of link clashes resulting from the large separation in frequency of the two communities.

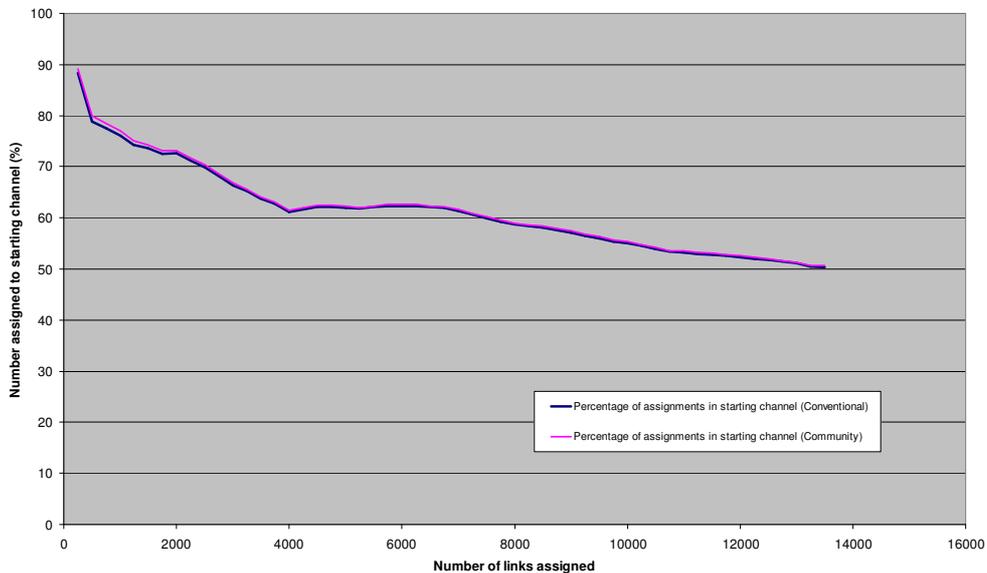


Figure 3: Two-community soft boundary assignment efficiency.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

This effect is also apparent using a (crude) statistic that attempts to capture the width of the community distributions. Defining,

$$B = \frac{1}{n} \sum_i |f_i - f_{start}| \quad (1)$$

then the variation of ‘average bandwidth offset’ with the number of links assigned is shown below.

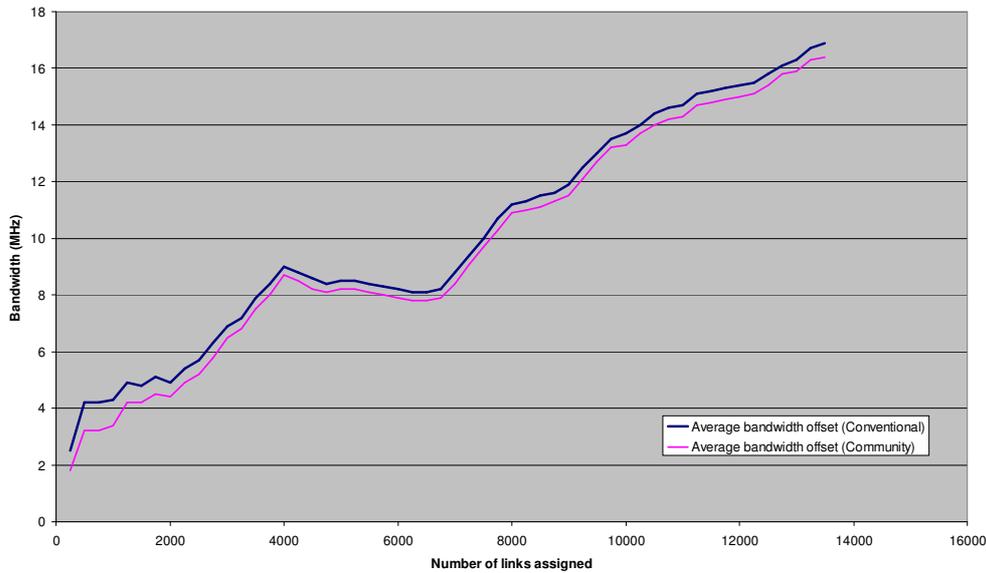


Figure 4: Two-community soft boundary assignment efficiency.

The aim of the investigation is to identify the point at which soft boundary assignment becomes more efficient than the conventional approach. This is clearly not possible with the existing statistics, since soft boundary assignment is *always* more efficient on all statistics thus far considered. Further investigations will look at the behaviour of the band as it becomes full, taking into account not only the number of links that are assigned but also the rate at which links *fail* to be assigned.

Assignment efficiency in reduced-band simulations

A series of simulations has been conducted in which the width of the 38 GHz band was artificially reduced in order to increase link congestion. The band was initially reduced to half its original size; however, this was not sufficient to cause the communities to overlap, so the band was then reduced to a quarter of its original size (i.e. to 280 MHz).

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

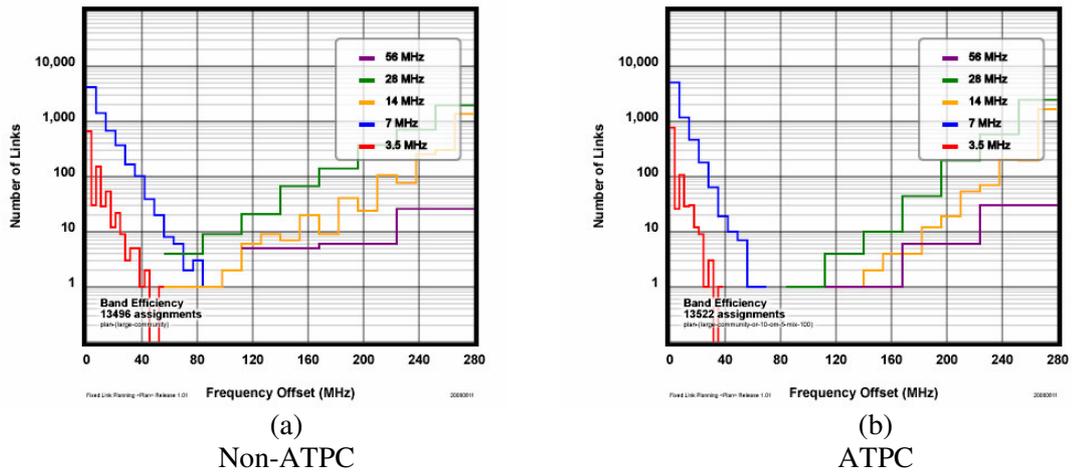


Figure 5: Two-community soft boundary assignment
 Community 1: DH, DJ, DI, DK, EK, DL, DM
 Community 2: DN, DP, DQ, DR, DS, DT, DU

In Figure 5a, the two communities can be seen to overlap, with empty channels appearing when ATPC is used to improve efficiency.

In the following example, a less balanced community definition is used in which the first community has many more links than the second. In this case, the first community occupies almost the entire band. Introducing the more efficient ATPC technology again separates the communities but, in this case, there are no unoccupied channels

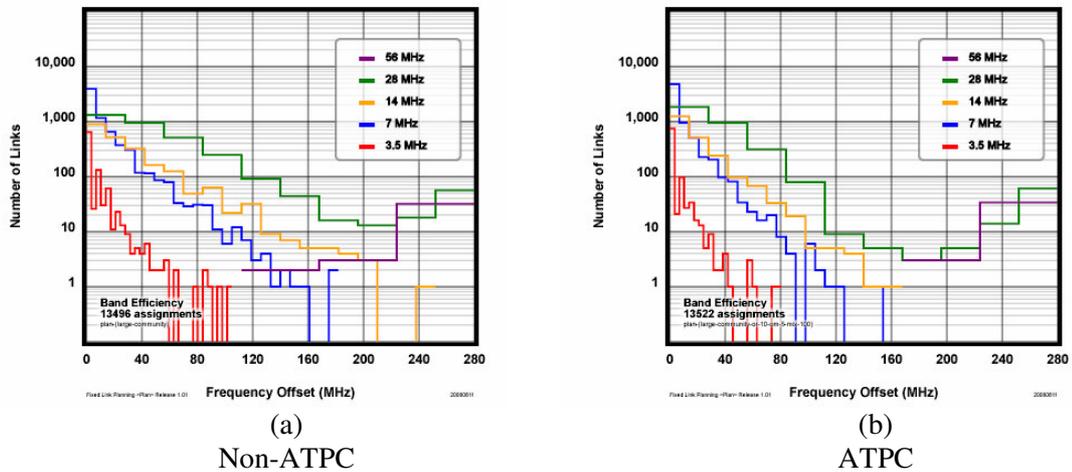


Figure 6: Two-community soft boundary assignment
 Community 1: DH, DJ, DI, DK, EK, DL, DM, DN, DP, DQ, DR, DS
 Community 2: DT, DU

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

The effect on planning efficiency, defined as the sum of the channels assigned to the starting channel for each community, is shown in the following figures. The trend is the same as for the full-band simulations: as measured by this statistic, the soft boundary approach is always more efficient than the conventional approach. As the communities become more evenly balanced, the difference in assignment efficiency increases.

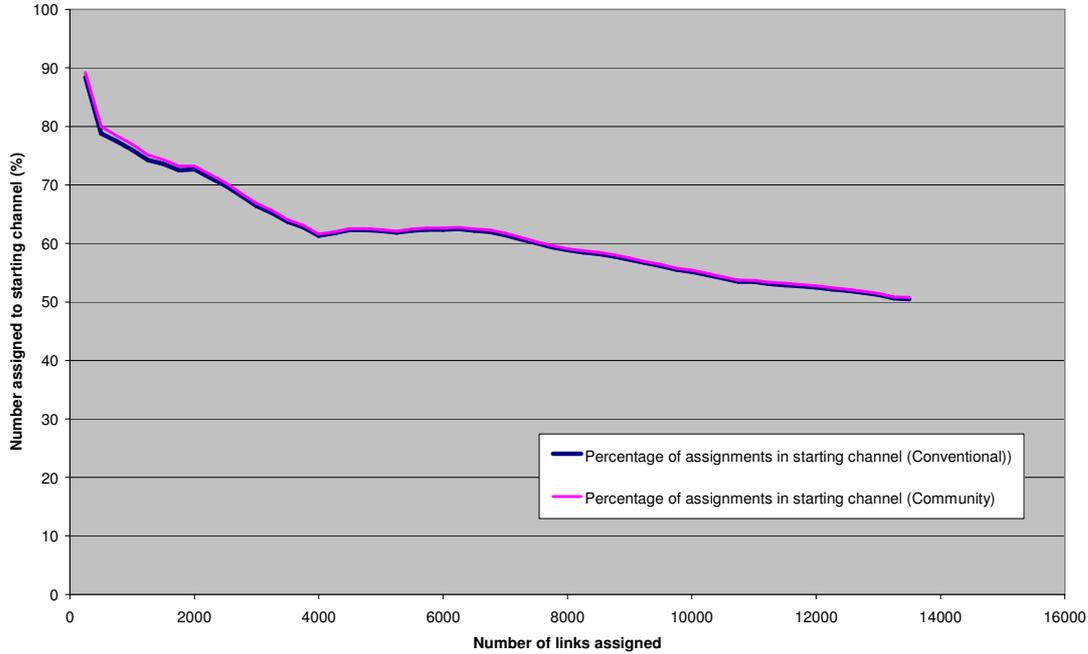


Figure 7: Two-community soft boundary efficiency, quarter band, 155 Mbps services in upper community (DT, DU).

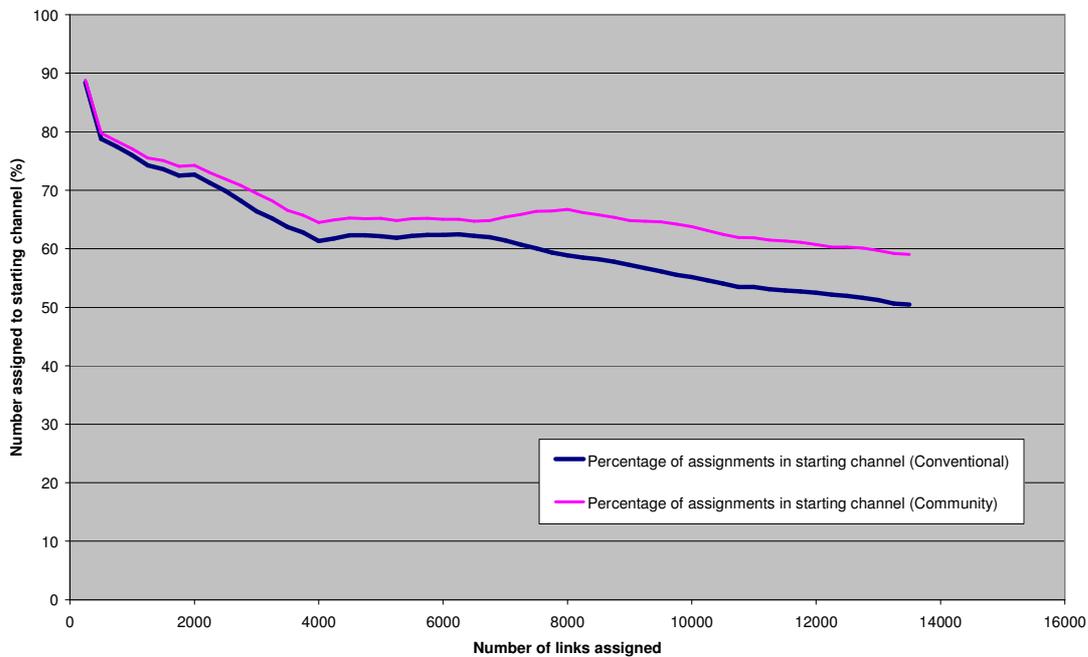


Figure 8: Two-community soft boundary efficiency, quarter band, ≥ 34 Mbps services in upper community (DR, DS, DT, DU).

