A Study for the Provision of Aggregation of Frequency to Provide Wider Bandwidth Services

Final Report

Dr Anil Shukla, Brian Willamson, John Burns, Eddie Burbidge,
Alan Taylor, David Robinson

QINETIQ/06/01773
August 2006

Requests for wider use or release must be sought from:
Intellectual Property Division
QinetiQ Ltd
Cody Technology Park
Farnborough
Hampshire
GU14 0LX

Administration page

Customer Information	
Customer reference number	83000127
Project title	Provision of Aggregation of Frequency to Provide Bandwidth Services
Customer Organisation	Ofcom
Customer contact	Ahmad Atefi
Contract number	83000127
Milestone number	Final Report
Date due	14th September 2006

Principal authors	
Dr Anil Shukla	+44 (0)1684 896461
QinetiQ Ltd, Room PB102, St Andrews Road, Malvern, UK.	akshukla@QinetiQ.com
John Burns	+44 (0)1932 860074
Aegis Systems Ltd, 30 Anyards Road, Cobham, Surrey UK.	john.burns@aegis-systems.co.uk
Brian Williamson	+44 (0) 207 3241807
Indepen Consulting Ltd, Diespeker Wharf, 38 Graham Street, London	brianwilliamson@indepen.co.uk
Edward Burbidge	eaburbidge@qinetiq.com
QinetiQ Ltd, Room PB315, St Andrews Road, Malvern, UK.	+44 (0)1684 895023
Alan Taylor	Atalylor2@qinetiq.com
QinetiQ Ltd, Room PB101, St Andrews Road, Malvern, UK.	+44 (0)1684 89488
Dr David Robinson	drobinson@qinetiq.com
QinetiQ Ltd, Room PC222, St Andrews Road, Malvern, UK.	+44 (0)1684 895908

Record of changes

Issue	Date	Detail of Changes
0.4a	July	Version discussed at Ofcom and notes made
0.5b	Aug	Version for internal comments
1.0	Sept	Delivered to Ofcom for comments
1.1	Oct	Minor changes incorporated after comments

Executive Summary

Introduction

This study has been commissioned by Ofcom to ensure that new methods of spectrum management do not result in the inefficient use of the spectrum. Ofcom's concern is that these new methods could result in spectrum fragmentation and thus lead to inefficiencies. This study, therefore, examines key issues such as: has fragmentation occurred and is it likely to occur in the future? What technical solutions could be used to overcome fragmentation? And what powers should Ofcom use to minimise any detrimental effects of fragmentation.

The report is structured with an Executive Summary, a main report body and a set of Appendices. The Executive Summary poses a set of key questions, reports the key findings and concludes with a number of Recommendations. The main body of the report contains the theoretical studies. The Appendices comprise supporting tables and calculations.

Study Background

The traditional 'command and control' approach to spectrum management is increasingly perceived as being economically inefficient and linked with restricting technical innovation. This has led Ofcom to improve efficiencies by subjecting the spectrum to increased market forces. Ofcom's strategic plan, therefore, is to move away from command and control techniques, currently applied to ~94% of the spectrum, to a 72% market based approach by 2010. As the control and management of spectrum moves towards being more market led, there are concerns that spectrum fragmentation could occur: through the adoption of more spectrally efficient technologies; through a combination of ill-defined usage rights and multiple trades; or as a product of the regulatory environment.

Current wireless communications systems are spectrum stove-pipes and designed to use spectrum from single dedicated bands. Future cognitive radio (CR) systems driven by software defined radios (SDR), may seek to utilise spectrum horizontally by sequentially and opportunistically hopping in and out of unused spectrum gaps. The CR operation, however, would be fundamentally different to a device that aggregates spectrum. An aggregating device would aim to exploit multiple, small spectrum fragments simultaneously to deliver a wider band service; i.e. a service not otherwise achievable using a single spectrum fragment. The ability to aggregate spectrum, therefore, could increase spectrum utilisation and potentially make available wider bandwidths for new communications services.

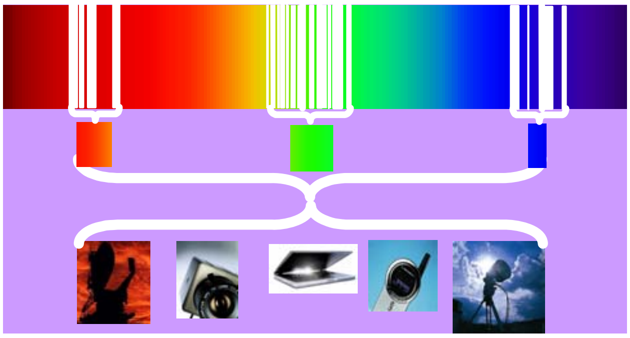

The concept of spectrum aggregation is to exploit spectrum fragments simultaneously to create wider bandwidths for communications systems.

Study Issues and Approach

This study investigates the issues surrounding potential spectrum fragmentation and aggregation by answering the following key questions:

1. *To what extent are there exploitable spectrum fragments?*
2. *What could cause fragmentation?*
3. *What powers should a regulator retain/use if fragmentation occurred?*
4. *What are the technical issues associated with using fragmented spectrum?*
5. *What is the "value" of fragmented spectrum, compared to non-fragmented?*
6. *Is there a minimum spectrum size beyond which the value of spectrum falls abruptly?*
7. *If spectrum becomes fragmented, is it a problem for a regulator?*

To answer these questions a team of RF technologists, spectrum engineers, and management economists reviewed data and performed theoretical studies.

The approach taken was to first conduct a spectrum review to identify any spectrum fragments that could be exploited by a new spectrum-aggregating device. The possible causes of fragmentation and the optimal extent of fragmentation versus aggregation, policy and retained regulatory powers were then investigated. Potential spectrum aggregating architectures and the technological constraints were then explored and hypothetical spectrum aggregating systems were designed. The hardware costs from these systems were then compared with conventional (non-aggregating) system architectures to investigate the economic aspects of developing the same service using a spectrum aggregating approach. The concept of using virtual aggregation for multiple services was also examined.

Key findings

1 – To what extent are there exploitable spectrum fragments?

There were very few exploitable fragments identified, but the unassigned Private Mobile Radio (PMR) frequencies could potentially be used by an aggregating device.

See Section 2

To be considered a spectrum fragment, exploitable by an aggregating device, fragments were required to be unallocated to a specific service, designated as guard-band spectrum or allocated to a service but not currently assigned to any user. A detailed analysis of current spectrum use in the range 100 MHz to 5 GHz was conducted by collating and interpreting a number of separate sources of data to produce frequency charts detailing users in particular bands.

The data sources used to derive the data included: the UK national frequency allocation table; Ofcom's published channel plans for services such as PMR and fixed links; and other Ofcom documents, in particular those relating to the spectrum framework review and dialogue with Ofcom's Business Radio Licensing team (which is responsible for managing PMR spectrum).

Data was also provided by Ofcom on current PMR (137 – 461MHz) assignments at four UK locations chosen to reflect different intensities of spectrum usage within 100 km of the designated location. At each location we identified unassigned spectrum that could be “available” to a spectrum aggregating device. The spectrum investigations identified and highlighted a few potential fragments. Although guard bands were included within the original scope of the study, following discussion with Ofcom licensing representatives it was concluded that they were not suitable for spectrum aggregation as they are effectively “owned by” the system they are protecting.

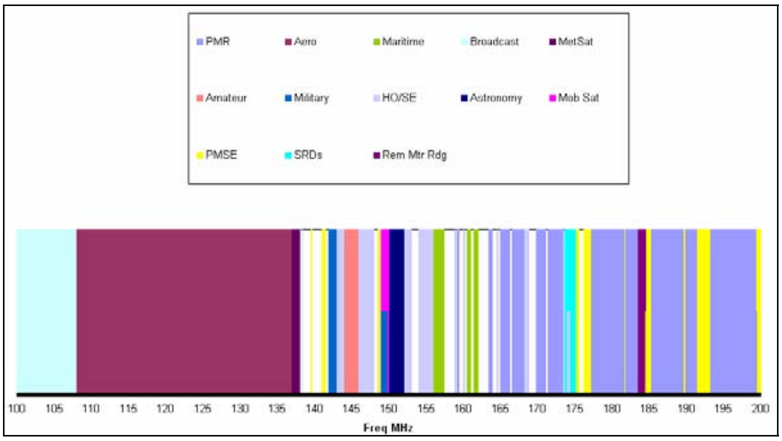

Example of a frequency chart produced showing spectrum usage between 100MHz to 200MHz

The smallest size of spectrum fragment considered suitable for aggregation was determined to be 25kHz, on the basis that smaller fragments would require excessive filtering and/or guard spectrum to protect adjacent spectrum. Hence two or more adjacent unused 12.5 kHz PMR channels could be considered as exploitable fragments. Analysis showed that the number and distribution of these fragments varied spatially. For example, London had 2.65 MHz of total unassigned spectrum fragments compared with 17.275 MHz for Ullapool. A total of 325 kHz of PMR spectrum was identified as available nationally. Our analysis also showed a number of other currently unallocated bands at frequencies above 862 MHz which appeared to be available on a national basis.

The analysis concludes that currently there is little fragmentation across the spectrum. A few fragments, other than guard bands were found. Unassigned frequencies in the PMR band could be considered as small spectrum fragments exploitable by an aggregating device.

One of the recommendations associated with this question is that Ofcom could significantly simplify the process of spectrum analysis to identify not only spectrum fragments but also under-utilised spectrum by providing and maintaining an effective spectrum database.

2 - What could cause fragmentation?

Spectrum markets, administration or the exploitation of spectrum saving technologies could be sources of fragmentation. There are no clear grounds, however, for thinking that spectrum markets will lead to excessive fragmentation. Rather the opportunity to trade, exploiting new technologies, could be expected to promote optimal levels of aggregation making use of any fragments created.

See Section 3

Trading could lead to further spectrum fragmentation and/or spectrum aggregation, depending on the economics of using spectrum fragments versus other spectrum.

Administrative decisions in relation to spectrum planning can result in large or small blocks of spectrum being allocated to different uses. Unused, unallocated or inefficient "fragments" of spectrum could then arise from a series of administrative decisions. For example, the original spectrum assigned to Cellnet (now O2) and Vodafone for analogue cellular telephones was augmented over the years and resulted in considerable interleaving between the two operators' assignments. Until these assignments were rationalised, this fragmentation had an adverse impact on the efficiency with which operators could utilise their spectrum assignments[1].

Fragments of spectrum may be allocated to a service but remain unused (e.g. the PMR frequencies identified in the previous section). This, however, would not necessarily imply that the allocation was sub-optimal, or that fragmentation was increasing as the fragments were not being used. The most efficient use of some spectrum at a given point in time may be to hold it in reserve for a prospective use. Idle spectrum, therefore, does not automatically imply inefficient use (or fragmentation). It may simply reflect the fact that the necessary technology or infrastructure is still under development, or that the optimal use is uncertain, or that the best current option is to wait to see how alternative prospective uses develop.

Spectrum fragments or unused spectrum may also involve sunk or irreversible investment either in terms of associated capital investment, or in relation to the acquisition of spectrum which is unlikely to involve transaction costs. It may, therefore, be efficient to hold on to unused spectrum until uncertainty is sufficiently resolved to allow a decision about whether to proceed with investment in a new service, or to dispose of the unneeded spectrum.

In other markets, such as the land market, fragmentation and aggregation is common. For example, a property developer may aggregate land to sell housing in lots. The non-use of land may also relate to the planning process whereby potential competing uses of land and questions over external impacts are resolved, in part, administratively: a process that can take some time before a final decision is made. For example, the planning decision over Terminal 5 at Heathrow took approximately a decade to resolve.

Spectrum fragments and the non-use of blocks of spectrum including fragments may reflect economic fundamentals and is not *per se* a cause for concern. The underlying circumstances of fragmentation must be examined to see if there is an issue in relation to spectrum fragments that raises public policy concerns.

---

[1] Radiocommunications Agency. 1998. "Report on Modifiers to be used in determining Administrative Pricing Fee Charges for Mobile Services." http://www.ofcom.org.uk/static/archive/ra/topics/spectrum-price/spec-pric/1998/report.htm

3 - What powers should a regulator retain/use if fragmentation occurred?

Ofcom should only retain powers to re-aggregate spectrum if the benefits, in terms of more economically efficient spectrum use, exceed the cost in terms of reduced regulatory certainty and the knock-on effects on licensee behaviour. We note that in other markets such powers do not in general exist, which suggests they may be unnecessary. In addition, whilst various forms of "market failure" can arise they would not necessarily lead to fragmentation, in which case powers to re-aggregate spectrum would not be an appropriate remedy.

If Ofcom retains powers to aggregate, licensees have weaker incentives to fragment the spectrum themselves (for good market reasons) in case Ofcom intervenes and changes their licence rights (Ofcom has proposed a five year notice period for changes to licences on spectrum management grounds.)

Ofcom already has powers to block trades for spectrum management reasons and we see no reason at this time to increase or to change these to deal with fragmentation of the spectrum.

4 - What are the technical issues associated with using fragmented spectrum?

The single most significant challenge to realising a spectrum aggregator (within-band) is the mitigation of intermodulation distortion (IMD), especially if fragments share a transmit/receiver chain, or chains need combining to share an antenna or amplifier, to reduce component count and overall size. ADCs dictate the exploitable fragment width and the lowest frequency of the fragment is determined by filter technology. Aggregating over separate bands raises the additional significant issue of antennas.

See Section 4

Broadly, two aggregation design options are available to the engineer; a receiver chain per spectrum fragment provided that only a few fragments are to be aggregated, or a single wideband receiver for many fragments. The former is achievable using narrow band technologies but increasing component count may be a problem as the number of fragments increases. The latter, although more elegant, is more difficult due to technological limitations of wideband components, antenna sharing and the challenge of managing intermodulation products. Based

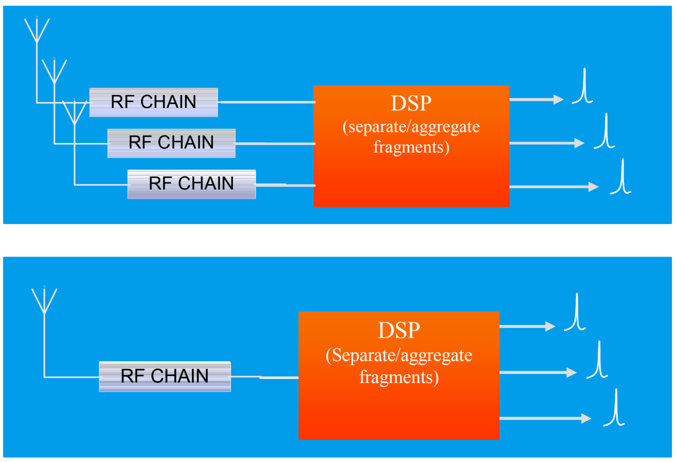

Design options available: a) single RF chains for a few fragments with a single tuneable RF chain for each fragment, and (b) wideband RF chain for many fragments on RF hardware design criteria, narrow band technologies aggregating fewer than five fragments are predicted to be the most effective solution for an aggregator being built today.

The single most significant challenge to realising a spectrum aggregator is the mitigation of intermodulation distortion (IMD), especially if fragments share a transmit/receiver chain, or chains need combining to share an antenna or amplifier, to reduce component count and overall size. Amplifiers and mixers are evolving to perform more linearly to combat IMD. Unfortunately there is little evidence to suggest combiner/divider technology is developing in the same manner. Instead, architectures employing combiners or dividers will have to rely on system linearisation techniques, which at present are immature. There is a drive, however to develop such linearisation techniques to improve RF efficiency.

There is a long term IMD solution based on ADCs, DACs and digital processing to dispense with the need to use any analogue frequency conversion. This will enable fragments and chains to be combined digitally without distortion.

Using today's hardware technology, it is possible to aggregate fragments over a limited number of bands, each band being at most 50MHz wide. The centre frequencies of these bands can be anything from a 100MHz up to 1.5GHz, and it is possible to have tuneable bands in a single aggregating device.

There is no fundamental technical limit on the minimum width of fragments, within the limits of 25kHz minimum and an arbitrary 1MHz maximum. If a spectrum aggregating device was built today, it would need to implement state-of-the-art DSP (Digital Signal Processing), ADC (analogue to digital conversion), DAC (digital to analogue conversion) and amplifier components. Using DSP automatically adds a cost premium.

In the scenarios analysed, the increased RF hardware costs for a 2-fragment solution ranged between 70% and 600% depending on the technology and type of service involved. The cost increase, however, is dominated by DSP costs, particularly when otherwise low cost RF hardware components are used.

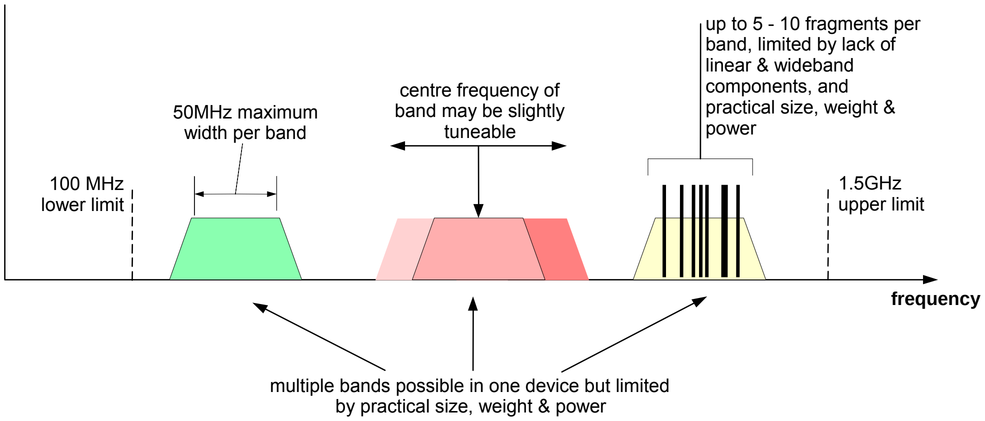

Summary of the capabilities of a spectrum aggregator built today

Typically the RF Research & Development (R&D) hardware costs represent a small portion of the overall development and life cycle budget of a radio service. In most cases, therefore it is anticipated that the RF R&D hardware costs alone are unlikely to hinder the development of new business models exploiting fragments. It is anticipated, however, that similar-performing systems built in five or ten years' time would cost significantly less due to reducing technology costs of the additional componants such as multiway splitters and DSPs.

It should be noted that the way in which fragments are distributed within a single band (close to the centre frequency) and the size of bands or fragments, does not affect the costs significantly. An in-band (i.e. using fragments near the centre frequency) spectrum aggregator could be built today.

If building aggregators to operate over separate bands, three significant issues need addressing:

- large, bulky antennas at low/widely separated frequencies
- band-limiting speeds of ADCs, DACs and DSP
- intermodulation distortion resulting in poor linearity

While these issues may not be immediately resolved, technological developments and trends underway will help to mitigate and eventually eliminate them.

5 - What is the value of fragmented spectrum, compared to non-fragmented?

The value of radio spectrum (fragmented or otherwise) is very dependent on how the spectrum can be used. There is some evidence that non-fragmented spectrum is valued more than smaller spectrum segments. This leads to the hypothesis that aggregation, rather than fragmentation will be the trend.

See Sections 3, 5 and 6

The “value” of spectrum fragments could be interpreted by considering the auction price of the spectrum as a function of spectrum size. This relationship, derived from GSM spectrum auctions in Austria and the Netherlands, suggest that bidders value larger blocks of spectrum more than smaller blocks (on a per-MHz basis), but it should be noted that other factors such as timing also have an impact on the value. The data support the hypothesis that there will be a market trend to aggregate rather than fragment spectrum because larger spectrum fragments are valued more than smaller fragments.

It should be noted that the monetary value of smaller segments may be lower if the spectrum is being purchased by an operator who already has a large amount of spectrum. Smaller segments, therefore, may be of more value to a provider who has no spectrum.

Fragments of spectrum might at times be allocated but unused, but that would not necessarily imply that the allocation was sub-optimal (or of little value). In a spectrum trading environment, spectrum might be fragmented or aggregated over time in response to changes in economic value of different configurations. For example, Nextel in the US aggregated spectrum in order to launch a mobile telephony service.

As discussed in other sections, radio equipment that can use fragmented bands is feasible and its potential capability and economics can be expected to improve over time. This will make it easier to use fragments, which in turn implies that fragments will become relatively more valuable. The optimal amount of fragmentation could therefore increase as the economics of using fragments improves. Alternatively, if new blocks of spectrum are released into the market, the relative value of spectrum fragments may fall.

To further investigate the value of fragmented versus un-fragmented spectrum services, the hardware costs involved for satellite, public mobile radio, wireless local area network and a microwave fixed link system have been compared. The hardware costs of an (in-band) fragmented system are, as would be expected, higher than a single band conventional system. The cost equation is dominated by the requirement for a DSP. Naturally, as the number of fragments increases, the costs tends to increase. It is interesting to note that for a WLAN system, however, a step change increase in cost is observed between 6 and 7 fragments due to the cost of higher-way splitters.

Another factor that may affect the value of fragmented spectrum is the consideration of the management overhead associated with frequency planning. Planning will be required between users either side of the spectrum fragments and it will be a disincentive to exploit fragments.

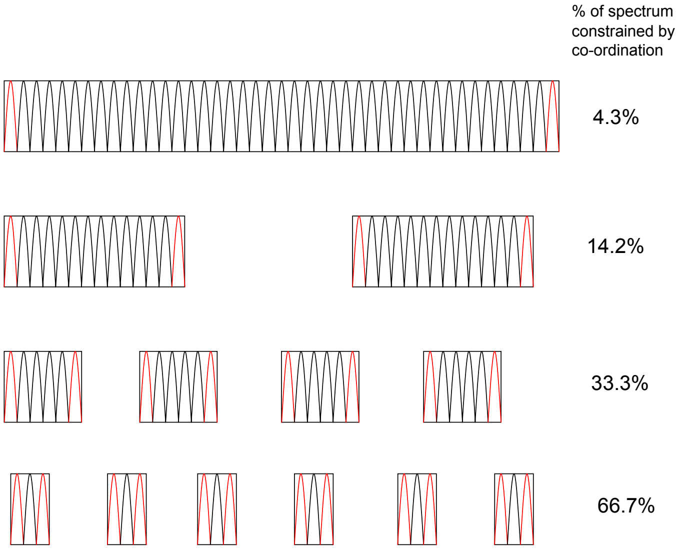

Illustration of how fragmentation could increase co-ordination requirements

A supplementary value question is: *had there been no contiguous spectrum would a service utilising fragments of spectrum been commercially viable*. To answer this the hardware costs derived in the study are further analysed as part of a larger investment scenario as the design/concept cost is only one small part of the total cost which must be passed on to potential consumers. By considering these other costs, we illustrate how a system using aggregated spectrum compares to a non-aggregating system and whether such an aggregating system would have been developed (assuming contiguous spectrum was not available). For the 2-fragment Satcom, Fixed Link and WLAN scenarios the NPV is positive within 5-years and these solutions may warrant a more detailed investment case analysis. Further analysis of the Satcom scenarios showed that by reducing any spectrum costs by 14% and 28% for the 2-fragment and dual band 4-fragment aggregating devices respectively, a similar business outcome could be achieved to a non-aggregating solution.

6 - Is there a minimum spectrum size beyond which the value of spectrum falls abruptly?

It is difficult to say at what point the spectrum value falls off abruptly as the true value of fragmented spectrum will largely depend on how it is to be used, the type of service that can be developed and frequency. Small spectrum fragments may be just as valuable as large segments (e.g. guard bands).

See Sections 3, 5 and 6

An analysis of spectrum value can lead to some surprising results. For example, it might appear intuitive that TV broadcast spectrum would be worth more than FM broadcast spectrum, which in turn would be worth more than AM broadcast spectrum, reflecting the relative revenues generated from these services. However this overlooks the amount of spectrum required to deliver the service to a large audience and the costs associated with delivery. A national analogue TV service requires up to eleven 8 MHz radio channels, i.e. a total bandwidth of 88 MHz, whereas a national FM service can be delivered with only 2 MHz and a national AM service can be delivered with only 18 kHz.

These differences are reflected in the opportunity cost estimates developed recently[2], which estimated the value of TV spectrum at £1.12M per MHz, FM radio spectrum at £1.8M per MHz and AM spectrum at £90M per MHz

To highlight the value aspect further it is interesting to note that auction prices paid for 12 licences, for 6.6MHz at 1781.7 – 1785MHZ and 1876.7 – 1880MHz, ranged from ~£50k to over £1M. This shows that separate groups of people value the same spectrum differently, which may be influenced by the amount of spectrum they currently hold, the sensitivity of their model to spectrum access etc.

7 - If spectrum becomes fragmented, is it a problem for a regulator?

Our work has concluded that spectrum fragmentation is not currently a significant problem. It is unlikely to become one in the future and our findings imply that a liberalised trading environment may encourage spectrum aggregation rather than spectrum fragmentation.

Ofcom also already has powers to block trades for spectrum management reasons. This could in principle include blocking certain trades if they are thought likely to lead to “excessive” fragmentation of the spectrum.

Spectrum aggregation seems technically feasible. The overall cost of developing a system which is capable of aggregating spectrum is unlikely to be significantly different from that of designing any other new radio system.

If spectrum becomes fragmented, Ofcom needs to allow the opportunity for spectrum users to exploit the fragments. If the opportunity exists then some manufacturers are likely to develop and market spectrum-aggregating-devices.

[2] “An economic study to review spectrum pricing”, prepared for Ofcom by Indepen, Aegis Systems and Warwick Business School, February 2004

Conclusions

Our analysis has shown that the spectrum utilisation gains from the use of spectrum aggregation techniques and technologies are currently relatively small and economic drivers are insufficient for this situation to change in the short term.

The technology to enable the aggregation of spectrum fragments is largely available. Development costs are, mostly, not anticipated to be different from any other kind of new radio design, but will require more RF components.

In summary, spectrum fragmentation is not currently a widespread problem and it is considered unlikely that the situation will become worse in a liberalised trading environment. Should fragmentation occur then technology may be brought to bear to provide effective communications services through use of spectrum fragments, and market economics will determine whether this is beneficial. Thus we conclude that there is no need to investigate any further regulatory actions at this time.

Recommendations

It is recommended that:

1. Ofcom considers the implications of making the spectrum fragments, and a number of apparently unallocated national-spectrum bands identified below nationally available.

Freq Range (MHz)	Bandwidth (kHz)
158.73125 – 158.80625	75
159.11875 – 159.14375	25
159.15625 – 159.19375	37.5
163.23125 – 163.30625	75
163.61875 – 163.64375	25
163.65625 – 163.69375	37.5
164.14375 – 1664.16875	25
454.41255 – 454.43755	25
Total	**325**

Spectrum fragments identified.

Frequency band	Bandwidth (MHz)
862 – 863 MHz	1 MHz
1375 – 1389 MHz	14 MHz
1399 – 1400 MHz	1 MHz
2290 – 2302 MHz	12 MHz
3440 – 3442 MHz	2 MHz
3475 – 3480 MHz	5MHz

Nationally available unallocated spectrum bands.

2. Spectrum data and users be made easily accessible. The exploitation of fragments (should they occur) will necessitate exchange of information about spectrum users adjacent to each fragment for coordination purposes. Data on spectrum users is not currently readily accessible, and there are many reasons in addition to coordination of fragmented spectrum, such as general spectrum trading, why this data should be available and managed in a centralised way.
3. If spectrum fragmentation occurs then technologies that drastically reduce the cost of their exploitation (e.g. DSPs, splitters) should be promoted.
4. The Cave Audit recommended a review of progress in relation to spectrum management and the development of spectrum trading in five years time. Such a review would provide an opportunity to examine developments in terms of spectrum fragmentation and aggregation, and test whether such developments raise public policy concerns. The impact of existing Ofcom powers to intervene in the market could also be examined to identify any possible impact on the development of a forward looking market which promotes efficiency in the short and long term.

List of contents

1 Introduction

This study was commissioned by Ofcom (under contract No. 83000127) to ensure that new methods of spectrum management do not result in the inefficient use of the spectrum. Ofcom's concern is that new management methods could result in the fragmentation of the spectrum and lead to spectrum inefficiencies. This study, therefore, examines the key issues regarding fragmentation by investigating if fragmentation has (or would) occur, what technical solutions could be used to overcome fragmentation and what powers Ofcom could use to minimise any detrimental effects of fragmentation.

Contributors to the report have comprised; QinetiQ Ltd (a Research Technology Organisation with expertise in RF communications technologies, propagation and resource trading), Aegis Ltd (an independent provider of specialist radio spectrum advice to users and regulators) and Indepen Ltd (a management consultancy specialising in economic and regulatory aspects).

1.1 The Need and Objectives

The traditional 'command and control' approach of spectrum licensing is seen as economically inefficient and is linked with restricting technical innovation. This has led Ofcom to look at new measures to improve efficiencies by subjecting the spectrum to increased market forces. Ofcom's strategic plan, therefore, is to move away from command and control, currently applied to ~94% of the spectrum, to a 72% market based approach by 2010.

As the control and management moves towards being more market led, spectrum fragmentation could occur through the adoption of more spectrally efficient technologies to deliver the same service; a combination of ill-defined usage rights and multiple trades; or as a product of the regulatory environment.

Current wireless communications systems are spectrum stove-pipes and designed to use spectrum from single available bands. Future cognitive radio (CR) systems driven by software defined radios, may seek to utilise spectrum horizontally by hopping in and out of unused spectrum gaps. The CR operation, however, would be fundamentally different to a device that aggregates spectrum. Such an aggregating device would aim to exploit multiple, small, spectrum fragments simultaneously to deliver a wider band service; i.e. a service not otherwise achievable using a single spectrum fragment. The ability to aggregate spectrum, therefore, could increase spectrum utilisation and potentially make available wider bandwidths for new communications services.

1.2 Objective

The objective of this study is to investigate the issues surrounding spectrum aggregation. This is achieved by answering the following key questions:

1. Are there any exploitable spectrum fragments?
2. What could cause fragmentation?
3. What powers should a regulator retain/use if fragmentation occurred?
4. What are the technical issues associated with using fragmented spectrum?
5. What is the "value" of fragmented spectrum, compared to non-fragmented?

6. Is there a minimum spectrum size beyond which the value of spectrum falls abruptly?
7. If spectrum becomes fragmented, is it a problem for a regulator?

1.3 Methodology

To study and answer the above questions a team of RF technologists, spectrum engineers, and management economists reviewed available data and performed theoretical studies.

The approach taken was to first conduct a spectrum review to identify any spectrum fragments that could be exploited by a new spectrum-aggregating device. The possible causes of fragmentation and the optimal extent of fragmentation versus aggregation, policy and retained regulatory powers were then investigated. Potential spectrum aggregating architectures and the technological constraints were then explored and hypothetical spectrum aggregating systems were designed. The hardware costs from these systems were then compared with conventional (non-aggregating) system architectures to investigate the economic aspects of developing the same service using a spectrum aggregating approach. The challenge of using the aggregated spectrum for multiple services (virtual aggregation) was also examined.

1.4 Structure of the Document

Section	Comment	Question
Executive Summary	Key study findings and recommendation designed for all readers.	All questions
Section 1 Introduction	Outlines the background, aims and report structure.	-
Section 2 Review of Fragmentation of Spectrum.	This reviews allocations between 100MHz and 5GHz and would of interest to spectrum managers.	1,7
Section 3 Policy and regulatory issues	The causes of fragmentation are discussed and would be of interest to spectrum economists.	2, 3, 5,6
Section 4 A Technical Solution for Using Fragmented Spectrum	The technical challenges of aggregation devices are detailed and would be of interest to radio technologists.	4
Section 5 Fragmented spectrum scenarios	Four hardware designs are detailed and costs estimated as the number of fragments increase. The section would be of interest to RF engineers.	5, 6
Section 6 Investment Impact Analysis	The investment case for developing an aggregating service is analysed and would be of interest to business modellers.	5, 6
Section 7 Virtual aggregation solution	This section looks at how multiple services may use fragments of spectrum to increase its value.	5
Appendices A, B, C, D, E and F	Provide equations and tables related to the sections above.	-

Summary sections, aimed at the general reader are included throughout the document.

2 Review of the Fragmentation of the Spectrum

2.1 Introduction

A short review was conducted to highlight fragmented spectrum that could be exploited by a spectrum aggregating system.

In identifying blocks of spectrum fragments, it was decided that a minimum contiguous bandwidth of 25 kHz (i.e. equivalent to two analogue PMR channels) would specify the minimum bandwidth required. No specific maximum bandwidth was chosen. Bands, however, that are currently identified in Ofcom's Spectrum Framework Implementation Plan for future release were excluded. To be considered a spectrum fragment (for spectrum aggregation), spectrum was required to be either unallocated, designated as guard band spectrum or allocated to a service but not currently assigned to any user. These three categories are explained in more detail below.

Unallocated spectrum: This means spectrum that is not currently allocated to some specific use, such as cellular mobile, fixed links, private mobile radio, etc. Spectrum allocated for government, military or aeronautical use has not been considered, nor has spectrum identified for other applications where the rights to use the spectrum lie outside Ofcom's remit.

Guard Band Spectrum: The terms of reference for the study require consideration to be given to whether guard band spectrum could be included as part of the aggregation process. Our analysis revealed a number of specific guard bands that are designated to protect adjacent band services, but may be useable under certain conditions (e.g. indoors or with restricted power levels) and providing that other services are not already using them.

Unassigned spectrum: A number of frequency bands are allocated to Private Mobile Radio (PMR) services and are assigned to individual users in the form of 12.5 kHz channels, which may be assigned nationally or (more generally) on a local or regional basis. This means many of these channels may not be in use in particular parts of the country (especially away from London and other major conurbations). In some cases several channels are available, creating the potential for a sizeable contiguous block of spectrum to be utilised at that location. Other allocated uses are more problematic in that either the full band is used on a national basis (e.g. for cellular mobile) or the co-ordination distances and interference protection criteria are stricter than for PMR (e.g. fixed links, broadcasting), hence these are not considered suitable for aggregation.

The methodology is described below and followed by the summary of the findings. The detailed analysis in reported in Appendix A.

2.2 Approach to identifying available spectrum

The approach involved a comprehensive analysis of the current use of spectrum in the range 100 MHz to 5 GHz. This was based on a number of sources, including the UK national frequency allocation table, Ofcom's published channel plans for services such as PMR and fixed links, other Ofcom documents, in particular those relating to the spectrum framework review, and dialogue with Ofcom's Business Radio Licensing team.

Data were also provided by Ofcom on current PMR assignments at four locations in the UK. These were intended to represent different intensities of spectrum usage. In each case all frequencies in use within 100 km of the designated location were considered to be unavailable; other frequencies were assumed to be available for aggregation purposes.

The four locations chosen were:

i) Central London– dense urban environment

ii) Newcastle – urban environment

iii) Brough (North Yorkshire) – small town with nearby urban areas

iv) Ullapool (Scottish Highlands) – remote rural area

The detailed analysis (Appendix A) comprises charts (for example, Figure 2-1) to illustrate the allocation of spectrum between 100 MHz and 5 GHz to various services and users. For clarity, the spectrum has been divided into smaller sub-bands detailed in the Table below.

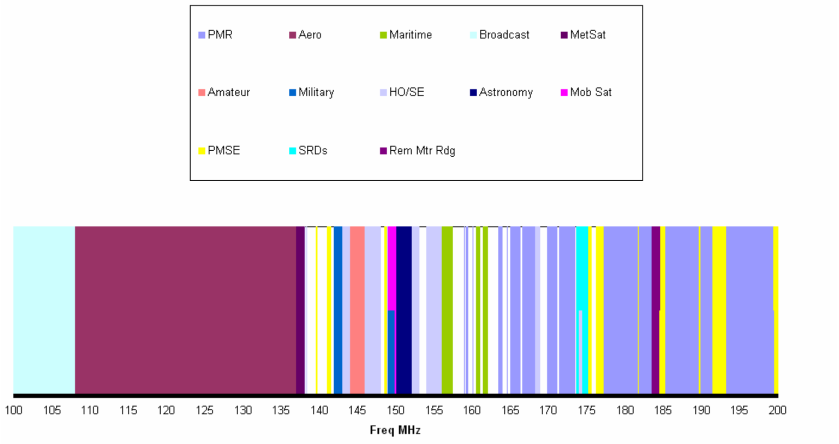

Figure 2-1: Frequency chart showing users between 100MHz to 200MHz

2.3 Availability of Fragmented spectrum for Aggregation Purposes

Based on the review of current allocations (Appendix A) the following spectrum has been identified as potentially available for a system that could aggregate spectrum. The following spectrum has been identified as potentially available on a national basis as it was (at the time of the study) not currently allocated to a specific use:

Frequency band	Bandwidth (MHz)
862 – 863 MHz	1 MHz
1375 – 1389 MHz	14 MHz
1399 – 1400 MHz	1 MHz
2290 – 2302 MHz	12 MHz
3440 – 3442 MHz	2 MHz
3475 – 3480 MHz	5MHz

Table 2-1: Frequencies not currently allocated to a specific use on a national basis.