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

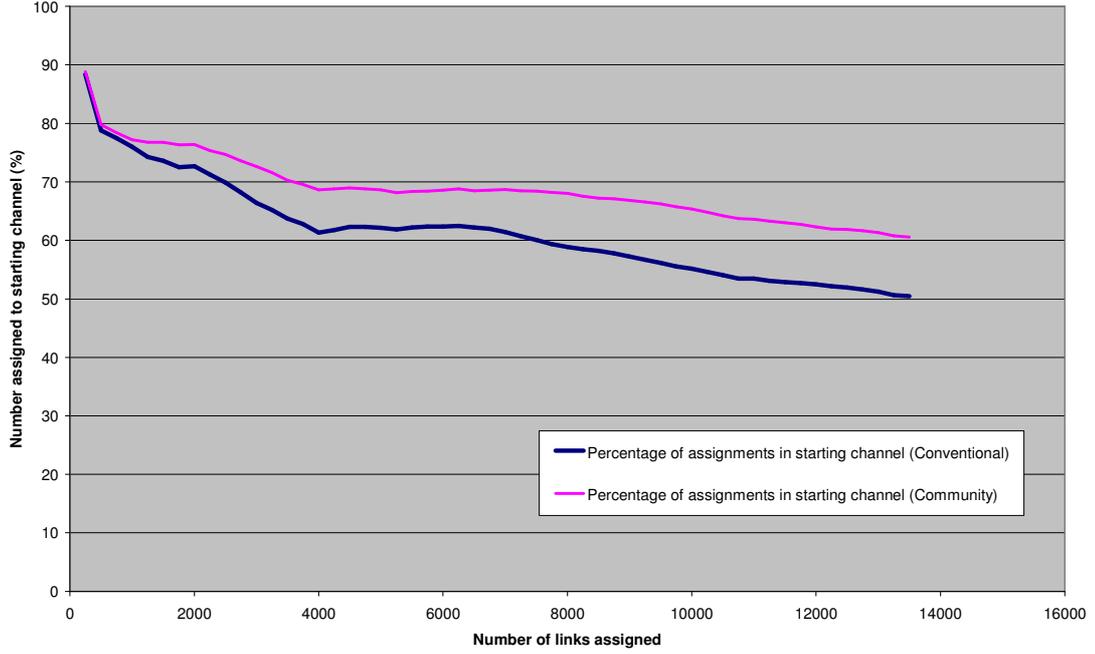


Figure 9: Two-community soft boundary efficiency, quarter band, ≥ 14 MHz services in upper community (DN, DP, DQ, DR, DS, DT, DU).

If the membership criteria for a community are well chosen, then links in the community will be compatible (i.e. less likely to fail a clash test). An assignment attempt in the planning process is abandoned at the first clash test failure and therefore no overall statistic of link compatibility is available. Moreover, deriving the total number of clash tests in developing a plan makes simulation run-times too long, as the clash test is the ‘expensive’ part of the computation. However, the following ‘channel offset’ statistic can be used as a measure of the number of clash test failures, since an assignment moves up or down a channel because at least one clash test has failed.

$$C = \frac{1}{n} \sum_i |c_i - c_{start}| \quad (2)$$

So, if the community membership criteria are effective, then the ‘C’ statistic should be minimised. The following figures show this statistic for the three communities considered previously. Again, the qualitative conclusions are the same: soft boundary assignment is more efficient than single community assignment, and the efficiency gains are highest when the communities are equally numerous.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

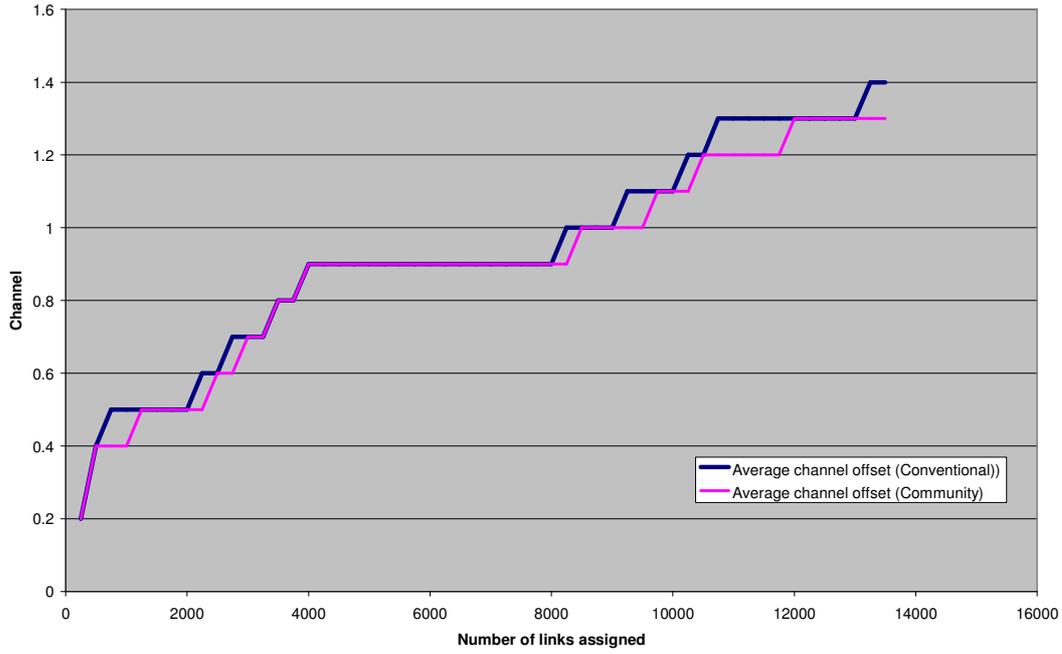


Figure 10: Two-community soft boundary channel offset, quarter band, 155 Mbps services in upper community (DT, DU).

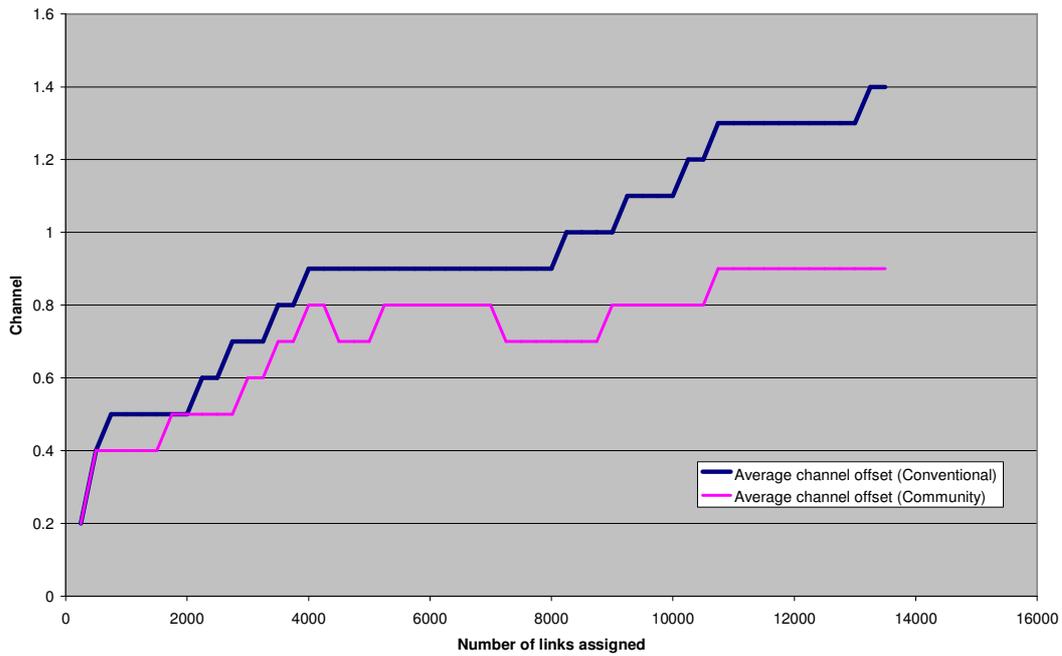


Figure 11: Two-community soft boundary channel offset, quarter band, ≥ 34 Mbps services in upper community (DR, DS, DT, DU).

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

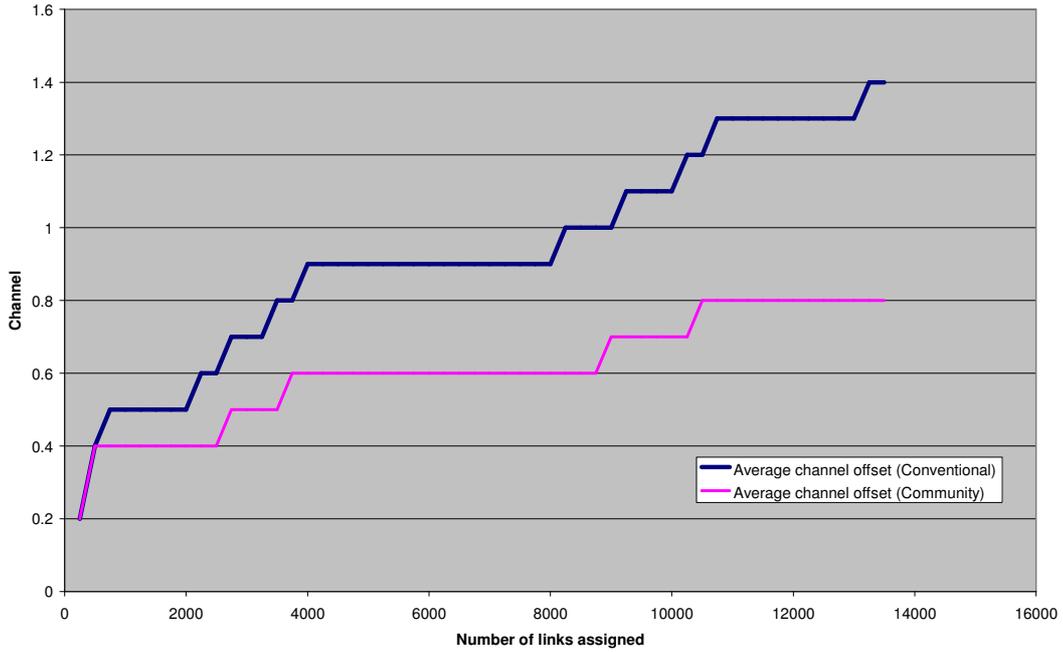


Figure 12: Two-community soft boundary channel offset, quarter band, ≥ 14 MHz services in upper community (DN, DP, DQ, DR, DS, DT, DU).

Further reduction of the band size to approximately one eighth of the original (three 56 MHz channels) results in similar effects. The conclusions remain unchanged.

Assignment failure rates in reduced-band simulations

An alternative approach is to examine the properties of those links that *are not* assigned, rather than deriving statistics from those links that *are* assigned. The following figure shows assignment failure rates as a function of the number of assignments.

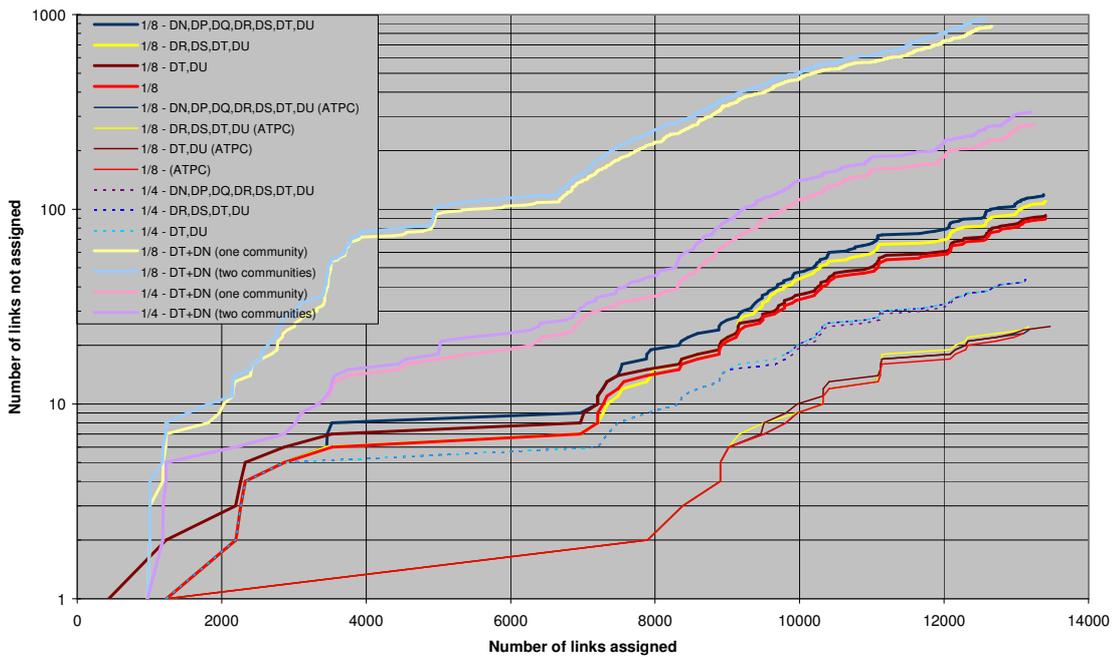


Figure 13: Assignment failure rates.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

Compressing the band to 1/8 of its original size does not significantly change the efficiency statistics, but does increase the number of link assignment failures.

The basis for soft boundary assignment—that ‘like’ communities are ‘compatible’—is not sustained by the results for failure rates found here, which show substantially more link assignment failures when multiple communities are used.

A further community definition was examined, in which all links were assigned in equal numbers either to the DT or the DN service code. The failure rates for this arrangement are ten times higher than for the mixed case and the two-community case has a higher failure rate than the single community case.

Work Package 2: Adaptive Coding & Modulation

Review of techniques

ACM is a technique commonly used in other applications, such as 2.4 / 5 GHz RLANs and WiMAX, where there are wide variations in received signal strength because of propagation behaviour and/or distance from the transmitter. The same technique can be applied to a fixed link operating where a variation in received signal strength is experienced due to rain attenuation on the propagation path.

By way of example, Ericsson has used the following range of modulation coding schemes in illustrative material:

Modulation (Coding)	$C/(N+I)^2$ – dB	Data rate in 28 MHz channel – Mbps	Example unavailability due to propagation - % and mins/year	
QPSK (2/3)	8	25	0.003	16
QPSK (1/1)	12	38	0.005	23
16-QAM (1/1)	20	77	0.008	42
64-QAM (1/1)	26	116	0.017	87

Table 4: Service code characteristics.

It can be seen from the target $C/(N+I)$ levels indicated above that a range of 18 dB distinguishes the lowest from the highest order modulation. A link having a fade margin of 18 dB (or more) would be able to take full advantage of the flexibility offered by the set of modulation/coding states illustrated above. A link with a fade margin less than 18 dB would be able to position itself in one of two ways:

1. With a lower EIRP using the lower order modulations
2. With a higher EIRP using the higher order modulations

Applying the current licence fee algorithm to these two options would result in the same charge, notwithstanding a possible implication for minimum link length in particular cases. It may therefore be necessary to either review the fee algorithm to reflect the different EIRPs or institute a rule that anchors the range used to the lowest order modulation.

Two other characteristics associated with this technique can be noted:

- A. In order to prevent the system from oscillating between two modulation/coding states, a small hysteresis (± 1 dB) is introduced between the threshold for switch-over and switch-back. This has not been taken into account in the simulation work reported in the following sections of this report.
- B. In order to determine the points at which the change between modulation / coding states take place, it is necessary to adjust the thresholds by an amount that defines the maximum degradation due to external interference tolerated by the system. This adjustment is the interference offset.

² These values refer to a BER of 10^{-11} and include implementation losses. They are greater than thresholds that would relate to link outage.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

ACM can also be used in association with ATPC whereby the transmit power is increased up to its maximum as the link path loss increases before the modulation / coding state is switched to a more robust one. When the path loss decreases, the ATPC maintains the transmit power at its maximum until the modulation / coding state is switched to a more efficient one. In order to examine the implications of ACM for fixed link planning the following analysis addresses links with ACM but without ATPC, since the principal objective is to determine what efficiency gains might be possible from introducing ACM and whether ACM might suffer from interference in realistic planned deployments.

Implementation approach

A number of changes were made to the fixed-link planning simulator in order to implement ACM. The planning simulator now operates as follows:

- Using the community-based assignment methods developed for soft boundary assignment, each service is allocated to the ACM or non-ACM community.
- The description of each ACM core service is expanded to include the higher-rate services that the link could use if propagation conditions allow. The higher-rate service description may also include an ‘interference offset’; adding this offset to the minimum RSL for the higher-rate service defines the point at which switching occurs.
- The determination of outages has been refined, so that the appropriate service descriptions are used for members of the ACM community, depending on whether the link has switched away from a higher-rate service as a result of a wanted link fade. A direct rain-induced outage is recorded for a link when all higher-rate services have fallen back and the fade margin for the core service is exceeded.

In practice, a simple simulation scenario was developed to explore the broad implications of the use of ACM:

- Changes in ACM scheme may only occur between bandwidth-compatible services. In practice, there is only ever one switching option for any service in the existing 38 GHz database. Real-world equipment may, of course, have more than one level of fall-back but, since the fading mechanism modelled here is rain, there is no significant benefit in capturing the additional methods (the simulator supports multiple additional service descriptions for each assignment).
- A baseline set of candidate assignments was established by demoting all assignments in the ACM community to the lowest bandwidth-compatible data rate service.
- An ‘enhanced’ set of candidate assignments was then created by adding switching capabilities to all assignments in the ACM community.
- It is assumed that switching between coding/modulation schemes is based on the properties of the wanted link (i.e. the rain fade) and does not require estimation of interference as well; antenna characteristics and service capabilities are also assumed to be the same as in the existing 38 GHz database.
- Hysteresis in the ACM switching is ignored. As the simulated rain fields do not necessarily represent a time series, the fade slope cannot be determined in order to implement different fall-back and return thresholds.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

- The interference offset was generally assumed to be 0 dB, in order to make the simulations as sensitive as possible to interference. A typical practical value might be 3 dB and simulations were also run for this value.

Three cases have been identified for simulation.

Case 1: Full fade margin

In this case, bandwidth-compatible services are demoted, with the higher bitrate service added as a switchable option. The purpose of the simulations is to quantify the effect of interference on services exploiting unused fade margin to operate at higher data rates.

Note that planning is based on the lower data rate service.

Case 2: Reduced fade margin

Planning is as for Case 1, but fade margins are reduced. The purpose of the simulations is to determine the spectrum efficiency gains that result from planning at reduced fade margin and to quantify the effect of rain fading and interference on these services.

Note that planning is based on the lower data rate service.

Case 3: Interference protection

In this case, the standard fade margin is allocated, but increased protection from interference is provided by using the W/U ratio of the higher-rate service. Again, the purpose of these simulations is to determine the effect on spectrum efficiency of the increased interference protection and the corresponding reduction in susceptibility to interference.

Simulation results

Effect on data rates

There are, of course, many possible scenarios that could be used to model the effect of introducing ACM into a band. The method used here is intended to produce a significant ACM penetration, so that any adverse effects can be detected by the simulation; the method does not attempt to determine a probable or economic service distribution. Two service codes are assigned for each service for which there is a lower data rate, bandwidth-compatible service—a ‘core’ service code and a ‘switched’ service code. The ‘baseline’ simulation cases correspond to all services operating with the lower data rate, core service; in the ‘switched’ simulation cases all services are operating—when possible—using the higher data rate, switched service. The following table shows the aggregate data carried^{3,4}.

³ For a full fade margin plan with 3 dB interference offset.

⁴ The 38 GHz fixed-link database does not include network associations, so summation of all link data rates will overestimate the actual data carried by the entire network.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

Service Code	Baseline Data Rate (Mbps)	Switched Data Rate (Mbps)	Total Number of Links	Baseline Data Rate (Mbps)	Switched Data Rate (Mbps)
DH	2	n/a	82	164	164
DH-DI	2	4	1,905	3,810	7,620
DJ	2	n/a	452	904	904
DJ-DK	2	4	2,238	4,476	8,952
DJ-EK	2	4	24	48	96
DK-DL	4	8	4,590	18,360	36,720
DK-DM	4	8	6,688	26,752	53,504
DN	8	n/a	1,730	13,840	13,840
DN-DP	8	17	2,618	20,944	44,506
DN-DQ	8	17	118	944	2,006
DR	34	n/a	1,874	63,716	63,716
DR-DT	34	155	160	5,440	24,800
DS	34	n/a	4,522	153,748	153,748
DU	155	n/a	74	11,470	11,470
Total			27,075 <i>67.7% ACM</i>	324,616	422,046

Table 5: Aggregate data rates.

Note that services operating at a lower data rate are not ‘promoted’ to a higher data rate, even if a bandwidth-compatible service is available. If such promotion were to be allowed, then the aggregate data rate would rise from 422 Gbps to 665 Gbps. Since the simulation method relies only on demotion of bandwidth-compatible services, the aggregate switched data rate is the same as the data rate for plans without ACM. The aggregate data rate will, of course, be reduced somewhat by the effects of rain: however, on an annual basis the change is very small.

Interference Offset

The effect of varying the interference offset can be seen in the table below, which shows the number of outages induced by a frontal rain event in a ‘full fade margin’ simulation containing a mix of ACM and non-ACM assignments.

Model	Interference Offset	Direct Outages	Interference Outages ⁵
Baseline (no ACM)	n/a	924	22
Switched	0 dB	828	61
Switched	1 dB	884	68
Switched	2 dB	924	70
Switched	3 dB	924	70

Table 6: Variation of outage rates with interference offset.

When there is no interference offset (i.e. 0 dB), the number of direct outages reduces when a switching capability is added. This happens because some higher-rate services have a lower minimum RSL than the lower-rate fall-back service; if there is no interference offset then the

⁵ These outages only occur in links that are already in outage because of rain fading. They do not therefore increase unavailability.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

higher-rate service will require a larger-than-planned fade to induce a direct outage—and has an effective uplift in its fade margin, which reduces the number of outages. For the service types that occur in the 38 GHz database, the largest difference in minimum RSL between bandwidth-compatible pairs of services is 1.5 dB, so the fade margin uplift effect disappears for an interference offset of 2 dB or greater.

To correct this anomaly an additional assumption about the ACM equipment was implemented in the simulation software: links were assumed to have switched to their fall-back service at or beyond the point at which a direct outage occurs. The response to varying the interference offset is then described in the following table.

Model	Interference Offset	Direct Outages	Interference Outages
Baseline (no ACM)	n/a	924	22
Switched	0 dB	924	44
Switched	1 dB	924	48
Switched	2 dB	924	53
Switched	3 dB	924	54

Table 7: Variation of outage rates with interference offset (corrected).

The number of direct outages is now unaffected by any variation in interference offset. The number of interference outages is reduced somewhat because links in direct outage are now operating with their fall-back service, which will usually have a less-constraining W/U requirement than the corresponding higher data rate service.

Case 1: Full fade margin

In this case, bandwidth-compatible services are demoted, with the higher bitrate service added as a switchable option. The purpose of the simulations is to quantify the effect of interference on services exploiting unused fade margin to operate at higher data rates.

Note that planning is based on the lower data rate service.

Annual percentage outage rates for simulations with varying interference offsets are listed below.

Baseline	0.00575%
Switched (int. offset 0 dB)	0.00575%
Switched (int. offset 3 dB)	0.00575%

Table 8: Variation of annual unavailability with interference offset.

The results show that there is no detectable increase in unavailability resulting from the exploitation of unused fade margin. A number of links that are in outage because the rain fade exceeds the fade margin will also fail the interference criterion, and this number increases as ACM is introduced: however, no link is in outage solely because it has failed the interference criterion. This is largely due to the switching options that are possible within the set of services that populate the 38 GHz fixed-link database. In particular, most bandwidth-compatible service changes do not affect the required W/U ratio and therefore do not increase the chance of an interference-induced outage: actual ACM-capable equipment may differ. The following table shows the variation in required W/U ratio and minimum RSL.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

Service Code	Service Description	W/U ⁶ (dB)	Min. RSL (dBW)
DH	2 Mbps in 3.5 MHz	27	-112.5
DI	4 or 2 x 2 Mbps in 3.5 MHz	27	-109.5
DJ	2 Mbps in 7 MHz	27	-108.0
DK	4 or 2 x 2 Mbps in 7 MHz	27	-105.0
EK	4 or 2 x 2 Mbps in 7 MHz	27	-105.0
DL	8 or 4 x 2 Mbps in 7 MHz	27	-106.5
DM	8 or 4 x 2 Mbps in 7 MHz	27	-106.5
DN	8 or 4 x 2 Mbps in 14 MHz	27	-102.0
DP	17 or 2 x 8 Mbps in 14 MHz	27	-103.5
DQ	17 or 2 x 8 Mbps in 14 MHz	27	-103.5
DR	34 or 2 x 17 Mbps in 28 MHz	27	-100.5
DS	34 or 2 x 17 Mbps in 28 MHz	27	-100.5
DT	155 Mbps in 28 MHz	38	-90.0
DU	155 Mbps in 56 MHz	36	-92.5

Table 9: Service planning parameters.

Case 2: Reduced fade margin

Planning is as for Case 1, but fade margins are reduced. The purpose of the simulations is to determine the spectrum efficiency gains that result from planning at reduced fade margin and to quantify the effect of rain fading and interference on these services.

Note that planning is based on the lower data rate service.

Annual percentage outage rates for simulations with varying fade margin reductions and interference offsets are listed below.

	Fade margin reduction			
	None	3 dB	5 dB	7 dB
Baseline	0.00575%	0.00688%	0.00761%	0.00826%
Switched (int. offset 0 dB)	0.00575%	0.00688%	0.00761%	0.00826%
Switched (int. offset 3 dB)	0.00575%	0.00688%	0.00761%	0.00826%

Table 10: Variation of annual unavailability with fade margin reduction.

As expected, the outage rate increases as the fade margin is reduced. This effect is almost entirely attributable to the change in the number of direct outages resulting from reduced protection against rain fades: the contribution from extra interference-induced outages is negligible. The only simulations in which any extra interference-induced outages occur are when the interference offset is reduced to 0 dB, in which case a maximum of 16 outages is recorded compared with 130,448 direct outages (for a fade margin reduction of 7 dB). When the interference offset has a more realistic value of 3 dB, then there are no extra interference-induced outages at all. The assumed interference offset therefore has no significant effect on outage rates.

⁶ Co-channel W/U used by the MWM library.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

The spectrum efficiency gains are listed below, for the ‘number of links assigned to channel 1’ statistic.

	Fade margin reduction			
	None	3 dB	5 dB	7 dB
Baseline	56.6%	59.4%	60.6%	61.9%
Switched (int. offset 0 dB)	56.6%	59.4%	60.6%	61.9%
Switched (int. offset 3 dB)	56.6%	59.4%	60.6%	61.9%

Table 11: Variation of spectrum efficiency with fade margin reduction.

The effect of the fade margin reduction is to reduce the required link EIRP, which allows more links to be assigned. The statistic therefore shows an increase as the fade margin reduction becomes larger. The assumed interference offset has no effect on spectrum efficiency.

From this analysis, ACM is a benign technology as far as interference is concerned. If higher-than-planned data rates are permitted in exchange for a fade margin reduction, then an increase in unavailability results, but that increase is determined by the direct effect of rain fading on the wanted link and not by interference from other links.