The following guard bands have been identified nationally. The status (ownership) of guard bands varies as some may be included as part of the licensed system they are designed to protect. They may be useable subject to appropriate interference mitigation measures, and provided that they are not already being used by services that have identified them as available. The guard bands of interest are:

- 915 – 917 MHz (GSM cellular)
- 1350 – 1350.5 MHz (fixed links)
- 3600 – 3605 MHz (FWA)
- 3641 – 3650 MHz (FWA)
- 3875 – 3925 MHz (fixed links)
- 3961 – 3970 MHz (fixed links)
- 4195 – 4200 MHz (fixed links)

Analysis performed on unassigned PMR frequencies in four locations across the UK is summarised in Table 2-1 (and detailed in Appendix A). The results show that the potential aggregated bandwidth (exploitable by an aggregating system) varies based on location.

Location (centre)	Unassigned spectrum
London	2.65 MHz
Newcastle-upon-Tyne	11.275 MHz
Brough	4.925 MHz
Ullapool	17.275 MHz

Table 2-2: Total bandwidth available in PMR band for differing locations in the UK.

The frequencies below may be useable in some areas (subject to discussion with Ofcom):

- 606 – 614 MHz (in areas where not used for Radioastronomy)
- 1389 – 1399 MHz (used by low power video links on a non-interference, non-protected basis).

2.4 Impact of fragmentation on spectrum efficiency / utility

It should be noted that the number and size of spectrum fragments affects the utility of the spectrum. For each individual fragment to be useful, it must be used in such a way that it does not harm users operating on frequencies adjacent to each spectrum fragment. For a large block of spectrum, such as the current 2G or 3G mobile assignments, this does not present a significant problem, as it is only the extreme edges of the assigned blocks that are constrained by the need to protect adjacent users. In the case of GSM for example, this generally involves a need to co-ordinate the 200 kHz channels at either edge of the operator's assignment. In the case of wider band systems, such as UMTS, the typical approach is to allow an additional 200 – 400 kHz spacing between carriers used by different operators than would be required for carriers used by the same operator.

Clearly, for a given technology, the narrower the assigned block, the greater the impact this will have on the co-ordination effort required. In the figure below we illustrate this by deriving a figure for the percentage of spectrum constrained by co-ordination activities. Taken to extremes – the assignment of a single GSM carrier between two other operators' assignments – the spectrum is likely to be of very little worth indeed.

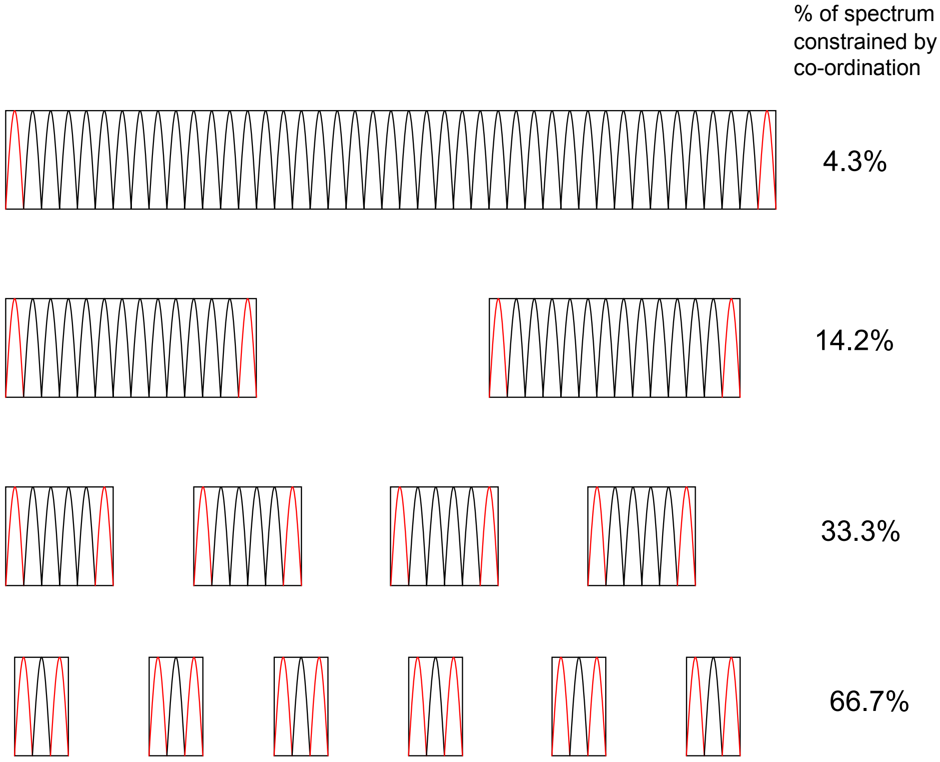

Figure 2-2: Illustration of impact of fragmentation on co-ordination requirements

Of course, it may still be possible to utilise very small spectrum fragments if the use of the adjacent spectrum is also based on narrow band channels – this would be the case in the PMR bands for example, but this will constrain the nature of the services that can be provided which in itself will tend to reduce the value of the spectrum. The adverse impact of spectrum fragmentation has been previously observed, for example in the context of cellular radio networks. The original (analogue) cellular networks in the UK were assigned spectrum on a progressive basis over several years and as a result there was substantial interleaving of the two operators' assignments, rather than two single, contiguous blocks. The impact of this was recognised at the time as introducing capacity limitations and planning and infrastructure logistical problems. This led to a proposal to reduce the fee applied to this spectrum under administrative pricing by 20%, pending the resolution of the fragmentation. [1]

In conclusion, it is probably the case that any fragmentation of spectrum below, say, 5 MHz will have an adverse impact on its value because of the more limited choice of technologies available (particularly for broadband applications) and the need for additional co-ordination. The minimum practical size of spectrum fragment will depend on the service and technology the user wishes to deploy and the way in which the adjacent spectrum is used, however based on typical guard band requirements for existing services such as GSM and TETRA a value of around 200 kHz would seem appropriate.

If fragments are to be exploited then the co-ordination effort at the edges of the band must be minimised. This may require spectrum usage rights data to be easily accessible and a negotiation/co-ordination process to be well understood and dynamic.

2.5 Analysis of Bandwidth Requirements

The bandwidth requirements of different services that could exploit a spectrum aggregation system are illustrated below. Three scenarios have been considered, namely:

- Home networking
- Office networking
- Wide area wireless access

These three scenarios are described in more detail below.

2.5.1 Home Services

A typical home may require wireless connectivity to entertainment, communication and IT devices.

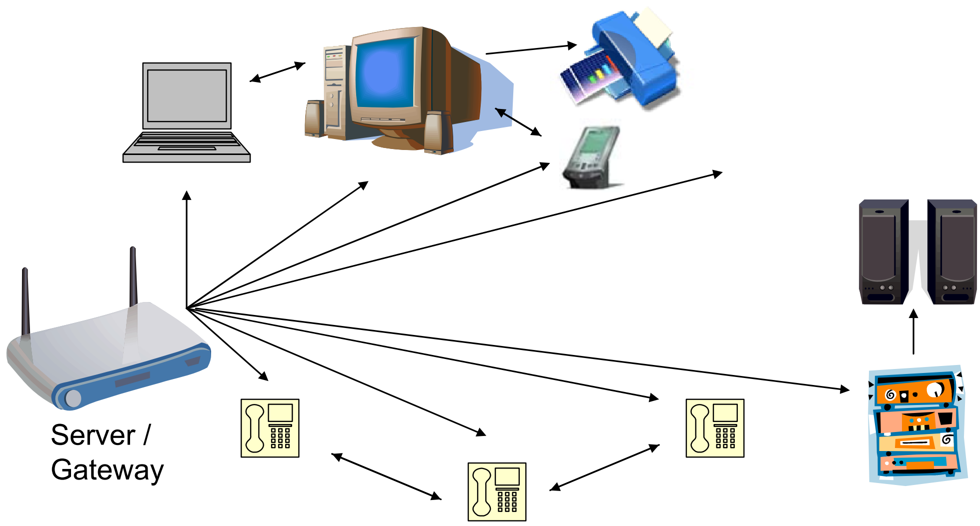

Figure 2-3: Example of Wireless Home Network

The typical bandwidths of wireless services that might be delivered in the home are:

Service	Bandwidth (kbps)
Voice telephony	48
Web browsing	512
Streamed video	1500
Streamed audio	128
E-mail	512
Data transfer (e.g. PC-PC or PC - printer	10000

Table 2-3: Bandwidth and or the Home Network Scenario

The operational range in a typical home environment is likely to be no more than 50 metres. Consequently frequencies in the GHz range would be favoured by system developers.

2.5.2 Office services

The Office services are in many ways similar to the home network, except that there will be a lower emphasis on entertainment services (e.g. audio / video streaming) and the operational range could be somewhat greater (up to several hundred metres in a large company premises.

The sort of services that might be delivered over an office wireless devices include:

Service	Bandwidth (kbps)
Voice telephony	48
Video telephony	144
Web browsing	512
Fax	10
E-mail	512
Data transfer (e.g. PC-PC or PC - printer	100000

Table 2-4: Bandwidth and or the Office Network Scenario

2.5.3 Wide Area Wireless Access

This scenario is quite different from the other two in that the geographic range could extend to several km rather than a few tens or hundreds of metres. The wireless link is essentially between the network and the user terminal, rather then between specific devices within a home or office network. It is therefore not appropriate to prioritise specific services, since the network will not generally be aware of what the data being transferred over the network is being used for. Instead, it is likely that the network will specify a particular peak bit rate and contention ratio, which may vary according to the tariff. For example, there might be a basic tariff offering 512 kbps with 40: 1 contention, a mid-range tariff offering 1 Mbps with 20:1 contention and a premium tariff offering 1 Mbps and 10:1 contention. In terms of apportioning spectrum resources, it will be necessary to give priority to the premium service when demand is highest, followed by the mid-range and finally the basic service. The preferred spectrum will be below 1 GHz, to maximise the operational range.

2.5.4 Conclusions and Recommendations

If we consider the typical services above then the bandwidths required could be obtained using the spectrum fragments highlighted. For example, short range home streamed video requiring 1.5Mb/s could be achieved for some users in London by low power devices aggregating the 2.65MHz spectrum available in the PMR band or by using the 2MHz that does not appear to be assigned in the 3440 – 3442 MHz band.

2.6 Impact of spectrum fragmentation on value (based on auction payments)

The “value” of spectrum fragments can be interpreted by considering the auction price of the spectrum as a function of spectrum size. This relationship is derived from GSM spectrum auctions in Austria and the Netherlands. The results suggest that bidders value larger blocks of spectrum more than smaller blocks (on a per-MHz basis), though it should be noted that other factors such as timing also have an impact on the value (in this regard the Netherlands results are probably more pertinent as these auctions were all held simultaneously. The data support the hypothesis that because larger fragments are valued more than smaller fragments, there will be a market trend to aggregate rather than fragment spectrum. This aspect is discussed further in the next section.

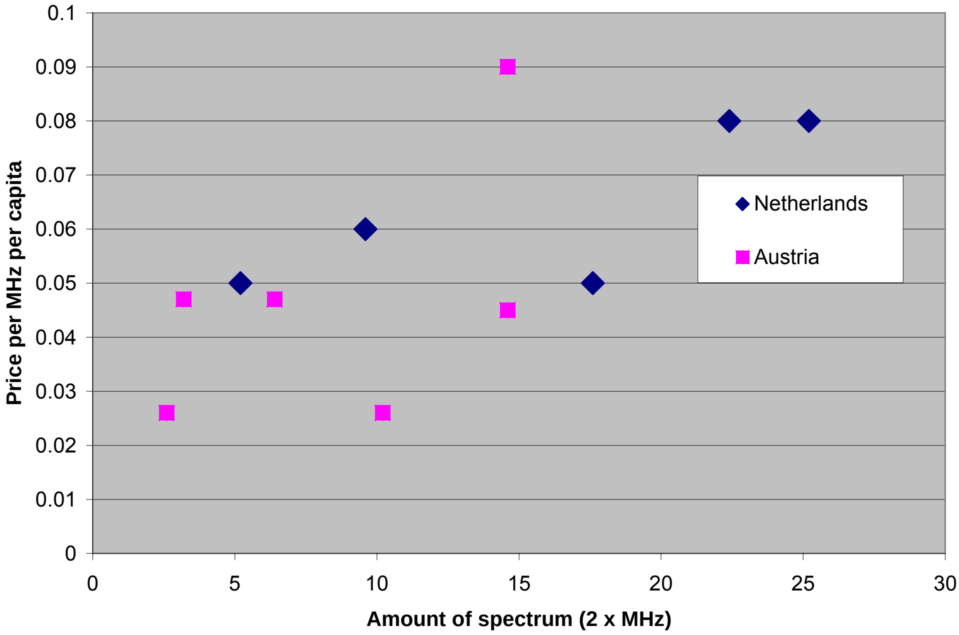

Figure 2-4: Amounts paid for GSM spectrum (€) in Austria and the Netherlands as a function of bandwidth acquired (normalised to bandwidth and population)

It should be noted that value of smaller segments may be lower if the spectrum is being purchased by an operator who already has a large amount of spectrum already. Smaller segments, therefore be more valuable to a provider who has no spectrum already.

2.7 Conclusions

The primary conclusion of the literature search was that if the PMR unlicensed bands were considered as exploitable fragments, the extent of these fragments varied significantly by geographic location. For example, London had only 2.65 MHz of total unassigned spectrum fragments compared with 17.275 MHz for Ullapool, 11.275 MHz for Newcastle-upon-Tyne, and 4.925 MHz for Brough.

A total of 325 kHz of nationally unassigned PMR frequencies were also identified, as detailed in the following table:

Freq Range (MHz)	Bandwidth (kHz)
158.73125 – 158.80625	75
159.11875 – 159.14375	25
159.15625 – 159.19375	37.5
163.23125 – 163.30625	75
163.61875 – 163.64375	25
163.65625 – 163.69375	37.5
164.14375 – 1664.16875	25
454.41255 – 454.43755	25
Total	**325**

Table 2-5: Frequencies not currently allocated to a specific use on a national basis

There were also a number of apparently unallocated spectrum bands that would appear to be available nationally, These are detailed below:

Frequency band	Bandwidth (MHz)
862 – 863 MHz	1 MHz
1375 – 1389 MHz	14 MHz
1399 – 1400 MHz	1 MHz
2290 – 2302 MHz	12 MHz
3440 – 3442 MHz	2 MHz
3475 – 3480 MHz	5MHz

Table 2-6: Nationally available unallocated spectrum bands

With respect to future fragmentation of the spectrum and its exploitation, it has been shown that even if the fragmented spectrum could be used, the management overhead in co-ordinating users at either end of the spectrum fragment will be a disincentive to exploit the fragment and will also result in spectrum inefficiencies because of the guard band needed at either end of the useful fragment. Data from spectrum auctions also shows that larger bandwidths are “valued” more than smaller bandwidths. This supports the hypothesis that in a spectrum-trading environment there will be more of a trend to aggregate spectrum than to fragment spectrum. This aspect is considered further in the next section.

2.8 Recommendations

It is recommended that Ofcom consider the implications of making available each of the spectrum fragments that have been identified as available nationally to one or more national service providers for use in conjunction with an appropriate aggregation mechanism, or to provide a national control channel that could be used in conjunction with such a mechanism, to identify locally-available traffic channels for aggregation purposes.

The exploitation of fragments (should they occur) will require data on users of spectrum adjacent to each fragment for co-ordination. These data should be made readily available to ease co-ordination issues.

2.9 References

[1] See http://www.ofcom.org.uk/static/archive/ra/topics/spectrum-price/spec-pric/1998/report.htm for details.

3 Policy and regulatory issues

3.1 Introduction

Fragments of spectrum already exist under the current administrative management of spectrum, and trading could lead to further fragmentation or aggregation of spectrum, depending on the economics of using spectrum fragments versus other spectrum. In turn, the level of spectrum aggregation will depend on the availability of other spectrum and the availability of technology that can utilise spectrum fragments efficiently.

In this section we consider what could cause fragmentation, comment on the "value" of fragmented spectrum and consider the powers that a regulator should retain.

3.2 What could cause fragmentation?

Fragmentation may arise due to administrative decisions or as a market outcome under spectrum trading. However, fragmentation may be optimal (in the sense that it maximises social surplus or the sum of producer and consumer surplus) depending on the underlying economics of using available spectrum in different ways.

Administrative decisions in relation to spectrum planning can result in large or small blocks of spectrum being allocated to different uses. Unused, unallocated or inefficient "fragments" of spectrum could then arise from a series of administrative decisions. For example, the original spectrum assigned to Cellnet (now O2) and Vodafone was augmented over the years and resulted in considerable interleaving between the two operators' assignments. Until these assignments were rationalised, this had an adverse impact on the efficiency with which the operators could utilise their spectrum assignments (the spectrum inefficiency was also reflected in a modifier applied to the spectrum fee) [1]

There are economic circumstances in which excessive fragmentation or aggregation might arise and/or persist and involve inefficiency. In general trading should lead to the optimal allocation of spectrum unless:

- particular uses of spectrum have public benefits that are not reflected in private values due to the existence of external costs or benefits of spectrum use or the services permitted by spectrum use (the public benefits associated with public broadcasting for example). In this instance other interventions such as the reservation of spectrum for particular uses and/or public funding of the associated output, may be utilised to bring private and public value into alignment.
- transaction costs, including the costs of trade and information asymmetry, block otherwise efficient trades.
- Market power; where aggregation of spectrum is motivated by a desire to acquire dominance in a market by controlling an essential input.

We note that licence holders might conceivably divest small amounts of spectrum they no longer need leading to fragmentation, for example, in particular geographical location. However, we consider this both unlikely (due to the option value associated with holding onto currently unused spectrum) and not necessarily inefficient if it did occur since the fragment would only be divested if it was of value to someone else.

The second and third points above could also result in inefficient instances of idle spectrum.

3.2.1 Transaction costs

Transaction costs that limit the scope for otherwise efficient trade in markets arise for two fundamental reasons:

- there are costs involved in locating a trading partner and carrying out a transaction
- information asymmetries in thin markets may result in strategic behaviour which results in some forgone trades

Markets will themselves adapt and adopt mechanisms to minimise the extent to which these problems limit efficient trade since there is economic surplus to be shared by doing so. However, in some instances inefficiency may persist – though there may or may not be a public policy solution.

Many services using spectrum require investment in relatively long-lived assets (10 years and more) and use of the service can persist for many decades with relatively few changes in technology, in part because of the large costs of replacing consumer equipment (e.g. radio and TV broadcasting, radars). This means that trading in spectrum markets could be sporadic: mostly coinciding with times when technologies and/or services change radically. This, together with the multi-dimensional nature of radio spectrum, will mean that much of the time the market may be rather thin in the sense that there are relatively few buyers and sellers for a given spectrum “product” [6].

Experience from Australia suggests licensees trade about once every seven years. In this situation, many transactions are likely to involve bi-lateral negotiations between a buyer and a seller rather than a competition for the sellers’ spectrum licence. Inefficiency, however, can occur in bilateral trade due to information asymmetries. If one spectrum user values spectrum more highly than another in a bilateral market trade may nevertheless not occur, since both parties to the trade misrepresent the value they place on spectrum in the hope of influencing the terms of trade. Under particular assumptions it has been demonstrated that 25 per cent of otherwise efficient trades would be foregone under a bilateral setting [7]. An administrative approach would not overcome this source of inefficiency, since the same information asymmetries would arise.

The possibility of post assignment inefficiency if markets are thin points to two public policy lessons:

1. ensure that the initial allocation is as close as possible to the efficient allocation (though this allocation is unlikely to remain efficient without subsequent trade as technology and tastes change over time)
2. ensure that the definition of property rights and the institutional environment minimises transaction costs

To see that (1) above can improve efficiency relative to an inefficient initial allocation and post allocation trade, consider the case (for simplicity) of a Vickrey auction where the winning bidder pays the price of the second highest bid. The dominant strategy in this auction is to bid your actual valuation (in the absence of collusion) which will result in an efficient allocation, even though bilateral trade in the secondary market would not necessarily result in efficiency. Careful design of the initial allocation process is therefore important taking into account issues such as the appropriate size of spectrum packages to be allocated, and whether complementarities might exist between packages.

In relation to the secondary market, a clear definition of spectrum assignments and a transparent register of assignments are likely to facilitate efficient trade (and efficient investment in complementary technology to make use of spectrum). In addition, predictability in terms of future allocations of spectrum can be expected to facilitate efficient trade.

Ofcom's Statement on Spectrum Trading (2004) indicates that there will be a database of assignment information which potential buyers of spectrum will be able to access to identify entities with whom they may wish to trade. This database is not expected to indicate the extent of spectrum use and in particular whether there are parts of assignments that are currently lying fallow. Licensees will be expected to indicate, say through a broker or exchange, via direct contact with potential buyers or via the trade press that such spectrum is available for sale.

Market based information exchange processes should in principle be sufficient to allow licensees to identify blocks of spectrum suitable for aggregation. Potentially more problematic is access to spectrum that has not been released by Ofcom but where there could be opportunities for aggregation.

3.2.2 Market power

The potential for market power to be exercised in a final goods or service market due to control over a key scarce input such as land or spectrum is another potential problem. Unchecked, scope to exploit market power in the final goods or service market could lead to excessive aggregation (rather than fragmentation).

Part of the policy response where concern over market power arises is to ensure that the initial spectrum allocation is not too concentrated by designing spectrum rights bundles and auction rules which promote an outcome with multiple holders of spectrum, and therefore potential competitors in the market. Competition from other platforms, or from services using other spectrum bands, may also alleviate concern over the exercise of market power based on control of particular spectrum (indeed, the examples cited earlier of swaps and aggregation to provide WiMax services were, in part, motivated by a desire to compete effectively with existing cellular operators).

Beyond the initial allocation it is usual to rely on general competition law to vet proposed mergers that could result in excessive dominance of a market. An example of a market where holding of a key input (land) has led to concern by the Office of Fair Trading is the grocery market. In May 2006, the Office of Fair Trading referred the groceries market to the Competition Commission under section 131 of the Enterprise Act 2002 on a number of grounds, including the following [9]:

"There are reasonable grounds for suspecting that the land holdings of the large supermarket multiples may reinforce their existing market position in some local areas. The OFT has also found evidence of prices that could have an anti-

competitive effect, including the use of restrictive covenants in relation to sites sold by the big supermarkets."

This case suggests that existing powers can be used in relation to concern over control of an essential input, in addition to the exercise of merger controls in the final goods or service market.

3.2.3 Summary

Whilst a range of market imperfections can arise in relation to a market allocation of spectrum, they would not necessarily lead to fragmentation. In addition, there are a range of policy mechanisms for addressing market imperfections that do not involve intervention in relation to fragmentation, or aggregation per se.

3.3 What powers should a regulator retain/use if fragmentation occurred?

For markets to operate most effectively, policy commitment and predictability is required, and this may be at odds with holding too much discretion to address perceived (short term) market problems. This is one reason for preferring, where possible, general frameworks such as competition law for addressing possible problems rather than sector specific powers such as the power to re-farm spectrum to achieve efficient use and to block trades for spectrum management reasons. Clearly, there is a balance to be struck, and a clear and transparent spectrum policy may provide assurance to market participants where specific powers to intervene are held.

Auctions, other trades or Administrative Spectrum Prices (AIP) may also help to facilitate efficient secondary trading and the optimal degree of aggregation by providing information about the opportunity cost of spectrum to potential parties to a trade. In the US, in relation to trading of Clean Air permits, annual auctions of a small percentage of permits is considered to have facilitated bilateral trading by publishing common values for permits [8].

Ofcom should only retain powers to re-aggregate spectrum if the benefits, in terms of more economically efficient spectrum use, exceed the cost in terms of reduced regulatory certainty and the knock-on effects on licensee behaviour. We note that in other markets such powers do not in general exist, which suggests they may be unnecessary. In addition, whilst various forms of "market failure" can arise they would not necessarily lead to fragmentation, in which case powers to re-aggregate spectrum would not be an appropriate remedy.

If Ofcom retains powers to aggregate, licensees have weaker incentives to fragment the spectrum themselves, in case Ofcom intervenes and changes their licence rights. (Ofcom has proposed there will be a five year notice period for changes to licences on spectrum management grounds.)

At present Ofcom has powers to refarm spectrum so as to achieve efficient use of the resource (such as re-aggregations), though in practice incumbent users are given relatively long notice periods, 5-10 years, before the refarming can occur. Under new licences currently being issued by Ofcom five year notice periods must be given before Ofcom can reclaim or otherwise modify the licences. Similarly 5 years notice of changes must be given before the end of longer duration licences.

Ofcom also has powers to block trades for spectrum management reasons. This could in principle include blocking certain trades if they are thought likely to lead to “excessive” fragmentation of spectrum. We see no reason at this time to increase or to change these to deal with fragmentation of the spectrum.

There are no clear grounds for thinking that spectrum markets will lead to excessive fragmentation. Rather, the opportunity to trade could be expected to promote optimal levels of aggregation.

3.4 What is the “value” of fragmented spectrum, compared to non-fragmented?

The value of radio spectrum to a user is very dependent on how the spectrum can be used and this can lead to some surprising conclusions. For example, it might appear intuitive that TV broadcast spectrum would be worth more than FM broadcast spectrum, which in turn would be worth more than AM broadcast spectrum, reflecting the relative revenues generated from these services. However this overlooks the amount of spectrum required to deliver the service to a large audience and the costs associated with delivery. A national analogue TV service requires up to eleven 8 MHz radio channels, i.e. a total bandwidth of 88 MHz, whereas a national FM service can be delivered with only 2 MHz and a national AM service can be delivered with only 18 kHz.

These differences are reflected in the opportunity cost estimates developed recently by Indepen and Aegis for Ofcom [10], which estimated the value of TV spectrum at £1.12M per MHz, FM radio spectrum at £1.8M per MHz and AM spectrum at £90M per MHz.

To highlight the value aspect further it is interesting to note that auction prices paid by 12 licences, for 6.6MHz at 1781.7 – 1785MHZ and 1876.7 – 1880MHz, ranged from ~£50k to over £1M.

Idle spectrum, spectrum fragments and aggregation do not imply inefficiency per se. In a spectrum trading environment spectrum might be fragmented or aggregated over time in response to changes in economic value of different configurations. For example, Nextel in US aggregated spectrum in order to launch a mobile telephony service. Further, spectrum swaps have been utilised in the US and Australia to build to build up blocks of spectrum for WiMax deployment. [2]

Fragments of spectrum might at times be allocated but unused, but that would not necessarily imply that the allocation was sub-optimal (or of little value), since the most efficient use of some spectrum at a given point in time may be to be hold it in reserve for a prospective use. Idle spectrum does not therefore automatically imply inefficient use, and may simply reflect the fact that the necessary technology or infrastructure is still under development, or that the optimal use is uncertain and waiting to see how alternative prospective uses develop is the best current option.

In other markets such as the land market, the optimal extent of aggregation (which maximises value) can increase or decrease over time as technology and demands change, for example, in terms of the economies of scale in agriculture or in terms of the style of residential housing. Resources may also at times be held in reserve for future use, for example, a property developer might aggregate land for development, and then sell smaller lots for residential or commercial use. Furthermore , land may be unused for a time, but in general one does not observe a tendency towards the creation of unused fragments or excessive aggregation. Non-use of land also relates to the planning process whereby potential competing uses of land and questions over external impacts are resolved, in part, administratively; a process that can take some time before a final decision is made.

For example, the planning decision over Terminal 5 at Heathrow took approximately a decade to resolve.

Idle spectrum may be efficient (and have value) when its future use is uncertain, a degree of irreversibility in terms of the costs incurred to use spectrum is involved and delay is possible. Formally decision problems of this kind are said to involve "real options" associated with the value of waiting for new information [3]. New applications of spectrum invariably involve a degree of uncertainty, for example in relation to the technical and market potential of mobile TV. They also involve sunk or irreversible investment either in terms of associated capital investment, or in relation to the acquisition of spectrum which is unlikely to involve transaction costs. It may therefore be efficient to hold on to unused spectrum until uncertainty is sufficiently resolved to allow a decision about whether to proceed with investment in a new service, or to dispose of unneeded spectrum. Lumpy investment and market demand growth profiles, for example in relation to 3G rollout, will also mean that spectrum may be underutilised in the early stages of a new service.

Whilst there are some costs in terms of the technical efficiency with which spectrum can be used as the size of spectrum fragments decreases (due to the fixed size of guard bands), in a spectrum trading environment spectrum rights holders can be expected to take this into account when deciding whether to aggregate or fragment spectrum.

Differences in the value of difference sized fragments are reflected in auction proceeds. For example, GSM spectrum auctions in Austria [4] and the Netherlands [5] suggest that bidders value larger blocks of spectrum more highly than smaller blocks (on a per-MHz basis). However, the costs and benefits of aggregation and fragmentation will change over time as technology and the availability of spectrum change.

As discussed in other sections of this report, radio equipment that can use fragmentation bands is feasible and its potential capability and economics can be expected to improve over time. This will make it easier to use fragments, which in turn implies that fragments will become relatively more valuable. The optimal amount of fragmentation could therefore increase as the economics of using fragments improves. Alternatively, if new blocks of spectrum are allocated or reallocated from non-commercial uses, the relative value of spectrum fragments may fall. These interactions between economic and technical considerations are illustrated in Figure 3-1.

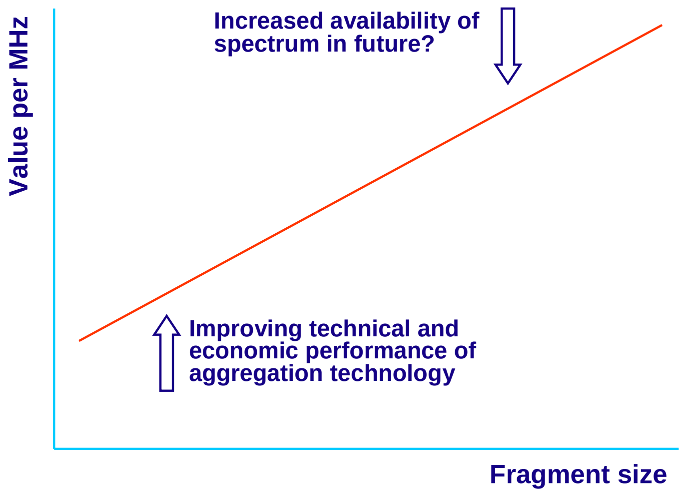

Figure 3-1: Interaction between economic and technical considerations

3.4.1 Summary

In summary, the value of spectrum fragments and therefore their extent will change in response to changes in the availability of spectrum and the technical and economic feasibility of utilising spectrum fragments. Any observed changes in fragmentation would therefore need to be evaluated relative to any such changes in fundamentals.

3.5 Conclusion

Spectrum markets would be expected to increase the likelihood that spectrum is aggregated to an optimal extent, by opening up market incentives for aggregation and fragmentation in response to changes in spectrum availability, the technology for using spectrum fragments and changes in demand for different services including new services.

Fragmentation per se should not necessarily raise concern since a degree of fragmentation may be optimal, and the optimal extent of fragmentation can be expected to change over time. Any observed trends would therefore need to be interpreted with caution in terms of their policy implications.

If policy intervention were considered appropriate then existing Ofcom powers are sufficient, possibly more than sufficient, to address issues in relation to fragmentation. Existing powers may in fact reduce the potential for the market to take a long term view of the optimal degree of aggregation, if they increase insecurity over spectrum rights.

We conclude that spectrum trading per se would not be expected to result in an inefficient increase in spectrum fragmentation; rather trading would be expected to facilitate aggregation wherever this is economic. The need for new policy approaches and powers to address spectrum aggregation is not therefore anticipated. Rather, clarifying and limiting existing powers to intervene in spectrum markets may facilitate optimal aggregation by encouraging traders to take a long-term view of the relevant costs and benefits of aggregation.

3.6 Recommendation

The Cave Audit recommended a review of progress in relation to spectrum management and the development of spectrum trading in five years time. Such a review would provide an opportunity to examine developments in terms of spectrum fragmentation and aggregation, and to test whether such developments raise public policy concerns.

The impact of existing Ofcom powers to intervene in the market could also be examined from the perspective of their possible impact on the development of a forward looking market which promotes efficiency in the short and long term.

3.7 References

[1] Radiocommunications Agency. 1998. "Report on Modifiers to be used in determining Administrative Pricing Fee Charges for Mobile Services."

http://www.ofcom.org.uk/static/archive/ra/topics/spectrum-price/spec-pric/1998/report.htm

[2] The Register. January 2006. “Sprint Nextel in Clearwire spectrum swap.” http://www.theregister.co.uk/2006/01/11/sprint_clearwire_spectrum_swap/

[3] Dixit and Pindyck. 1993. “Investment under uncertainty.” Princeton.

[4] http://www.rtr.at/web.nsf/englisch/Telekommunikation_Frequenzvergabe_Bisherige+Auktionen

[5] Eric van Damme. 1999. “The Dutch DCS-1800 Auction.” http://arno.uvt.nl/show.cgi?fid=3943

[6] A secondary market for trading spectrum: promoting market liquidity, M Bykowsky, Telecommunications Policy 27, 2003.

[7] Robert B Myerson and Mark A Satterthwaite. 1983 “Efficient mechanisms for bilateral trading”. *Journal of Economic Theory*, 28:265—281.

[8] Leese, Levine and Rickman. January 2002. “The economic effects of spectrum trading.” http://www.econ.surrey.ac.uk/staff/plevine/channeltrade7%5B1%5D-4.pdf

[9] Office of Fair Trading. May 2006. “The grocery market.” http://www.oft.gov.uk/NR/rdonlyres/1A2D7FA2-FEA3-4459-9B25-4A737A20023D/0/oft845.pdf

[10] Study into the potential application of Administered Incentive Pricing to spectrum used for Terrestrial TV & Radio Broadcasting, October 2005

QinetiQ Proprietary

4 A Technical Solution for Using Fragmented Spectrum

4.1 Introduction

Should spectrum fragmentation occur, the challenges in designing and building a practical aggregation-enabled system must be considered. These challenges are addressed in this section by investigating the technical feasibility of designing and building a system capable of aggregating spectrum. The technological and economic challenges of producing such a system using current state-of-the-art hardware are considered along with possible waveforms suitable for spectrum aggregation. The technological development trends that may affect the design and/or manufacture of future aggregating systems are also summarised.

In this section, four main technical questions are addressed:

- Over what bands can fragments be aggregated?
- What is the smallest and the largest fragment width possible?
- How many fragments can be aggregated in one go?
- What waveforms are suitable for spectrum aggregation?

Several assumptions have been made during this section. For example it is envisaged that spectrum aggregation's main application will be a static one, probably in the home or office. Therefore we have considered the feasibility of fixed installations only, ignoring (within reason) constraints on power consumption or the size of aggregating technology. Fragment sizes (or width of a fragment) are assumed to be in the range 25kHz to 1MHz. Anything smaller than 25kHz is considered too narrow to be of use, while an individual fragment wider than 1MHz should really be considered as a block of unused continuous spectrum in its own right. In spectrum aggregating transmission systems, a power spectral density (PSD) of +37dBm/MHz (5W per 1 MHz) is assumed sufficient.

No assumptions have been made about the number of fragments that might be present, or where those fragments might sit in the spectrum (for example evenly-spaced or clustered together).

4.2 RF Architectures

Radio frequency system-level designs, or RF architectures, can be split into three broad categories; receivers, transmitters and transceivers. For the purposes of this report, a transceiver architecture is one that shares common RF analogue components (other than the antenna) in both transmit and receive operations. Unsurprisingly then, transceivers are invariably the most complex and the most difficult systems with which to achieve good performance. This is particularly so for wideband systems and systems requiring full duplex operation (simultaneous transmission and reception). To simplify things, transceiver architectures are dismissed as too complicated for consideration in this report. But this does not mean that spectrum aggregating transceivers are a future impossibility.

4.2.1 Supersonic Heterodyne Principle

Regardless of how it is implemented, virtually every communications receiver architecture in existence today uses the supersonic heterodyne (superhet) principle of converting from RF to IF. Other receiver design principles such as tuned radio frequency (TRF), homodyne (also called direct conversion or zero-IF) and regenerative/super-regenerative circuits have become obsolete, suffering from inferior frequency stability, selectivity, and (in the case of TRF) sensitivity. Basic functionality of the superheterodyne architecture is shown below, depicted in what is often termed a component chain or receive chain.