Case 3: Interference protection

In this case, the standard fade margin is allocated, but increased protection from interference is provided by using the W/U ratio of the higher rate service. Again, the purpose of these simulations is to determine the effect on spectrum efficiency of the increased interference protection and the corresponding reduction in susceptibility to interference.

Note that the ‘number of assignments in channel 1’ for a re-plan in which no service demotions occur is 50.3%. This represents the worst-case decrease in efficiency that would result from planning the band based on the higher-rate ‘switched’ service. In practice, the Case 3 plans should have a better efficiency because the assigned EIRP will be that required by the core service and not the switched service.

Baseline	56.6%
Switched (int. offset 0 dB)	55.6%
Switched (int. offset 3 dB)	55.6%

Table 12: Variation of spectrum efficiency with interference offset.

Replacement links

An alternative method of assessing the efficiency of ACM as a means of adding data capacity is to compare the ACM case with a case in which each ACM-capable link is replaced by a link carrying the core service and an additional link to carry the extra capacity.

A worst-case simulation has therefore been run in which a Case 3 plan (i.e. all services protected) is compared with a plan in which additional links are added over the same paths, with the core service availability. In practice, the enhanced service links would have a lower availability than the core service.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

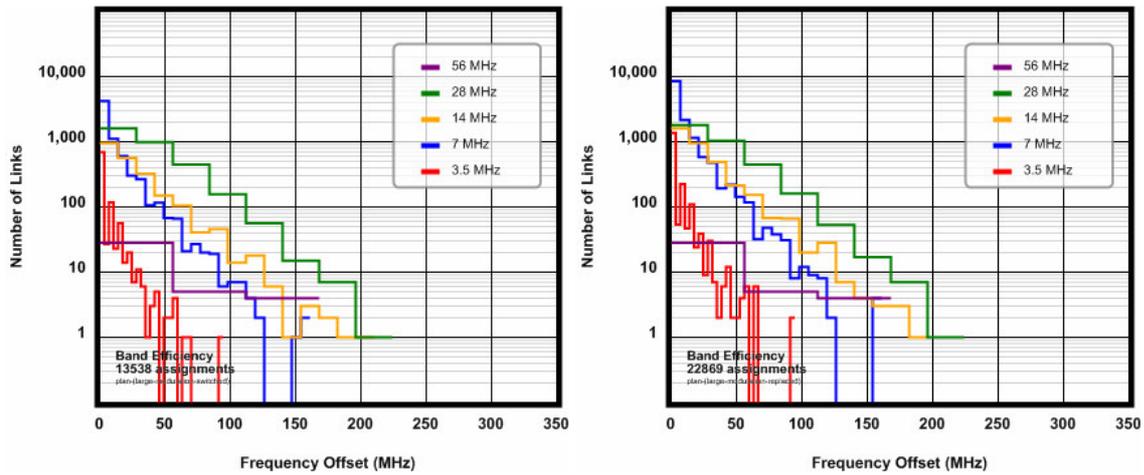


Figure 14: Plans with aggregate data rate from ACM and replacement links.

It is apparent from the graphical presentation of the plans that the shapes and extents of the assignment distributions are similar, though the number of assigned links has clearly increased in order for the ‘replacement’ plan to carry the same data as the baseline case (from 13,538 to 22,869). In other words, the efficiency of the replacement plan is not greatly altered—even with a large number of additional parallel links. In fact, the ‘replacement’ plan has a higher efficiency as measured by number of assignments made to the starting channel (58.1% vs 55.6%).

Work Package 3: Rain Field Simulator

The rain field simulation work done in this project expands on the previous work already done as part of the ATPC project [Callaghan et al, 2005]. The rain field model described in the final report produces two dimensional rain fields of a stratiform/convective type, with each two-dimensional field a separate realisation of a rain event. In order to test the long-term statistics of the behaviour of a system operating with ATPC or ACM, simulated rain field datasets are required which match the annual rain statistics for the average year. Hence, a method of converting from a set of single events into a set capable of reproducing annual statistics effectively is the desired aim of this work package.

The following figures show example realisations of single convective and stratiform rain fields resulting from the rain field simulator described in [Callaghan et al, 2005].

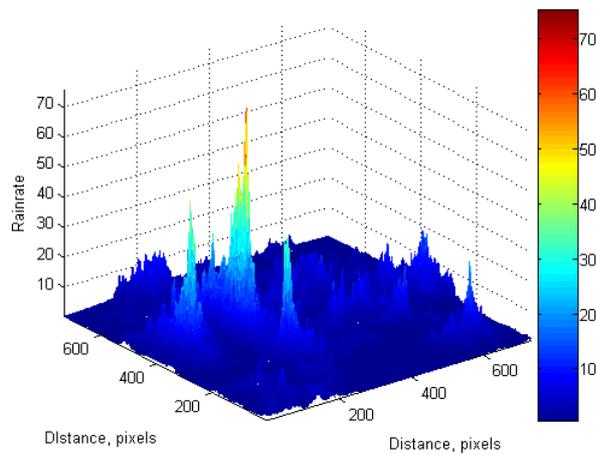


Figure 15: Simulated convective rain field.

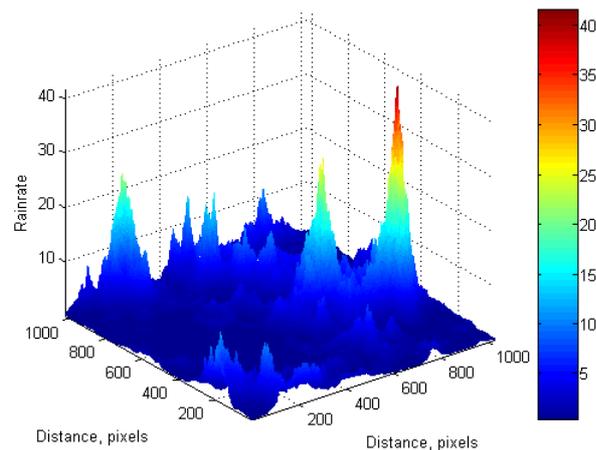


Figure 16: Simulated stratiform rain field.

Modifications made to the simulator

Scaling according to $R_{0.01}$

The discrete, additive cascade process used to produce the simulated rain fields results in arrays that are equivalent to log rain rate fields. Hence, some processing is required to convert to equivalent rain rates.

This can be done by using the following formula:

$$R_{sim} = 10^{(V + \log_{10}(R_{0.01}) - V_{0.01})} \quad (2)$$

where V is the simulated array values, $V_{0.01}$ is value in V exceeded by 0.01% of the points in V , $R_{0.01}$ is the rain rate exceeded in 0.01% of an average year and R_{sim} is the resulting simulated array of rain rates. ($V_{0.01}$ must be greater than $\log_{10}(R_{0.01})$.)

This offset alters the mean of the simulated distribution but will not affect any of the higher order moments. $R_{0.01}$ is chosen as the parameter to which to scale the simulated distribution, as it is used in link budget planning, and hence is easily calculated from ITU-R Recommendation P.837-4.

Distribution of stratiform and convective events through an average year

Stratiform events typically occur during the winter and spring months. They are characterised by long periods of low intensity rain, which is widespread and spatially homogeneous (see Figure 17). Vertical meteorological radar scans through the rain event reveal a horizontal bright band several kilometres about the ground (Figure 18). This is the melting layer, which corresponds with the zero degree isotherm. Above the melting layer, atmospheric water is found as ice crystals, while below the melting layer, the temperature is such that the ice melts, forming liquid rain drops. The bright band is the layer of the atmosphere where liquid water and ice crystals coexist.

Convective rain events typically occur during the summer and autumn months. They are intense rain showers with short durations, and are localised (see Figure 19). Vertical meteorological scans show that there is no melting layer observable during convective events, but the convective events themselves are very turbulent (Figure 20).

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

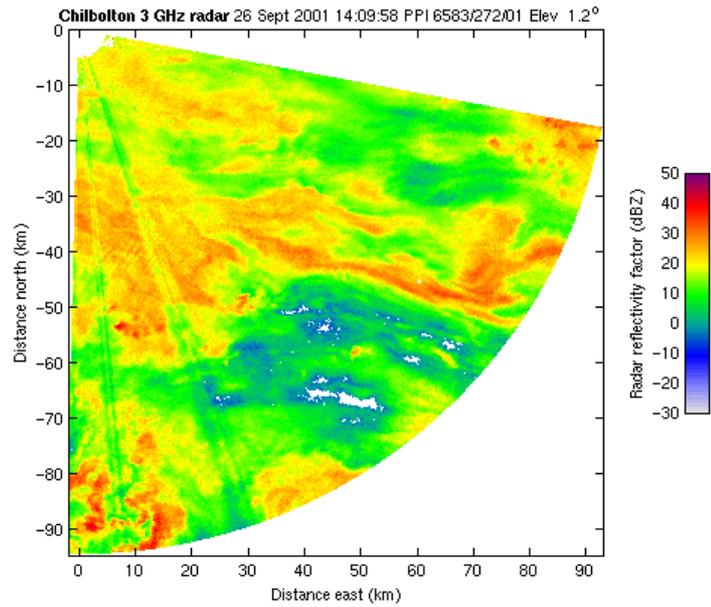


Figure 17: Example horizontal cross section through a stratiform rain field

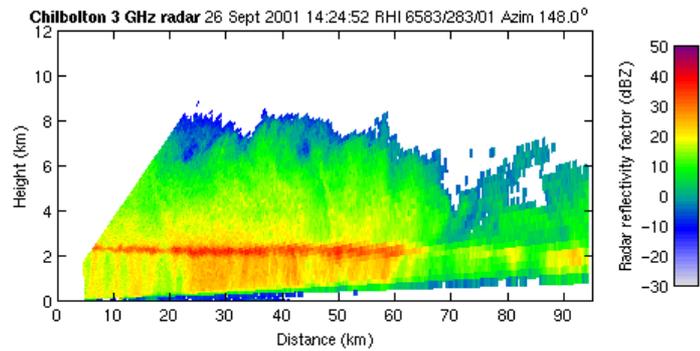


Figure 18: Example vertical cross section through a stratiform rain field.

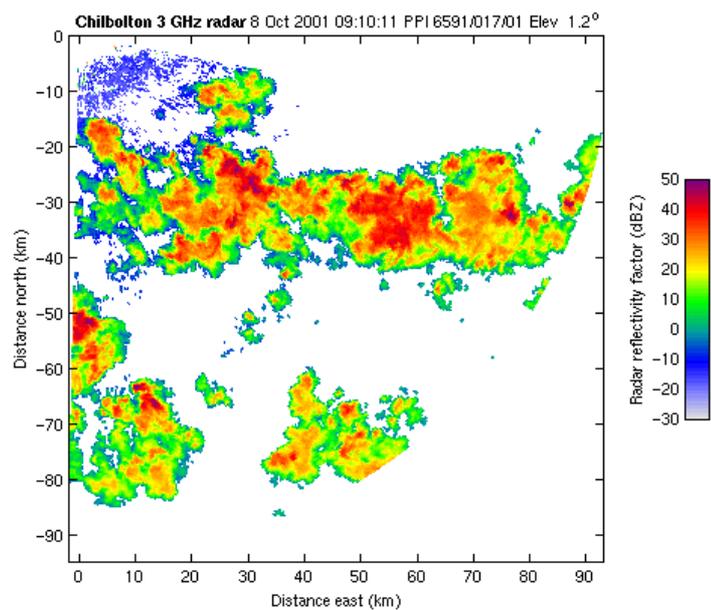


Figure 19: Example horizontal cross section through a convective rain field.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

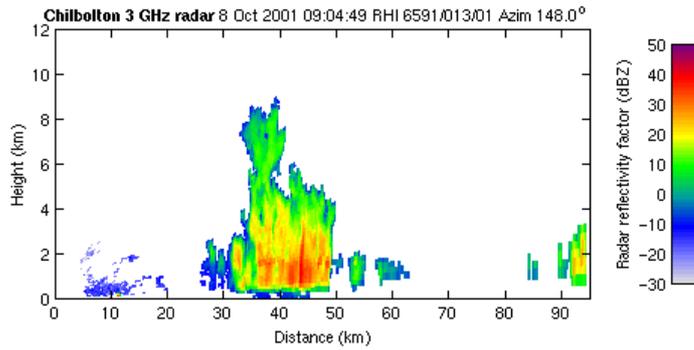


Figure 20: Example vertical cross section through a convective rain field.

A thorough literature search has failed to find any studies that investigate what proportion of rain events can be classified as stratiform or as convective in an average year. As can be seen from the graphs above, meteorological radar provides the best information on whether a rain event has a melting layer (i.e. is stratiform) but long-term continuous radar data sets are rare, due to the expense involved in collecting them.

Long term continuous rain gauge data sets exist, however they only provide rainfall accumulations, and do not provide any information about the vertical structure of the rain field.

An indication of the proportions of stratiform and convective rain in an average year can be found using the variables M_c and M_s as given in Rec. 837-4, where:

$$M_c = \text{annual average convective rainfall amount (mm)} \quad (3)$$

$$M_s = \text{annual average stratiform rainfall amount (mm)} \quad (4)$$

Creation of a simulated annual dataset

Each realization of the simulated rain fields is independent of each other, hence simulating multiple arrays and taking the same point from multiple arrays will not produce realistic simulated time series. However, it is possible to produce a large number of simulated arrays and taking them as a group, scale them to $R_{0.01}$. These simulated databases are then capable of reproducing long-term statistics.

As convective and stratiform arrays are produced separately, it is necessary to determine what proportion of rain events in the simulated database should be stratiform and what convective. For a simulated “annual” database, these proportions can be found using the variables M_c and M_s as given in Rec. 837-4.

Comparisons with ITU-R Recommendation P.837-4

Figure 21 shows the cumulative distributions of the simulated rain fields, in comparison with measured rain gauge data recorded at Sparsholt, Hampshire, Rec. 837-4, and the CDF of the total of all the simulated fields. The simulated fields are scaled to $R_{0.01}$ given by Rec. 837-4 for the latitude and longitude of Sparsholt, hence the crossing point where all the simulated curves and the model curve coincide.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

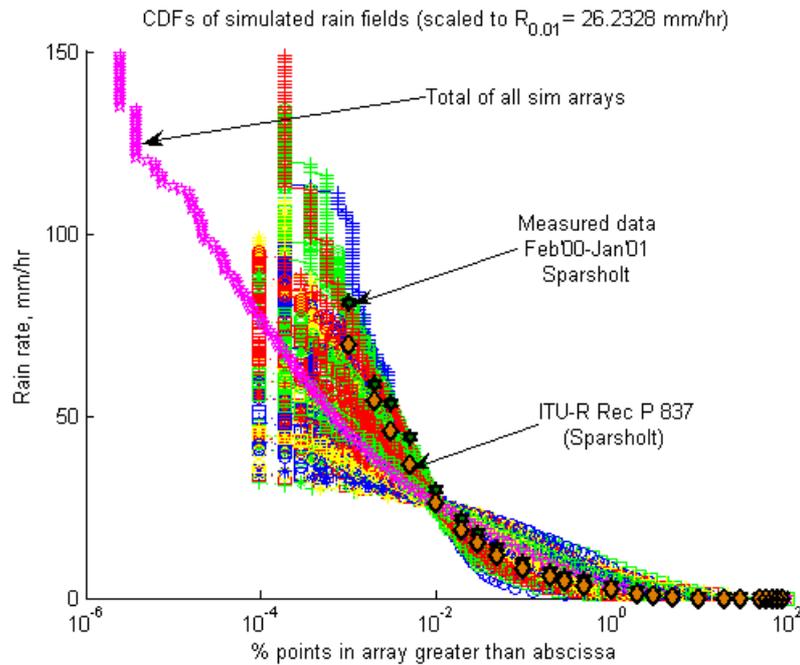


Figure 21: Cumulative distributions of each simulated rain field, in comparison with measured rain gauge data recorded at Sparsholt, Hampshire, ITU-R Rec. P.837-4, and the CDF of the total of all the simulated fields.

In order to investigate the effects of changing the proportion of stratiform and convective events to simulate a year, spatio-temporal equivalence is assumed, and the simulated and measured data sets are required to have the same number of data points.

For example:

Calendar year: February 2000 to January 2001 = 31,622,400 data points (seconds in a year—note that 2000 was a leap year).

1 simulated convective field = 532,900 data points (730*730 pixels in each field).

1 simulated stratiform field = 1,050,625 data points (1025*1025 pixels in each field).

Therefore, a dataset can be created where 30 stratiform fields are equivalent to a year, 60 convective fields are equivalent to a year, or 30 convective fields and 15 stratiform fields are equivalent to a year.

Figure 22 shows the results for these different combinations of simulated fields in comparison with measured rain gauge data recorded at Sparsholt and Rec. 837-4. From these curves it can be seen that the simulated data sets overestimate slightly at low rain rates, and underestimate at high rain rates. However, given annual variability and the uncertainty of what proportions to mix the stratiform and convective fields, the simulated curves are in reasonable agreement with the measurements and ITU-R model.

It is also interesting to note that a 50-50 mix of stratiform and convective fields results in a curve that is very close to the all-convective curve.

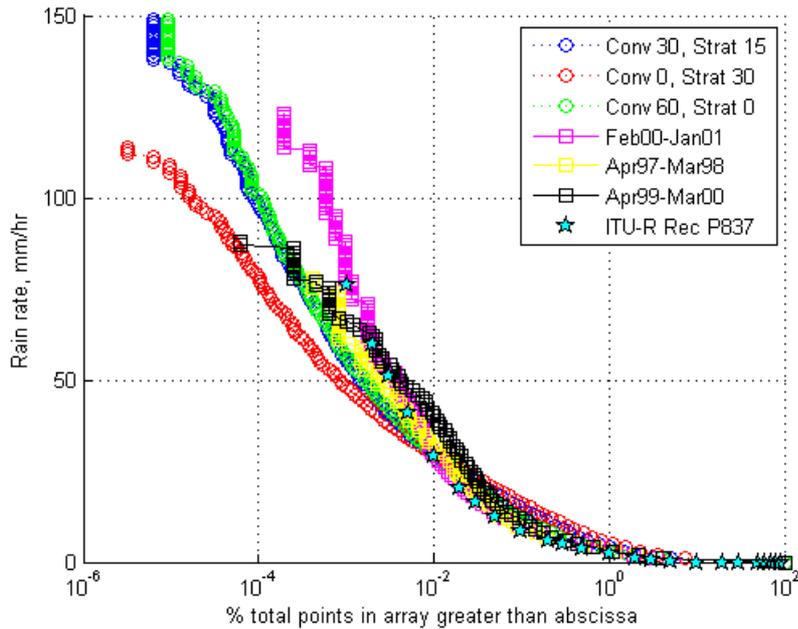


Figure 22: Different combinations of proportions of stratiform and convective fields in comparison with measured yearly statistics derived from Sparsholt rain gauge measurements and ITU-R Rec. P.837-4.

Figure 23 shows the simulated dataset for Chilbolton, assuming spatio-temporal equivalence, and using the stratiform and convective rain fields in proportion as given by M_c and M_s . The simulated curve is still scaled to $R_{0.01}$. It can be seen that using M_c and M_s gives good agreement along most of the length of the two curves.

This also suggests that if an entire year's worth of two-dimensional rain fields is simulated in the proportions given by M_c and M_s and scaled them to $R_{0.01}$ then the resulting statistics would be consistent with the ITU model.

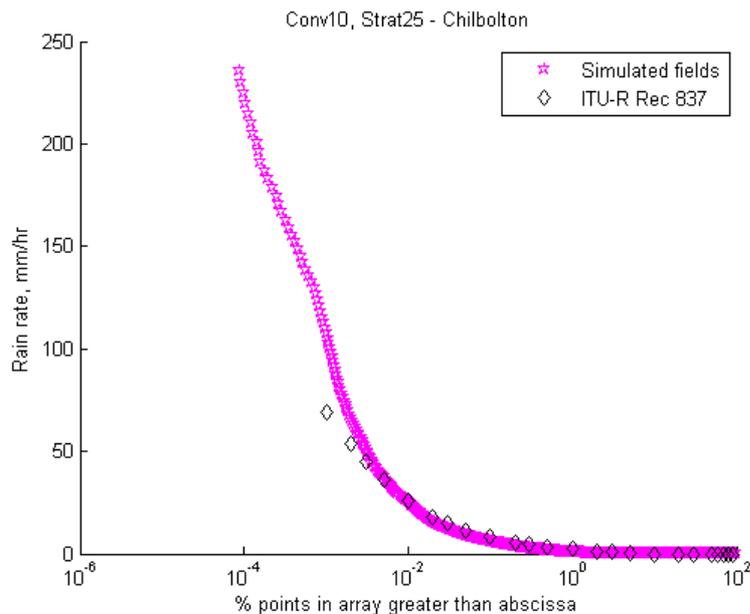
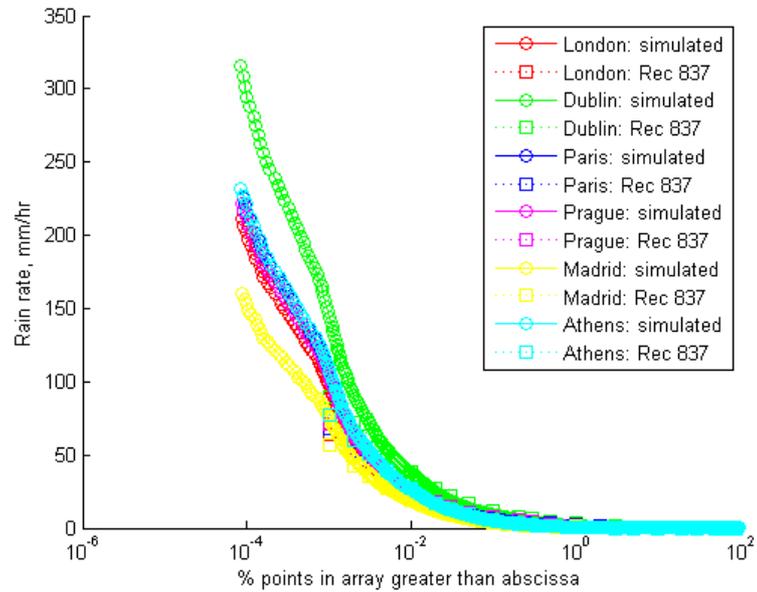


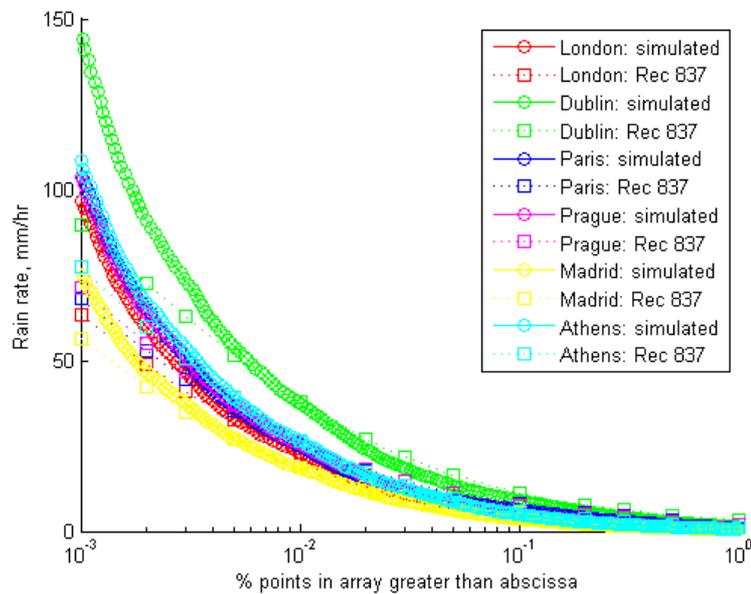
Figure 23: Simulated annual curve (using convective and stratiform events in proportions given by M_c and M_s) in comparison with the ITU-R Rec. P.837-4 model for Chilbolton.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

For different European cities, Figure 24 shows the Rec. 837-4 model curve in comparison with the simulated annual statistics, scaled to $R_{0.01}$ taken from Rec. 837-4, again using convective and stratiform events in proportions given by M_c and M_s . In general, the agreement between the simulated statistics and the ITU-R model curves is good. This also demonstrates that the simulator is suitable for other locations, not just the UK. Further work is suggested to determine if the simulator is appropriate for use in geographic locations outside Europe.



(a) Complete curves



(b) Restricted range of most interest to system engineers

Figure 24: Simulated annual curve in comparison with the ITU-R Rec. P.837-4 model for various European cities.

Figure 24b shows that the simulated data seems to overestimate the ITU predictions for percentage values less than 0.01%. However, for percentage values greater than 0.01%, the fit between ITU model and simulation is good.

Simulated two-dimensional annual database

If it is assumed that each simulated field is equivalent to a meteorological radar snapshot at any given minute, then 43,200 arrays will be needed to create a simulated month (30 days) and 525,600 arrays for a year. The memory requirements to store the simulated arrays are such that a year's worth of arrays would occupy ~6,150 GB of computer memory and would take ~40 days to create.

It is known that it only rains for a small percentage of a year; however, in the case of the simulated arrays, the rain doesn't cover the entire array area, i.e. there are areas of rain rate which approach zero (equivalent to no rain). Hence there is a desire to create a subset of events which will accurately reproduce the tail of the annual distribution as given by Rec. 837-4.

Simulating the tail of the ITU annual model

If it is assumed that each simulated array is equivalent to 1 radar snapshot/minute, then a database covering 0.1% of a year requires 526 arrays. This smaller simulated dataset was therefore fit to the Rec. 837-4 values in the range 0.1% to 0.001%.

To do this scaling, two new parameters are needed, a and b —note that these two parameters *supersede* the use of $R_{0.01}$ as a scaling parameter.

$$R_{sim} = 10^{\frac{V}{a} + b} \quad (5)$$

where:

R_{sim} is the simulated rain array (mm/hr)

V is the data array produced by the simulator

The parameter a is given by:

$$a = \frac{std(R_{GA} \geq m_{GA})}{std(\log_{10}(I_{R \leq 0.1}))} \quad (6)$$

where:

R_{GA} is the simulated rain gauge data extracted from V (equivalent to log values)

$I_{R \leq 0.1}$ is the rain rate exceeded for percentages of time less than 0.1% (from Rec. 837-4)

m_{GA} is the mean of R_{GA}

std is the standard deviation

The parameter b is given by:

$$b = \left| \frac{m_{GA}}{a} \right| + \left| \log_{10}(I_{R=0.1}) \right| \quad (7)$$

R_{GA} is normally distributed and has a zero mean, hence:

$$b = \left| \log_{10}(I_{R=0.1}) \right| \quad (8)$$

where $I_{R=0.1}$ is the rain rate exceeded for 0.1% of the time (from Rec. 837-4).

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

For the ITU model for Chilbolton, $a=3.2266$ and $b=0.9106$.

The step-by-step methodology for determining the values of a and b can be found in appendix B.

Figure 25 shows the Rec. 837-4 curve for percentage times less than 0.1%, in comparison with the measured Chilbolton and Sparsholt annual curves (Jan 05–Dec 05) and the CDFs for simulated rain gauges A and B. The simulated gauges are separated by a distance equivalent to 7.5 km, approximately equal to the distance between Sparsholt and Chilbolton.

The CDF of the simulated gauge A has been scaled using parameters a and b to fit it to the ITU curve (making it equivalent to Chilbolton). The CDF of B uses the same a and b , but isn't as closely tied to the ITU curve.

There is good agreement of the simulated curves with the ITU curve for percentage values less than 0.05%.