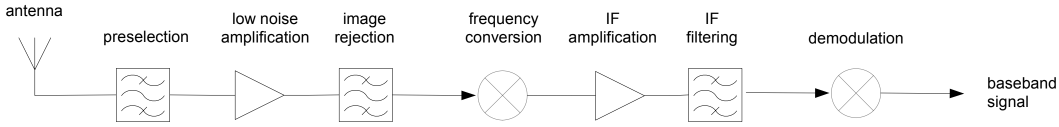

Figure 4-1: Receive chain using superheterodyne architecture

At the very front end of the receiver is the antenna, followed by a preselect filter (preselector) for preventing the undesirable reception of out-of-band signals, and a low noise amplifier (LNA) for boosting weak wanted signals. This is usually followed by image rejection (filtering) and frequency conversion (mixing) to an intermediate frequency (IF). Sometimes two stages of frequency conversion are used, for various reasons (including improved selectivity and/or image rejection), constituting a double superheterodyne receiver. Once at a suitable IF, further amplification and filtering can be achieved using cheaper and better-performing components/methods. Finally, the receiver architecture ends with the demodulation of the IF into a baseband signal, or modulation waveform. Depending on the type of demodulation this can sometimes be considered as an additional frequency down-conversion, from IF to baseband/DC, but is never classed as a heterodyning action in itself.

For relatively low RF bands that are close to the IF, double superheterodyning will probably be required to give adequate image rejection, IF rejection and LO isolation. The RF is first heterodyned to a relatively high frequency (perhaps around 1GHz), where the image and LO will be well displaced from the IF band and thus can be filtered with relative ease. However, this first IF must be selected carefully because if it is too high it will be difficult to find good performance components (after all, the whole point of converting to an IF is to use cheaper, better components at a lower frequency). Therefore the analogue component count at this first IF should be minimised but nevertheless some filtering is essential for the second stage image rejection. Filtering the output of the first mixer may also be required, and probably some intermediate amplification to overcome the insertion losses of the filter(s).

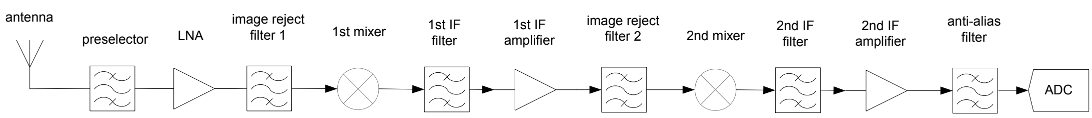

Figure 4-2: Double superhet architecture for low RF bands close to the IF

Drawbacks to this architecture are the increased component count, the greater care needed over frequency planning when choosing LO and IF, and the fact that two LO sources are required to drive the mixers which could pose self-interference problems. These issues might outweigh the desire to operate at lower RF bands.

Transmitters can also make use of the superheterodyne architecture, modulating up to an IF before signal processing and further up-conversion to RF. There are some obvious differences in the transmit chain, such as the types of filters and amplifiers used. The biggest difference between receive and transmit chains is the front-end amplification; usually in a transmit architecture a pre-amp is installed to boost signals before entering a high power amplifier (PA), which feeds the antenna. A low-loss filter may be placed after the power amplifier, to attenuate out-of-band emissions. Nevertheless the superheterodyne principle is the same.

4.2.2 Chain Options for Few Fragments

In general, wideband chains are harder to design and build than narrowband ones. They are more expensive and may not perform so well. Therefore, if only a few fragments need aggregating, it is desirable to use individual narrowband chains – one per fragment – than to try and use one wideband chain to share amongst the fragments.

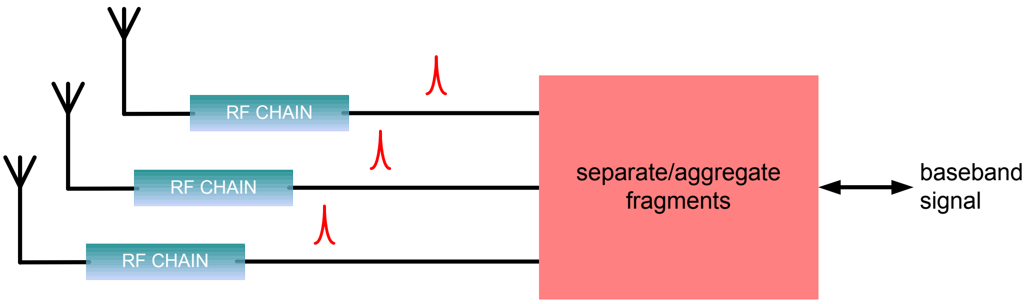

Figure 4-3: One chain per fragment

4.2.3 Chain Options for Many Fragments

In cases where there are more than a few fragments (say more than five), it is impractical to have separate chains dedicated to each fragment. The number of individual components (the 'component count') would become too high, resulting in an unacceptably large, heavy, expensive device with high power consumption even for a fixed installation.

Therefore chains must be shared by multiple fragments within a given band. This greatly reduces the number of components to a feasible level, but means that each chain must be wideband. Wideband performance is difficult to achieve and multiple-tone third-order intermodulation products caused by applying more than one fragment at a time to an analogue component (e.g. amplifier) will cause linearity problems. It may be possible to overcome such problems by using wideband, highly linear, high two-tone third order intercept (IP3) devices designed for multi-carrier applications. These kinds of devices are expensive and also quite large and power inefficient, and are therefore only suited to fixed installation scenarios.

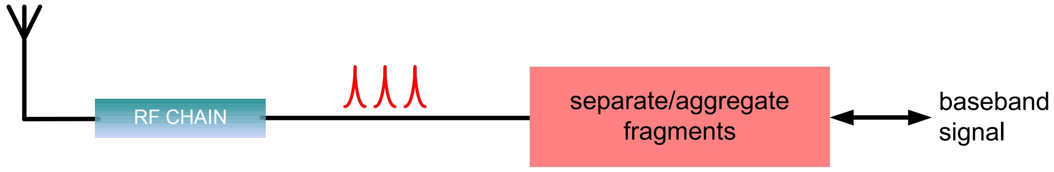

Figure 4-4: Wideband chain shared amongst multiple fragments

Of course there are limits on how wide a wideband chain can be, depending on the available technology. This will be discussed in more detail later, but the requirement for a chain to be wideband depends on how fragments are positioned in the spectrum. For example, if fragments are clustered together, only one wideband chain may be needed to 'capture' all them. But if fragments are broadly scattered, especially beyond the coverage of a single wideband chain, then more than one chain will be needed to capture all the fragments. The number of chains must be limited to a few, otherwise the same problems of high component count that are associated with narrowband chains are encountered. However, multiple wideband chains will still be able to capture many more fragments then a multiple narrowband chain architecture.

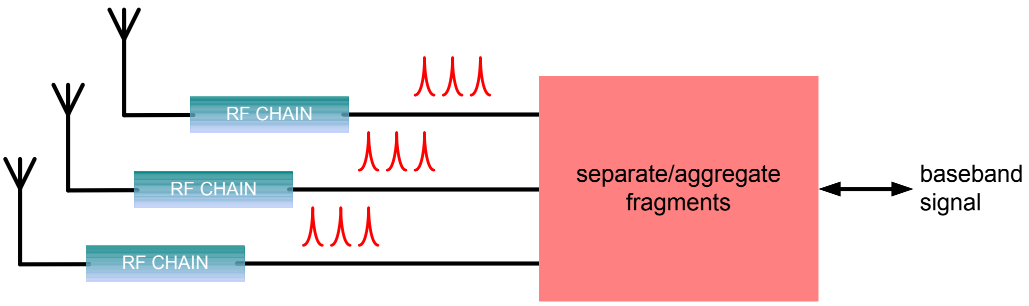

Figure 4-5: Multiple wideband chains could pass many fragments

4.2.4 Antenna Sharing

Whether they are narrowband or wideband, multiple chains still exhibit a high component count. One of the bulkiest components is the antenna. For this reason, and also because of potential electromagnetic mutual coupling effects (which are hard to predict) that can occur between closely located antennas, it is undesirable to have a forest of antennas. To prevent this, one or two wideband antennas could be shared amongst chains.

For receive chains, an architecture somewhat like that shown in Figure 4-6 could be used to share an antenna. An important characteristic of any receive chain is a low noise figure, which can be achieved by minimising signal loss while maximising gain at the front end of the chain. While a splitter will add loss, placing one overall wideband low noise amplifier (LNA) between the antenna and the splitter should compensate for this and keep the overall noise figure of the chain to an acceptably low level. Sharing the LNA itself does not really add benefit to the design, since additional LNAs will still be needed in each individual chain (remember that the shared LNA is there only to compensate for the splitter loss).

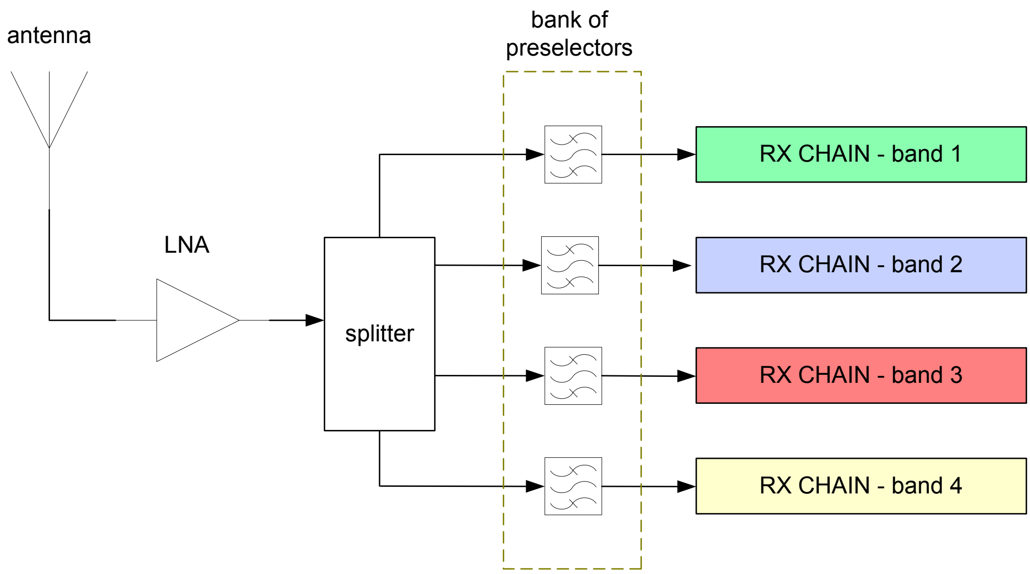

Figure 4-6: Antenna sharing by receive chains

4.2.5 Power Amplifier Sharing

In transmission systems, the power amplifier (PA) is usually a bulky (and power hungry) component as well as the antenna. Therefore it would be desirable to share both the antenna and the PA in multiple transmit chain designs. A combiner can be used to do this but it introduces even more unwanted intermodulation products in addition to those already generated if using multiple fragments within a wideband chain. A broad band low-loss band-pass filter (BPF) after the combiner can suppress the out-of-band products, but not those in-band. Thus a very high performance linear PA is required. The current and future possibilities of linear power amplifiers and linearisation techniques are discussed later, but the basic architecture behind PA sharing is shown in Figure 4-7.

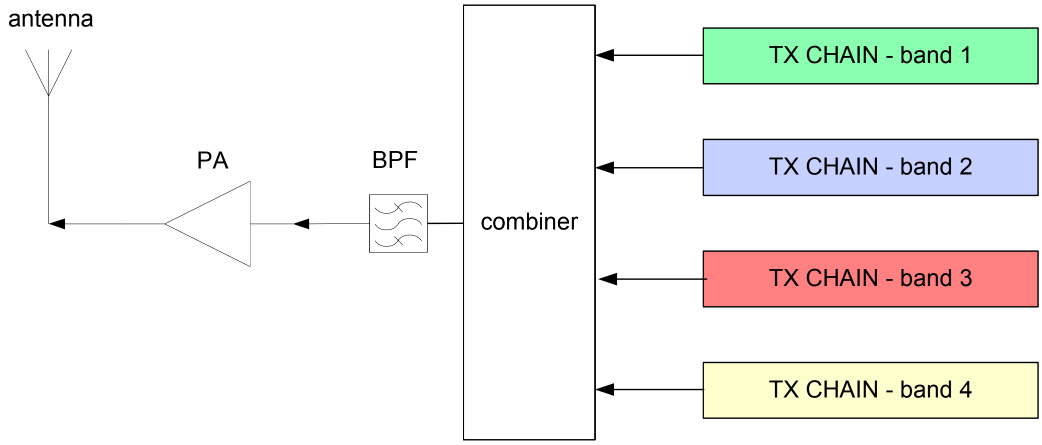

Figure 4-7: Antenna and PA sharing by transmit chains

4.2.6 Summary

In considering the RF architecture needed to exploit multiple spectral fragments, a design is clearly dependent on the size of fragments and how they are scattered in the spectrum. A single chain of transmit and receive components could be expected to handle several fragments clustered within a frequency band. However, the components must be highly linear to minimise the generation of inter-modulation products, some of which could fall in band and be difficult to filter out.

Because of the limits of current radio technology, multiple parallel component chains are needed to handle fragments scattered over a number of frequency bands. There is a desire to share bulky and costly components such as the antenna and transmitter power amplifier, but this requires a sophisticated combiner, splitter and multi-band antenna, which are not readily available.

4.3 RF Components

4.3.1 Digital Signal Processing (DSP)

For the purposes of this chapter DSP means DSP in general, be it processors, FPGAs, ASICs, algorithms, etc. Regardless of whether a chain is narrowband, wideband, transmit or receive, an aggregating device will at some point need to use DSP. Even multiple fragments within a chain must be individually filtered and separated from each other. This is not practical using analogue components. But in the digital domain, signals can be copied, conditioned and otherwise processed with much greater ease. For example, it is easy to create a high-Q 25kHz band-pass filter centred around several hundred MHz using DSP; to get the same performance using analogue hardware components would be extremely difficult, and result in very bulky, lossy and expensive filters.

An important functional block of DSP for receivers is the digital down converter (DDC), comprising complex multipliers, oscillators, low pass filters and decimators to provide what is essentially a digital equivalent of a mixer, LO and filter used for heterodyning (Figure 4-8). Regardless of the various techniques available to implement digital down conversion, the digital portion of each chain needs to perform two things; separate fragments into individual carriers, then demodulate them. To do this, a number of narrow band DDCs are required, one per fragment. Depending on the analogue carrier frequencies being digitised, a single, shared wideband DDC may also be implemented to shift the entire band down to a lower centre frequency beforehand.

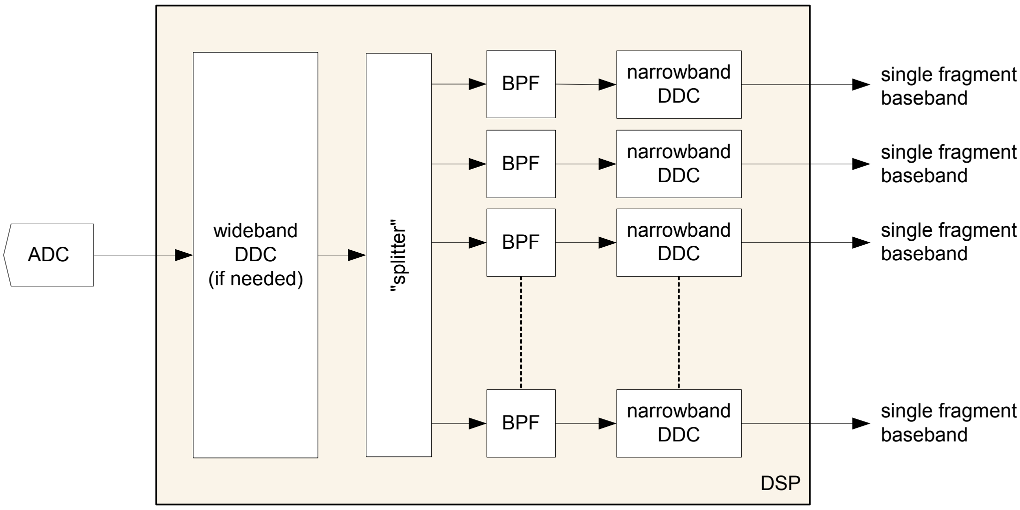

Figure 4-8: Digital Down Conversion for a receiver chain

State-of-the-art DSP platforms have in the order of 200,000 logic cells which are quite capable of managing multiple fragments and DDCs; exactly how many fragments per microchip is difficult to predict as the processing power required depends on many factors including the width of each fragment, the type of modulation/demodulation scheme in use, up/down conversion parameters, filtering demands, and so on.

Whereas DDC blocks are used in the receive chain, for the transmit chain digital up converters (DUC) are required. DUCs will accept baseband signals and up-convert them to an IF carrier. The bandwidth of the DUC is fundamentally limited by the clock speed of the digital processor. State-of-the-art DSP platforms can operate at clock speeds up to ~500MHz, therefore limiting current DUC performance to 250MHz bandwidth (Nyquist criterion applies).

So, although DSP has developed enough to allow implementation of digital counterparts that are superior in some ways to analogue component functions (e.g. filtering, splitting, combining), all such digital functions are, at present, limited to low frequency. Therefore analogue components are also a necessity for all but the most limited/lowest RF chains. State-of-the-art DSP platforms, such as Xilinx's Virtex-4 series, cost in the region of £400 per device.

4.3.2 ADCs & DACs

If both analogue and digital components are required in a chain, then there is a need for an analogue/digital interface. Such an interface is provided by analogue-to-digital converter (ADC) devices in receivers and digital-to-analogue conversion (DAC) devices in transmitters.

As mentioned in the previous section, multiple fragments within a chain must be individually filtered and separated from each other, and this is only practical if implemented in the digital domain. This gives rise to an important implication; digitisation of the entire band (width of a chain) is required before aggregating or separating fragments.

Research indicates that the majority of wireless communication receivers require a dynamic range equivalent to 16 bits, so any ADC with lower bit resolution is deemed unsuitable. Sampling speed is also an important factor, to maximise the chain bandwidth.

The fastest ADC devices use flash conversion employing 2n comparators for an n-bit resolution. They are available with sample rates up to 1.5Gsps, but are limited to resolutions of 8 - 10 bits. Other drawbacks include high cost and power consumption.

Sigma-Delta converters offer the advantage of providing noise-shaping and hence the anti-aliasing filtering requirements may be significantly relaxed. They offer high resolution, but do not currently provide the sample rates required for wideband receivers, with sample rates in excess of 10Msps not commonly available. However, this technology (in particular Continuous-Time (CT) sigma delta converters, which offer low power and size) could provide ADCs for future radio systems. CT sigma delta technology capable of 80Msps is currently under development.

The most suitable ADC architecture for wideband receivers at present is pipelined conversion, which employs a cascaded series of flash converters. The current state of the art offers 125Msps at 16-bit resolution. Therefore Nyquist-sampled digitisation is limited to around 60MHz maximum IF. However, under-sampling techniques could be used to digitise higher IF whilst preserving a 50MHz signal bandwidth. The cost of 16-bit 125Msps ADCs is around £100.

There is a wider choice of suitable DAC devices for two reasons. First, the bit resolution required for wireless transmission is lower at only 12 bits. Second, DAC performance exceeds that of contemporary ADCs. There are 12-bit DACs with 1Gsps speeds available at reasonable cost in today's market (few £10s). These can generate carrier frequencies up to the Nyquist rate, i.e. up to 500MHz.

It is worth noting that the analogue/digital interface is gradually moving toward the antenna in modern RF design, as ADC/DAC speed and resolution improve with technological progress (Figure 4-9). Thus an ever growing portion of modern RF architectures rests in the digital domain, and less in the analogue domain. The ultimate goal is to directly digitise the entire frequency band of interest, eliminating the need for filtering and frequency conversion in the analogue domain beforehand. This is, of course, the aim of software radio and will require very fast DACs and ADCs. As mentioned, DAC performance leads ADC performance, so the main limit to analogue/digital interface progression is the ADC. Since the late 1980's the rate of increase in ADC performance has followed a trend; sampling speed has approximately tripled every 6-8 years. Or put another way, dynamic range increases by ~1.5 bits for any given sample rate over the same period.

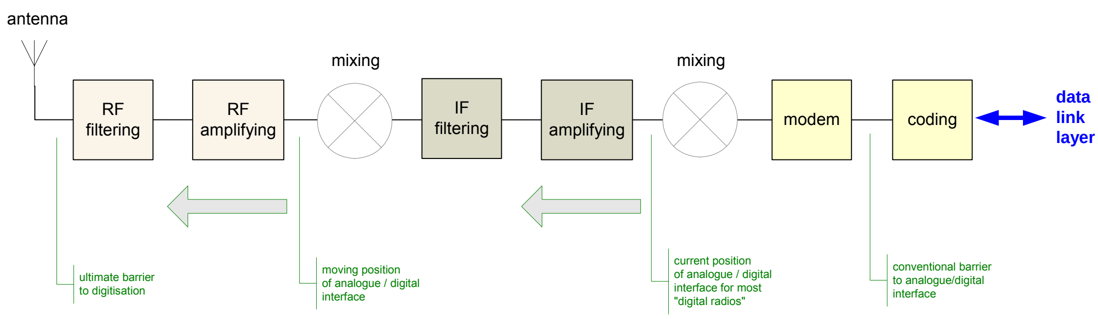

Figure 4-9: The analogue/digital interface is gradually progressing toward the antenna

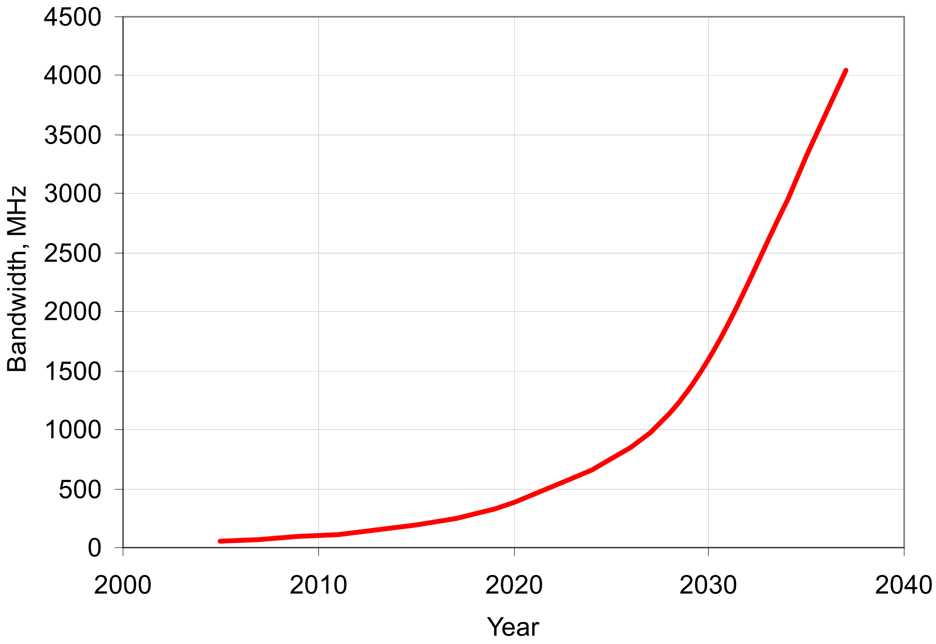

Figure 4-10: Future bandwidth able to be digitised if ADC developments continue as they have over the last 20 years (for 16-bit resolution).

4.3.3 Antenna

In the aggregating device the antenna will need to be wide band and, in the case of a transmitter, be able to radiate power. There should be no problem in obtaining an antenna with adequate power rating: most transmit antennas have a power rating of 5W or more.

Obtaining acceptable wideband properties is more of a challenge. Ideally, one broadband antenna would serve all bands across the spectrum. For receive-only applications this is possible using scanner type antennas, such as discone designs. Acceptable performance between 75 – 3000MHz can be achieved. The drawback is that they are large and bulky.

For transmit applications it is much harder to make good wideband antennas. Unfortunately the laws of physics means there is always going to be a compromise between the frequency range, electrical efficiency and size. Conventional antenna designs have made it possible to have practical antennas capable of covering a single band (e.g. HF, VHF, UHF) or at best two or three bands. Available wideband transceiver antennas can cover approximately 100 – 500MHz, 470MHz – 800MHz, and 800MHz – 2.5GHz. Therefore it is feasible to expect that two or three antennas will cover most of the spectrum required for spectrum aggregation. However, it should be noted that the 470MHz – 800MHz band is covered by very large very high power TV broadcast antennas; these are not suitable for spectrum aggregation, but it proves that the technology does exist and custom-made antennas at smaller, lower power applications are possible.

Recent developments are aiming to solve the problem of wideband antenna operation. There are numerous multi-band designs that achieve efficient operation at a number of prescriptive frequencies spread over a wide range. While they are not truly wideband, multi-band antennas are an important step forward in wideband antenna technology.

Some multi-band antenna designs include mechanical (as opposed to electrical) tuning by switching on or off elements of an antenna. Micro-electromechanical system (MEMS) technology seems likely to have a key role to play in making such antennas compact and practical (Figure 4-11).

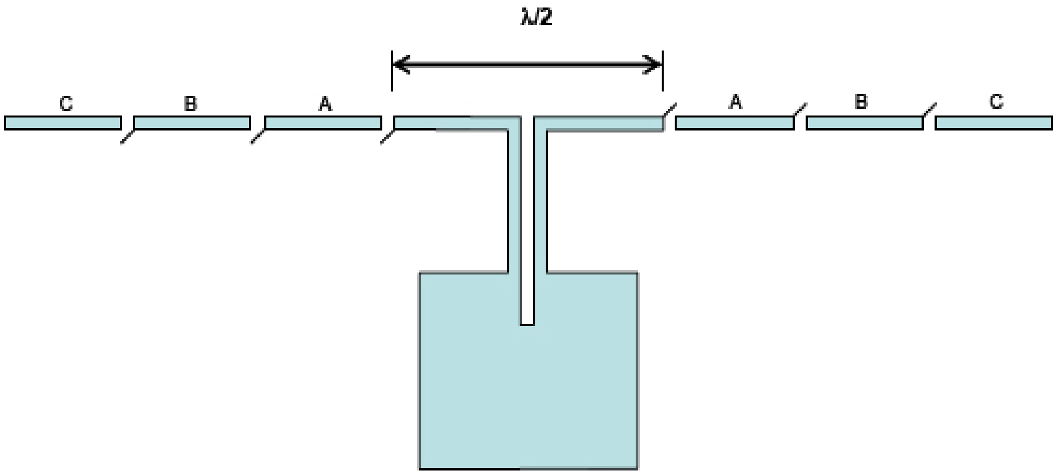

Figure 4-11: Mechanical Tuning of an antenna using switchable elements

Another multi-band design is to have a fixed physical structure with many resonances, such as a spiral antenna or sinuous antenna. This type of antenna can operate efficiently at each of the resonant frequencies. A major advantage over the mechanically tuned design is the ability to operate on more than one frequency at the same time (as well as being more robust). The natural progression from this is the fractal antenna, whose physical form is composed from a simple geometric shape which is repeated many times to make a complex pattern. The self-similar nature of fractals means that a high number of resonances are available over a wide range of frequencies. Also, fractal antennas (Figure 4-12) should go a long way in helping solve the 'bulky antenna' problem, because they are much smaller than their conventional equivalents. Currently there is only one company that has built and marketed fractal antennas, called Fractal Antennas Inc., based in America. They have proved the technology is feasible and economical.

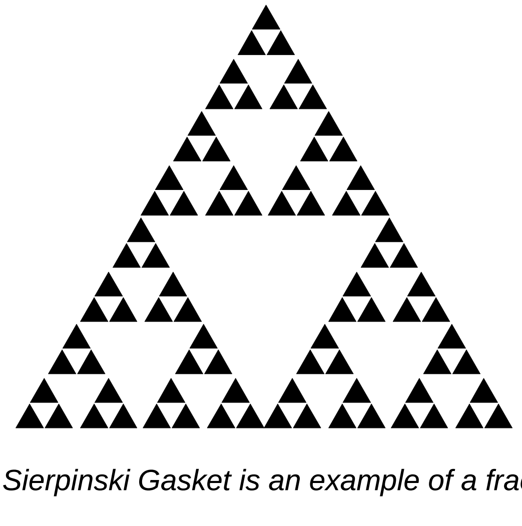

Figure 4-12: Sierpinski Gasket is an example of a fractal antenna geometry

Also worthy of note is the plasma antenna, which uses dielectric lenses and concentrations of plasma to electronically steer and focus RF beams. Because this technology is based on lenses and reflectors, it is inherently wideband and is not limited to particular resonances. The plasma antenna is a new invention that has yet to see mass production.

4.3.4 Divider/Combiner

To enable antenna sharing and aggregation, an analogue divider/combiner is needed at the front-end of the chain. Unfortunately such devices are not very wideband, with good performance devices operating from ~500MHz up to 2GHz with the necessary power rating. There are other ranges, such as 1-500MHz, 2-4GHz and 4-8GHz, but it is not possible to use multiple divider/combiners to extend frequency coverage in this way whilst sharing one antenna. The only way to achieve greater bandwidth by using a second divider/combiner would be to have an additional, separate antenna and front-end which defeats the point of combining chains in the first place.

Although a resistive splitter circuit could be used for the receiver to give an operating range of DC – 12GHz, this is extremely lossy (half the signal power is lost in dissipation) and the circuit does not work in reverse, i.e. it cannot be used as a combiner for the transmit architecture.

Standard configurations of divider/combiners include 2-way, 4-way and 8-way devices. Wideband, higher-way types can cost £1000’s. For a higher price, custom-made devices with any number of ways are feasible (although a sensible limit would be around 16-way).

4.3.5 Mixers

The key parameters of a mixer are the RF (input) and IF (output) frequency ranges, intermodulation performance, conversion loss and inter-port isolation. In addition, the mixer may be designed to provide a degree of image rejection in order to reduce filtering requirements and improve performance. Intermodulation performance is a measure of the level of spurious signals in the mixer output caused by mixing of harmonics of the input signals to the mixer. These harmonics are generated by non-linearities in the mixer, but some of them may be cancelled by employing a single, double, or triple balanced mixer. Additional benefits of double and triple balanced mixers include improved isolation, and wider band operation.

Mixers are currently available in a variety of sizes, types, packages and prices, although generally they are all quite cheap at just a few pounds per device. High IP3 broadband models can typically cover frequencies ranging from 500MHz to 5GHz. Mixers are not therefore considered to be a significant technological or economic hurdle to spectrum aggregation.

Frequency conversion is a well-established technology and there is not much in the way of mixer development. However there does seem to be a gradual trend toward improved linearity and lower intermodulation distortion mixers, especially at wider bands. Such improvements are relevant and will help in designing and building spectrum aggregators.

Mixer requirements are the same for receive and transmit chains. Analogue mixing occurs far enough away from both ends of both chains so that there is little different between the two.

4.3.6 Filters

A major advantage of using digital techniques is the potential to perform complex filtering techniques in the digital domain with relative ease. However, the analogue receiver and transmitter chains still require a number of filters operating at RF and IF frequencies. There are various analogue filter technologies including lumped element, combline, interdigital, helical, tubular, ceramic resonance, surface acoustic wave and waveguide. All of them have limits to what they can pass and reject, and there is no solution that will provide a perfect “brick wall” response.

For a spectrum aggregating system, filters must be sufficiently wideband to cover the required fragments, while still providing a steep roll-off so that unwanted signals are adequately suppressed. Due to the limits on filter percentage bandwidth, this can present a significant challenge, especially at lower frequencies. For example, in §4.3.2 it was established that ADC limitations restrict a receive chain’s bandwidth to ~50MHz. Filtering 50MHz is easy to do at higher frequencies, but at lower frequencies it pushes filter requirements to the limit of current technology. On top of this there are the usual considerations for filters such as insertion loss, phase distortion (critical for some modulation schemes), sharpness of roll-off at edge of band (size of transition band), and flatness of pass-band response (ripple).

If antenna and/or PA sharing is being considered, then multiple bands will need separating through filters. Because of the unavailability of perfect “brick wall” filters (Figure 4-13), guard bands above and below each chain band are required, lest signals from adjacent bands leak into the wrong chain. This imposes limits on the separation between different bands. The size of these guard bands determines the minimum separation between chain bands, and may be different for different band centre frequencies, due to the variation in filter technology and performance.

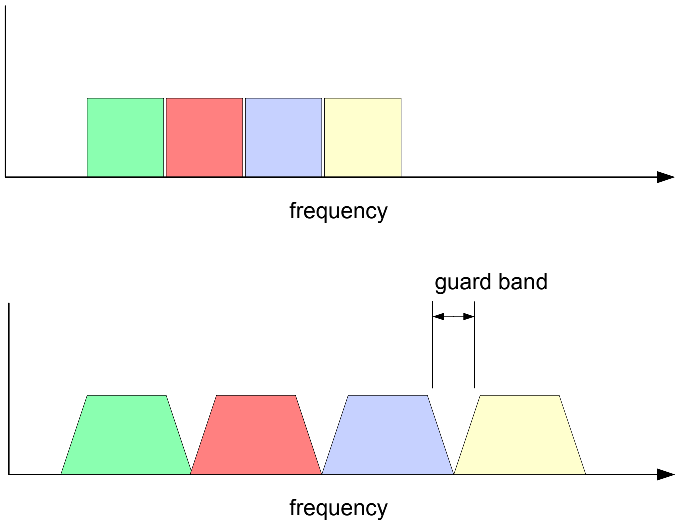

Figure 4-13: Only perfect “brick wall” filtering would allow adjacent bands to touch (top). In reality, guard bands are needed to protect adjacent bands from inadequate filtering (bottom).

Filtering in the receive chain has slightly different requirements to that in the transmit chain. For example image rejection is important in receivers, but in transmitter IMP rejection is also important. Anti-alias filtering before an ADC may be more difficult to do than reconstruction filtering after a DAC, because anti-aliasing usually requires a band-pass filter (BPF) while reconstruction could be performed using a low pass filter (LPF). For a given filter technology/size/cost it is easier to implement a LPF than a BPF. Preselect filters, or preselectors, are a usual requirement in receive chains but are not needed in the transmit chain. Instead, the transmit chain may require an overall broadband filter just prior or after the power amplifier. If it is placed after the amplifier, then there are addition complications such as obtaining a filter with a high enough power rating to pass the transmission. Filters normally degrade in phase distortion and insertion loss performance as their power rating increases. However, a very high performance power amplifier may eliminate the need for such a filter.

Leading filter manufacturers include BSC, Filtronic and Spectrum Microwave. Really good filters are expensive (£100's each) and usually bulky, especially at frequencies lower than 500MHz. However there are two emerging technologies that should help reduce the size of high performance, high-Q filters: Micro-electromechanical system (MEMS) technology has already been mentioned in §4.3.3 as a potential enabler of reconfigurable antennas through use of miniature RF switches. But MEMS also offers miniature capacitor and inductor components, thus very small lumped-element filters could become a reality in the near future (Figure 4-14).

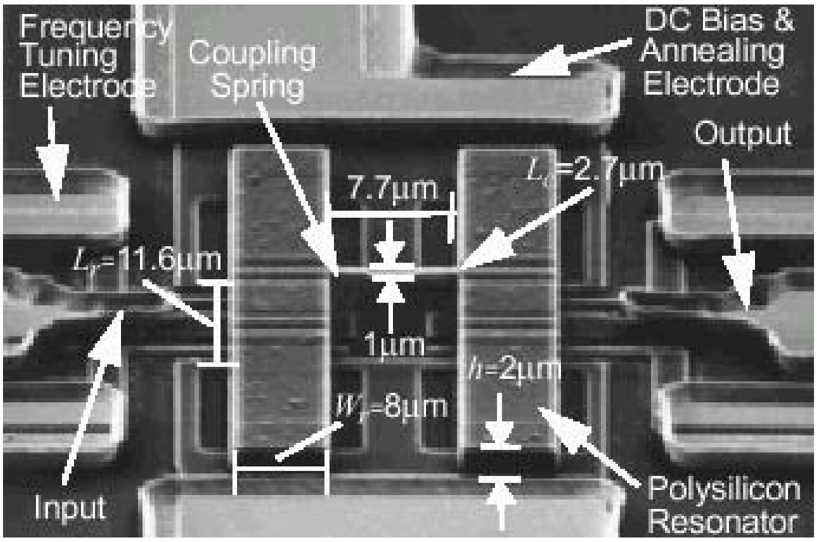

Figure 4-14: MEMS technology should enable extremely small lumped-element filters

Film acoustic bulk resonance (FBAR) is another technology with the potential to produce compact yet high-Q filters. FBAR may not be as useful as MEMS for filters though as its potential is limited to higher frequency filters only.

4.3.7 Low Noise Amplifier (LNA)

Near the front end of a receiver, it is important to achieve amplification (or gain) whilst introducing as little noise as possible to the system. A poor noise figure at the beginning of the chain will have a detrimental effect on the rest of the chain, and cannot be traded for good performance in following IF stages. For spectrum aggregation applications an amplifier with good linearity and high third order intercept (IP3) is also required. Fortunately there is plenty of choice in the form of low noise amplifiers (LNA), especially made for the purpose of front-end receiver gain. Partly due to the booming business of mobile telephony and short-range wireless (such as Bluetooth), LNAs are very cheap and come with a wide range of gain and IP3 characteristics. It is very easy to obtain inexpensive surface mount device (SMD) low noise gain blocks which can singly operate between DC and 6GHz. Many have output IP3 (OIP3) well above +30dBm (some approaching +40dBm), and offer sufficient gain and low enough noise for most radio receiver applications.

Some typical manufacturers of suitable LNAs include Sirenza, Watkins-Johnson, and Agilent. Typical price per device is a few pounds.

4.3.8 Power Amplifiers

The two most important characteristics of the power amplifier (PA) for spectrum aggregation are wideband performance and linear performance. Unfortunately, high linearity does not usually go hand in hand with wideband performance. Demands on wideband and linear performance are even greater if sharing the PA amongst several chains: the PA must be wideband enough to cover all the chains' bands, while the combining of signals at the input of the PA results in a multi-tone signal that will create its own non-linear intermodulation products (IMPs). It is inevitable that some of these products will fall in-band, i.e. at frequencies at which the PA is operating, making them impossible to filter out. So the PA has to be relied upon to give good linear performance and suppress these in-band IMPs. Fortunately, third generation mobile telephony and W-CDMA have driven the development and manufacture of multi-carrier power amplifiers (MCPA). Highly linear PA devices do exist, although there is not as large a choice as there is for LNAs. They are available for wideband operation and adequate powers: 10W amplifiers exist that can cover 20MHz – 1000MHz, 800MHz - 2GHz, 2GHz - 4GHz, and even 800MHz – 4.2GHz in a single device. These type of low distortion amplifiers rely on what is called class A or class A/B design. The drawback of class A or class A/B amplifiers is that they tend to be bulky, heavy and extremely power inefficient (5% efficiency is not uncommon). Even for the purposes of static spectrum aggregation system, these drawbacks could pose a major problem. They also tend to be very expensive, in the order of £1000's.

Another potential technology is metal-oxide substrate (MOS) field effect transistor (FET), more specifically vertical double diffuse MOSFET (VDMOS) and lateral double diffuse MOSFET (LDMOS). The main differences between these technologies are in the design and positioning of the gates, drains, sources and material used in the substrate. MOSFET devices are common and have enabled compact wideband PAs with relatively high power in a monolithic package. VDMOS is a progression from the standard MOSFET that has improved the frequency range and power rating of monolithic PAs, enabling performance up to 1GHz at relatively high powers of 8W – 12W. VDMOS has been on the market since the mid-1980's although recent advances in the technology have improved the gain, power output and efficiency of VDMOS devices.

LDMOS is the newer technology, having emerged in the mid-1990's. It offers improved gain, efficiency and linearity over VDMOS but has not improved the power rating of PAs. However, most significantly to spectrum aggregation LDMOS has higher frequency operation: it can deliver 7-12W monolithic PA's working at 1.5GHz. The cost of VDMOS and LDMOS amplifiers is typically in the region of a few hundred pounds.

Despite the ultra-linear performance of today's MCPA, some transmission spurii may still leak through to the antenna at unacceptably high levels, especially when one considers the large number of tones in a spectrum aggregating RF signal. To combat this one can look to another benefit of the 3G and other cellular industries, that is the advances made in linearisation techniques. These are discussed in the next section.

4.3.9 Linearisation

In recent years there has been a drive to develop compact, wideband and effective linearisation techniques. This is mostly due to wideband multi-carrier problems encountered in 3G mobile telephony, but also in part due to the growing popularity of re-using PAs from one application in another application. Re-use avoids designing and building new PAs for different applications. Some calibration and characterisation work is usually required to adapt the PA's performance to the new application.

All PA linearisation techniques work in the same basic way; the amplitude and phase of the input signal envelope is compared with that of the output, so that differences can be determined and the appropriate corrections made. Most of these techniques concentrate around the fundamental concept of feedback, and (generally speaking) there are two types; closed-loop and open-loop.

Closed-loop feedback includes direct negative feedback and various forms of envelope feedback such as Polar Loop and Cartesian Loop methods, providing high levels of linearisation. Initially one might think of direct negative feedback to be an effective approach. While this works for lower frequency analogue circuits, however, for RF applications there are stability and time causality problems which means other types of feedback must be looked at. But there is another problem that makes closed-loop methods as a whole unsuitable for linearisation; they exhibit very narrow modulation bandwidths. This makes them totally useless for multi-channel (fragmented spectrum) applications, and even for single channel applications closed-loop techniques are severely limited when one considers modern digital modulation schemes, which can have baseband frequencies of the order of MHz. For these reasons, closed-loop linearisation is not suitable for spectrum aggregation purposes.

On the other hand, open-loop feedback can cope with bandwidths wide enough for use in spectrum aggregation applications and has no stability problems. However, it does not have the accuracy of closed-loop feedback methods. The most common open-loop feedback technique is pre-distortion, which is (usually) simple and quite cheap to implement(Figure 4-15).

Both analogue and digital pre-distortion techniques are available, although limitations in DSP performance mean that digital pre-distortion is rare, at least for the moment. There is a third linearisation technique which does not fit into either of the above categories, called feed-forward. In theory it exhibits the precision of closed-loop and the bandwidth of open-loop methods. The drawbacks include inefficient use of power, difficult gain and phase tracking requirements, and (significantly) the potential to add further non-linearities to the desired signal.

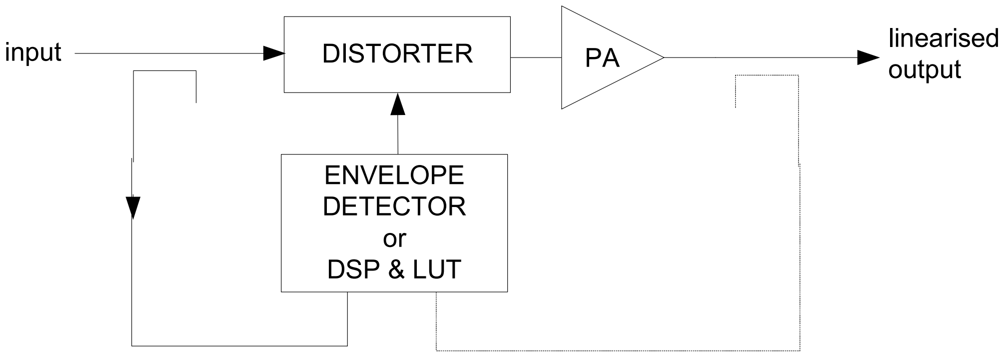

Figure 4-15: Basic Concept of Pre-Distortion

It should be noted that linearisation is not a substitute for good, linear design of amplifiers and component chains. It cannot convert a non-linear system into a linear one, but it can help lower the levels of distortion caused by non-linearities that already exist in components and systems. There are various techniques to choose from, some of which can also be combined to give better results. Linearisation can be performed in the digital or analogue signal domains.

Traditionally linearisation techniques have focused on RF power amplifiers, whose non-linear behaviour can severely limit a system, especially with increased demand for multi-channel applications in modern times. But linearisation can also be used for IF stages, and not necessarily just for amplifiers. Mixers, combiners and even chains of components could have linearisation applied to them. This latter method of linearising multiple components at once is known as system linearisation. It is very difficult to implement successfully and is still an immature development. Nevertheless in a transmit architecture where several chains are sharing a PA and antenna, like that shown below (Figure 4-16), some form of system linearisation technique is likely to be required, even if a highly linear good performance PA is used. This is because it is not enough for the linearisation technique to simply look at the input and output of the PA, because it will not know which tones at the input are wanted fragments and which are unwanted intermodulation products due to the combiner. Therefore some form of control over the linearisation is required, with information regarding the intended signals coming from the fragment planning part of the system (probably implemented in DSP or other processing technology).

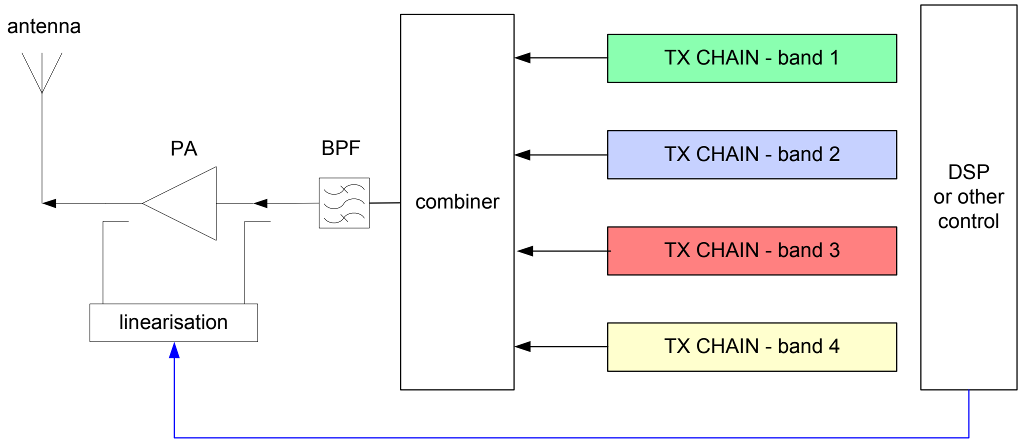

Figure 4-16: Combining transmit chains to share PA and antenna will require a difficult system linearisation technique.

4.3.10 Summary

Digital processing of multiple spectral fragments is superior to analogue components when considering the need for linearity and precision in the mixing and separation of multiple signals. However, digital components, including processors, analogue-to-digital and digital-to-analogue converters, are currently not fast enough to process RF signals over all spectrum fragments of interest. The expectation is that, within a few decades, more powerful digital components will enable an affordable 'software-defined' radio to satisfy most needs in spectrum aggregation.

A spectrum-aggregating radio ideally needs a single antenna capable of receiving and radiating sufficient power over a wider frequency range than is currently achievable. Several novel designs are being developed using MEMS technology and fractal techniques, but it remains a significant challenge to produce a sufficiently wide-band antenna to complement the software-defined radio.

4.4 Waveforms and Channel Coding

4.4.1 Digital Modulation

Digital modulation techniques are replacing or have replaced conventional analogue modulation in a wide range of systems including radio and television broadcasts and cellular communications systems, and have enabled a new range of technologies based on wireless data networks. The advantages of digital modulation include greater noise immunity and robustness to channel impairments, easier multiplexing of various forms of information and multiple users and greater security. Further, the potential to use error detection and correction techniques and signal conditioning and processing techniques such as source coding, encryption and equalisation enable the performance and reliability of the system to be increased. For an aggregating system, careful choice of an appropriate digital modulation scheme is essential to enable the most efficient use of the small spectrum fragments, while preventing interference to the existing band users and offering the flexibility to adapt to the available fragments.

Digital modulation involves representing the information in a baseband data signal as changes in an RF carrier or set of carriers. As with analogue modulation, this can be achieved by shifting the carrier frequency, phase or amplitude, or some combination of these parameters. Key forms of digital modulation include Phase Shift Keying (PSK), Frequency Shift Keying (FSK) and Quadrature Amplitude Modulation (QAM). Modulation schemes with a constant envelope (amplitude) do not require a linear PA and so offer an increase in power efficiency at the cost of lower spectral efficiency. Higher spectral efficiency can be obtained by using a modulation scheme with variable amplitude, but constructing a PA with sufficiently high linearity to prevent a high level of spurious out-of-band emissions increases cost and reduces power efficiency. This trade-off can be compared to the choice of AM or FM in an analogue system. Obtaining a PA which is linear over a wide bandwidth, as would be required in an aggregation enabled system, presents an even greater challenge. It should be noted that pulse shaping techniques, which are normally used on baseband data signals before modulation in order to reduce inter-symbol interference, can remove the constant envelope property.

QinetiQ Proprietary

4.4.2 Common Digital Modulation Schemes

In evaluating digital modulation schemes, the key trade-off is between power efficiency and spectral efficiency. Spectral efficiency is a measure of the data transmission bandwidth available per Hz, and power efficiency is a measure of the SNR required for a given BER, and hence the power required for a certain level of link performance. In general, increasing the spectral efficiency will require a higher SNR to maintain the same BER.

In addition to the intrinsic power efficiency of the modulation scheme determined by the required SNR, power efficiency of the PA is affected by the modulation scheme. Modulation schemes with a variable envelope require linear amplification to prevent spectral spreading. The higher the Peak-to-average power ratio (PAPR) the tougher the requirements for PA linearization. In the following discussion, the power efficiency refers to the SNR required for a given BER, i.e. the intrinsic mathematical efficiency of the modulation.

In Binary Phase Shift Keying (BPSK), the carrier phase is switched between two values (usually separated by 180°), representing binary 0 and 1. This method enables transmission of 1 bit per symbol. The two signals representing the two binary states have the same frequency and amplitude. Quadrature Phase Shift Keying (QPSK) is an extension of BPSK in which there are 4 possible levels represented by 4 values of carrier phase, separated by 90°. Therefore, each QPSK symbol transmits 2 bits. PSK modulation schemes with pulse shaping to prevent inter-symbol interference generally require linear amplification, but offer good spectral efficiency at a cost of low power efficiency. Another similar scheme is Offset QPSK (OQPSK), which is a modified version of QPSK in which the largest instantaneous phase change is limited to 90°, thereby relaxing the requirement on PA linearization to some extent (Rappaport, 2002). PSK can be generalised to M-ary PSK, enabling n-bits to be transmitted per symbol using M=2n different possible phase values. As n is increased, the spectral efficiency increases but the power efficiency decreases since resolving the multiple closely-spaced levels requires a higher SNR.

Modulation schemes exploiting variations in both the carrier amplitude and phase (or frequency) can also be used, for example Quadrature Amplitude Modulation (QAM). QAM is an extension of M-ary PSK, in which the carrier amplitude is allowed to vary in addition to the phase. This allows n-bits to be transmitted per symbol using M=2n different possible combinations of amplitude and phase. As with PSK, increasing the number of levels increases spectral efficiency at the expense of power efficiency (Table 4-1). QAM therefore offers high spectral efficiency at the expense of high power requirements (in order to provide the SNR to required to resolve a large number of levels), but it offers higher power efficiency for a given spectral efficiency compared to PSK. QAM requires linear amplification.

Frequency Shift Keying (FSK) is a constant envelope modulation scheme, and is akin to a digital form of FM. The simplest form of FSK is Binary FSK (BFSK), in which the frequency of the carrier is switched between two different values, representing 0 and 1. The low SNR requirements for a given BER and ability to use a non-linear amplifier make FSK a very power efficient modulation scheme, however it offers poor spectral efficiency. Gaussian Minimum Shift Keying (GMSK) is a variant of FSK which provides enhanced spectral efficiency while maintaining the constant envelope property of FSK. The constant envelope make GMSK ideal for battery limited handheld applications and it is the modulation scheme used in 2G (GSM). M-ary FSK transmits n-bits per symbol using M=2n possible carrier frequencies. The carrier frequencies are often chosen to be orthogonal, by relating their spacing to the symbol rate. As the number of frequencies (M) increases, the spectral efficiency decreases, but the power efficiency increases. The orthogonality characteristic between adjacent carriers at a suitable spacing is exploited in Orthogonal Frequency Division Multiplexing (OFDM), which is described below.

Commonly used M-ary modulation schemes are contrasted in Table 4-1 (Rappaport, 2002). In each case, the spectral efficiency ηB = Rb/B (where Rb is the data rate and B is the first null-to-null bandwidth) is compared to required power, measured as the ratio (dB) of the energy-per-bit to the channel noise spectral density required to give a Bit Error Rate (BER) of 10-6. The signal-to-noise ratio required is given by the product of ηB and Eb/N0.

Since an aggregating system may need to make use of small fragments of spectrum, spectral efficiency is particularly advantageous. Therefore M-PSK or QAM are obvious candidates for this purpose, with QAM offering slightly higher power efficiency. For the aggregating applications currently considered, power consumption is not a critical factor, so transmitter power could potentially be increased to provide a higher SNR for a given noise environment. However, implications of the higher SNR requirements must be taken into account when considering co-existence with the existing band users.

	PSK (VARIABLE ENVELOPE IF PULSE SHAPING)		QAM (VARIABLE ENVELOPE)		COHERENT FSK (CONSTANT ENVELOPE)	
M	η_B	E_b/N_0 (BER=10^{-6})	η_B	E_b/N_0 (BER=10^{-6})	η_B	E_b/N_0 (BER=10^{-6})
2	0.5	10.5			0.4	13.5
4	1	10.5	1	10.5	0.57	10.8
8	1.5	14			0.55	9.3
16	2	18.5	2	15	0.42	8.2
32	2.5	23.4			0.29	7.5
64	3	28.5	3	18.5	0.18	6.9
256			4	24		
1024			5	28		
4096			6	33.5		

Table 4-1: Comparison of spectral efficiency and power requirements of common digital modulation schemes (Rappaport, 2002).

4.4.3 Polar Modulation

Polar modulation is a recently developed modulation technique designed to enable the use of spectrally-efficient modulation schemes with an AM component (such as QAM) while using a power-efficient non-linear PA. The principle is to separate the modulation into an AM component and a constant envelope component, and modulate them separately. The constant envelope component is modulated onto the carrier in the normal way, and this constant envelope signal is fed to a PA which has its gain controlled by the AM component of the original signal. Since the signal at the input to the PA is constant-envelope, a highly power-efficient PA may be used without causing distortion of the signal and spectral spreading, while the addition of the AM component modulating the PA gain enables highly spectrally efficient schemes such as QAM to be generated.

A driver for polar modulation was the need to develop dual mode 2G/3G cellular transceivers, capable of coping with the very different demands of GSM (Global System for Mobile communications – 2G), a constant envelope system, and UMTS (Universal Mobile Telecommunication System – 3G), an envelope-varying system requiring linear amplification (Sowlati et al, 2004). An outline of polar modulation architecture for use in a UMTS device is shown in Figure 4-17 (McCune, 2003).

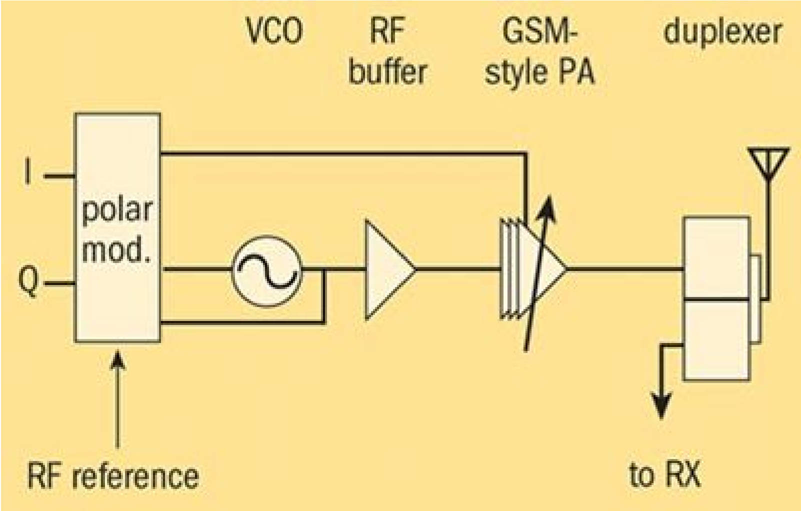

Figure 4-17: Polar modulator for use in UMTS transceiver (McCune, 2003)

A more sophisticated polar modulation architecture includes a feedback loop to combat variations in PA gain due to temperature and aging as well as any non-linearities in the magnitude feed forward path (Ditore, 2003). The polar modulator architecture with feedback is illustrated in Figure 4-18.

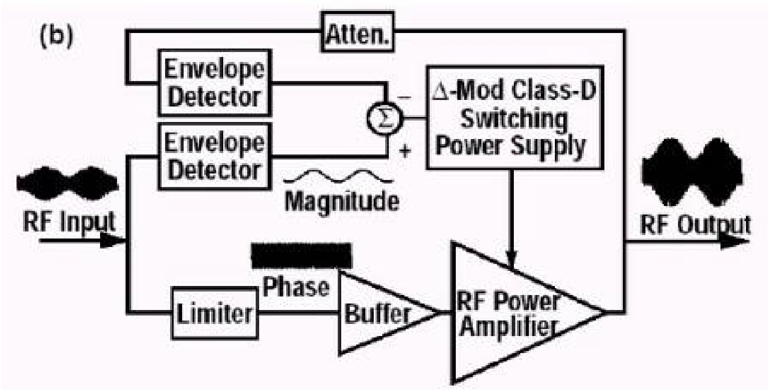

Figure 4-18: Polar modulator featuring feedback loop (Ditore, 2003)

Polar modulation has proved popular for EDGE (Enhanced Data-rates for Global Evolution – 2.5G) systems, and several chipsets are in production for this application.

Although the polar modulation technique has been primarily developed for use in multiple-standard cellular devices where space and power consumption are at a premium, polar modulation could have applications in an aggregating system which is required to operate over a wide bandwidth and may use modulation schemes requiring linear amplification.

4.4.4 Aggregating Fragments and Multiple Access

The problem of allowing multiple users to share the same spectrum can be solved by division of spectrum in the frequency domain (Frequency Division Multiple Access – FDMA), time domain (Time Division Multiple Access – TDMA) or power domain (Code Division Multiple Access – CDMA). In a spectrum aggregation context, the simplest method of allowing multiple access is to aggregate fragments of spectrum into chunks of suitable size for single users, say sufficient to provide 64kbps-1Mbps bandwidth. This is an implicit form of FDMA, so it could be envisaged that an aggregation-enabled system could cover a wide range of frequencies but each user would only be allocated a sufficient number of fragments to fulfil their bandwidth requirements.

Multiple access could alternatively be achieved by time or code separation. TDMA is not thought to be promising for aggregating systems, since it requires synchronisation between all users. TDMA is suitable for a number of users linking to a single base station, but is not suitable for multiple users operating autonomously. CDMA is the other option for allowing multiple access, and this is explored in the discussion of spread spectrum systems in §4.4.6.

The development of FDMA techniques has led to multi-carrier systems such as OFDM, which could prove very well suited to spectrum aggregation due to the potential to cover non-contiguous blocks of spectrum.

4.4.5 Non-Contiguous OFDM

Orthogonal Frequency Division Multiplexing (OFDM) is a multi-carrier system in which the carrier spacing is chosen such that each subcarrier is located at the first null in the spectrum of the adjacent subcarrier, and this orthogonality eliminates inter-carrier interference. OFDM has good spectral efficiency for a large number of carriers, and may be generated fairly easily in the digital domain using an IFFT. Additionally, using suitable guard intervals, OFDM has low Inter-Symbol Interference (ISI) allowing lower complexity receivers to be used. However, drawbacks are high Peak-to-Average Power Ratio (PAPR) and hence requirement for highly linear amplification, requirements for accurate frequency and time synchronisation to retain orthogonality and data capacity loss due to guard intervals to prevent ISI. OFDM has been chosen for ADSL, DAB, DVB and DRM systems.

OFDM could be very suitable for aggregating systems due to the ability to “switch off” unwanted subcarriers, and hence produce a signal with a non-contiguous frequency spectrum which may be tailored to transmitting in available spectrum fragments. This is illustrated in Figure 4-19.

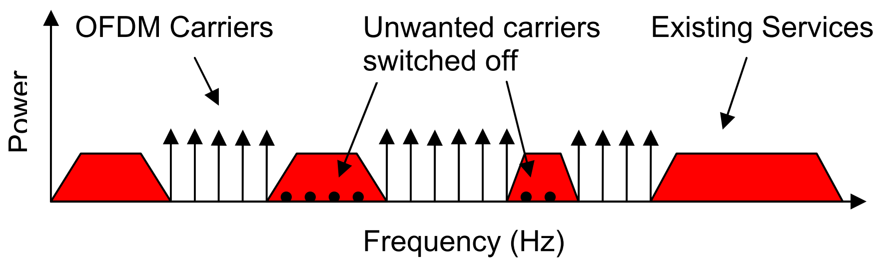

Figure 4-19: Discontiguous OFDM for using fragmented spectrum

In the frequency domain, assuming a rectangular pulse shape, each subcarrier has the form of a sinc function (sin[x]/x) and appropriate choice of subcarrier spacing and symbol rate of the modulation ensures the orthogonality condition is fulfilled. However, although the OFDM subcarriers are orthogonal with one another, orthogonality with the signal from the existing band user cannot be assumed and hence the side-lobes from the subcarriers close to the edge of each OFDM block may interfere with the adjacent existing user.

For a large number of contiguous OFDM subcarriers, the bandwidth of the out-of-band side-lobes represents a small proportion of the total bandwidth occupied by the OFDM system and so guard bands may be used to prevent interference. However, in an aggregating system designed to utilise many small fragments of spectrum, the bandwidth-cost of guard bands becomes high and suppression of side-lobes becomes an important task. In addition to the use of guard bands (which may be created flexibly by switching off a variable number of subcarriers at the edges of the OFDM blocks), additional techniques are required in order to use the bandwidth efficiently while not causing interference to the existing band user. Two techniques for side-lobe suppression are Cancellation Carriers (CCs) and use of signal windowing in the time domain (pulse shaping).

Windowing the signal in the time domain smoothes the transition between symbols and this reduces the out-of-band radiation in the frequency domain, at a cost of increasing the symbol interval and computational complexity (Weiss 2004, Brandes et al, 2005). A common choice of window is the Raised Cosine (RC) window, as illustrated in Figure 4-20 (Weiss, 2004). In Figure 4-20, the total symbol time is the sum of the usable time for data (TU), the prefix time (Tprefix) and the postfix time (Tpostfix). However, since the symbols are allowed to overlap in the roll-off region, the symbol interval (TS) is βTS shorter than the total symbol time, where β is the roll-off factor of the windowing function. The prefix time includes the guard interval which is required to prevent Inter-Symbol Interference (ISI). The postfix time must be longer than βTS in order to preserve orthogonality between subcarriers. This demonstrates the overhead of windowing in terms of reduced data throughput. Windowing in the time domain is relatively easy to implement since it simply requires multiplying the signal by the windowing function. However, windowing does not completely solve the side-lobe problem, as shown in Figure 4-21 (Weiss, 2004). Even at high roll-off factors, only about 6 dB reduction is achieved in the first sidelobe, and so additional side-lobe suppression is required.

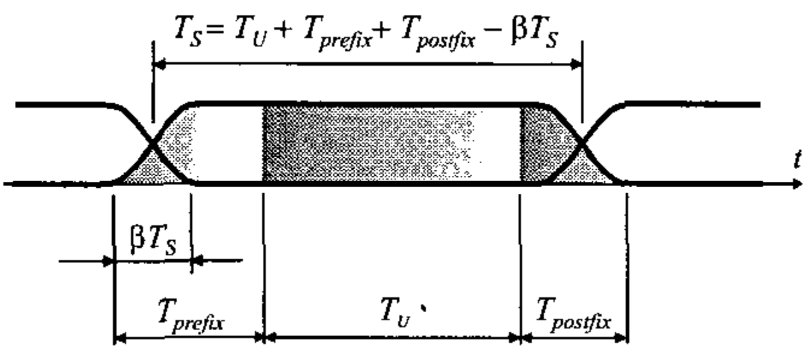

Figure 4-20: The raised cosine window (Weiss 2004)

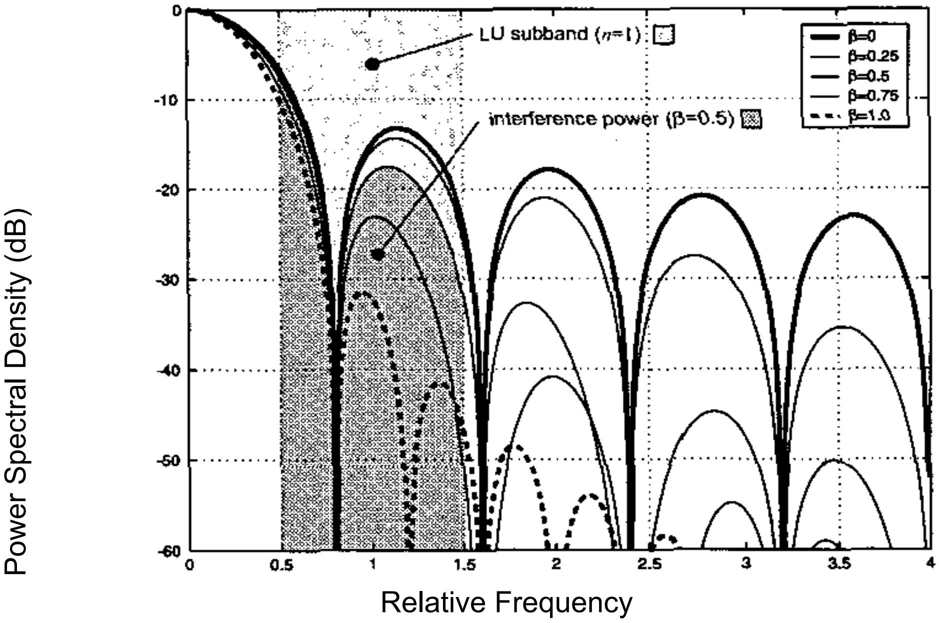

Figure 4-21: Effect on sidelobes of RC window with various roll-off factors (Weiss 2004)

Cancellation Carriers (CCs) are subcarriers added at the edges of the OFDM signal block and modulated in such a way to cancel the side-lobes from the data-transmission subcarriers (Brandes et al, 2005). The cost of adding these subcarriers is a loss of Bit Error Rate (BER) due to the fact that a certain amount of power is required for the CCs and is therefore not available for data transmission, as well as increasing the complexity of the system. However, CCs can achieve significant side-lobe suppression. CCs are implemented after serial-parallel conversion of the data before it is modulated on to the multiple carriers by IFFT (Inverse Fast Fourier Transform) and the guard interval is added.

In order to implement the CCs, the symbol vector consisting of the data symbols and the weighting factors of the CCs is normalised so that the total transmit power is the same with and without CCs. The symbol vector is then modulated onto all subcarriers (data and CC). The weighting factors for the CCs can be determined by a minimisation process, whereby the lowest value of side-lobe power is determined. This optimisation may be set up as a least-squares problem, and a constraint is normally added limiting the maximum fraction of transmitter power which may be used by the CCs.

A simulation has been performed in (Brandes et al, 2005) showing how use of CCs and windowing can enable alternate 25 kHz channels to be used by the existing users and the aggregating system while out-of-band emissions remain within acceptable limits. For the purposes of the simulation, each channel was filled with 12 subcarriers (2.0833 kHz spacing) consisting of 8 data-transmission subcarriers surrounded by 2 CCs each side. The total power in the CCs was constrained to 25% of total transmitter power, and the guard interval was 4% of the symbol interval. In addition to the CCs, an RC window was used for additional side-lobe suppression, with a roll-off factor of $\beta = 0.2$. The modulation used in the simulation was QPSK. Five contiguous 25 kHz channels were chosen for the simulation, and the spurious signals in channels #1, #3 and #5 produced by the aggregating system operating in channels #2 and #4 were calculated. The results are shown in Figure 4-22. Addition of CCs and RC windowing has reduced the average side-lobe power to -54.4 dB in channels #1 and #5 and -50.8 dB in channel #3. The side-lobe power has been reduced by more than 30 dB compared to a conventional OFDM signal.

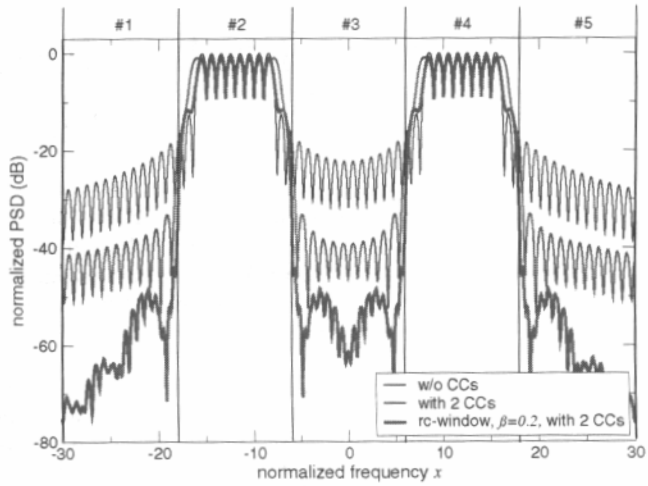

Figure 4-22: PSD of a QPSK OFDM signal, showing addition of 2 CCs per OFDM block edge (max 25% of transmitter power) and RC windowing (β = 0.2) [1]

4.4.6 Spread Spectrum Techniques

Spread spectrum techniques in this context are taken to refer to systems using a Pseudo-Noise (PN) spreading code in order to spread the signal in the frequency domain in some fashion. All the modulation techniques described so far are based on transmitting a signal using the minimum amount of bandwidth possible. Spread spectrum techniques use at least an order of magnitude more bandwidth than required to transmit the data signal, and allow multiple users to access the band simultaneously. In Code Division Multiple Access (CDMA), each user is assigned a unique PN code allowing all users to transmit simultaneously on the same frequency. The advantages of spread spectrum include robustness to narrowband interference and multipath and high scalability, but the disadvantages are low peak data rate and limited capacity due to multiple access interference (Fazel and Kaiser, 2003).

In Direct Sequence Spread Spectrum (DS-SS), the data is spread over a wide bandwidth by mixing with a spreading code before transmission. The spreading code is at the chip rate which is at a considerably higher rate than the data signal rate. To allow multiple access, the spreading codes of the users are chosen to have low cross correlation properties, so that the receiver can “lock on” to the wanted signal and reject the other users signals. Spreading the data over a very wide bandwidth, using a Pseudo-Noise (PN) spreading code increases security of the data being transmitted (since the receiver must know the PN code used) and provides a high tolerance against interference. The Processing Gain, P_G, produced by the spreading process is defined as the ratio of the bandwidth used by the spread signal to the bandwidth of the data signal (i.e. the ratio of the data rate to the chip rate). This gives a measure of the interference resistance – if a system requires a given SNR to perform correctly, S, after spreading this requirement is reduced to $S - P_G$. The longer the PN sequence, the more noise like the signals appear, and the more users which may be supported, however acquisition of the signal becomes more difficult.

In Frequency Hopping Spread Spectrum (FH-SS), the spreading code is used to determine the sequence of frequencies over which the signal is hopped. Each channel has the same bandwidth as the original signal, and the PN codes are chosen to maximise the use of the available bandwidth for all users. As with DS-SS, each user receiver must know that users PN code. FH-SS may be either 'fast' or 'slow' depending on whether less or more than one symbol is transmitted per hop. It is also possible to use a hybrid DS/FH spreading system, enabling very wide spreading using two cheaper, slower code generators.

It is possible to combine multi-carrier and spread spectrum, the two principle systems being Multi-Carrier Code Division Multiple Access (MC-CDMA) and Multi-Carrier Direct-Sequence Code Division Multiple Access (MC-DS-CDMA). MC-CDMA is also referred to as OFDM-CDMA, since it relies on transmitting the chips of the spreading code on multiple carriers, thus giving frequency diversity. User separation is carried out in the code domain, so multiple users can transmit using the same frequencies simultaneously. MC-DS-CDMA relies on transmission of different chips of the spread signal on different carriers, and therefore transmits the data in parallel. This corresponds to time diversity, since the data rate is lower on each subcarrier due to transmission of data in parallel.

Spread spectrum techniques could be used for spectrum aggregation, but it should be noted that they generally only provide an advantage in the case of multiple access scenarios, so for spectrum aggregation this would assume that a sufficient amount of fragments had been aggregated to serve several users. In this case, the trading algorithm could issue each user with a PN code along with the frequency assignment, to prevent mutual interference. FH-SS could prove ideal for this purpose, since the narrowband channels over which the signal is hopped could readily form a non-contiguous set. Clearly, the limit on the bandwidth of each of the channels over which the signal is hopped would be given by fragment size, and this may limit each FH-SS channel to less bandwidth than required by the user. In this case, each user could be allocated more than one channel, but this would then present the same aggregation problem of splitting the data into multiple aggregated channels, and considerably increase hardware complexity.

DS-SS seems less suitable for spectrum aggregation. A spread signal could be transmitted covering a large bandwidth, including occupied channels, at such a low power density that the existing services saw it only as a slight increase in the noise floor. The signal could even be transmitted below the noise floor provided sufficient processing gain and a pilot carrier in an unoccupied fragment to allow the receiver to acquire the signal. Multiple users could share by code division, although care must be taken to ensure the existing user does not suffer interference, since as more users share the frequency the noise floor will progressively rise. Provided the block of bandwidth containing the fragments was sufficiently wide, frequency division could also be used to increase the number of users. However, this is not actually spectrum aggregation, as the fragments of spectrum are only being used for establishing the link, and the extra bandwidth is achieved through the spreading and does not use the available fragments.