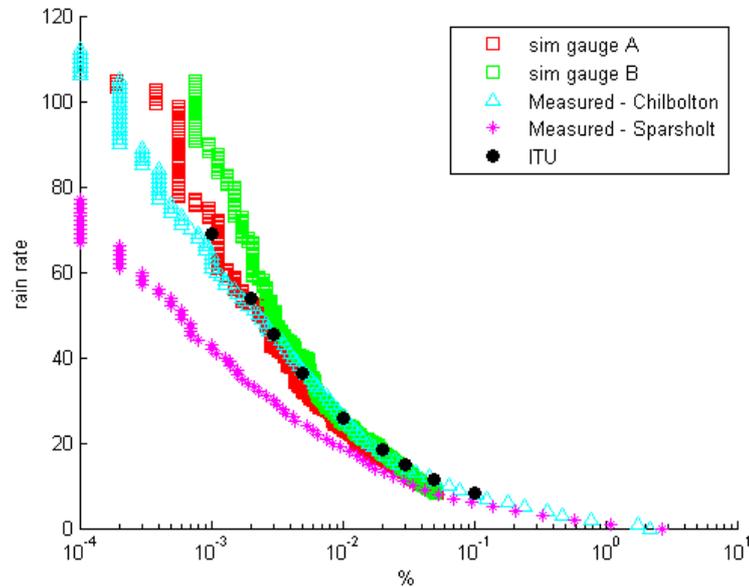


Figure 25: ITU Rec. P.837-4 curve for percentage times less than 0.1%, in comparison with the measured Chilbolton and Sparsholt annual curves (Jan 05–Dec 05) and the CDFs for simulated rain gauges A and B.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

The following figures show example simulated rain fields, scaled according to the Chilbolton values for a and b .

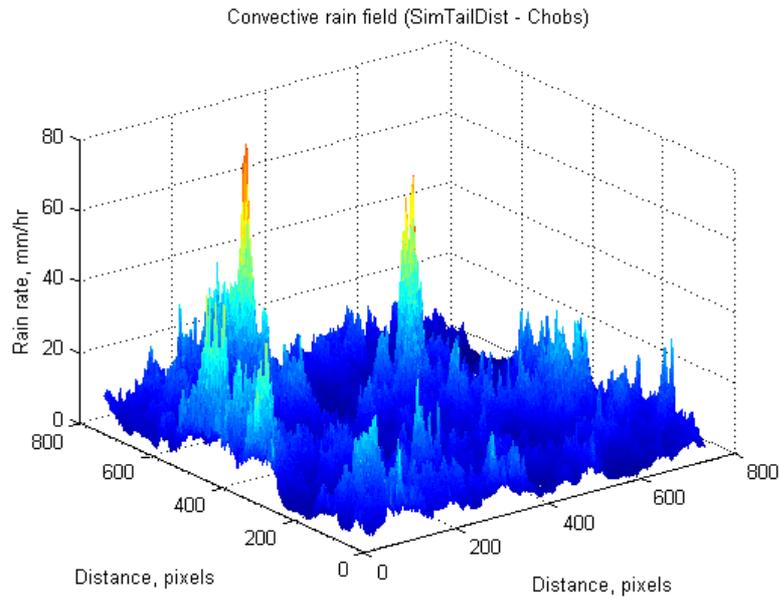


Figure 26: Example convective rain field scaled according to Chilbolton data ($a=3.2266$ and $b=0.9106$).

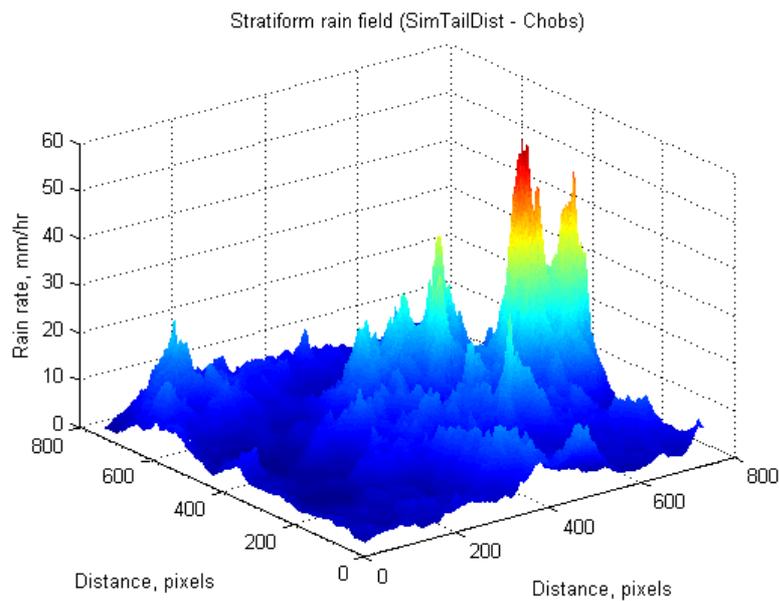


Figure 27: Example stratiform rain field scaled according to Chilbolton data ($a=3.2266$ and $b=0.9106$).

Conclusions and caveats regarding the rain field simulator

- The results presented in this report suggest that simulating an entire year's worth of 2-D rain fields in the proportions given by M_c and M_s and scaling them to $R_{0.01}$ results in statistics that are consistent with the ITU model.
- Unfortunately, this is not practical, due to computer memory and processing time constraints.
- The parameters a and b allow the scaling of the CDFs of simulated gauges (derived from simulated rain fields) to the tail of the ITU model for the rain rate CDF, specifically for time percentages less than 0.05%.
- The simulated database required to do this has 526 arrays, in the stratiform/convective proportions given by M_s and M_c .
- Comparison of the scaled simulated and ITU modelled CDFs are in good agreement for the tail of the ITU distribution. Results so far have only been produced for the single site of Chilbolton. This method should be tested for other geographic locations.

The results presented here suggest that this method of scaling works. However, there are a few caveats that should be noted:

1. Tails of distributions are never that well defined, due to the limited numbers of data points that create them. Hence, fitting simulated data to the tail of the ITU model carries the risk that the ITU model's tail might be inaccurate.
2. To fit the simulated data to the tail of the ITU model, the mean of the simulated data distribution is being adjusted, but so is the standard deviation. This is unlikely to affect the spatial correlation, as all the points in all the simulated data arrays are treated in exactly the same way. However, this tweaking of the statistical parameters is not mathematically rigorous. Further work is suggested to give this method a more solid statistical foundation.
3. Each point in a simulated rain field, when scaled by a and b , will have a rain rate greater than zero. Given the areas that the simulated fields cover, this is physically very unlikely.
4. After scaling, the simulated CDF percentage values only map to the ITU curve for percentage values less than 0.05%. This may not be that much of an issue to systems designers, depending on what outage probabilities they require.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

References

- Bacon, D., "Investigation into effect of modulation order on spectral efficiency in fixed-link assignment: 38 GHz simulation results", PPWG (03-03)/053, 2003.
- Callaghan, S.A., Hansell, P., Inglis, I., "Impact of introducing Automatic Transmit Power Control in P-P Fixed Service Systems operating in bands above 13 GHz", Final report for Ofcom, 19th December 2005.
- AXT-PtP Planning Approach. Ericsson. 7 February 2006.
- Deidda, R., "Multifractal analysis and simulation of rainfall fields in space", *Physics and Chemistry of the Earth Part B—Hydrology, Oceans and Atmosphere*, 24 (1-2), 73–78, 1999.
- Ferraris, L., Gabellani, S., Parodi, U., Rebori, N., von Hardenberg, J., Provenzale, A., "Revisiting multifractality in rainfall fields" *J. Hydromet.* 4, pp 544–551, 2003.
- Gupta, V.K. and Waymire, E., "A statistical analysis of mesoscale rainfall as a random cascade", *Jnl. Applied Meteorology*, Vol. 32, 251–256, 1993.
- Harris, D., Menabde, M., Seed, A.W., Austin, G.L., "Multifractal characterization of rain fields with a strong orographic influence", *Jnl. Geophys. Res.* Vol. 101, D21, 1996.
- Lovejoy, S. and Schertzer, D., "Multifractals, universality classes and satellite and radar measurements of clouds and rain fields", *Journal of Geophysical Research*, Vol. 95, 2021–2034, 1990.
- Olsson, J., and Niemczynowicz, J., "Multifractal analysis of daily spatial rainfall distributions", *Journal of Hydrology*, 187, 29–43, 1996.
- Schertzer, D. and Lovejoy, S., "Physical modelling and analysis of rain and clouds by anisotropic scaling multiplicative processes", *Journal of Geophysical Research*, Vol. 92, No. D8, 1987.

Appendix A: Mathematical basis for the simulator

The rain field simulator used in this research is based on an additive cascade process, operating in the logarithmic domain. The resulting fields are monofractal, and are hence characterised by a single fractal dimension. Other cascade models presented in the literature [Deidda, 1999, Lovejoy and Schertzer, 1990] are multiplicative, and produce multifractal fields, which are characterised by a whole spectrum of fractal dimensions.

There is some controversy in the literature over whether or not rain fields can be adequately simulated using monofractals. The following sections attempt to address this issue.

The $K(q)$ function for characterisation of multifractals

It is commonly taken that the $K(q)$ function, sometimes known as the moment scaling function, can be viewed as a characteristic function of multifractal behaviour. The shape of the $K(q)$ function specifies the type of scaling involved for a given dataset. A curved $K(q)$ function indicates a multifractal structure, whereas if the $K(q)$ function is straight, this is indicative of a monofractal structure.

Schertzer and Lovejoy [1987] present a method of calculating $K(q)$ through investigation of the variation of statistical moments with scale. Fields produced by multiplicative cascade processes may have scaling behaviour that can be expressed by different scale independent scaling relationships. One of these relationships is given by:

$$\langle \varepsilon_\lambda^q \rangle \approx \lambda^{K(q)} \quad (\text{A.1})$$

where $\langle \varepsilon_\lambda^q \rangle$ is the ensemble average q th moment of the field studied at (i.e. averaged over) a scale specified by λ . λ , sometimes called the scale ratio, is defined as the ratio of the outer, maximum scale of the field to the averaging scale.

The following general forms have been proposed for the $K(q)$ function [Lovejoy and Schertzer, 1990, Olsson and Niemczynowicz, 1996]:

$$K(q) = \frac{C_1}{\alpha_L - 1} (q^\alpha - q) \quad 0 \leq \alpha_L < 1, 1 < \alpha_L \leq 2 \quad (\text{A.2})$$

$$K(q) = C_1 \log(q) \quad \alpha_L = 1 \quad (\text{A.3})$$

where C_1 is the co-dimension of the mean process and α_L , also known as the Lévy index, is related to the type of multifractal process involved. It should be mentioned that the generality of these forms has been questioned [Gupta and Waymire, 1993], as they are based on certain assumptions about the cascade structure of the fields.

The moment scaling function was used to characterise the multifractal nature of the measured radar data in preference to other methods of multifractal analysis, as this enabled us to directly compare results with others published in the literature [Olsson and Niemczynowicz, 1996, Harris et al, 1996].

Multifractal analysis of radar derived rain fields

The multifractal behaviour of the radar derived rain fields was investigated through the variation of statistical moments with scale using Equation A.1.

The fact that $\langle \mathcal{E}_\lambda^q \rangle$ is an ensemble average means that it is not certain that an individual field (i.e. radar raster) is properly described by the above equation. It is only after averaging over a number of fields that the expected multifractal behaviour is likely to appear. It is also assumed that the data has temporal stationarity for this analysis, as temporal stationarity is required to get accurate ensemble averages from multiple fields. This temporal stationarity, however, has not been confirmed by measurements.

To analyse the radar data, the area of the field is successively divided into non-overlapping squares of side length λ . For each λ the average rainfall volume in each square \mathcal{E}_λ is obtained as the average rain rate in the square. To obtain \mathcal{E}_λ^q , the average values are raised to the power of q , and $\langle \mathcal{E}_\lambda^q \rangle$ is obtained by averaging \mathcal{E}_λ^q firstly over all squares, and secondly over all fields (radar rasters).

The validity of Equation A.1 is tested by plotting the average moments $\langle \mathcal{E}_\lambda^q \rangle$ as a function of λ on a log-log diagram. If the points fall on a straight line, then the value of $K(q)$ can be estimated as the slope of line. The entire $K(q)$ function can be estimated by performing the above procedure for different values of q .

Figure A.1 shows a plot of the average moments $\langle \mathcal{E}_\lambda^q \rangle$ against λ for values of q between 0.5 and 3.5. As can be seen, Equation A.1 holds, as plots for various q are straight lines.

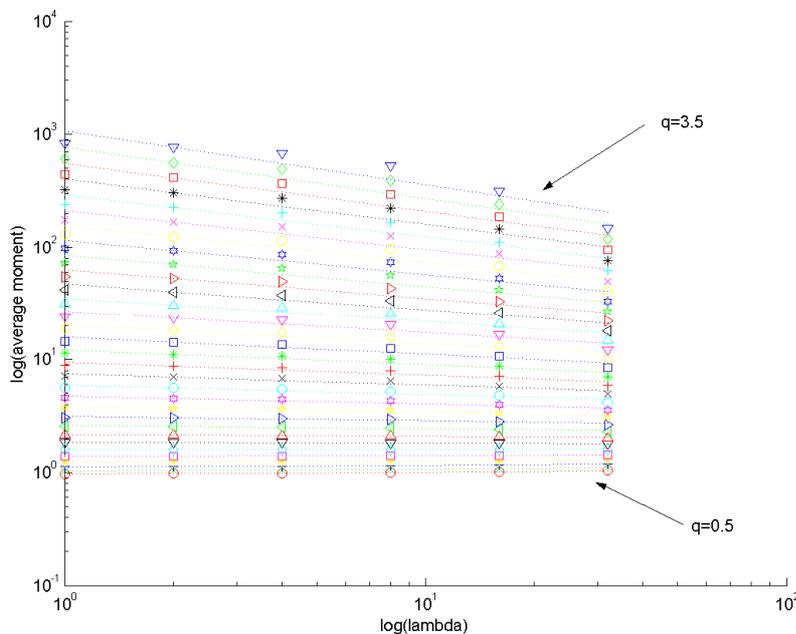


Figure A.1. Estimation of $K(q)$ values for different q .
Frontal rain event, recorded on the 1st May 2001.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

Figure A.2 shows the full $K(q)$ functions for the three different rain event types studied. All rain types show a curving $K(q)$ function, indicating multifractal behaviour, though there is dramatic variation between the rain types, with the frontal event being the most curved and the stratiform the closest to a straight line. This corroborates other results presented by Olsson and Niemczynowicz [1996], though their results show a convex curve. Here, the results curve in the opposite direction because when the rain is averaged across the radar rasters, the resulting ensemble average is generally very small. When these average values are raised to exponents greater than 1, their value decreases, resulting in the concave curve seen here. Table A.1 shows the number of points in each data set for the three event types, and gives the percentage of those points with a value greater than 1 mm/hr.