The alternative DS-SS mode would be to choose appropriate spreading codes in order to use the available fragments of spectrum, the advantage being different size fragments could be occupied by simply varying the choice of spreading code, as in the Wideband CDMA (W-CDMA) standard used in 3G systems. The advantage of this technique is flexibility to fill arbitrary-sized fragments, but in practice this type of wideband technique is not suitable for the small fragments used by an aggregating system.

4.4.7 Link Budgets

A spectrum aggregation system must co-exist with the existing users of the spectrum without causing interference or itself being adversely affected by interference. The ability to operate without causing interference to the existing band users is determined by the out-of-band emissions of the new system, which are controlled by the waveform in use and the performance of the transmitter PA. If a wideband implementation is used whereby many fragments are aggregated by digitising a large block of spectrum directly, a key factor for withstanding interference from the existing users will be the dynamic range of the ADC. An implementation digitising each fragment separately or using only a contiguous chunk of spectrum would not have such stringent dynamic range requirements, since pre-digitisation filtering could be used to remove unwanted signals.

Some link budget calculations based on example aggregating systems and the existing systems with which they would have to co-exist are detailed in the Appendix B and summarised below. The calculations enable an estimate of dynamic range required for the case of digitising a large bandwidth (including many unwanted signals) and, taken together with out-of-band emission data for both systems, they identify possible inter-system interference issues.

Table 4-2 shows a summary of the link budgets calculations, including indicative dynamic range estimates which do not include the requirements of the modulation scheme in use. Aggregating transmitter power is 125 mW for ‘short’ (30 m) links, 1 W for ‘long’ (1 km or more) links. Aggregating system transmitter and receiver have antennas with 0 dBi gain for <1 GHz and 6 dBi gain >1 GHz. EIRP of unwanted signal is based on a typical existing user at each frequency. Propagation losses are calculated using an exponent for range of 6 for the wanted intra home link (30 m), representing a typical empirical value for worst case non-line-of-sight indoor, and 2.7 for the interferer and long range links, representing a best-case urban scenario.

Frequency	Range of wanted link	Distance to unwanted TX	Wanted Signal Strength	Unwanted Signal Strength	Min Dynamic Range
(MHz)	(m)	(m)	(dBm)	(dBm)	(dB)
173	30	30	-85	-10	75
173	20,000	30	-103	-10	93
460	30	30	-93	-19	74
460	10,000	30	-104	-19	85
650	-	1000	-	-19	-
1400	30	30	-91	1	92
1400	3000	30	-87	1	88
3400	30	30	-98	3	101
3400	1000	30	-82	3	85

Table 4-2: Summary of the co-existence challenges facing an aggregating receiver

The table above shows that a target dynamic range would be around 90-100 dB for detection of the signal, plus the additional requirements imposed by the modulation of choice. Operation along boresight of high-power FWA link transmitters around 1.4 and 3.4 GHz would prove very challenging if attempting to either operate close in frequency or digitise the entire band including wanted and unwanted signals (to

QinetiQ Proprietary

enable multiple fragments per chain). The trading algorithm has the facility to define geographical 'buffer zones' around high-EIRP existing systems when allocating fragments to users and this mechanism could be used to prevent interference from FWA links and TV transmitters.

4.4.8 Out-of-Band Emissions

For PMR transceivers operating in 25 kHz channels, adjacent channel emissions must be reduced to -70 dBc (-60 dBc for 12.5 kHz channelisation) or 0.2 μW, whichever is greater (ETS 300 086, 1991). For PMR transceivers with output power of 25 W, these limits correspond to 2.5 μW (-26 dBm) at 25 kHz and 25 μW (-16 dBm) at 12.5 kHz.

Consider an aggregating system operating in 25 kHz channels, with a PMR system as an existing user. Considering Table 4-2, it can be seen that the adjacent channel unwanted signals at the aggregating receiver are -10 dBm and -19 dBm at 173 MHz and 460 MHz respectively. This means that in the channel in which the aggregating receiver is operating, the signals will be a maximum of -80 dBm at 173 MHz and -89 dBm at 460 MHz, stronger than the wanted signal strengths in Table 4-2. This could therefore place a considerable limitation on the achievable range of the links unless the power of the aggregating system was increased, or guard bands were left between the existing and new systems. Guard bands are extremely undesirable when considering small spectrum fragments, since they reduce the spectral efficiency considerably.

It is assumed that the aggregating transmitter must meet the same adjacent channel power limit as the existing user, in order to comply with standards. Rather than using the figure of -70 dBc as per the PMR specification, a value of 2.5 μW (-26 dBm) is used, as calculated above based on -70 dBc relative to a 25 W PMR transmitter. This allows the aggregating system to meet the specification due to having a lower power output. The non-contiguous OFDM described in §4.4.5 had adjacent channel suppression of about 50 dB for a 25 kHz channels, so this would enable a transmitter power of 24 dBm or 250 mW for the aggregating system per 25 kHz channel (10 W/MHz), whilst still meeting the requirements for an adjacent channel power of no more than -26 dBm. 10 W/MHz is twice the maximum power proposed for the aggregating system.

These calculations show that interference from existing PMR users on adjacent channels to the aggregating receiver would likely pose more of a challenge than meeting the requirements for interference to existing PMR users, provided that the new system is allowed to operate at lower power in order to meet the absolute out-of-band limit (while not meeting the relative limit).

Considering the fragments around 1.4 and 3.4 GHz, the principal users operating at high EIRP in this range are FWA links, so co-existence with these users must be considered. The specifications for FWA links at these frequencies state that they must be capable of accepting interference levels of around -120 dBm (1.4 GHz) / -109 dBm (3.4 GHz) from co-channel interferers and -96 dBm (1.4 GHz) / -74 dBm (3.4 GHz) from adjacent channel interferers, with a wanted signal of around -90 dBm (1.4 GHz) / -70 dBm (3.4 GHz) (OfW 30, 2004; OfW 46, 2004). This gives an idea of the performance requirements for an aggregating receiver operating adjacent in frequency to existing FWA links.

4.4.9 Error-Control Coding

Throughout the discussion of digital modulation and waveforms, the trade-off between spectral efficiency and power efficiency to obtain a sufficiently low BER has been considered. In practice, to enable communications systems to transmit information with a sufficiently low error rate while maintaining feasible bandwidth and power requirements, some form of error control coding is required. The basic principle of error-control coding is to add overhead to the data to enable errors introduced by the propagation channel to be detected (and optionally corrected). The addition of overhead implies a lowering of the available data rate for a given bandwidth, and so channel coding effectively lowers the spectral efficiency in order to reduce the effective BER. Provided the Shannon limit of channel capacity is not exceeded, it is theoretically possible to obtain a coding scheme giving error-free information transfer. Although error-control coding adds significantly to the complexity of the receiver and transmitter, it is necessary in order to achieve sufficiently low error-rates while remaining within bandwidth and power consumption limits.

Channel coding can be broadly divided into two main categories, namely Forward Error Correction (FEC) and Automatic Retransmission Request (ARQ). FEC relies on adding sufficient overhead to the data so that errors can be corrected by the receiver, and can be used on 1-way and 2-way links. ARQ requires only the capability to detect errors causing a repeat request to be issued, but clearly requires a 2-way link.

The main types of FEC are known as block codes and convolution codes. Block codes split the data into blocks of a fixed number of bits, and add a fixed number of parity bits to each block to provide the robustness required. Some examples of block codes are the Hamming code, BCH codes and Reed-Solomon codes. The Hamming code has a block length of 2m - 1, consisting of 2m - m - 1 message bits and m parity bits, where m ≥ 3. The Hamming code can correct one error per block. The BCH codes can correct t errors per block, where the block length is 2m - 1, consisting of 2m - mt - 1 message bits and mt parity bits, where m ≥ 3. Reed-Solomon codes are a class of non-binary codes, that is they operate on m-bit symbols where m ≥ 1 rather than individual bits. A t-error correcting Reed-Solomon code has a block length of 2m – 1 symbols, consisting of 2m - 2t - 1 message symbols and 2t parity symbols (Haykin, 1994).

Block codes operate on the data on a block-by-block basis, therefore they are known as memory-less coding schemes since the code for each block is not dependent on the data which has gone before. Convolution encoders are said to have memory, since the encoder operates on a continuous data stream using a sliding window approach. An advantage of this technique is that the encoder does not need to buffer an entire block at a time during the encoding. One of the highest performance coding techniques is the family of turbo codes, which are used in the 3G standard. Turbo codes are akin to a type of convolution code, but their implementation requires a receiver capable of determining the SNR of the link in real time. Turbo codes have the potential to obtain levels of performance close to the Shannon limit.

The choice of error-control coding will be a key consideration for development of a practical aggregating system in conjunction with choice of modulation scheme. It is not, however, likely to be a limiting factor in determining the performance of an aggregating system, or determining the limits on maximum number of fragments which may be aggregated, etc. The coding scheme chosen will have to be adaptive to allow operation at various bandwidths. Using non-contiguous spectrum does not necessarily pose additional challenges for channel coding, since the coding is applied before modulation, so encoding the data stream will use the same techniques irrespective of the modulation process.

4.4.10 Summary

Channel coding and modulation schemes are needed to provide a reliable data channel over multiple spectral fragments that may differ greatly in terms of propagation delay and interference. Adaptive error control may be needed to cope with RF channels of variable quality, but implementation should have little impact on the radio design, unlike the choice of modulation scheme.

The effect of various modulation schemes has been considered. Constant envelope modulation (e.g. FSK variants) allows the use of an efficient non-linear power amplifier, whereas more spectrally efficient schemes (e.g. QAM variants) need highly linear amplifiers that are much less power efficient. The latest techniques, in particular polar modulation, should enable the use of the most spectrally efficient schemes while minimising distortion when amplifying multiple signals.

A prime candidate waveform for fragmented spectrum is Orthogonal Frequency Division Multiplex, which is now highly developed for digital radio and TV consumer products. The challenge is to develop a variant whereby individual carriers in the multi-carrier waveform can be sufficiently suppressed outside the available spectral fragments. An analysis indicates that the main problem is in suppressing interference from services operating in spectrum adjacent to spectral fragments.

4.5 Proposed Example Aggregator

4.5.1 Receiver

As mentioned in §4.3.2, fragments must be individually filtered and separated from each other in the digital domain, thus requiring the entire band to be processed through the analogue and ADC parts of the chain beforehand. Suitable ADCs can sample signals of up to 50MHz width, and this is the most band-limiting device of the entire receive chain. It is extremely difficult to adequately filter alias, image and other unwanted frequency bands at this low frequency, in no small part because of the need to have a band-pass filter to avoid DC offset problems at the input to the ADC. Therefore in practice, to maximise the ADC's sampling capability, under-sampling at an IF that is higher than 50MHz is recommended. Choosing an IF higher than the ADC's sampling rate relaxes the requirements on filtering and enables better rejection performance at both IF and RF stages. Certain conditions must be met in order for under-sampling to work, putting some strict limitations on the choice of IF, signal bandwidth and sampling rate. However it has been checked that an ADC sampling at 125Msps could successfully digitise a bandwidth of 50MHz centred around a 100MHz IF.

Converting the low IF to a higher RF will achieved by using a superheterodyne architecture for each chain. It is feasible (for a certain range of frequencies within the limits of filters and LO technology) that a variable LO could be used to tune the band somewhat, offering some flexibility. This is desirable in practice because at higher carrier frequencies, 50MHz bandwidth is relatively narrow and a large number of bands would be needed to cover the spectrum. As a compromise, fewer bands could be used but offer some tuneable range to concentrate on clusters of fragments, see Figure 4-23 below. Rather than full fragment coverage, most fragments could be covered with a saving in the number of hardware chains needed. This does, of course, rely of the nature of fragmentation; the more evenly-spaced the fragments are, the less advantage there is to using tuneable filters. If all fragments were evenly spaced then no advantage would be gained.

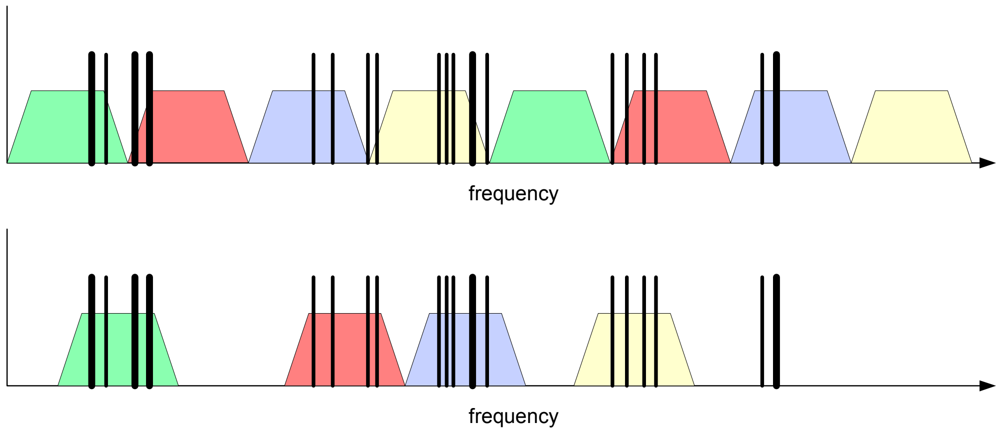

Figure 4-23: Many fixed bands of 50MHz width are needed to capture all the fragments in a broad spectrum (top picture). Fewer tuneable bands could be used to focus on clusters of fragments, at the cost of missing out on other more scattered fragments (bottom picture).

Because of the static nature of the spectrum aggregation application, it is feasible to have changeable plug-in filters and a simple reprogramming of a variable LO to tune a chain to a certain band. It is important to note that tuning the bands to focus on clusters of fragments could not be made adaptive at this point in time, because of the unrealistic demands that would be incurred on the RF filters.

It is feasible to share one wideband antenna and use a divider/combiner to separate out bands into individual chains, with one overall wideband LNA prior to the divider/combiner. Because of the current limitations of divider/combiners, the lowest and highest carrier frequencies that the receiver could operate at would be 500MHz and 2GHz respectively. To operate outside this range, separate chains with their own antennas would have to be used. The carrier limit then becomes ~6GHz, as dictated by LNA limitations, although specialist narrowband antennas would be needed which might also restrict chain bandwidths.

Components have a good enough linear performance to allow several fragments to be shared by each chain, which would be subsequently separated using digital filtering. Thus the number of fragments possible is greater than the number of chains. Of course, because each chain is fairly limited in bandwidth (50MHz) then fragments must be clustered together in order for chain sharing to occur. Digital filtering will easily allow any fragment size between 25KHz and 1MHz, and there is no restriction on the variance of fragment size (i.e. fragments do not have to be all the same size).

An overall view of the kind of spectrum aggregating receiver that could be built today is shown in Figure 4-24.

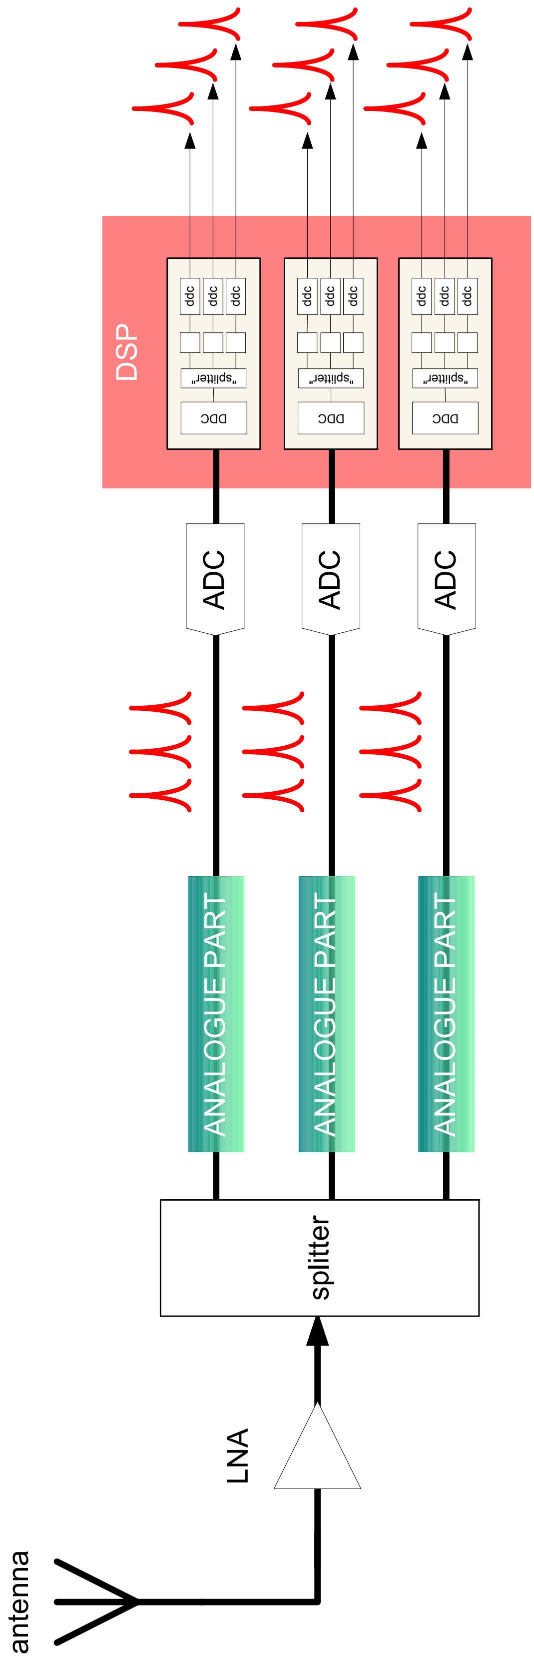

Figure 4-24: A possible Spectrum Aggregation Receiver

QinetiQ Proprietary

4.5.2 Transmitter

Whereas ADCs are the band-limiting device for receive chains, today's DAC speeds are not the limiting factor for transmit chains. Instead, it is the clock speed of DSP that dictates the restriction. As mentioned in §4.3.1, the fastest DSP speeds available today are ~500MHz, resulting in a maximum signal bandwidth of 250MHz that can be processed by DUCs. However, in practice the centre IF must be greater than half the IF bandwidth. This is because of the imperfect roll-off response of filters used further down the chain to suppress image frequencies. Unlike the receiver design, this problem is not caused by the reconstruction filters at IF, since a low-pass filter can be used. The limiting factor is the RF band-pass filter image rejection. The greater the ratio of centre IF to IF bandwidth, the farther removed is the image from the wanted signal, thus the easier it is to filter. This is less of a limit for image filters in a transmit chain than it is for anti-alias filters in a receiver chain. To get good image rejection using today's filter technology, a realistic maximum IF bandwidth of 200MHz centred around 150MHz is recommended (i.e. 50 – 250MHz IF).Similar superheterodyne architectures as used for receive chains can be used to equal effect for implementing transmit chains, converting the low IF to much higher RF bands. As with receive chains, it is feasible to have some tuneable ability by varying the LO drive, but because of the wider chain bandwidth the filter requirements are stricter and thus the range of tuning is likely to be reduced. But for the very same reason (wider chains) it may not be that important to have tuneable chains in a transmitter system.

The carrier-limiting factor for transmit chains is the PA. Up to 1.5GHz carrier RF is possible at adequate power using today's LDMOS technology. Higher frequency operation will require bulkier and less efficient PA technologies.

Antenna sharing is not feasible for transmitter systems, for two main reasons. First, it has been shown that transmit antennas are not as wideband as receive antennas, and to cover the range 100MHz – 1.5GHz would require at least three different antennas. Therefore, even with antenna sharing, multiple antennas cannot be avoided. The second reason is more significant: adequate linear performance cannot be achieved using today's PAs and linearisation techniques if multiple transmit chains were to be combined at their outputs prior to the antenna. It demands system linearisation techniques which are extremely difficult, if not impossible, to implement at the moment.

The inability to share the antenna or the PA means that each transmit chain must contain its own antennas and PA. Therefore the number of chains is quite limited, in order to keep the overall system component count below an acceptable limit and to mitigate mutual coupling effects between collocated antennas. However, it is still possible to share several fragments per chain because there are high-performance PAs and single-component linearisation techniques adequate to do this without introducing unacceptable levels of intermodulation distortion. Exactly how many fragments per chain will depend on individual components and techniques used, but it is likely to be a low number, say two or three fragment per chain.

An overall view of the kind of spectrum aggregating transmitter that could be built today is shown in Figure 4-25.

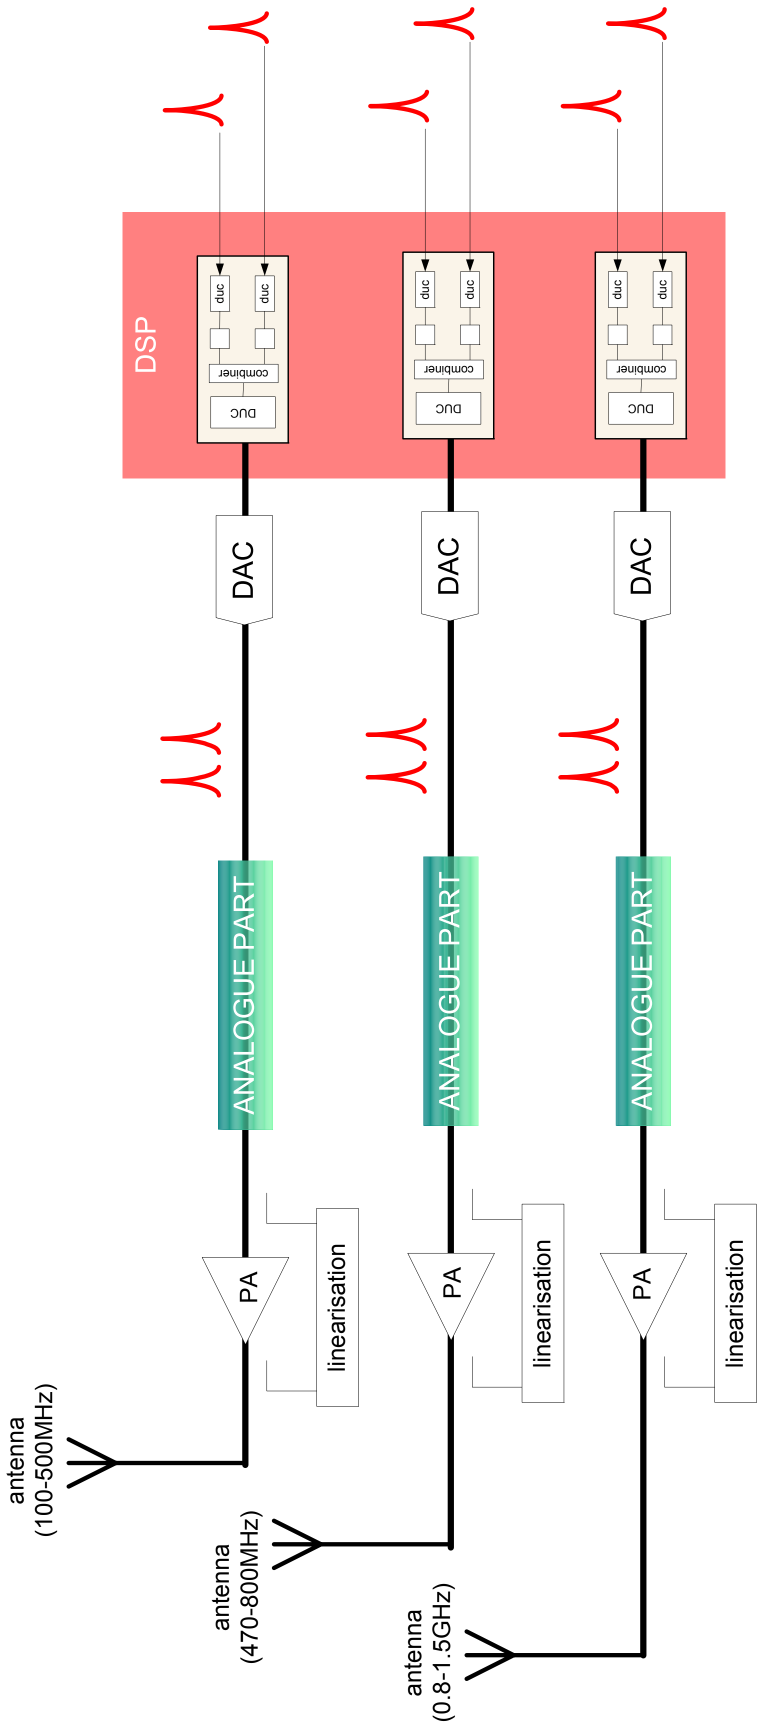

Figure 4-25: A possible spectrum aggregator transmitter

4.5.3 Cost

The prices quoted in this chapter for various devices are typical prices for one-off or low production scales. These are the kinds of prices that would be paid to build one or two prototype spectrum aggregators systems. Based on QinetiQ's previous experience in building prototype RF systems, it is estimated that a single prototype receiver and transmitter pair would cost in the region of £15K - £20K in parts. The labour cost required for detailed design, build and test of a prototype might be between £120K - £180K, but this would be a one-off initial cost as subsequent system builds would only cost parts and build, not design.

For mass production there would be no design costs and the cost per component would be dramatically reduced. It is estimated that the cost per aggregator could fall to ~£1,000 for mass production. But if such devices proved popular on a national or international scale, the cost could fall even lower, perhaps to that comparable to a satellite TV receiver box found in the home.

The theoretical costs of developing a specific aggregation system are considered in more detail in the next section.

4.5.4 Summary

A spectrum-aggregating radio has been outlined which would consist of several separate transmitter and receiver chains, each able to exploit several spectral fragments within a bandwidth of about 50MHz. The cost of such a radio, although expensive to prototype, should fall dramatically as the design is refined and brought into mass production.

4.6 Conclusion

Using technology that is available today, is it possible to aggregate fragments over a limited number of bands, each band being at most 50MHz wide. This 50MHz limit is imposed by the sampling speeds of today's analogue-to-digital converters. As sampling speeds increase this band limit will widen. The centre frequencies of the bands can be anything from 100MHz up to 1.5GHz. The performance of power amplifiers is limited beyond 1.5GHz, where high transmit power means impractical size, weight and power supply requirements. It is possible to have tuneable bands in a single aggregating device using the superheterodyne principle.

There is no limit on the width of fragments, other than our arbitrary boundaries of 25kHz minimum and 1MHz maximum.

Between five or ten fragments could be aggregated in one go using a single device, without that device becoming too bulky, inefficient or costly. A system aggregating a higher number of fragments will probably have poorer performance than one that is only aggregating a few fragments, because of the need to share fragments over chains. If only one fragment per chain is aggregated, performance will be superior; around five chains is a practical limit before a system becomes too bulky and expensive.

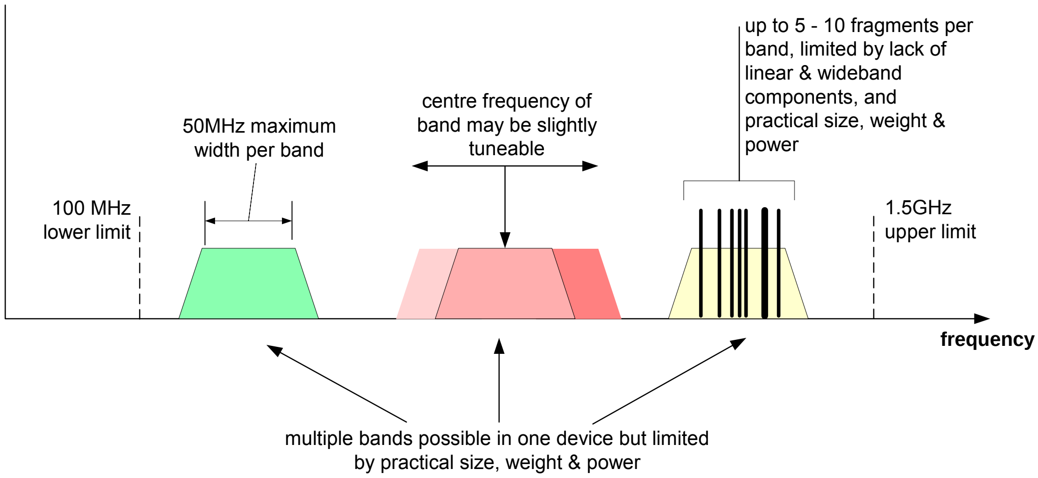

Figure 4-26: What a spectrum aggregator built today might look like

Key considerations for waveforms are spectral efficiency, flexibility to generate waveforms which can adapt to the available fragments, hardware requirements (e.g. PA linearity required determined by peak-to-average power ratio) and co-existence with the existing band users. Since spectral efficiency is more important than power efficiency for the fixed systems considered at this stage, multi-level PSK or QAM seem ideal modulation schemes for an aggregating system. Although spread spectrum techniques have some advantages for flexible use of spectrum, a multi-carrier system without use of a spreading code is considered most suitable for aggregation of narrowband fragments, and a flexible non-contiguous OFDM scheme is proposed. With suitable measures for side-lobe suppression this could meet the criterion for co-existence with existing band users. Calculations indicate that interference to the aggregating system from higher power existing band users would be more likely than the reverse, and a trading algorithm could define geographical 'buffer zones' around high EIRP existing users to prevent interference to the aggregating system. A disadvantage of OFDM is the high PAPR and therefore PA linearity requirements, but this could be mitigated by the use of polar modulation. The aggregating system would also need an adaptive channel coding system to provide the required error detection and correction.

If such a spectrum aggregating device was built today, it would need to implement state-of-the-art DSP, ADC, DAC and amplifier components. This would make it more expensive than perhaps it should be, for example if a similar-performing system was built in five or ten years' time, the cost should have fallen.

Although a spectrum aggregator could be built today, there are three significant issues that need addressing. These are

- large, bulky antennas at low/wide frequencies
- band-limiting speeds of ADCs, DACs and DSP
- intermodulation distortion resulting in poor linearity

While these issues many not be immediately resolved, technological developments are underway that will help to mitigate and eventually eliminate them. Fractal antennas are rapidly emerging as multi-band, compact alternatives to bulky antennas. Looking further into the future, plasma antennas will also significantly reduce the size of antennas and offer the benefit of true wideband operation.

Clock speeds of DSPs and the sampling speeds of ADCs and DACs are ever increasing, and if the trend of the past 20 years continues then digitisable bandwidth is expected to triple every 6 – 8 years. Thus, over the next two to three decades, most of the frequency coverage discussed in this chapter will be directly digitised and no longer need to be converted using analogue components beforehand. As well as expanding the band of each chain, this will result in smaller and cheaper aggregator devices.

The single most significant challenge to realising a spectrum aggregator is the mitigation of intermodulation distortion (IMD), especially if fragments share a chain, or chains need combining to share an antenna or amplifier, to reduce component count and overall size. Amplifiers and mixers are evolving to perform more linearly to combat IMD. Unfortunately there is little evidence to suggest combiner/divider technology is developing in the same manner. Instead, architectures employing combiners or dividers will have to rely on system linearisation techniques, which at present are immature but are under development.

However, in the long term there is another solution to preventing IMD. When ADCs, DACs and digital processing becomes fast enough to directly digitise RF, there will no longer be a need to use analogue frequency conversion. This will enable fragments and chains to be combined digitally. In the digital domain, IMD can be eliminated because signals can be copied and combined without the need for impedance matching (in effect, digital functional blocks have perfect matching properties). In the analogue domain a simple filter, LNA or PA, and a wideband antenna will still be needed. Of these components only the PA has the potential to create significant IMD, which could then be overcome by single component linearisation techniques (if needed at all).

4.7 Recommendation

If fragmentation occurs Ofcom or its licensees should endeavour to make sure that the fragments are not evenly distributed across the whole spectrum, i.e. clusters of fragments around the band are preferable to widely-spread fragment to facilitate the development of any potential aggregating systems.

Ofcom should highlight any occurrences of spectrum fragmentation so that hardware suppliers can develop specific devices that can facilitate aggregation such as splitters, wideband amplifiers, DSPs etc.

4.8 References

[1] Brandes, S, I Cosovic and M Schnell, 'Reduction of Out-of Band Radiation in OFDM Based Overlay Systems', Proc. IEEE DySpAN 2005 Conference, Baltimore, MD, USA, November 2005.

[2] Ditore, F, 'Polar Modulation Ups Efficiency in Mobile PA Designs' CommsDesign, 4 Sep 2003, available at:

http://www.commsdesign.com/story/OEG20030904S0013

[3] ETSI Standard, ETS 300 086, 1991.

[4] Fazel K and S Kaiser, 'Multi-Carrier and Spread Spectrum Systems', John Wiley & Sons, 2003.

[5] Haykin, S, 'Communication Systems', 3rd Ed., John Wiley & Sons, 1994.

[6] McCune, E, 'Polar modulation resolves multimode-mobile problems', Institute of Physics Wireless Web, Feb 2003, available at:

http://wireless.iop.org/articles/feature/4/2/2/1

[7] OfW 30, 'Fixed Point-to-Point Radio Services with Digital Modulation Operating in the Frequency Ranges 3.6 to 3.875 GHz paired with 3.925 to 4.2 GHz', Ofcom Technical Frequency Assignment Criteria, Sep 2004, available at:

http://www.ofcom.org.uk/radiocomms/ifi/tech/tfacs/tfac_ofw30.pdf

[8] OfW 46, 'Fixed Point-to-Point Radio Services with Digital Modulation operating in the Frequency Ranges 1350 to 1375 MHz paired with 1492 to 1517 MHz', Ofcom Technical Frequency Assignment Criteria, Sep 2004, available at:

http://www.ofcom.org.uk/radiocomms/ifi/tech/tfacs/tfac_ofw46.pdf

[9] Rappaport, T S, 'Wireless Communications Principles and Practice ', 2nd Ed., Prentice Hall, 2002.

[10] Sowlati, T et al, 'Quad-Band GSM/GPRS/EDGE Polar Loop Transmitter', IEEE Journal of Solid-State Circuits, Vol. 39, No. 12, Dec 2004, pp 2179-2189

[11] Weiss, T, A Krohn, J Hillenbrand, and F Jondral, 'Mutual Interference in OFDM-based Spectrum Pooling Systems', Proc. of the Vehicular Technology Conference VTC2004-Spring Milan, Italy, May 17-19, 2004.

5 Fragmented spectrum scenarios

5.1 Introduction

If spectrum fragmentation occurs we have shown that there are no significant technical barriers to aggregating spectrum to deliver a service. The question arises, however, on the detailed costs between a system that employed aggregation with one that did not and would an aggregating solution been viable. There are many variables to consider for this latter question; market size, technology advances matching the market need etc.

In this section we will look at the additional hardware cost incurred in building a system capable of spectrum aggregation, with respect to a non-aggregating system capable of the same level of performance and service. This extra cost only relates to the additional hardware components required to build each spectrum-aggregating transceiver unit. Other hardware costs, such as casing, masts, circuit board, controls, power supply, batteries, etc. are not considered as it is assumed these would be similar regardless of whether the transceiver was conventional or aggregating.

This section looks at the component cost of existing non-aggregating transceivers, by considering the number and type of key components required in a generic non-aggregating radio system architecture.

The component cost of a theoretical spectrum-aggregating transceiver will then be estimated in the same way, by considering the additional and different components required in order to achieve the same level of performance as the existing non-aggregating transceiver. This can be repeated for a varying number of spectrum fragments. The data are then used to perform a basic analysis to indicate if a spectrum aggregating solution would have been viable.

Devices from four applications are considered:

- Satellite communication (SATCOM)
- Public Mobile Radio (PMR)
- Wireless local area networks (WLAN)
- Microwave fixed link.

Two forms of spectrum aggregation are considered. The first is when aggregation uses fragments within the current spectrum band. The second is when aggregation uses spectrum in the adjacent or near by band.

It should be noted that the designs and costs were developed to inform this specific study and should not be used for any other purpose.

5.2 Satellite Communication Terminal

In this section we will look at a UHF SATCOM voice terminal (Table 5-1). The current, conventional SATCOM design does not allow for multiple channel use as a spectrum-aggregating system would (for example two fragments of 12.5kHz each). There are, then, 1200 channels over the 30MHz band but only one can be selected at a time.

UHF SATCOM is used primarily by the military for two-way voice communications. The number of users is therefore limited to several hundred. However, similar principles can be applied to other SATCOM services which deliver commercial satellite voice communications, of which there are estimated to be tens of thousands of users in the European region alone. For example, the world merchant navy is estimated at around 50,000 vessels, of which one third (in tonnage) is owned by European countries [1]. There are many more smaller vessels, such as fishing boats and pleasure/private craft; the size of the European fishing fleet in 2000 was estimated at over 90,000 vessels [2].

Transmit Band	290 – 320 MHz
Receive band	240 – 270 MHz
T/R mode	Full duplex (FDD)
Channel size	25 kHz
EIRP	40dBm (10 Watts) minimum

Table 5-1: Basic specification of a conventional UHF SATCOM voice terminal.