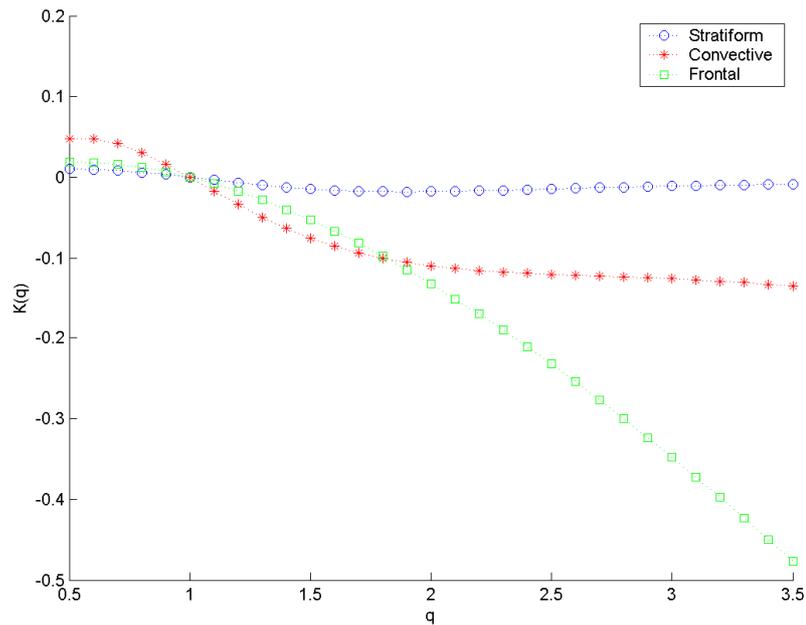


Figure A.2. $K(q)$ functions for different rain event types.
Values averaged are rain rate, in mm/hr.

Data set	Total number of points in data set	% of points in data set with values greater than 1 mm/hr
7 th December 2000 (stratiform event)	9,224,784	90.4%
1 st May 2001 (frontal event)	8,164,464	66.6%
16 th May 2001 (convective event)	4,630,064	96.1%

Table A.1: Number of data points in each data set, and percentage of those points with a rain rate value of less than 1 mm/hr.

Figure A.3, similarly to Figure A.2, shows the full $K(q)$ functions for the three different rain event types studied. In this case, however, it is the *log rain rate values* that are analysed, rather than the rain rate values. This is a crucial distinction, as the $K(q)$ functions for the log rain rate

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

values are not curved, and may be considered to be straight lines, *indicating a monofractal structure*.

From this demonstration, it can be seen that even though rain fields can be considered multifractal, log rain rate fields can equally validly be considered to be monofractal. This provides a justification for using a method of simulating rain fields that is based on a monofractal, additive process in the logarithmic domain.

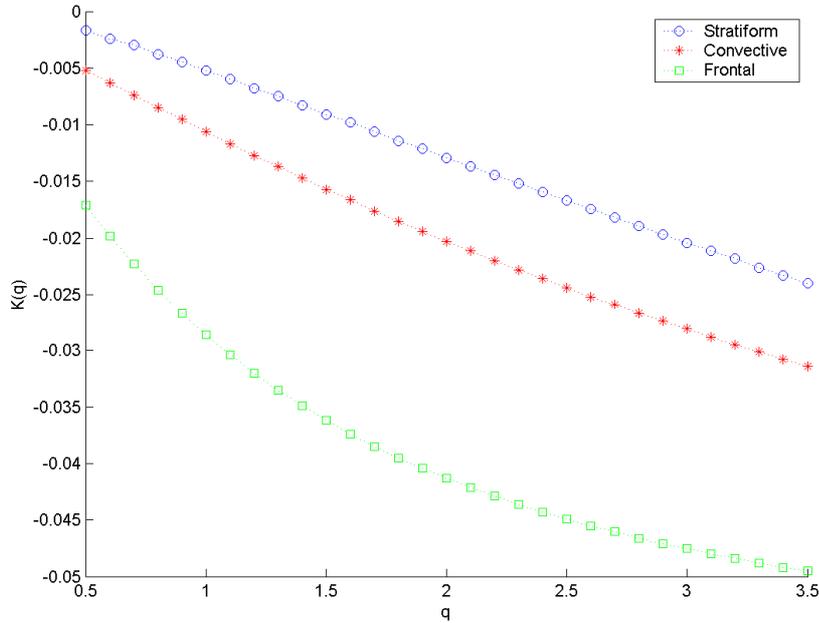


Figure A.3. $K(q)$ functions for different rain event types.
Values averaged are log rain rate.

Converting from a linear additive process to a non-linear multiplicative process

Ferraris et al. [2003] discusses the relationships between additive processes and multiplicative processes in detail. The following discussion uses the terminology used in [Ferraris et al, 2003].

In an additive process, a random variable S (the simulated rain field) is obtained as the sum of many individual components:

$$S = \sum_i X_i \tag{A.4}$$

If the individual elements X_i have a finite variance and are independent of each other, this additive process leads to a Gaussian amplitude distribution for S , because of the central limit theorem⁷.

In a multiplicative process a random variable Q is obtained as the product of a set of factors Y_i :

$$Q = \prod_i Y_i \tag{A.5}$$

⁷ The central limit theorem states that if the sum of the variables has a finite variance, then it will be approximately normally distributed.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

A multiplicative process of this type can be expressed as an additive process in log coordinates:

$$\log Q = \log \prod_i Y_i = \sum_i \log Y_i \quad (\text{A.6})$$

In the simplest multiplicative process, the factors Y_i are independent and have a finite variance. In this case, again due to the central limit theorem, the multiplicative process leads to a log-normal distribution. (It is commonly accepted that rain is lognormally distributed in space and time.) The corresponding additive process in log coordinates is also the sum of independent terms with finite variance and can be associated with a linear stochastic process.

By transforming the variables, it is possible to linearise the problem. Here the non-linear variable Q is obtained by a non-linear transformation of a linear additive process given by the sum of independent terms. In this case, the higher order spatial correlation and the Fourier phase correlations are entirely generated by the a posteriori non-linear transformation as they are absent in the original additive process. This type of process is called “Meta-Gaussian”.

Not all multiplicative processes are so simple. In some cases, taking $\log Q$ produces a sum of correlated terms as well as a mixture of additive and multiplicative terms that can't be fully resolved by a simple log transformation. For such a process, Q cannot be expressed as a non-linear transformation of a linear additive stochastic process and the dynamics cannot be linearised by a change of variable.

Results discussed in the paper [Ferraris et al, 2003] show that the multifractal behaviour of the GATE radar rain fields can be originated by a non-linearly filtered, linearly correlated stochastic process. As a consequence, multifractal cascades and non-linear stochastic processes are not necessary to generate the anomalous scaling observed in large intensity GATE rainfall fields. This agrees with the results from the multifractal analysis performed on the CAMRa measured rain field data.

Appendix B: Method used to determine the scaling parameters a and b .

As mentioned in the main body of the report, 0.1% of a year is 526 arrays, assuming each array is equivalent to 1 radar snapshot/minute. This smaller dataset needs to be fit to the ITU values in the range 0.1% to 0.001%.

Figure B.1 plots the whole distribution of a simulated rain gauge in the dataset of 526 arrays, without scaling, in comparison with Rec. 837-4 values. A simulated rain gauge is used because Rec. 837-4 is formulated using rain gauge measurements, rather than rain field measurements. From this simple comparison it can be seen that the simulated curve needs significant scaling.

Note that during these comparisons the raw values in the simulated fields V are being compared with log10 rain rate as given by the Rec. 837-4 model.

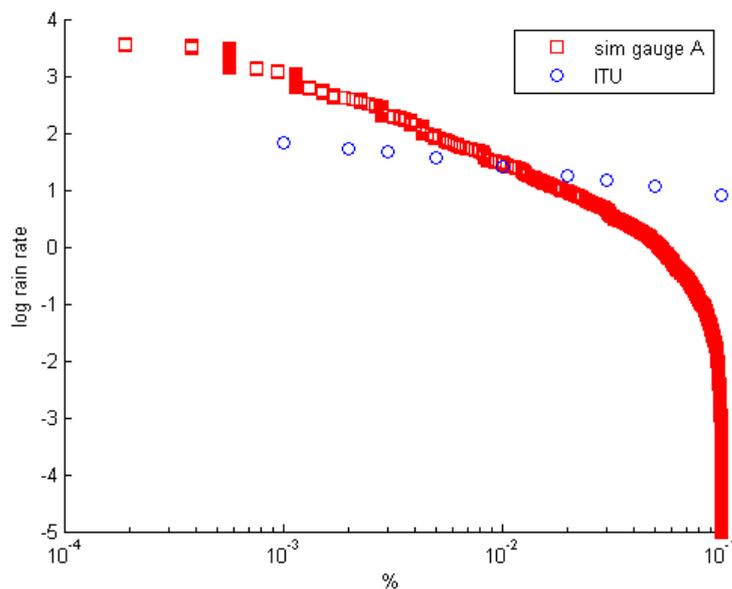


Figure B.1: Comparison of ITU Rec. P.837-4 curve with the complete curve from a simulated rain gauge taken from the reduced (526 arrays) simulated dataset.

As can be seen from Figure B.1, it's not practicable to fit this dataset to the ITU values in the range 0.1% to 0.001%, as this is trying to fit a whole distribution into the tail of another. Hence half the simulated curve is selected and scaled—i.e. the part of the simulated curve with log rain rate values greater than the mean. Figure B.2 shows this comparison, where the simulated gauge values plotted are still the original values as extracted from the simulated rain fields, i.e. there is no new scaling applied.

Selecting half the simulated curve allows a fit it to the ITU percentage values less than 0.05%.

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

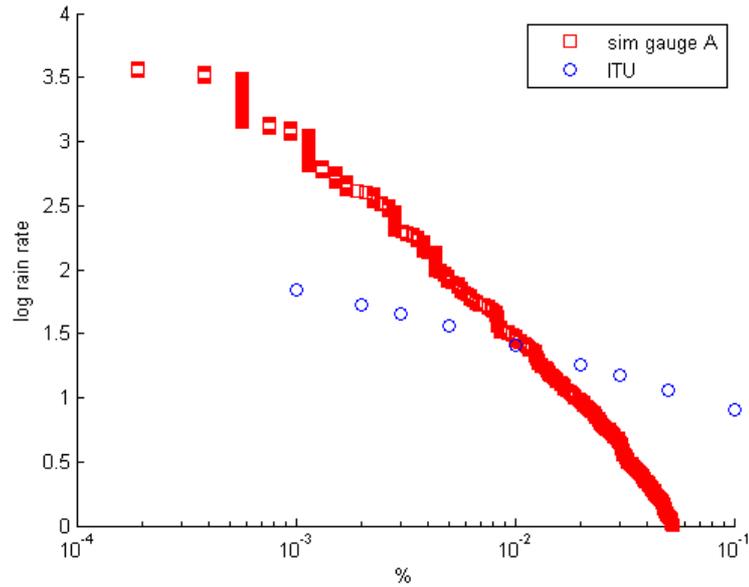


Figure B.2: Comparison of ITU Rec. P.837-4 curve with the part of the simulated curve with log rain rate values greater than the mean.

The parameter a modifies the standard deviation of the simulated distribution to fit in with that of the ITU (tail of the) distribution, and to give the simulated curve a similar slope to the ITU distribution. Figure B.3 shows the results of modifying the simulated data with the parameter a . The simulated curve and ITU curve are a better fit, shape-wise, but the simulated curve now needs to be offset to match the ITU curve.

Hence, the parameter b is introduced, which governs this offset.

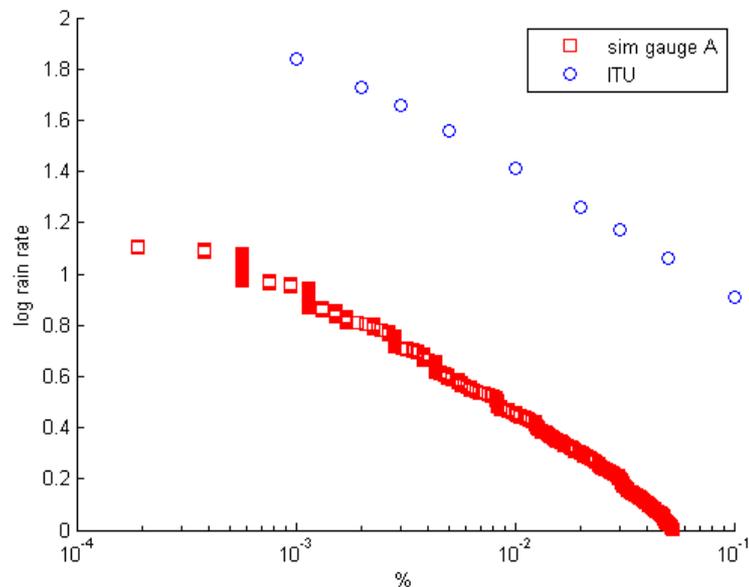


Figure B.3: Comparison of ITU Rec. P.837-4 curve with the section of the simulated curve after being scaled by a .

Appendix C: Implementation test of rain model.

The report has described the generation of a database of rain fields which can be used to expose fixed link plans to simulated annualised rain. An important step in gaining confidence in the annualised outage percentages obtained from runs of the fixed link planning simulator is to show that the simulator does indeed produce outages at the planned rate for a plan constructed under standard planning assumptions (i.e. no novel technologies).

An exercise has therefore been conducted to examine the performance of the simulator. Firstly, the simulator was tested with a series of uniform rain fields to check whether rain attenuation is calculated the same way in the outage assessment software as in the planning software (which uses library software supplied by Ofcom). Secondly, a standard plan was constructed and the outage rate measured in response to a sample of annualised rain fields.

Comparison of propagation models

A series of 41 uniform rain fields was created, in which the rain rate varied from 10 mm/hr to 50 mm/hr in steps of 1 mm/hr. A standard plan was then analysed using each of the rain fields and the number of outages counted. The outage rate should then show an increase as the rain rate increases, until all the links are in outage.

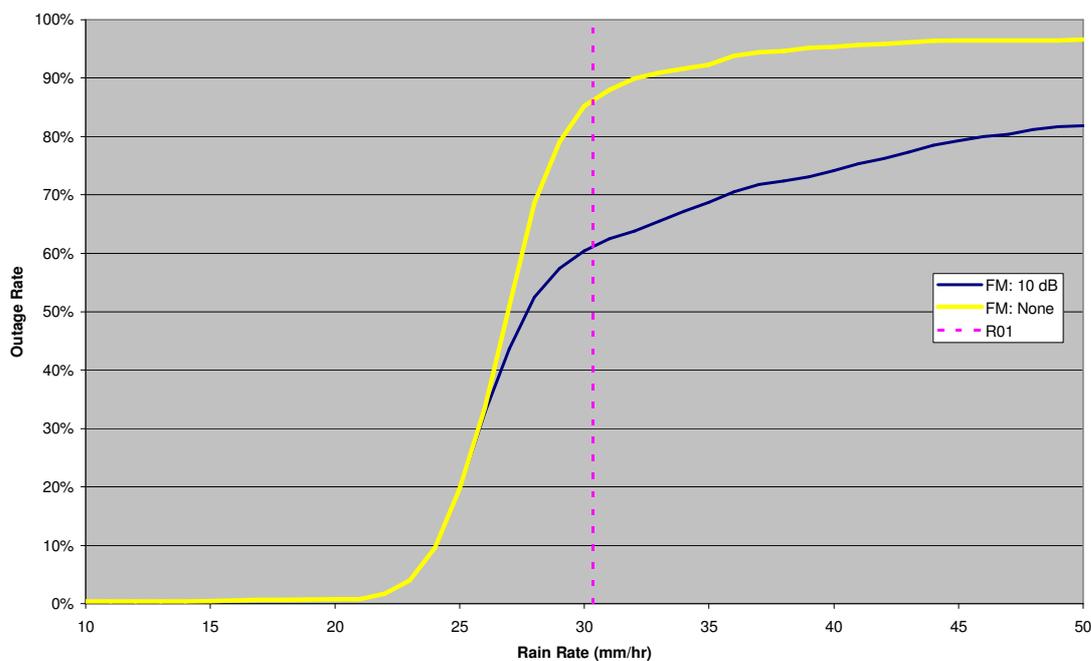


Figure C.1: Outage rate in a uniform rain field for a standard fixed link plan.

The figure above shows the result of a series of simulations in which the rain rate of a uniform rain field (i.e. one in which each pixel has the same rain rate) is varied. The figure is annotated with the 0.01% rain rate for the location of the centre of the simulated area (i.e. 30.3604 mm/hr at TQ250850). Curves are plotted for simulations with and without a minimum fade margin. The significant difference between these curves shows that there are substantial distortions in outage rates caused by planning and other assumptions.

Some assumptions are obvious. The minimum fade margin policy allocates excessive fade margin to some links, allowing those links to continue operating well beyond the point at which the increasing rain rate passes the rain rate appropriate to the required availability (mostly

Soft Boundary Frequency Assignment Techniques and Adaptive Coding and Modulation: Final Report

99.99%). Eliminating the minimum fade margin policy then produces a curve in which most links have indeed failed as the 0.01% rain rate is reached. The second, obvious, distortion is that the real 38 GHz plan includes links with a range of required availabilities: links with a higher than 99.99% availability will continue to operate past the 0.01% rain rate.

There are other assumptions:

- Rain rate varies across the UK and the planning software uses a rain rate that varies with position. For links to fail at the 0.01% rain rate, the simulated rain field at each link must match the planned rain rate at that link; the simplest way to achieve this is to use a single rain rate for planning (i.e. the rain rate at the centre of the test area).
- The determination of fade margin splits the required unavailability between clear air and rain effects. For links to fail at the 0.01% rain rate, all the unavailability must be assigned to rain attenuation.
- The amount of rain attenuation considered in calculating the fade margin is based on the rain rate at the centre of the path and an effective path length. For links to fail at the 0.01% rain rate, the link fade margin must be calculated assuming that the entire path is subject to fading at the same rate.

Corrections to the planning software were implemented in a test version to eliminate the biases caused by these assumptions: the results are shown in Figure C.2, where it is apparent that all links now fail as expected at the planned rain rate. This test therefore establishes that the fixed link simulator and the propagation library have compatible implementations of rain attenuation—it is then possible to continue with a more interesting test in which a standard plan is confronted with a set of annualised rain fields.

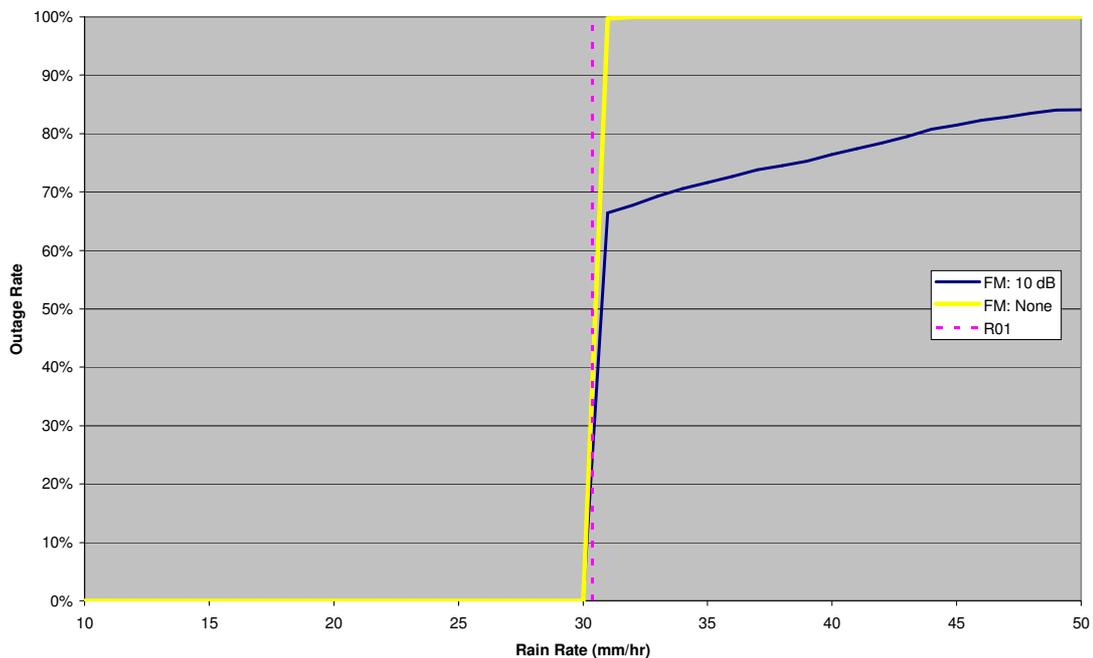


Figure C.2: Outage rate in a uniform rain field for a standard fixed link plan (with corrections).

Outage rate in a standard plan exposed to annualised rain

It is now apparent that a user of the database of annualised rain fields must interpret the numbers of induced outages with care. Insofar as the database accurately represents annual rainfall then simulated outage rates will indeed be annual rain-induced outage rates; however, outages will not generally occur at exactly the same rate as the planned unavailability.

A second test has been conducted to establish what outage rates might be simulated when a standard fixed link plan is exposed to simulated annualised rain fields. As with the previous test, biases were progressively removed in order to determine the degree of difference between actual outage rates and planned unavailability.

Bias Removed	Outage Rate (annual)	Comment
<i>Standard Plan</i>	0.00334%	This simulation reports the annual outage rate from a mix of services subject to standard planning assumptions (including minimum fade margin).
Availability	0.00329%	All services are now planned to have a 0.01% unavailability. Relative to the standard plan, links with higher availability requirements will now have a lower availability than previously and therefore a lower fade margin.
Rain Rate	0.00573%	All services are now planned with the same 0.01% rain rate as is assumed in scaling the simulated rain fields. This results in a substantial increase in outages because the simulated rain was scaled using a lower rain rate than that used in planning ⁸ —so the simulated rain was relatively less potent.
Fade Margin	0.00798%	The minimum fade margin policy has been removed. The removal of ‘excessive’ fade margin increases the outage rate.
Clear Air	0.00845%	No unavailability is allocated to clear air fading. The outage rate should now reflect the planned rate as closely as possible.

Table C.1: Outage rate in a simulated rain field for a standard fixed link plan (with corrections).

The results show that a plan constructed with the objective of achieving a 0.01% unavailability has, when exposed to simulated annualised rain, a measured unavailability of 0.008%. This is close enough to demonstrate the general method. However, it is probable that further improvements to the simulated rain (or additional data sets) might improve the ability of that rain to provoke the expected unavailability. Any user of the data should also be aware that the various biases in the planning process mean that link outages may occur at only half the planned rate. In other words, the planning process at 38 GHz is, in this limited respect, conservative.

⁸ The rain rate was taken from Rec. 837-4.

Rain rate distributions applicable to fixed link planning in the United Kingdom

A consideration in the use of the simulated rain is that a single 0.01% rain rate is required for scaling, whereas in practice the planned rain rate varies with position, as shown below. Moreover, the UK fixed link planning software uses a rain rate database that differs in some areas from Rec. 837-4.

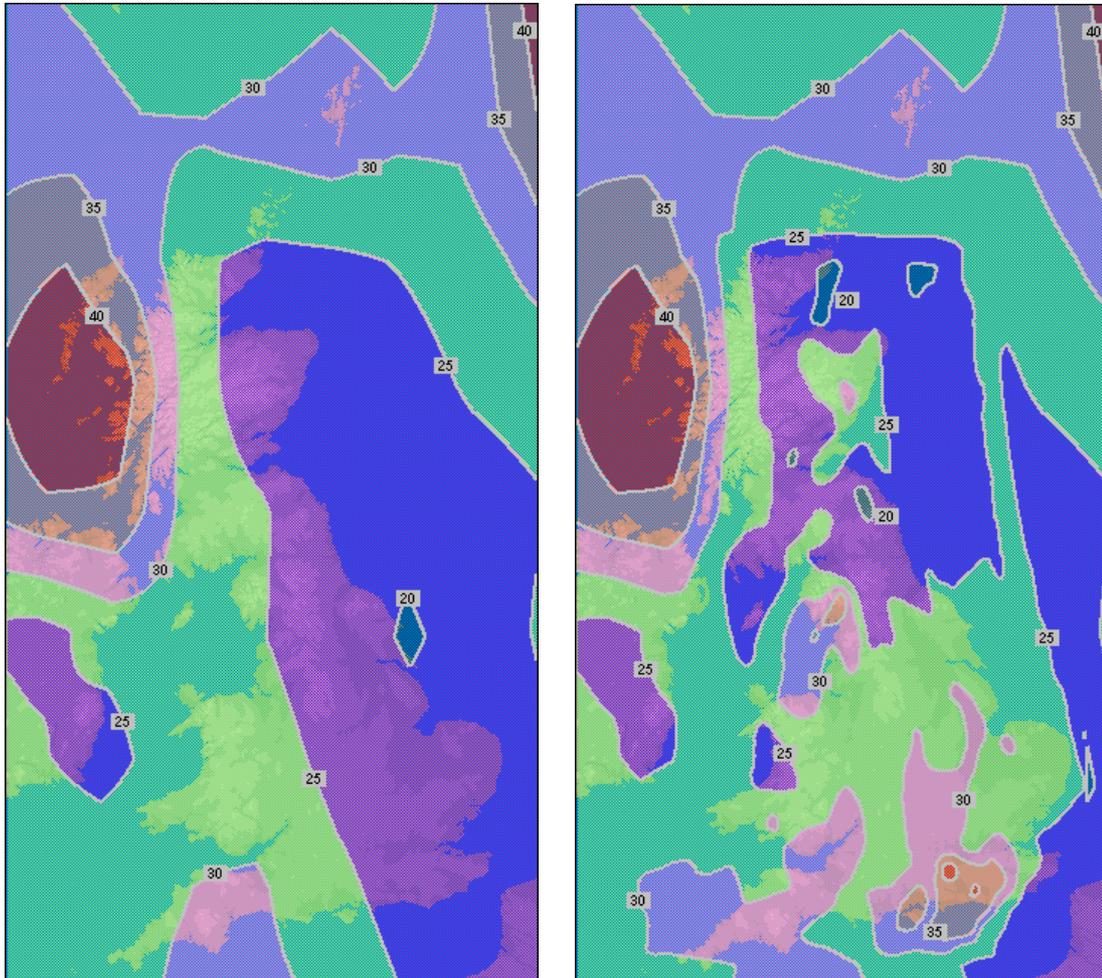


Figure C.3: Rain rate contours from ITU-R and UK databases.