5.2.1 Conventional Design

The outline design for this SATCOM terminal is given in Figure 5-1. It consists of five main parts:

- Receiver Chain
- Transmit Chain
- Transmit/Receive Switch
- Phase Locked Loop (PLL) Stage
- Analogue-to-Digital and Digital-to-Analogue Conversion

From a design point of view it is useful to split the design into five stages, namely: RF, Conversion, IF, AF and PLL.

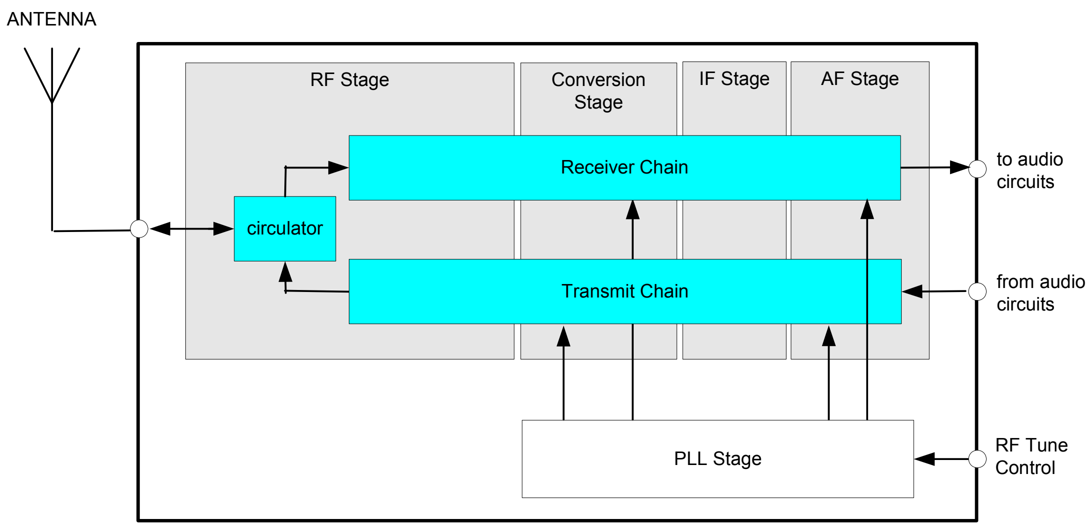

Figure 5-1: Outline design for UHF SATCOM terminal

The receiver chain contains the analogue components such as low noise amplifiers, filters and mixers required to boost the weak received signal and down-convert this to an IF suitable for further analogue signal conditioning and final down-conversion to audio band. Conversely, the transmit chain contains similar components such as mixers and filters, and higher power amplifiers, for converting the AF to IF, then IF to RF, with enough gain in each stage to achieve the required EIRP.

The receive and transmit chains stretch across the RF, conversion, IF and AF stages, and are driven by local oscillators (LO) sourced from the PLL stage.

Because the mode of operation is frequency division duplex (FDD), a circulator (also known as an orthomode transducer in SATCOM engineering) is used to make sure that the transmit signals go out through the antenna and not back into the receive chain. It also ensures that the receive signals coming from the antenna go into the receive chain and not the transmit chain.

More detailed RF architecture showing individual RF components of the two chains that lie within the RF, conversion, IF and AF stages are detailed in Appendix C.

5.2.2 Spectrum-Aggregating Design

To turn the SATCOM terminal into a spectrum aggregating device, multiple receive and transmit chains are needed for each fragment. This could result in the component count escalating dramatically, save for the fact that the channel bandwidth is narrow at 25kHz so it is unlikely that it would be split into more than five fragments (of 5kHz each, perhaps).

As the conventional design stands, only the IF filters need to have their pass-band narrowed in order for the terminal to operate on a single fragment. Then it is a simple matter of adding more chains, one pair (transmit and receive) for each additional fragment.

However, the entire chain does not have to be replicated. It is feasible for all of the components in the RF stage, and some in the conversion stage, to be shared by all fragments, since these components are already wideband enough to encompass the entire tuneable band. This assumes that all usable fragments will lie somewhere within the band, which is a realistic assumption for SATCOM as using out-of-band fragments would mean changing the space element (i.e. launching new satellites). In the receiver chain the preselector, LNA and image rejection filter can be shared. In the transmit chain, the harmonics suppression filter can be shared. The IMP2 suppression filter could also be shared, but the pre-amp and PA cannot, and it is not worth combining transmit chains together to share the filter just have them split back out to separate amplifiers, then combined back to share the harmonics suppression filter. In the RF stage the circulator can also be shared, and so can the antenna.

The rest of the components in the RF, conversion and IF stages will need to be replicated, one replication per additional chain/fragment. However, the AF stage can be replaced entirely by analogue-to-digital and digital-to-analogue conversion, since combining of the fragments into a data stream will require digital signal processing. This means that only one LO drive is required per chain, i.e. one dual synthesiser per transmit/receiver chain pair. Thus some saving can be made on AF and LO components.

Figure 5-2 shows an overview of design of a SATCOM terminal aggregating two fragments, both in receive and transmit modes.

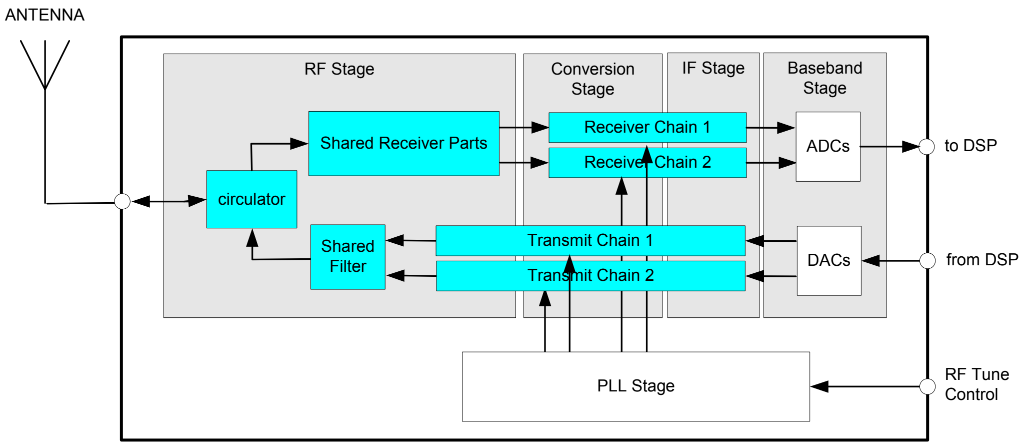

Figure 5-2: Outline design for spectrum-aggregating SATCOM terminal

The detailed spectrum aggregating design and costs are detailed in Appendix C. The results are summarised in Table 5-2 below which shows that as the number of fragments increases from the baseline (1 fragment) to two fragments the cost of the hardware increases by ~70%.

System type	**Price, £**	**extra cost %**	**Difference**
Conventional terminal	**2314**	**0**	
2-fragment aggregator	**4006**	**73**	**73%** more than conventional system
3-fragment aggregator	**4713**	**104**	**31%** more than 2-fragment system
4-fragment aggregator	**4998**	**116**	**12%** more than 3-fragment system
5-fragment aggregator	**5721**	**147**	**31%** more than 4-fragment system
6-fragment aggregator	**6317**	**173**	**26%** more than 5-fragment system
7-fragment aggregator	**7174**	**210**	**37%** more than 6-fragment system
8-fragment aggregator	**7808**	**237**	**27%** more than 7-fragment system
9-fragment aggregator	**8493**	**267**	**30%** more than 8-fragment system
10-fragment aggregator	**9318**	**303**	**36%** more than 9-fragment system

Table 5-2: Cost of spectrum aggregating devices, up to 10 fragments in-band

The major contributor to the cost increase is DSP which makes up approximately 50% of the additional cost. DSP is required to aggregate, but the amount of DSP cannot be correlated to the number of fragments: this is because the DSP workload depends on numerous factors, some related and some unrelated to fragments. Thus the significant DSP cost may only be encountered in the transition from conventional non-aggregating design to an aggregating design, regardless of the number of fragments. On the other hand, there will eventually be a limit on the number of fragments any single DSP can deal with and additional DSP – along with additional large cost – will be needed. The threshold number of fragments at which extra DSP would be needed cannot be estimated without a detailed design of a particular device. Therefore, in all the examples shown here, it has been assumed that only one DSP device is required, and the cost of this is included in the 2-fragment design: i.e in the transition between a conventional and an aggregating device.

Based on experience it should be noted that the development costs of hardware may only comprise some 3-5% of the total development costs of the system.

5.2.3 Additional Band Design

So far we have only considered adding fragments that lie within the original band of the SATCOM terminal (290-320MHz or 240-270MHz). To use fragments that are outside of this band, components in the RF stage need to be replaced with wider band parts or be duplicated to allow more than one band at a time. Duplication is the cheaper of the two, since wideband components are rare (and hence expensive), and this approach also allows for more flexibility (it may be possible to tune the bands, or at the very least have two bands that are not necessarily neighbouring each other in the spectrum).

Figure 5-3 shows the modified system overview required to enable fragment across multiple bands (in this case two bands). There are still some components in the RF stage that can be shared, namely the preselector and LNA on the receiver side (these are readily available in wideband versions) and the harmonics suppression filter on the transmit side. In fact, no changes need be made on the transmit components within the RF stage, as long as a combiner can be found that is wideband enough to accommodate all the transmit fragments.

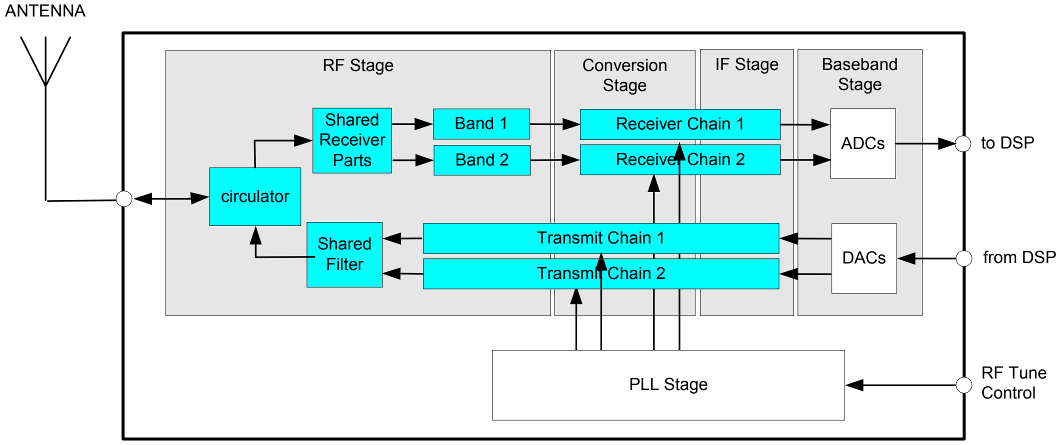

Figure 5-3: Outline design for spectrum-aggregating SATCOM terminal across bands

The costs are detailed in Appendix C and they show shows that dividing fragments across different bands does cost extra (when compared to aggregating fragments in the same band), but not very much. The additional cost is £177.66; only 7.7% of the original non-aggregating design cost. The dual-band five-fragment design incurs a similar 8.4% increase. Compared to their 4- and 5-fragment single-band counterparts, the extra cost is only 3.5%. This is to be expected since the majority of component duplication is caused by adding a fragment to the system, requiring extra components in the conversion, IF and PLL stages, not the RF stage. Because most of this duplication occurs in the conversion, IF and PLL stages, the value of the carrier frequency – and hence the upper and lower limits of the bands – will have little impact on the cost of each fragment. Also, the size of a fragment will not affect the hardware cost, since final fragment size is determined by digital filtering within the DSP which can be changed via software. Remember, we have assumed a single DSP is required for any number of fragments up to 10 and that subsequent DSP is not required. The same assumption has been made for the dual-band examples.

It is concluded, therefore, that the size of fragments and how they are distributed across multiple bands is not as significant as the total number of fragments and the total number of bands.

System type	Price, £	extra cost %	Difference
Dual-band 4-fragment aggregator	5176	156	3.5% more than 4-fragment single band system
Dual-band 5-fragment aggregator	5915	156	3.4% more than 5-fragment single band system

Table 5-3: Summary of SATCOM terminal costs for differing bands

5.2.4 Summary of SATCOM costing

A summary of the cost increases involved in building a system aggregator is given in Table 5-3. We can see that building a 2-fragment aggregator is dominated by the DSP costs. Aggregating more fragments is better value for money (in terms of the cost per fragment). The trend in cost per fragment is summarised in Figure 5-4 excluding the DSP elements.

Exploiting fragments across multiple bands increases the cost, although not significantly. The way in which fragments are distributed within one or more bands, and the size of bands as well as individual fragments, does not affect the cost.

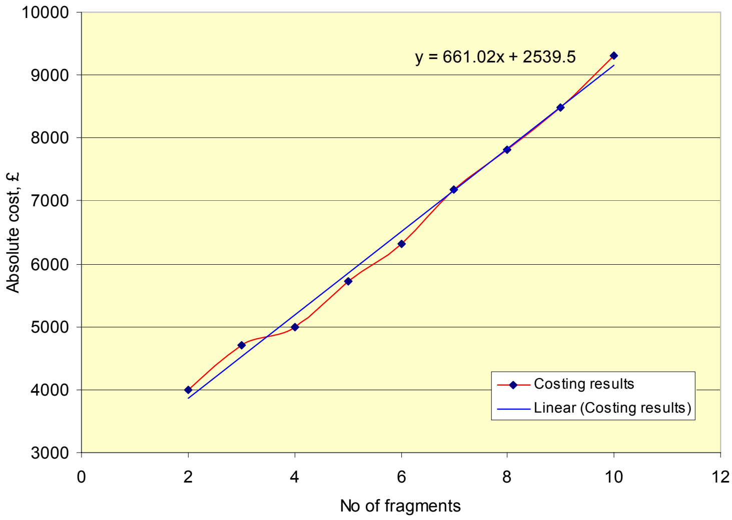

Figure 5-4: Trend of cost against number of fragments (discounting DSPs) for an in-band system.

5.3 Private Business Radio (PBR) Transceiver

In this section we will look at a VHF hand-portable or mobile PBR voice terminal with some limited data handling such as scrambling, 2/5 tone calling, and dual tone multi frequency (DTMF) and trunk networking capability.

Some basic specifications of the transceiver are given in Table 5-4. The current, conventional design does not allow for multiple channel use as a spectrum-aggregating system would. There are 48, 59 or 96 channels available (depending on the channel spacing setting) over the transceiver's band but only one can be selected at a time. Furthermore, not all devices will have the capability to select from the entire range of channels: for example some basic hand portables only offer a choice of 4, 12 or 16 channels.

PBR is used by small to medium sized private, local businesses, councils and community groups who require two-way voice communications. The limited range of PBR means many user groups can re-use the same frequencies across the UK.

There are approximately 3000 user groups in the UK, and it is estimated that between them around 800,000 individual devices are used [3]. Typical user groups include taxi firms, construction firms, national couriers and breakdown organisations.

Transmit Band	146 – 174 MHz
Receive band	146 – 174 MHz
T/R mode	Half-duplex (simplex PTT operation)
Channel size	12.5KHz, 20KHz or 25 kHz
EIRP	37dBm (5 Watts) minimum 44dBm (25W) maximum

Table 5-4: Basic characteristics of a PBR

5.3.1 Conventional Design

The outline design for a common hand-portable transceiver is given in Appendix C. All PBR radios will constitute this basic design, with more advanced devices having additional hardware for various features such as trunk networking, or simply better quality parts for improved robustness and performance.

VHF is a low enough frequency for PBR radios to be made of discrete electronic components such as transistors, inductors, capacitors as well as some integrated circuits, as opposed to dedicated RF analogue components. Thus PBR design is split into a number of circuit blocks, categorised into three groups, as follows:

- Receiver Circuits:-
    - Antenna Switching Circuit
    - RF Circuit
    - 1st Mixer & IF Circuit
    - 2nd IF & Demodulator Circuit
    - AF Amplifier Circuit
    - Squelch Circuit
- Transmitter Circuits:-
    - Microphone Amplifier
    - Modulation Circuit
    - Drive Amplifier Circuit
    - RF Power Amplifier Circuit
    - APC Circuit
- PLL Circuits:-
    - Reference Oscillator Circuit
    - Programmable Divider & Phase Detection Circuits

- Charge Pump & Loop Filter Circuits
- VCO Circuit

From a design point of view it is useful to split the design into four units: RF unit, Main unit, VCO board, Logic unit

The design of the conventional PBR system is detailed in Appendix C which predicts the hardware costs for such as system to be £266.

5.3.2 Spectrum-Aggregating Design

To turn a PBR hand-portable into a spectrum aggregating device, additional receive and transmit circuit blocks are required, just as additional components were required for a spectrum-aggregating UHF SATCOM terminal. The number of additional blocks will depend upon the number of fragments being aggregated.

To aggregate fragments within the 146 – 174MHz band, some components and circuits can be shared across all fragments since they already operate over the entire band. They are: the antenna; low pass filter (LPF), antenna switch, and automatic power control (APC). The reference oscillator part of the PLL circuitry can also be shared, as long as its output is boosted. All these shared blocks reside in the RF unit.

The overall design changes required to change the conventional standard PBR into a two-fragment aggregator system is detailed in Appendix C.

The summary of the costs as the number of fragments is increased within the band are given in Table 5-5 below. As in the UHF SATCOM terminal example, adding DSP to the design constitutes the majority of the additional cost (i.e £1200). Even the smallest DSP devices are likely to be able to cope with processing of more than two fragments, so the DSP cost does not need to be duplicated for three, four, five or more fragment designs (although eventually additional DSP would be needed as the number of fragments goes up). Also the type (and hence price) of RF analogue splitters and combiners will change depending on the number of fragments.

System design	Price, £	extra cost %	difference
Conventional terminal	265.94		
2-fragment aggregator	1848.78	595	595% more than conventional system
3-fragment aggregator	2122.31	698	103% more than 2-fragment system
4-fragment aggregator	2378.92	795	96% more than 3-fragment system
5-fragment aggregator	2683.53	909	115% more than 4-fragment system
6-fragment aggregator	2734.14	928	19% more than 5-fragment system
7-fragment aggregator	2904.75	992	64% more than 6-fragment system
8-fragment aggregator	3435.36	1192	200% more than 7-fragment system
9-fragment aggregator	3665.57	1278	87% more than 8-fragment system
10-fragment aggregator	3686.48	1286	8% more than 9-fragment system

Table 5-5: Cost of PBR aggregating devices, up to 10 fragments in-band

5.3.3 Additional Band Design

All of the above PBR examples assume that the fragments lie within the PBR band of 146 – 174MHz. But what about using fragments outside of this band? Because the original PBR band is comparatively narrow anyway, it is easy to widen the spectrum-aggregating PBR design to exploit out-of-band fragments without too much cost or extra design effort. Components that are already shared: antenna, antenna switch, low pass preselection filter, APC circuitry, LNA, RF splitters and RF combiners, are available or can be made to operate over wider bandwidths than in the designs considered above.

Therefore by using wider band shared components, the system architecture does not need to change in order to accommodate additional band fragments. There is, of course, a limit to this. As high and higher frequency fragments are considered, there will be a point where discrete components are no longer suitable and must be replaced by dedicated RF analogue devices. This will significantly increase the cost per fragment of the RF and VCO parts of the system, up to levels comparable with the UHF SATCOM cost (roughly £660 per fragment). Also, the system design will have to change as the transmit carrier can no longer be modulated directly without an IF stage. The frequency threshold may be somewhere in the region between 200-300MHz, i.e. to exploit fragments above this will require dedicated analogue components. However, there will be a dip in the cost of parts in the frequency region 400 – 500MHz, where PBR radios are also prevalent and thus parts are readily available and cheap.

The detailed multi-band fragment design and costs are detailed in Appendix C. The costs are summarised below.

System type	Price, £	extra cost %	difference
Multi-band 2-fragment aggregator (1 distant fragment)	2423.78	811	31% more than 2-fragment single band system
Multi-band 4-fragment aggregator (2 distant fragments)	3528.92	1227	48% more than 4-fragment single band system
Multi-band 6-fragment aggregator (3 distant fragments)	4459.14	1577	63% more than 6-fragment single band system

Table 5-6: Cost of PBR aggregating devices, in multiple bands

From these costing it can be deduced that the distribution of bands will affect the cost, as will the number of fragments in a given band. But the width and distribution of fragments within a band will not affect the cost, nor will the number of bands.

For example, if an aggregating device only has bands distributed in or close to the PBR region, the number of bands will not really change the cost and the total additional cost will depend only upon the number of fragments. But introducing one or more bands far away from the PBR region will add significantly to the cost, and this cost will vary depending on the number of fragments in these far-out bands.

This is different to the situation for UHF SATCOM, where the total number of bands affected cost more than the distribution.

QinetiQ Proprietary

5.3.4 Summary of PBR costing

We can see that building a 2-fragment aggregator is dominated by the DSP costs, even more so than for UHF SATCOM because in PBR all the conventional parts are very cheap.

As the number of fragments increases, so does the cost. There is not a strong linear correlation between the two, but Figure 5-5 shows how, as the number of fragments increase, the cost starts to swing around a linear generalisation. The non-linear trend is due to the cost and combination of splitters used to achieve the design. This swing means that aggregating more fragments is not necessarily better value for money, but aggregating 4, 6, 7 or 10 fragments is better value than other fragment quantities. Remember our assumption that no extra DSP is required beyond the 2-fragment aggregator design.

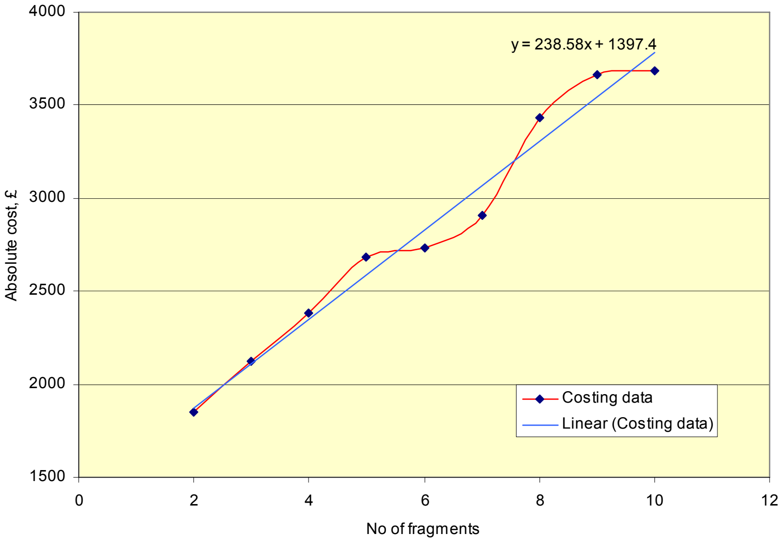

Figure 5-5: Linear trend of cost against number of fragments (discounting extra DSP)

The cost of exploiting fragments across multiple bands will depend upon the distribution of the bands, but not the number. Bands that are close to the PBR frequencies will not have much impact, but any band greater than, say, 100MHz away will significantly increase the cost. Moreover, it is the total number of fragments in these distant bands that will increase cost, and not the number of distant bands.

The distribution and size of fragments within each band, and the size of each band, will have little or no effect on cost. In effect, bands have no meaning for PBR aggregating systems and the cost depends entirely on where each fragment lies in the spectrum.

5.4 Wireless Local Area Network (WLAN)

WLAN uses RF to link two or more computers together in a small area (for example, an office or a home). The most common type of WLAN is the infrastructure WLAN, in which computers connect to an access point (AP). The AP acts rather like a base station in cellular telephony and is wired to a server (and perhaps other APs and wired computers). The fundamental difference is that in a

WLAN, computers within a "cell", or basic service set (BSS) can talk directly to each other without having to go via the access point. This allows for the second type of WLAN, the peer-to-peer or ad-hoc network, which is essentially a BSS without an AP.

The frequency at which WLAN operates depends on the protocols used, and there are several of these. The main three groups are 802.11 (Wi-Fi), 802.15 (Bluetooth) and 802.16 (WiMAX). The most common frequency used by these protocols is 2.4GHz, so this is the frequency we shall concentrate on in this example. Of the 2.4GHz protocols, some use direct sequence spread spectrum (DSSS), which is incompatible with spectrum aggregation. The other protocols either use frequency hopping spread spectrum (FHSS) or orthogonal frequency division multiplexing (OFDM). FHSS tends to be used with lower data rates and OFDM for higher data rates. Some very basic specifications are given in Table 5-7.

Transmit Band	2.402 – 2.495GHz
Receive band	2.402 – 2.495GHz
Data rate	1, 2, 6, 9, 12, 18, 24, 36, 48, 54 Mbps
EIRP	100mW maximum

Table 5-7: Basic characteristics of WLAN

5.4.1 Conventional Design

A combination of low transmit power and abundance means that 2.4GHz WLAN technology is very cheap and advanced compared to many other radio types - to the point where complete 2.4GHz transceivers are available on a single microchip at reasonable prices. Quite simple microcontrollers can be used to run these transceivers in FHSS modes, otherwise for the more complicated OFDM and COFDM, DSP is required. Figure 5-6 shows the very simple architecture, for either a FHSS or OFDM type device.

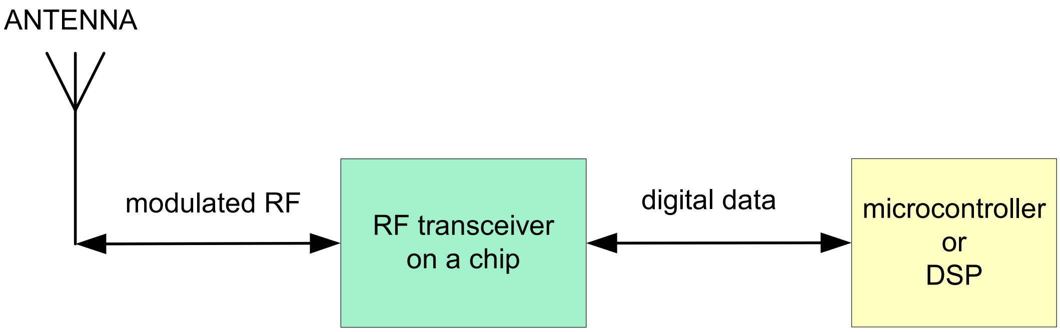

Figure 5-6: Architecture for a conventional 2.4GHz WLAN transceiver

The transceiver microchip does all the usual RF processing such as filtering, amplification, frequency synthesis, phase-locked loop circuits. More often than not they will have some flexibility to control the transmit power, centre frequency, channel width and hence data rate. They also contain all the necessary ADCs, DACs, coding, (de)modulation, buffering and packet handling required to enable a simple digital interface to the microcontroller or DSP.

It is very difficult to calculate the cost of a one-off WLAN device. Because they are so prolific, manufacturers and distributors will only sell in bulk, so the price of the transceiver chip is immediately reduced due to economy of scale. Normally, if only one or two units are required, a distributor will supply them as free-of-charge samples. As a guide, some 2.4GHz transceiver chips are sold for £3.60 per unit, for a minimum quantity of 1000 units. Microcontroller chips can be bought singly and might cost between £8 - £15 each, or less than £8 for quantities over 100. In contrast, a one-off DSP system tailored to radio applications will cost around £1,200. A WLAN antenna is required to completed the system. There is a wide variety of WLAN antennas on the market, the cheapest being for indoor use only at around £10, while larger, higher-gain outdoor antennas can cost as much as £130. The average price seems to be around £20.

The variation in prices and types of components makes is very difficult to estimate the one-off cost of a 2.4GHz WLAN system. In the table below, the average price for each component has been used, but the total cost could still vary greatly depending on the application intended. The estimated cost (detailed in the Appendix C) of one-off WLAN devices (non-aggregating) is 1,223.60 (for an OFDM solution) and 35 for a FHSS solution.

5.4.2 Spectrum-Aggregating Design

To turn a WLAN device into a spectrum aggregating device depends on whether it performs OFDM or FHSS. OFDM can be considered a form of spectrum aggregation without the need for further or different hardware. Only modifications in the control and coding software, to select which spot frequencies to use and to ensure channels are orthogonal with adjacent channels, are required. Thus devices that are already OFDM capable require no change in hardware, as long as fragments are within the usual frequency band of operation, see Figure 5-7 (to use fragments outside the band, see next subsection where additional band design system is considered).

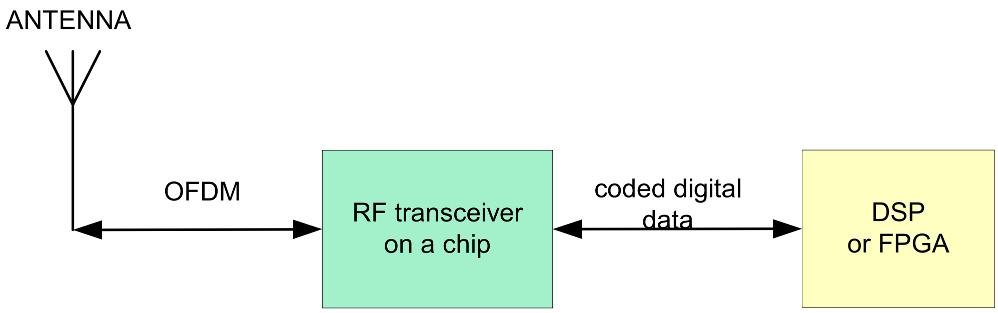

Figure 5-7: Spectrum aggregating WLAN system using OFDM

Like most wideband devices, FHSS is capable of operating on many different centre frequencies, but only one at a time. Each centre frequency could be considered as a fragment, and to get two or more fragments working simultaneously would require subsequent sets of transceivers chips, and a splitter/combiner to allow them to share one antenna, see Figure 5-8. A two-fragment design would, then, essentially be using two hopping schemes at once. This puts limitations on where they can hop to, to avoid fragments clashing. The more fragments, the greater the limitation. There are normally around 80 channels in a Bluetooth WLAN FHSS protocol, but not all channels are used in a given hopping scheme, since the original idea of hopping was to avoid poor channels caused by interference. Bluetooth can work satisfactorily with only 20 channels. Therefore, if all channels were available in the band, up to four hopping schemes – and thus fragments - could be used without fear of clashing. Higher number of fragments could still be used, but the schemes require more careful planning to time-share channels. This makes it unlikely that a spectrum-aggregating WLAN system would have a high number of fragments. The additional processing required to run multiple transceiver chips can be provided by a higher-end microcontroller unit (MCU), but does not warrant the use of more expensive DSP or FPGA.

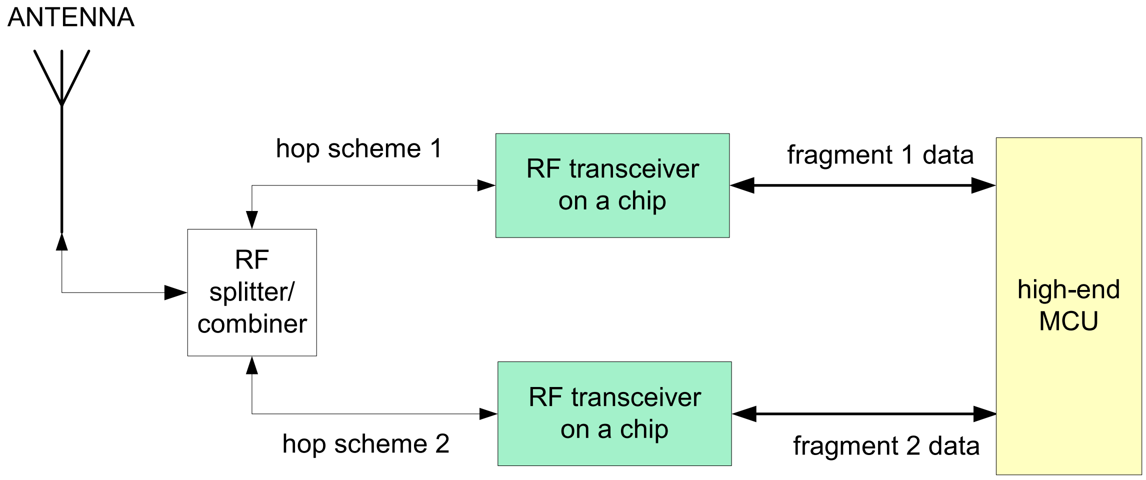

Figure 5-8: Spectrum aggregating WLAN system using FHSS

For a two-fragment aggregator, there are no additional parts and hence no additional cost if OFDM is already being used. Since the 802.11g standard can have up to 52 OFDM sub-carriers, it is highly unlikely that additional DSP will be needed to aggregate two fragments, or indeed much higher numbers of fragments. For a FHSS design, there is additional cost in the form of a splitter/combiner, an extra transceiver chip and more capable MCU. For higher numbers of fragments, in OFDM system the cost will be the same. In FHSS systems the cost will vary and are detailed in Appendix C and summarised in Table 5-8.

As DSP was the dominant expense in SATCOM and PBR aggregating systems, so the improved MCU is the initial and dominant expense in moving from a conventional to a spectrum-aggregating WLAN device. Once this initial increase is accounted for, however, the cost steadily increases as the number of fragments increases, roughly £5.70 per fragment, for up to seven fragments. Devices with eight or more fragments see another step increase, this time caused by the sudden increase in the cost of 8-way splitter/combiners with 2.4GHz performance.

System type	Price, £	extra cost %	difference
Conventional terminal	35.60		
2-fragment aggregator	95.95	170	170% more than conventional system
3-fragment aggregator	101.00	184	14% more than 2-fragment system
4-fragment aggregator	110.60	211	27% more than 3-fragment system
5-fragment aggregator	113.20	218	7% more than 4-fragment system
6-fragment aggregator	118.25	232	14% more than 5-fragment system
7-fragment aggregator	130.35	266	34% more than 6-fragment system
8-fragment aggregator	254.00	613	347% more than 7-fragment system
9-fragment aggregator	260.10	631	17% more than 8-fragment system
10-fragment aggregator	265.15	645	14% more than 9-fragment system

Table 5-8: Summary of FHSS WLAN costs for upto 10 fragments in -band

5.4.3 Additional Band Design

To aggregate fragments outside of the WLAN bands will incurr large costs for both OFDM and FHSS type systems. The cost will depend on the specific bands in which fragments reside, but in general a superhetrodyne approach using individual RF analogue components will be required. This means ADCs, DACs and possibly DSP – or more MCU processing power – will also be needed.

Thus the cost of aggregating outside the WLAN band will be similar in nature to that of the PBR examples, i.e. for WLAN additional band design, bands and fragments are interchangeable. Also, the distribution of bands will affect the cost, as will the number of fragments in a given band, but the width and distribution of fragments within a band will not affect the cost, nor will the number of bands.

System type	Price, £	extra cost %	difference
Multi-band 2-fragment aggregator (1 distant fragment)	2423.78	811	31% more than 2-fragment single band system
Multi-band 4-fragment aggregator (2 distant fragments)	3528.92	1227	48% more than 4-fragment single band system
Multi-band 6-fragment aggregator (3 distant fragments)	4459.14	1577	63% more than 6-fragment single band system

Table 5-9: Cost of WLAN aggregating devices, in multiple bands

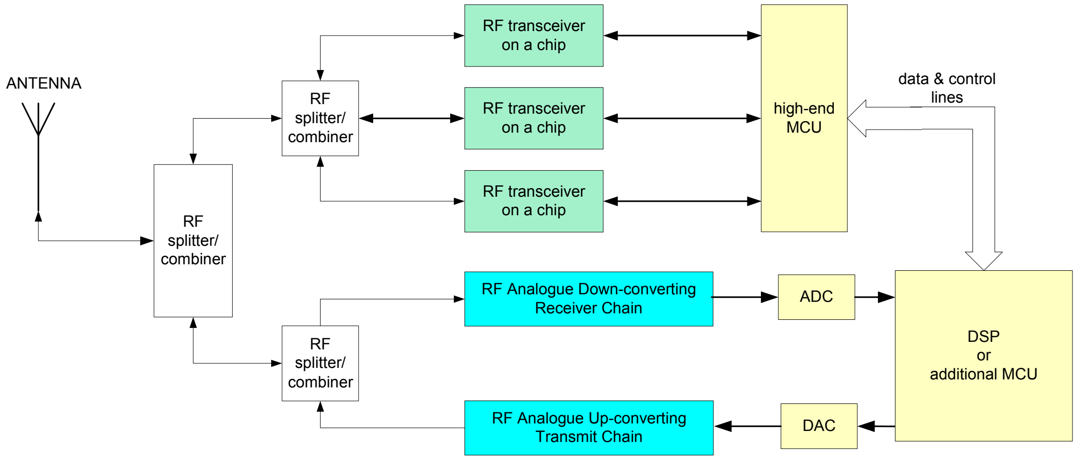

Figure 5-9: Aggregating from bands outside WLAN will be costly due to analogue parts & DSP

5.4.4 Summary of WLAN costing

To summarise the costs involved, see Table 5-8. We can see that building a 2-fragment aggregator is dominated by the more advanced MCU costs, although this increase is not on a par with the cost of DSP as has been the case for SATCOM and PBR.

Because of the popularity and abundance of WLAN devices, adding more hardware to enable fragment aggregation is generally very cheap. As the number of fragments increases, so does the cost, but only by about £5.70 per fragment. As Figure 5-10 shows, there is a good linear relationship between the absolute cost and the number of fragments, until the fragment quantity exceeds seven. A step increase is then seen, even more dominant than the initial MCU cost, due to a sudden increase in the cost of higher-way splitters (i.e. 8-way splitters). Beyond this, for 9- and 10-fragment devices, we see the linear trend continue once more. It is interesting to note that this step increase interrupting the linear trend is similar to what we would see in the SATCOM and PBR examples if additional DSP were added at a certain threshold of fragment quantity.

The cost of exploiting fragments across multiple bands will depend upon the distribution of the bands, but not the number. Any band outside the specified WLAN frequency region will significantly increase the cost. Moreover, it is the total number of fragments in these distant bands that will increase cost, and not the number of distant bands. The cost increase will be exacerbated if DSP is required to process these additional bands.

The distribution and size of fragments within each band, and the size of each band, will have little or no effect on cost. As is the case for PBR, outside bands essentially have no meaning for WLAN and the extra cost depends almost entirely on where each fragment lies in the spectrum.

These costs relate only to FHSS type WLAN. Exploitation of fragmented spectrum is already possible today using OFDM and COFDM. Devices that are not capable of OFDM but have an FHSS ability are unlikely to be converted into aggregating devices because of the easier, cheaper option of OFDM.

The significant cost in aggregating fragments beyond the WLAN bands is a major barrier and means that combining cheap WLAN technologies with other expensive technologies is unlikely to happen, if the reason is just to aggregate spectrum.

However, because it is so cheap, WLAN is being added to all sort of other devices – including other radios – and so converting other types of radio into aggregating devices by adding WLAN channels is a much more realistic proposition.

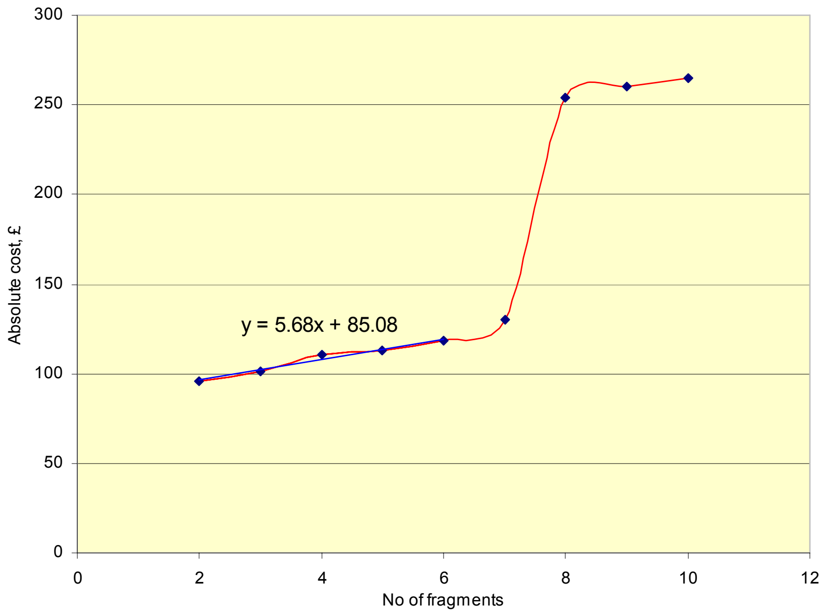

Figure 5-10: Linear trend interrupted by step increase (between 7-8 fragments)

5.5 Microwave Fixed Link

There are many fixed links operating at frequencies between 5 - 40GHz. They are generally used as large bandwidth "pipelines" between remote or hard-to-access areas, such as rural GSM base stations. Indeed, a major part of the fixed link market constitutes mobile phone operators. Fixed links are usually two-way, using frequency division duplex (FDD). Two common sets of frequencies are 18.82/18.87GHz and 21.2/23.6GHz. Therefore we will focus on hardware around the 20GHz region in this example.

Some basic specifications are given in Table 5-10. Note that in a fixed point-to-point system the location of the transceivers are known, therefore directional antennas can be utilised. Directional antennas have high gain, concentrating RF energy in one direction only so that only a small amount of power need to supplied to the antenna by the power amplifier in order to get a high EIRP. In microwave fixed links, antennas gains between 13dBi – 43dBi are typical.

Transmit Band	18.82 / 21.20 GHz
Receive band	18.87 / 23.60 GHz
T/R mode	Full duplex (FDD)
Channel size	2 – 34 MHz
ERP	20dBm (0.1 Watts) maximum
EIRP	up to 63dBm (2000 Watts)

Table 5-10: Basic specification of a conventional microwave fixed link

5.5.1 Conventional Design

The outline design for a fixed microwave link transceiver station is given in Appendix C. It consists of five main parts:

- Receiver Chain
- Transmit Chain
- Antenna and antenna sharing device
- Phase Locked Loop (PLL) Stage
- Analogue-to-Digital and Digital-to-Analogue Conversion

Some microwave fixed links may be acting as relays, in which case they will be using two antennas (pointing in different directions), one for receive and one for re-transmit. In this case, an antenna sharing device is not needed as either chain connects directly to an antenna. But in most cases, one antenna will be shared by the chains.

The design of the system is detailed in Appendix C. along with the design and cost of a conventional system which is predicted to be £8421.

5.5.2 Spectrum-Aggregating Design

To turn the microwave link transceiver station into a spectrum aggregating device, multiple receive and transmit chains are needed for each fragment. To share the antenna amongst the chains, a single circulator is placed at the antenna feed which is then connected to splitters that combine signals to/from the individual chains (the alternative is to use one splitter and a circulator for each chain fragment, but circulators or more expensive than splitters so it is cheaper to share one circulator).

A detailed architecture of the two-fragment example is given in the Appendix C and shows that, with just two fragments, the component count and complexity is starting to become significant. The breakdown in cost of the additional parts as the number of fragments increases is also detailed in the same section.

Compared to the cost of a non-aggregating conventional microwave transceiver station, the additional 2-fragment cost is equivalent to a 58% increase. Almost all of this extra cost is due to the duplication of components, since there is little scope for sharing common parts across chains/fragments.

Cost breakdowns for 3-fragment, 4-fragment and 5-fragment aggregating systems are given in the Appendix C. On average, a 58% increase cost per fragment pair is observed (approximately £4,884). This is to be expected due to the large duplication of components. The cost as a function of fragments is summarised in the table below.

System type	Price, £	extra cost %	difference
Conventional terminal	8,420.85	0	
2-fragment aggregator	13,314.4	58	58% more than conventional system
3-fragment aggregator	18,208.15	116	58% more than 2-fragment system
4-fragment aggregator	23,101.8	174	58% more than 3-fragment system
5-fragment aggregator	28,495.45	238	64% more than 4-fragment system
6-fragment aggregator	25,295.35	300	62% more than 5-fragment system
7-fragment aggregator	29,938.95	356	55% more than 6-fragment system
8-fragment aggregator	34,582.55	411	55% more than 7-fragment system
9-fragment aggregator	39,226.25	466	55% more than 8-fragment system
10-fragment aggregator	43,869.95	521	55% more than 9-fragment system

Table 5-11: Summary of microwave link costs for an in-band aggregator system

5.5.3 Additional Band Design

The cost for aggregating microwave fragments in different microwave bands will be similar to that of aggregating higher number of fragments, because multiple antennas (and therefore circulators) will be needed to cover the various bands.

Figure C-15 shows an example of a dual-band four-fragment aggregator, with two fragments in each band. The cost, ~£17,000 is broken down in Appendix C.

The additional cost of a dual-band aggregator with five fragments, two in one band and three in the other, are shown in the Appendix C. We have already seen that the approximate cost for a fragment pair is around £5,000. These dual-band examples show that this cost per fragment is still the same, and that the cost of adding a new band is largely due to the additional antenna and circulator. However, even though these are expensive parts in a microwave system (accounting for a third of the cost of our conventional system example), it does not significantly add to the overall cost – only 10% at most. Therefore, the number of fragments is the dominant factor in cost, while the cost can be pushed up slightly further depending on the distribution of fragments in relation to bands, or in other words the fragment density within a band. For a given total number of fragments, if the fragment density is low, then this means there are more bands and thus higher cost than if all the fragments were concentrated within just one or two bands.

The summary costs are shown in the table below.

System type	Price, £	extra cost %	difference
Dual-band 4-fragment aggregator	25,428.80	202	10% more than 4-fragment single band system
Dual-band 5-fragment aggregator	30,322.55	260	6.5% more than 5-fragment single band system

Table 5-12: Summary of microwave link costs for a dual band system

5.5.4 Summary of microwave fixed link costing

To summarise the costs of spectrum-aggregating microwave links, see Figure 5-11. There is a very strong linear relationship between the number of fragments and the cost of hardware. This is due to the large amount of duplication required for each fragment, with few parts that can be shared amongst fragments or component chains.

We can see that, even assuming that an extra antenna and circulator will be needed to aggregate more than five fragments, this does not affect the linear correlation between cost and number of fragments. Similarly in multiple-band designs, there is only a small increase in cost incurred due to adding a new band. The main cost is due to the number of fragments, not their distribution.

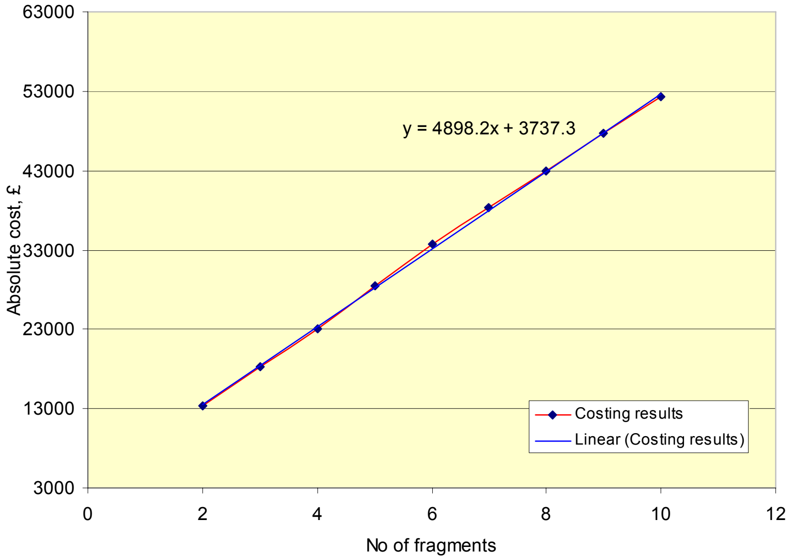

Figure 5-11: Linear trend of cost against number of fragments (discounting extra DSP)

5.6 Summary and conclusions

The hardware costs involved for a satellite, a public mobile radio, a wireless local area network and a microwave fixed link system have been considered as if they were designed using conventional spectrum (i.e. a single band). The hardware costs for the same system, but operating in a fragmented spectrum, in the same band, were then derived as the number of fragments exploited increases to a maximum of 10. The design of a fragmented system, but this time operating across multiple bands, has also been considered.

The hardware costs of a fragmented system are, as expected, more costly than a single band conventional system. As the number of fragments increases the cost tends to increase proportionally.

To summarise the costs involved for a satellite system we note that building a 2-fragment aggregator is dominated by the DSP costs. Aggregating more fragments is better value for money (in terms of the cost per fragment). Assuming no extra DSP is needed, the cost of an aggregator is fairly linear with respect to the number of fragments. Exploiting fragments across multiple bands increases the cost, although not significantly. The way in which fragments are distributed within one or more bands, and the size of bands as well as individual fragments, does not affect the cost.

For the PBR system we note that building a 2-fragment aggregator is again dominated by the DSP costs, even more so than for UHF SATCOM because in PBR all the conventional parts are very cheap. As the number of fragments increases, so does the cost. There is not a strong linear correlation between the two, but there is a sinusoidal swing around a linear trend. This swing means that aggregating more fragments is not necessarily better value for money, but aggregating 4, 6, 7 or 10 fragments is better value than other fragment quantities. The cost of exploiting fragments across multiple bands will depend upon the distribution of the bands, but not the number. Bands that are close to the PBR frequencies will not have much impact, but any band greater than, say, 100MHz away will significantly increase the cost. Moreover, it is the total number of fragments in these distant bands that will increase cost, and not the number of distant bands. The distribution and size of fragments within each band, and the size of each band, will have little or no effect on cost. In effect, bands have no meaning for PBR aggregating systems and the cost depends entirely on where each fragment lies in the spectrum.

Building a 2-fragment aggregator for a WLAN system is dominated by the more advanced MCU costs, although this increase is not on a par with the cost of the additional DSP as required for the SATCOM and PBR cases. Because of the popularity and abundance of WLAN devices, adding more hardware to enable fragment aggregation is generally very cheap. As the number of fragments increases, so does the cost, but only by about £5.70 per fragment. There is a good linear relationship between the absolute cost and the number of fragments, until the fragment quantity exceeds seven. A step increase is then seen, even more dominant than the initial MCU cost, due to a sudden increase in the cost of higher-way splitters (i.e. 8-way splitters). Beyond this, for 9- and 10-fragment devices, we see the linear trend continue once more. It is interesting to note that this step increase interrupting the linear trend is similar to what we would see in the SATCOM and PBR examples if additional DSP were added at a certain threshold of fragment quantity. The cost of exploiting fragments across multiple bands will depend upon the distribution of the bands, but not the number. Any band outside the specified WLAN frequency region will significantly increase the cost. Moreover, it is the total number of fragments in these distant bands that will increase cost, and not the number of distant bands. The cost increase will be exacerbated if DSP is required to process these additional bands. The distribution and size of fragments within each band, and the size of each band, will have little or no effect on cost. As is the case for PBR, outside bands essentially have no meaning for WLAN and the extra cost depends almost entirely on where each fragment lies in the spectrum.

To summarise the costs of spectrum-aggregating microwave links we note that there is a very strong linear relationship between the number of fragments and the cost of hardware. This is due to the large amount of duplication required for each fragment, with few parts that can be shared amongst fragments or component chains. We also note that, even assuming that an extra antenna and circulator will be needed to aggregate more than five fragments, this does not affect the linear correlation between cost and number of fragments. Similarly in multiple-band designs, there is only a small increase in cost incurred due to adding a new band. The main cost is due to the number of fragments, not their distribution.

5.7 Recommendations

The design studies in this section support the recommendations that if fragmentation occurs Ofcom or its licensees should endeavour to make sure that the fragments are not evenly distributed across the whole spectrum because aggregating spectrum across differing bands is expensive. Clusters of fragments are preferable to widely-spread fragment to facilitate the development of any potential aggregating systems.

Ofcom should highlight any occurrences of spectrum fragmentation so that hardware suppliers can develop specific devices that can facilitate aggregation such as splitters, wideband amplifiers, DSPs etc.

5.8 References

[1] Chamber of Shipping. URL: http://www.british-shipping.org/british/index.htm, 19th July 2005.

[2] ICT driving change in the fishing industry: A review of the years 1990 – 2002, European Foundation for the Improvement of Living and Working Conditions, 2003, page 15.

[3] Davies, C., The demand for Private Business Radio: A study for the Radiocommunications Agency, ref. quotient/report/qc2434, Quotient Communications Ltd.

6 Investment Impact Analysis

In the previous section, illustrative aggregating and conventional hardware-costs were determined for UHF Satcom, PBR, WLAN (FHSS) and microwave fixed link systems. Aggregation architectures included fragments within a band (i.e. close to the centre frequency) and fragments outside of the band. The hardware results showed (not-unexpectedly) that as the number of fragments increased the hardware cost also increased. In the WLAN example, the hardware costs exhibited a step-function increase.

The hardware design costs provide the first level of indicative costs to show if a potential market using an aggregating device could be achieved. The hardware concept/design (prototype) phase may typically account for 3-8% of the full production development budget [1] and if these costs are many orders of magnitude greater than a conventional system costs, a spectrum aggregating device could be uneconomic. To analyse the costs further an investment scenario is developed to investigate if a business case could be made for the aggregating service. The aim of this analysis is to answer the hypothetical question: would the service have been developed if only fragmented spectrum was available?

In considering the above question, many assumptions have been made. For example, if the hardware manufacturers knew that only fragmented spectrum could be exploited in certain bands then they may have developed cheap multi-way splitters, DSPs etc. Other assumptions such as number of units, timescale for investment etc have also been made. The results, therefore, should not be used for any other purpose.

To investigate the larger investment scenario, it is assumed that each service using a conventional (non-aggregating spectrum) terminal must achieve a positive Net Present Value (NPV) figure within a 5-year market with a 10% discount factor applied. The same constraints are then applied to a service built using a fragmented spectrum device (using within and out of band frequencies).

6.1 Using Target Costing to Establish a Price

Before we are able to carry out the NPV analysis, a selling price to base expected revenues must be established. This is achieved using a Target costing process [2] which builds upon a design-to-cost (DTC) approach with the focus on market-driven target prices as a basis for establishing target costs. In our scenario we have already calculated a build-cost, based on hardware costs, and can reverse the normal DTC model to produce a Manufacturers Selling Price which can then be used in the NPV appraisal. Table 6-1 shows a typical Target Cost calculation spreadsheet. The "%" factors are illustrative and will be market dependant and very commercially sensitive.

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£495.00
−	Standard Dealer Margin		30%	£ -	£148.50
=	Cost to Retailer				£346.50
−	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£331.50
−	Distribution Cost/Mark-up		15%	£ -	£49.73
−	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£264.78
−	Profit Margin		8%	£ -	£21.18
−	Warranty Cost		2%	£ -	£5.30
−	Corporate Allocations		10%	£ -	£26.48
−	Business Unit Selling, General & Administrative		12%	£ -	£31.77
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
−	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£174.05
−	Overhead		45%	£ -	£78.32
=	Direct Target Cost (Labour & Material)				£95.73

Table 6-1: Typical Target Cost Calculation Spreadsheet [2]

Table 6-1 shows how a manufacturers suggested retail price for a product of £495 is broken down to produce a direct target (i.e. cost to build the item) of £95.73. In this example this would be the cost (labour & materials only) that the relevant business unit (within the company structure) would aim to build each unit for.

An explanation of each cost/margin is shown below:

- **Overhead**. This is business unit specific and represent the overheads (e.g. heating, lighting, IT provision etc) for that unit.
- **Non-Recurring Development Cost**. Refers to the cost of creating a new product which is paid up front. For example, in the semiconductor industry, this is the cost of developing the circuit design and photomasks; the production cost is the cost to manufacture each wafer.
- **Business Unit Selling, General & Administrative**. This margin covers the business units administrative, marketing and sales costs which is not normally a direct cost but rather a cost that is shared between the various projects and products.
- **Corporate Allocations**. These are margins that may be allocated at Company board level to be levied across projects/product development to meet Corporate financial targets.
- **Warranty Cost**. A margin that may be levied across all projects/products for a central warranty account that services claims from across the company.

- **Profit Margin**. An assigned margin to meet expected profit target within the company.
- **Shipping/Logistics Cost to Distributor Centre**. A per unit cost to cover the costs involved in shipping and logistics to the distributor centre.
- **Distribution Cost/Mark-up**. This is an average mark-up a distributor might apply to goods they are handling.
- **Shipping/Distribution Cost to Retailer**. A per unit cost to cover the costs involved in shipping and distributing the goods to the retailer.
- **Standard Dealer Margin**. This is an average mark-up a dealer would expect to make on goods.

The model has been used in reverse to go from a direct target cost to a manufacturers selling price. We have assumed that the build costs from the previous Section directly map to the Direct Target Cost (labour & materials). In a real world situation we would expect these build costs to reduce considerably due to the economies of scale and design refinements. It is assumed, therefore, that the costs derived earlier also cover the labour allocation of the cost. The completed Target Cost calculations for the aggregating terminals are detailed in Appendix D.

In the analysis no distinction is made between the selling price determined using the above methodology, and the real-world standard (non-aggregating) terminal price. Any variances, however, may be attributed to a large number of overheads and allocations (e.g. corporate risk, product-line development). Comparisons, therefore, between our calculated price and the real-market price should not be made. The final (real-world) selling price, however, does not impact on the NPV analysis we have carried out.

Table 6-2 below summarises the results obtained from the Target Cost analysis done for each example selected (using example %-factor values). The figures are used within the NPV analysis performed below.

Service	**Terminal type**	**Direct target cost (build)**	**Manufacturers selling price (sell)**
Satcom	Standard	£2,315	£4,430
	Single band 2-fragment	£4,009	£7,674
	Dual band 4-fragment	£5,683	£10,877
PBR	Conventional	£266	£509
	Single band 2-fragment	£1849	£3539
	Multi band 4-fragment	£3529	£6754
WLAN	Conventional	£36	£68
	2-fragment	£96	£184
	10-fragment	£265	£508
Fixed Link	Conventional	£8421	£16118
	Single band 2-fragment	£13314	£25484
	Dual band 4-fragment	£25429	£48671

Table 6-2: Build Costs & Selling Price for the Chosen Terminals

6.2 Net Present Value Analysis

Using the figures from Table 6-2 an NPV analysis for the various terminals listed has been performed using the following assumptions:

- Initial investment to build 20,000 units (target size of market)
- For each standard terminal the market position will be built up over 5 years, i.e. 20,000 units sold evenly split per year over this period
- Discount rate of 10% to take account of inflation, corporate risk and opportunity cost (i.e. what the money could make if it were invested elsewhere)
- Due to the increased price of the aggregating terminal (very significant in some examples) it is assumed that the terminals would sell more slowly and in some cases would take more than the 5 years shown. There are numerous examples of this behaviour in real markets i.e. LCD TVs and DVD recorders where initial sales were to ‘early adopters’ because of the high price. In our examples this effect has been modelled through a % decrease in the number of units sold based on the magnitude of price increase. Details of this % decrease and how this applies to each service example are shown in Table 6-3 and Table 6-4.

Range of Price Increase (%)		Decrease in units sold (%)
Lower Limit	Upper Limit	
0	99	10
100	249	25
250	399	40
400	549	55
550	699	70
700	-	90

Table 6-3 : Sales Volume Percentage Decrease

Type	RX	Price Increase (%)	Decrease in units sold (%)
Satcom	Standard	0	0
	2-Fragment	73	10
	DB 4-Fragment	156	25
PBR	Conventional	0	0
	2-Fragment	595	70
	MB 4-Fragment	1227	90
WLAN	Conventional	0	0
	2-Fragment	170	25
	10-Fragment	645	70
Fixed Link	Conventional	0	0
	2-Fragment	58	10
	DB 4-Fragment	202	25

Table 6-4 : Sales Decrease Applied to Terminal Examples

For reference, an example of the NPV analysis for the standard Satcom terminal and the 2-fragment Satcom terminal is shown in *Table 6-5* and *Table 6-6*.

Year	Discount factor 10%	Units Built @ £2,315	Units Sold @ £4,430	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£46,300,000	£0	-£46,300,000	-£46,300,000	-£46,300,000
1	0.91		4000	£0	£17,720,000	£17,720,000	£16,109,091	-£30,190,909
2	0.83		4000	£0	£17,720,000	£17,720,000	£14,644,628	-£15,546,281
3	0.75		4000	£0	£17,720,000	£17,720,000	£13,313,298	-£2,232,983
4	0.68		4000	£0	£17,720,000	£17,720,000	£12,102,998	£9,870,016
5	0.62		4000	£0	£17,720,000	£17,720,000	£11,002,726	£20,872,742

Table 6-5 : NPV Analysis for the Satcom Standard Terminal

Year	Discount factor 10%	Units Built @ £4,009	Units Sold @ £7,674	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£80,180,000	£0	-£80,180,000	-£80,180,000	-£80,180,000
1	0.91		3600	£0	£27,626,400	£27,626,400	£25,114,909	-£55,065,091
2	0.83		3600	£0	£27,626,400	£27,626,400	£22,831,736	-£32,233,355
3	0.75		3600	£0	£27,626,400	£27,626,400	£20,756,123	-£11,477,232
4	0.68		3600	£0	£27,626,400	£27,626,400	£18,869,203	£7,391,971
5	0.62		3600	£0	£27,626,400	£27,626,400	£17,153,821	£24,545,792

Table 6-6 : NPV Analysis for the Satcom Single Band 2-Fragment Terminal

The 2-fragment Satcom terminal is 73% more expensive to build (*Table 6-4*) than a conventional terminal. The assumption, therefore, is made that 10% less units will sell in the 5 year market we are assessing. The impact is that although the NPV goes positive in the same year (Year 4), the return is almost £1.5M less at this point. All other examples are shown at Appendix D.

6.2.1 NPV Results

The NPV results for all four services are shown at Figure 1-1. The results are plotted as NPV versus Time (Year) curves using the data from the NPV Tables in Appendix D.

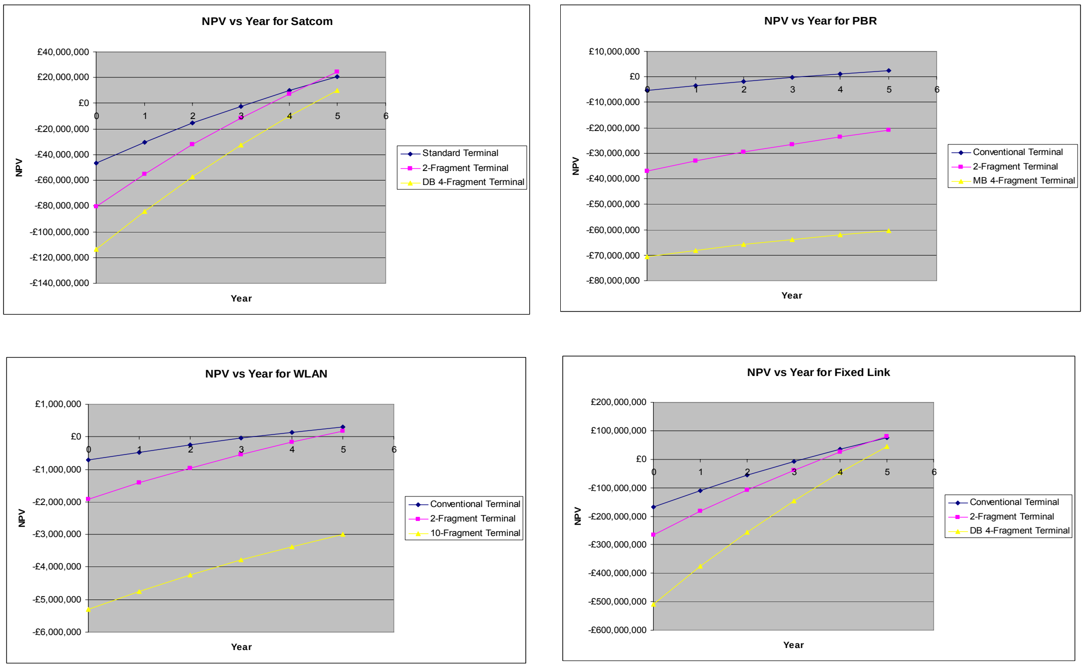

Figure 6-1 NPV Results for the Example Services

Analysis of the NPV results shows that the curves for Satcom and Fixed Link are very similar, as expected, as the price increases for the aggregating terminals are very similar (Satcom 73% & 156%; Fixed Link 58% & 202%). Technically, for both of these services an aggregating architecture is fairly straightforward to achieve and the increase in costs rose linearly as the number of fragments increased due to the component count. In pure NPV terms and within the constraints of our assumed market, the business model is positive for both of these services and further detailed business analysis with more meaningful and accurate cost models would be justified, i.e. an aggregating solution may have been developed had there been no conventional spectrum.

The results for PBR and WLAN (FHSS) both look a lot less likely as an investment case, with very long payback periods for all bar the WLAN 2-fragment terminal being predicted. Specifically, in the PBR scenario, the large cost increases in an aggregating solution are dominated by the high cost of DSP against the low cost of the components. The PBR radio looks like a very poor investment opportunity. For WLAN the increases are fairly linear as far as eight fragments and up to this point, when there is a step change due to the high cost of higher-way splitters.

6.2.2 Cost of Spectrum Impact.

The investment scenarios have ignored the cost of any spectrum and in this brief section we consider its impact on one of the examples. The question asked is: could any spectrum costs be reduced for the fragmented scenario to the extent that the NPV becomes positive on the same timescale as that for the standard terminal?

If it is assumed (in the Satcom NPV example) that in each fragmented case the investment figures (the expenses) include a cost element for the spectrum, and that the revenues streams remain unchanged, we can then calculate the reduction necessary on the investment side to achieve positive NPV at the same rate as the standard terminal. To do this in practise we must achieve an approximate 2:1 relationship between revenues and costs (expenditures). Therefore if we reduce the costs in the Satcom 2-fragment and dual band 4-fragment by 14% and 28% respectively we satisfy the 2:1 relationship and achieve the NPV trend we are looking for. The effects of this on the NPV analysis and results are shown below in Table 6-7 and Figure 6-2

Cash Flow 2- fragments			Net Present Value	NPV Payback & Total
Expenses	**Revenues**	**Balance**		
-£69,000,000	£0	-£69,000,000	-£69,000,000	-£69,000,000
£0	£27,626,400	£27,626,400	£25,114,909	-£43,885,091
£0	£27,626,400	£27,626,400	£22,831,736	-£21,053,355
£0	£27,626,400	£27,626,400	£20,756,123	-£297,232
£0	£27,626,400	£27,626,400	£18,869,203	£18,571,971
£0	£27,626,400	£27,626,400	£17,153,821	£35,725,792
(£69,000,000.00)	£138,132,000.00	**< Totals**		

Cash Flow 4 fragments			Net Present Value	NPV Payback & Total
Expenses	**Revenues**	**Balance**		
-£81,800,000	£0	-£81,800,000	-£81,800,000	-£81,800,000
£0	£32,631,000	£32,631,000	£29,664,545	-£52,135,455
£0	£32,631,000	£32,631,000	£26,967,769	-£25,167,686
£0	£32,631,000	£32,631,000	£24,516,153	-£651,533
£0	£32,631,000	£32,631,000	£22,287,412	£21,635,879
£0	£32,631,000	£32,631,000	£20,261,284	£41,897,163
(£81,800,000.00)	£163,155,000.00	**< Totals**		

Table 6-7: NPV Analysis for Satcom 2-fragment and Dual Band 4-Fragment Terminals with a 14% and 28% Reduction in Expenditure Respectively

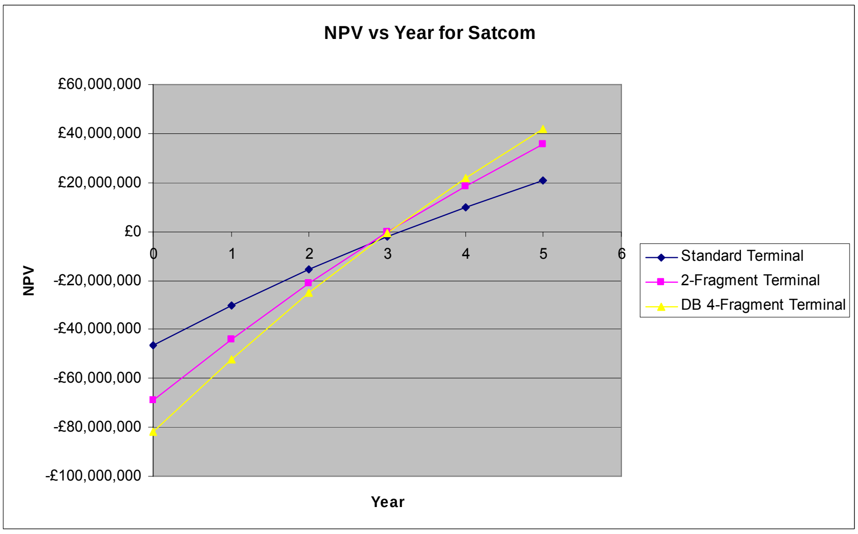

Figure 6-2: NPV Results Plotted for Satcom Terminals After Reductions in Expenditure

The analysis shows, that if our initial investment had to consider the cost of spectrum and that this was a reasonably dominant cost (between ~ 20-30% of the total), there would be opportunities within some of the service examples examined for closing that gap in the NPV analysis by reducing the cost of the fragmented spectrum. Note that in the worst case examples, i.e. PBR and WLAN (10-fragment) the reduction in costs would have to be 50% and more and it is unlikely that the cost of spectrum would be as dominant a part of the cost base to make this transformation likely.

6.3 Conclusions

The hardware costs for non-aggregating (standard) and aggregating terminals in four different services, UHF Satcom, PBR, WLAN (FHSS) and microwave fixed link have been previously derived. In this section these costs have been used to develop a larger investment case to predict if a fragmented solution would have been possible if no conventional spectrum was available. This has been achieved via NPV analysis of chosen examples and how they behave in this assumed market. The results should only be used for illustrative purposes for this study.

The analysis has shown that for fragmented PBR and WLAN (10-fragment terminal) scenarios, the investment case would not be made under our assumed market conditions. However, for Satcom, Fixed Link and WLAN (2-fragment terminal) NPV is positive within the 5-years and could, therefore, have attracted a more detailed investment analysis. The investment analysis has also shown that in circumstances where the cost of spectrum is a dominant part of the investment, an opportunity exists whereby if the fragmented spectrum were of a lower cost than its non-fragmented counterpart, positive NPV can be achieved on a timescale commensurate with that of the standard terminal. For the Satcom service, 2-fragment and dual band 4-fragment examples should have fragmented spectrum prices lower by 14% and 28% respectively than non-fragmented spectrum.

It should be noted that care should be taken in interpreting these results because of the assumptions made. The investment analysis does illustrate, however, that aggregated services could have provided a return on investment and hence may have developed in the absence of conventional spectrum.

6.4 References

[1] DUTTON, J. J., *Target Setting: Key to Successful NPD Outcomes,* PDMA Visions, Apr 1998. http://www.pdma.org/visions/apr98/dutton.html

[2] http://www.npd-solutions.com/target.html

7 Virtual aggregation solution

7.1 Introduction

If it is assumed that fragments of spectrum can be utilized (e.g. using spread spectrum such as frequency hopping spread spectrum, direct sequence spread spectrum or hybrid spread spectrum) in an IP packet type concept to transmit data, the challenge then is how to exploit the fragments of spectrum to ensure "optimum" use. One method investigated was on how you could load the fragments with as many services as possible to maximize use. The management of the services could be achieved by the application of an economic based trading model where a central server is the trading engine and the spectrum segments and the data (for the differing services) to be transmitted are traded. The application of a trading model applied in a hierarchical manner to the spectrum where trading is carried out at many levels was investigated.

The results of the investigation are summarised in the Conference paper reproduced in Appendix E and further detailed in Appendix F.

7.2 Conclusion

The virtual aggregation solution highlighted the type of data required if fragments were to be virtually aggregated and then traded. A number of potential applications were considered and an experimental scenario for testing the solution was proposed but not implemented.

The primary conclusion of this part of the study was that greater knowledge was required on how fragmentation may occur and potential single service hardware solutions that could be implemented.

A Fragmentation of the Spectrum

A.1 Introduction

This section details the findings of the fragmentation study in graphical format. Each graph shows the users in the particular band of the spectrum and provides an indication of multiple users of the spectrum.

A.2 100 – 200 MHz:

In this range, the biggest uses of spectrum are for aeronautical, mobile and broadcast applications. The only potentially available spectrum for aggregation comprises locally unassigned channels within the PMR bands around 150 and 170 MHz.

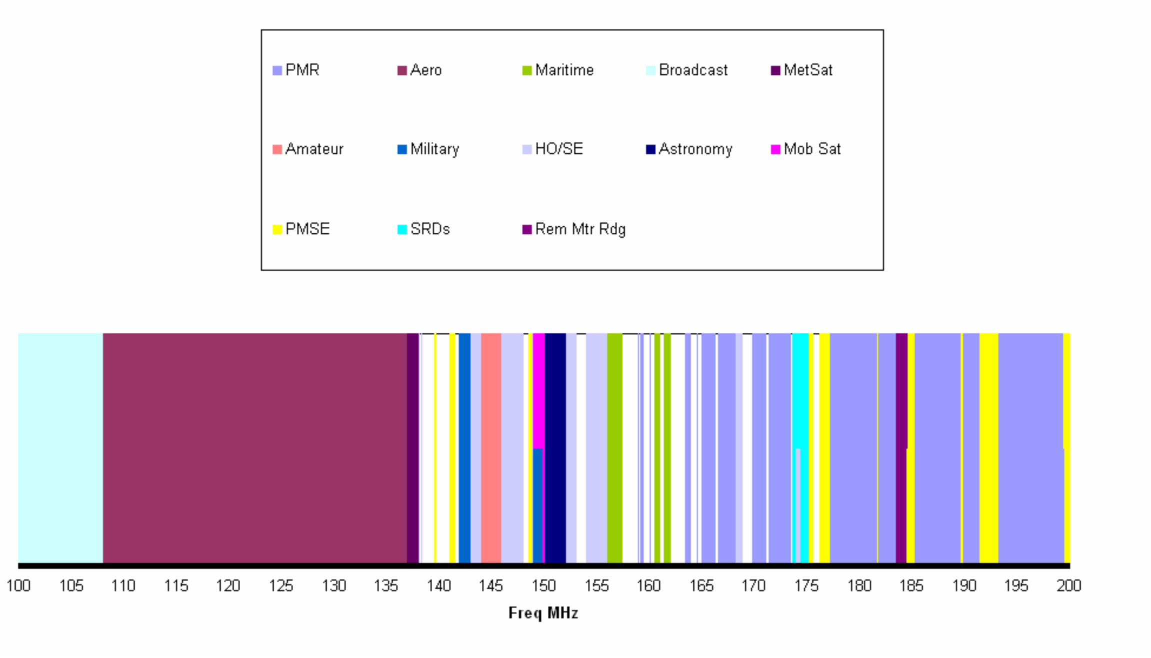

A.3 200 – 500 MHz

In this range, much of the spectrum is allocated to the military, with civil mobile (PMR) and broadcast (DAB at the bottom end and TV at the top) services also accounting for significant amounts.

Some of the spectrum in Band III is currently unused, however our understanding from Ofcom is that this should not be considered for aggregation currently, pending the outcome of the forthcoming ITU Regional Radio Conference (RRC).

PMR spectrum in the 420 MHz region (UHF1 band) is constrained by the presence of a military radar in the band and therefore cannot be used. Consequently the only spectrum suitable for aggregation in this range comprises unassigned PMR channels in the UHF2 band (450 – 470 MHz).

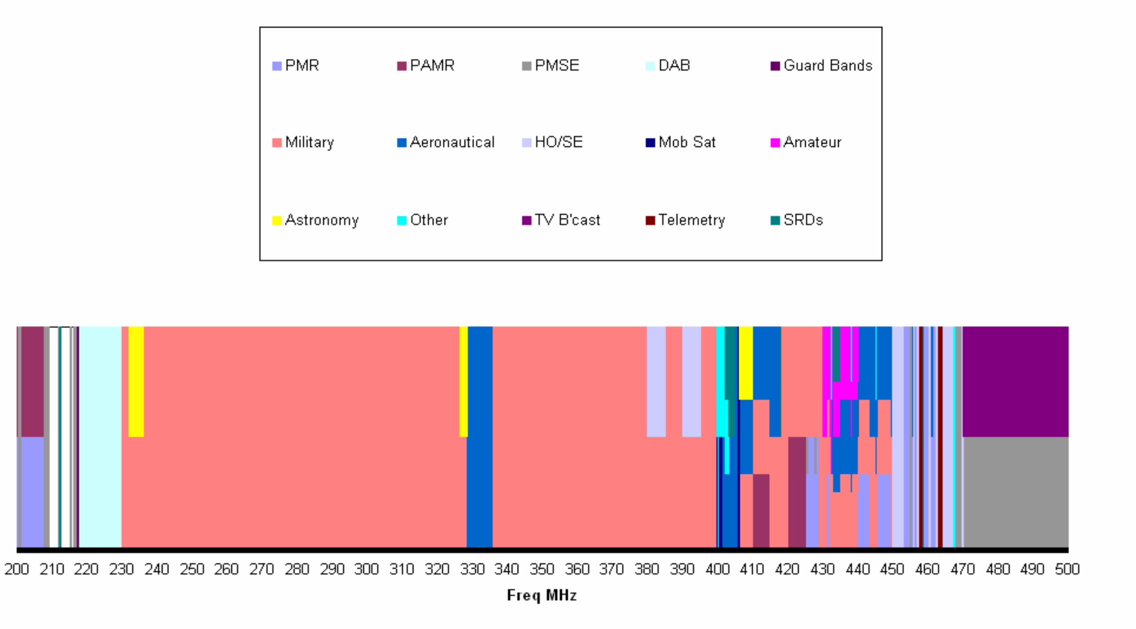

A.4 500 – 1000 MHz

In this range, most of the spectrum is allocated to analogue TV broadcasting and where this spectrum is not used for this application it is used for PMSE. Cellular mobile (GSM) and military applications are also significant users, as is aeronautical use (radars) at the top end of the band. The only currently available spectrum identified is a 1 MHz block between 862 and 863 MHz, although the GSM guard band at 915 – 917 MHz might be useable with the agreement of the military. TV channel 38 is reserved for Radioastronomy use and may be useable for other services in certain areas.

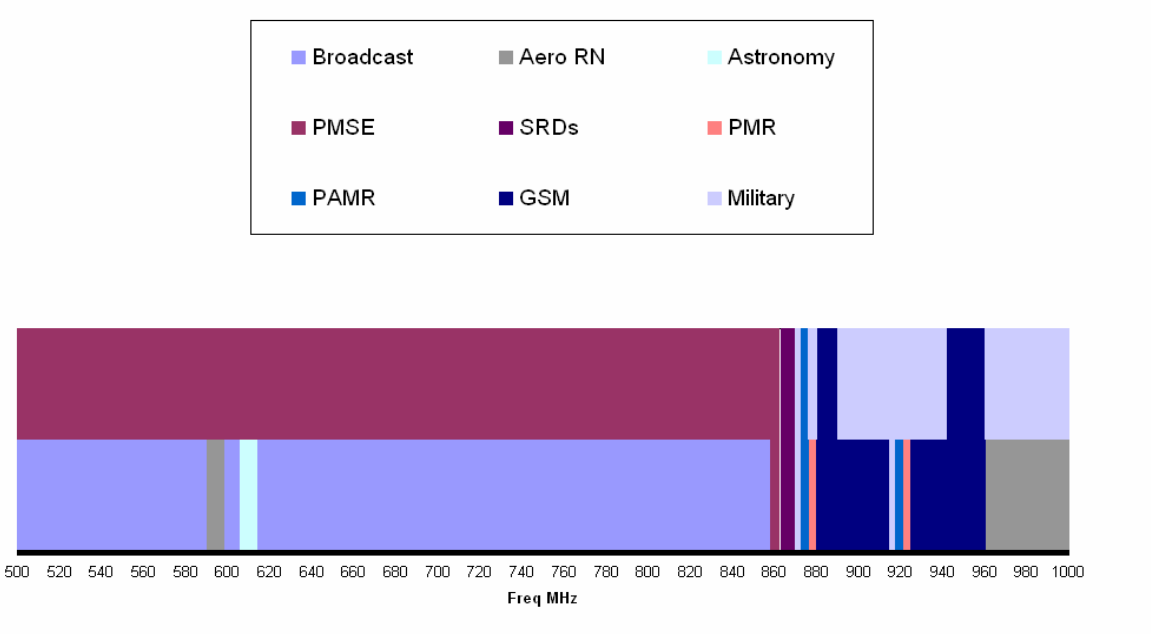

A.5 1000 – 1500 MHz

Most of the spectrum in this range is used by aeronautical radars (civil and military), though a sizeable chunk is also reserved for future DAB or similar applications. Potentially available spectrum has been identified between 1375 – 1389 MHz and 1399 – 1400 MHz. The intervening block (1389 – 1399 MHz) may also be available but is currently used by low power video links on a non-protected, non-interference basis. There are also two guard bands adjacent to the fixed link band: 1350 – 1350.5 MHz and 1374.5 – 1375 MHz.

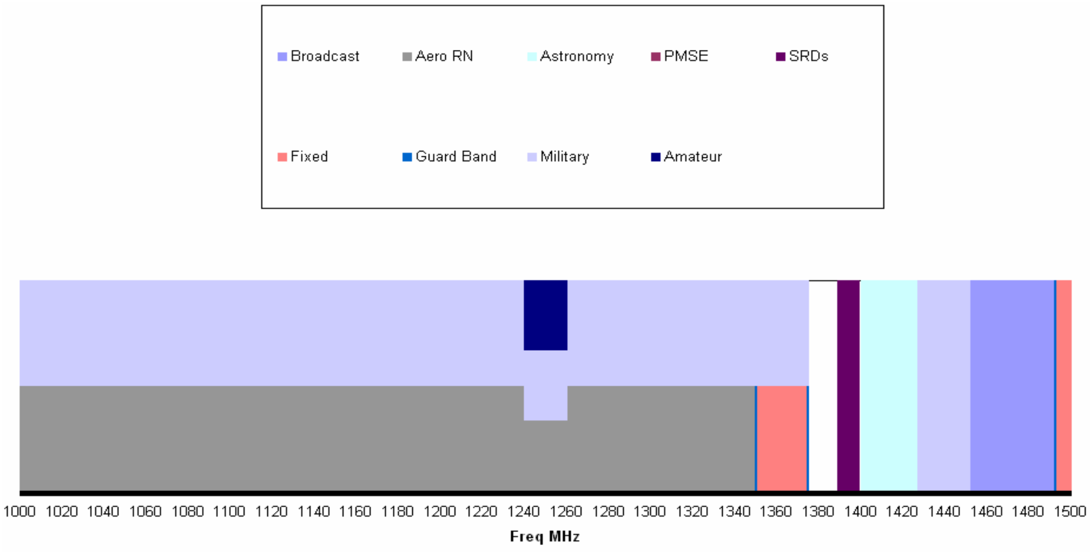

A.6 1500 – 2000 MHz

In this range, mobile cellular (2G and 3G) is the biggest allocation, though there are also sizeable allocations to mobile satellite and aeronautical services. The only spectrum identified as currently available is between the GSM base and mobile return bands, however this spectrum is already earmarked for release by Ofcom.

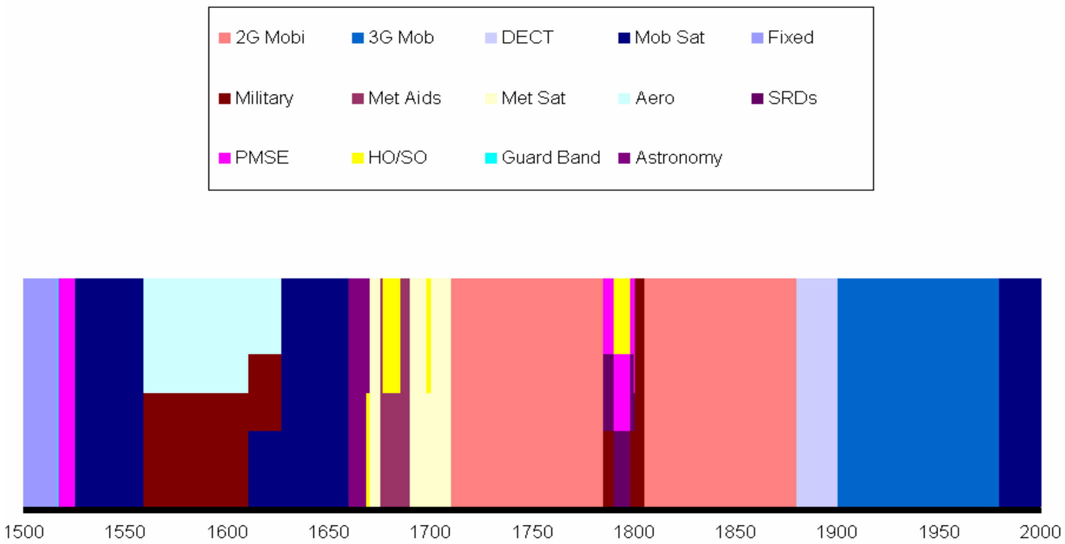

A.7 2000 – 2500 MHz

Military use dominates in this range, although there are also sizeable allocations to 3G mobile, PMSE and licence-exempt services (including WiFi at the top end of the band. Potentially available spectrum has been identified in the band 2290 – 2302 MHz, though this is used on a non-protected basis by radioastronomers at Joddrell Bank. The band 2302 - 2310 MHz is also unused currently but has been earmarked for release by Ofcom.

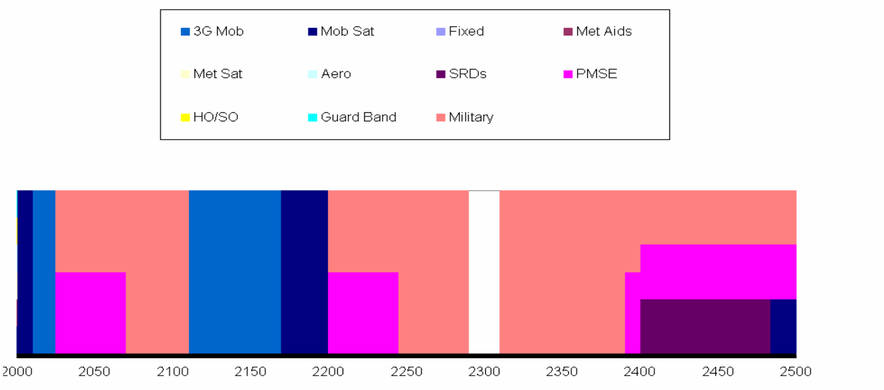

A.8 2500 – 3000 MHz

This band is almost entirely allocated to mobile and aeronautical services. The mobile spectrum is currently earmarked for future release by Ofcom, either as expansion spectrum for 3G mobile services or for auction on a technology and service neutral basis. The aeronautical spectrum is used internationally for aircraft radar and although it has been suggested in the past that sharing with other services may be feasible this is not considered an option at this stage. Hence no spectrum has been identified as available in this range.

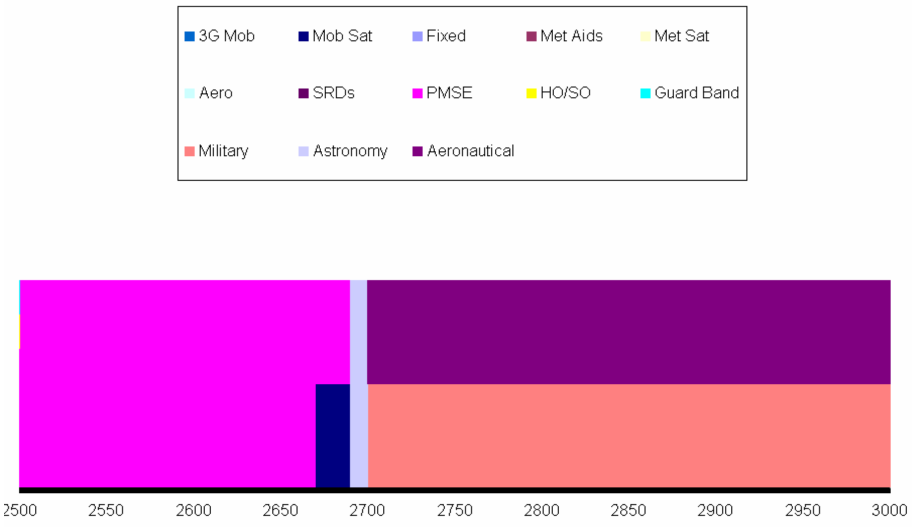

A.9 3000 – 3500 MHz

Spectrum below 3400 MHz is allocated to the military, above 3400 to public safety, PMSE and fixed wireless access. Unused spectrum is apparent at 3440 – 3442 MHz and 3475 – 3480 MHz.

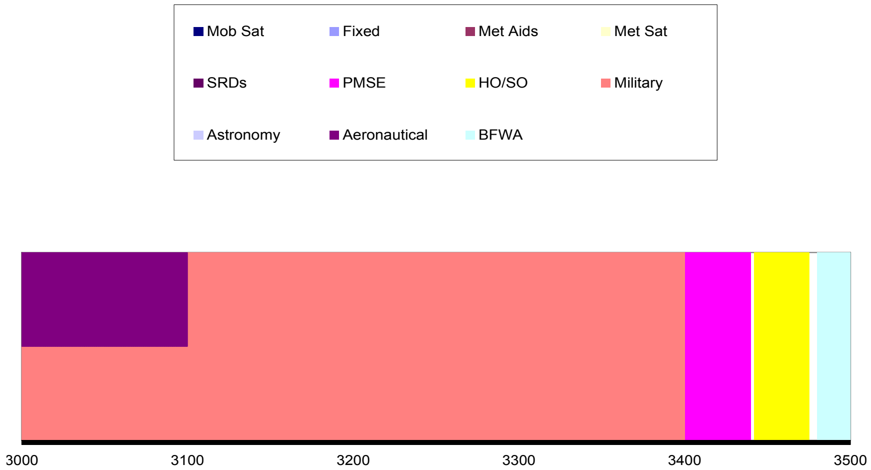

A.10 3500 – 4000 MHz

This band is mostly allocated to fixed links and fixed wireless access, with some PMSE use. There are number of designated guard bands, namely:

- 3600 – 3605 MHz (FWA)
- 3641 – 3650 MHz (FWA)
- 3875 – 3925 MHz (fixed links)
- 3961 – 3970 MHz (fixed links).

Although these are shared with satellite services, they could be used for aggregation purposes, subject to exclusion from to satellite earth stations.

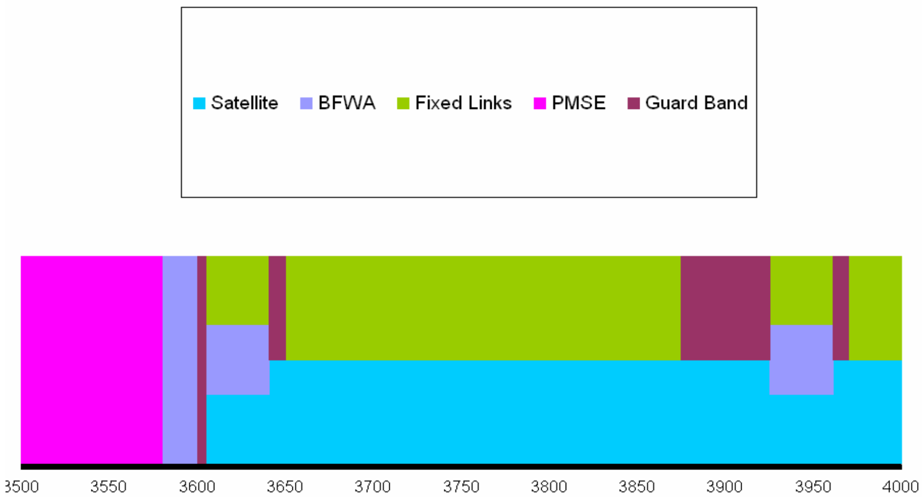

A.11 4000 – 4500 MHz

In this case, the lower part of the spectrum is allocated to fixed links and satellite earth stations, the upper part to aeronautical and military use. A guard band exists at the top of the fixed links band (4195 – 4200 MHz).

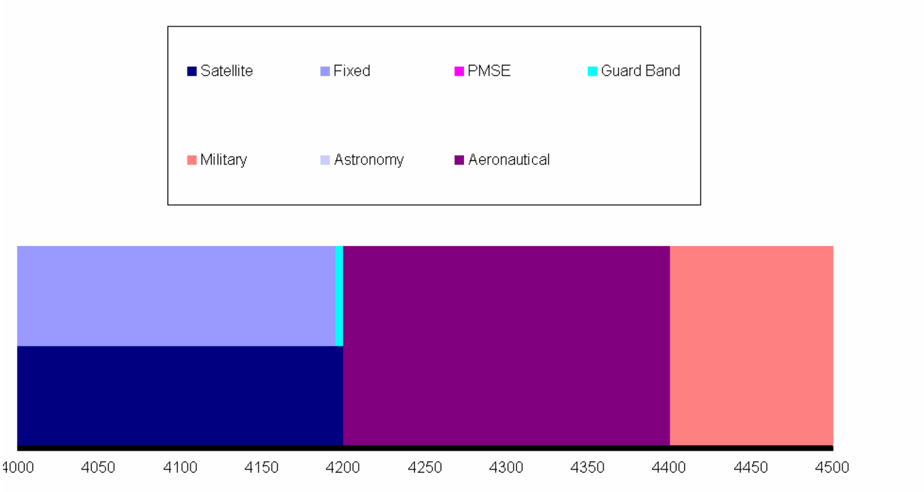

A.12 Unassigned spectrum within existing PMR allocations

The following frequencies have been identified as being unused within 100 km of the location specified (i.e. London, Newcastle-upon-Tyne, Brough and Ullapool) and are therefore, potentially available for aggregation purposes. Note that the “units” column in the tables refers to the number of 25 kHz spectrum aggregation units within each identified block.

A.12.1 London

Total unassigned spectrum identified in London: 2.65 MHz

Lower freq	Upper freq	Units	BW (MHx)
137.9625	138.0875	5	0.125
138.1125	138.2375	5	0.125
157.4625	157.4875	1	0.025
158.1875	158.2125	1	0.025
158.3875	158.4125	1	0.025
158.7375	158.7875	2	0.05
159.0625	159.1875	5	0.125
159.6875	159.7125	1	0.025
160.9875	161.1375	6	0.15
161.1875	161.2125	1	0.025
161.2625	161.2875	1	0.025
161.3125	161.3625	2	0.05
161.4375	161.4625	1	0.025
162.1125	162.1375	1	0.025
162.4375	162.4625	1	0.025
162.4875	162.5125	1	0.025
162.5625	162.5875	1	0.025
162.6125	162.6625	2	0.05
162.7125	162.8375	5	0.125
162.8625	162.8875	1	0.025
162.9875	163.0375	2	0.05

Lower freq	Upper freq	Units	BW (MHx)
163.1625	163.1875	1	0.025
163.2375	163.2875	2	0.05
163.5625	163.6875	5	0.125
164.1875	164.2125	1	0.025
165.8625	165.8875	1	0.025
170.6625	170.6875	1	0.025
173.0875	173.6375	22	0.55
173.6625	173.9875	13	0.325
452.9875	453.0125	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4625	454.4875	1	0.025
454.4375	454.4875	2	0.05

A.12.2 Newcastle-upon-Tyne

Total unassigned spectrum 11.275 MHz

Lower freq	Upper freq	Units	BW (MHz)
137.9625	138.0875	5	0.125
138.1125	138.2375	5	0.125
140.0625	140.1375	3	0.075
140.1625	140.3375	7	0.175
140.4125	140.5125	4	0.1
148.5125	148.5375	1	0.025
148.5875	148.6375	2	0.05
148.6875	148.8375	6	0.15
148.9125	148.9875	3	0.075
153.0875	153.3625	11	0.275
153.3875	153.4125	1	0.025
153.4375	153.5125	3	0.075
157.4625	157.4875	1	0.025
157.5125	157.5375	1	0.025
157.5625	157.8375	11	0.275
157.8625	157.9375	3	0.075
157.9625	157.9875	1	0.025
158.0125	158.1875	7	0.175
158.2125	158.4875	11	0.275
158.5125	158.9375	17	0.425
159.0375	159.1875	6	0.15
159.2625	159.5125	10	0.25
159.5375	160.2875	30	0.75
160.3125	160.5875	11	0.275
160.9875	161.1375	6	0.15
161.1625	161.2875	5	0.125

Lower freq	Upper freq	Units	BW (MHx)
161.3125	161.4125	4	0.1
161.4375	161.4625	1	0.025
162.0625	162.5375	19	0.475
162.5625	162.5875	1	0.025
162.6125	162.7875	7	0.175
162.8125	163.6875	35	0.875
163.7375	164.0125	11	0.275
164.0625	164.1625	4	0.1
164.1875	164.7875	24	0.6
164.8125	165.0375	9	0.225
166.1125	166.1375	1	0.025
168.4375	169.0125	23	0.575
169.5625	169.8375	11	0.275
170.7625	170.7875	1	0.025
173.0875	173.9875	36	0.9
452.9875	453.0125	1	0.025
454.3625	455.0125	26	0.65
455.5625	455.6125	2	0.05
456.3125	456.3375	1	0.025
456.3625	456.4625	4	0.1
456.4875	456.5125	1	0.025
456.5375	456.6625	5	0.125
456.6875	456.7625	3	0.075
456.7875	456.9625	7	0.175
458.8125	459.5125	28	0.7
460.9625	460.9875	1	0.025
461.0875	461.1125	1	0.025
461.1625	461.2875	5	0.125
461.4125	461.5125	4	0.1
461.8125	461.8375	1	0.025
461.8625	461.9375	3	0.075

A.12.3 Brough

Total unassigned spectrum: 4.925 MHz

Lower freq	Upper freq	Units	BW (MHx)
137.9625	138.0875	5	0.125
138.1125	138.2375	5	0.125
139.9625	139.9875	1	0.025
148.4625	148.4875	1	0.025
157.4375	157.6125	7	0.175
157.6875	157.7625	3	0.075
157.8375	157.9375	4	0.1
157.9625	158.0625	4	0.1
158.0875	158.2625	7	0.175
158.2875	158.3125	1	0.025
158.3375	158.3625	1	0.025
158.3875	158.4125	1	0.025
158.6625	158.6875	1	0.025
158.7375	158.7875	2	0.05
158.8875	158.9125	1	0.025

Lower freq	Upper freq	Units	BW (MHx)
159.0375	159.1875	6	0.15
159.2625	159.2875	1	0.025
159.3125	159.3375	1	0.025
159.6875	159.7125	1	0.025
160.9875	161.1375	6	0.15
161.1875	161.2125	1	0.025
161.2625	161.2875	1	0.025
161.3125	161.3625	2	0.05
161.4375	161.4625	1	0.025
162.0375	162.0625	1	0.025
162.0875	162.2125	5	0.125
162.2875	162.3625	3	0.075
162.4375	162.5375	4	0.1
162.5625	162.6625	4	0.1
162.6875	162.9125	9	0.225
162.9375	162.9625	1	0.025
162.9875	163.0375	2	0.05
163.1625	163.1875	1	0.025
163.2375	163.2875	2	0.05
163.3875	163.4125	1	0.025
163.5375	163.6875	6	0.15
163.7625	163.7875	1	0.025
163.8125	163.8375	1	0.025
164.1875	164.2125	1	0.025
165.8625	165.8875	1	0.025
165.9625	165.9875	1	0.025
169.4125	169.8375	17	0.425
170.6625	170.6875	1	0.025
173.0875	173.9875	36	0.9
452.9875	453.0125	1	0.025
454.1875	454.2375	2	0.05
454.4375	454.6875	10	0.25
454.7125	454.7625	2	0.05
454.8125	454.8375	1	0.025
456.2875	456.3375	2	0.05
456.3625	456.3875	1	0.025
456.4125	456.4375	1	0.025
460.6875	460.7375	2	0.05
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025
461.7375	461.7625	1	0.025

A.12.4 Ullapool

Total unassigned spectrum: 17.275 MHz

Lower freq	Upper freq	Units	BW (MHx)
137.9625	138.0875	5	0.125
138.1125	138.2375	5	0.125
139.9875	140.5125	21	0.525
148.5125	148.8625	14	0.35
148.8875	148.9875	4	0.1
153.0625	153.3625	12	0.3
153.3875	153.5125	5	0.125
157.4625	157.4875	1	0.025
157.5125	157.5375	1	0.025
157.5625	157.8375	11	0.275
157.8625	157.9375	3	0.075
157.9625	157.9875	1	0.025
158.0125	158.1875	7	0.175
158.2125	158.4875	11	0.275
158.5125	158.9375	17	0.425
159.0375	159.1875	6	0.15
159.2625	159.5125	10	0.25
159.5375	160.2875	30	0.75
160.3125	160.5875	11	0.275
160.9875	161.1375	6	0.15
161.1625	161.2875	5	0.125
161.3125	161.4125	4	0.1
161.4375	161.4625	1	0.025
162.0625	162.5375	19	0.475
162.5625	162.5875	1	0.025
162.6125	162.7875	7	0.175
162.8125	163.6875	35	0.875
163.7375	164.0125	11	0.275
164.0625	164.1625	4	0.1
164.1875	164.7875	24	0.6
164.8125	165.0375	9	0.225
165.2375	165.2625	1	0.025
165.4125	165.4375	1	0.025
165.8625	165.8875	1	0.025
165.9375	165.9875	2	0.05
166.1125	166.2125	4	0.1
166.6375	166.7875	6	0.15
166.8375	166.8875	2	0.05
166.9375	166.9625	1	0.025
166.9875	167.0125	1	0.025
167.0625	167.0875	1	0.025
167.1125	167.4125	12	0.3
167.5375	167.8625	13	0.325
167.8875	167.9375	2	0.05
167.9875	168.1875	8	0.2
168.2875	168.3125	1	0.025
168.8125	168.8375	1	0.025
168.8625	168.9375	3	0.075
169.4125	169.8125	16	0.4

Lower freq	Upper freq	Units	BW (MHx)
170.0375	170.0625	1	0.025
170.2125	170.2375	1	0.025
170.6625	170.6875	1	0.025
170.7375	170.7875	2	0.05
170.9125	171.0125	4	0.1
171.4375	171.5875	6	0.15
171.6375	171.6875	2	0.05
171.7375	171.7625	1	0.025
171.7875	171.8125	1	0.025
171.8625	171.8875	1	0.025
171.9125	172.2125	12	0.3
172.3375	172.6625	13	0.325
172.6875	172.7375	2	0.05
172.7875	172.9875	8	0.2
173.0875	173.9875	36	0.9
452.9875	453.0125	1	0.025
453.1375	453.1875	2	0.05
453.2375	453.3375	4	0.1
453.3625	453.3875	1	0.025
453.4375	453.4625	1	0.025
453.4875	453.5875	4	0.1
453.6625	453.8875	9	0.225
453.9125	454.3125	16	0.4
454.3375	455.0125	27	0.675
455.4875	455.6375	6	0.15
455.6625	455.7625	4	0.1
455.7875	455.8625	3	0.075
456.0125	456.3375	13	0.325
456.3625	456.4625	4	0.1
456.4875	456.5125	1	0.025
456.5375	456.6625	5	0.125
456.6875	456.7625	3	0.075
456.7875	456.9625	7	0.175
458.8125	459.5125	28	0.7
459.6375	459.6875	2	0.05
459.7125	459.9125	8	0.2
459.9375	459.9625	1	0.025
459.9875	460.0125	1	0.025
460.0625	460.0875	1	0.025
460.1625	460.3875	9	0.225
460.4125	460.4875	3	0.075
460.7375	461.0625	13	0.325
461.0875	461.2875	8	0.2
461.4125	461.8375	17	0.425
461.8625	461.9625	4	0.1
461.9875	462.1625	7	0.175
462.1875	462.2375	2	0.05
462.2625	462.4625	8	0.2

QinetiQ Proprietary

B Link Budgets

A spectrum aggregation system must co-exist with the existing users of the spectrum without causing interference or itself being adversely affected by interference. The ability to operate without causing interference *to* the existing band users is determined by the out-of-band emissions of the new system, which are controlled by the waveform in use and the performance of the transmitter PA. If a wideband implementation is used whereby many fragments are aggregated by digitising a large block of spectrum directly, a key factor for withstanding interference *from* the existing users will be the dynamic range of the ADC. An implementation digitising each fragment separately or using only a contiguous chunk of spectrum would not have such stringent dynamic range requirements, since pre-digitisation filtering could be used to remove unwanted signals.

Some link budget calculations based on example aggregating systems and the existing systems with which they would have to co-exist are summarised below. The calculations enable an estimate of dynamic range required for the case of digitising a large bandwidth (including many unwanted signals) and, taken together with out-of-band emission data for both systems, they identify possible inter-system interference issues.

In each case, a short range, intra-home link is considered (30 m) together with a long range link. The aggregating transmitter uses 125 mW for the intra-home link and 1 W for the long range link (Equivalent to 25 kHz and 200 kHz respectively at 5 W/MHz). Propagation losses are estimated using empirically derived link range exponents of 6 for the intra-home link (worst case indoor shadowed) and 2.7 for the long range and interfering links (best case urban). Antenna gains for the aggregating system are 0 dBi at VHF/UHF, and 6 dBi at microwave. Note that the dynamic ranges calculated are generalised figures for detection of the signal, the actual dynamic range requirements will be increased by a modulation-specific amount in a practical system.

<u>173 MHz:</u>

The first scenario considered is aggregating fragments identified around 173 MHz (in the VHF PMR band). For intra-home of link of 30 m, the wanted signal from the aggregating system suffers a propagation loss of:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^6\right) \approx 106dB$$

Assuming G_{TX} = G_{RX} = 0 dBi, a transmitter power of P_{TX} = 21 dBm (125 mW) would give:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_W = 21dBm + 0dB + 0dB - 106dB = -85dBm$$

A 'worst-case' interferer scenario is taken to be a taxi parked outside the house in which the aggregating system is operating. The taxi is taken to be 30 m distant from the aggregating receiver, and so propagation loss for interferer is given by:

$$L_I = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 57dB$$

The taxi transmits P_{TX} = 44 dBm (VHF PMR radios usually use 25W) into an antenna with G_{TX} = 3 dBi gain. This gives the received interferer signal strength:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_I = 44dBm + 3dB + 0dB - 57dB = -10dBm$$

Therefore a dynamic range of 75 dB plus the required SNR to demodulate the signal is required.

The long range link is D = 20 km at this frequency, giving:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 133dB$$

Again G_{TX} = G_{RX} = 0 dBi. If P_{TX} = 1 W = 30 dBm, this gives a received signal level of:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_{total} = 30dBm + 0dB + 0dB - 133dB = -103dBm$$

For the long range link, an estimate of the dynamic range is therefore 93 dB plus the required SNR to demodulate the signal.

<u>460 MHz:</u>

For intra-home of link of 30 m, the wanted signal from the aggregating system suffers a propagation loss of:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{6}\right) \approx 114dB$$

Assuming G_{TX} = G_{RX} = 0 dBi, a transmitter power of P_{TX} = 21 dBm (125 mW) would give:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_W = 21dBm + 0dB + 0dB - 114dB = -93dBm$$

Considering the available fragments at 452 – 462 MHz, the existing users are again PMR systems. Again the existing user is assumed to be located at a distance of 30 m and transmitting P_{TX} = 44 dBm into an antenna with G_{TX} = 3 dBi. This gives:

$$L_I = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 66dB$$

So the received power from the interferer is:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_I = 44dBm + 3dB + 0dB - 66dB = -19dBm$$

Therefore a dynamic range of 74 dB plus the required SNR to demodulate the signal is required.

The long range link is D = 10 km at this frequency, giving:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 134dB$$

Again G_{TX} = G_{RX} = 0 dBi. If P_{TX} = 1 W = 30 dBm, this gives a received signal level of:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_{total} = 30dBm + 0dB + 0dB - 134dB = -104dBm$$

For the long range link, an estimate of the dynamic range is therefore 85 dB plus the required SNR to demodulate the signal.

<u>650 MHz</u>

One of the most challenging spectrum environments in which to operate would be in the TV broadcast bands. This is presently not proposed, since spare TV channels are used for Programme Making and Special Events (PMSE), but as TV broadcasting changes and the switchover to digital occurs, it may be desirable for future aggregating systems to have this capability. TV transmitters operate at up to 1MW (90dBm) EIRP. At a distance of 1 km, the loss at 650 MHz (approximately the centre of the UHF TV broadcast band), using best case urban is:

$$L_I = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 109dB$$

Therefore, the received power assuming an isotropic receiving antenna would be -19 dBm, even at a distance of 1 km. This shows that the trading algorithm would have to define geographical exclusion zones around TV transmitters if aggregating systems were to operate at adjacent frequencies. In a rural area, the losses would be lower, and the exclusion zone would need to cover a larger area.

<u>1.4GHz:</u>

For intra-home of link of 30 m, the wanted signal from the aggregating system suffers a propagation loss of:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{6}\right) \approx 124dB$$

Assuming G_{TX} = G_{RX} = 6 dBi, a transmitter power of P_{TX} = 21 dBm (125 mW) would give:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_W = 21dBm + 6dB + 6dB - 124dB = -91dBm$$

The free spectrum identified at 1375 – 1389 MHz lies between allocations to Short Range Devices (SRDs) and Fixed Wireless Access (FWA). The FWA links will use considerably higher power, so they are considered as the existing user with which the aggregating system must co-exist. The maximum power which will normally be assigned to an FWA link at 1.4 GHz is 40 dBW EIRP, or 70 dBm EIRP (OfW 46, 2004). At D = 30 m and f = 1.4 GHz, the loss for the interferer is:

$$L_I = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 75dB$$

So the received power from the interferer is:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_I = 70dBm + 6dB - 75dB = 1dBm$$

Therefore a dynamic range of 92 dB plus the required SNR to demodulate the signal is required.

The long range link is D = 3 km at this frequency, giving:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 129dB$$

Again $G_{TX} = G_{RX}$ = 6 dBi. If P_{TX} = 1 W = 30 dBm, this gives a received signal level of:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_{total} = 30dBm + 6dB + 6dB - 129dB = -87dBm$$

For the long range link, an estimate of the dynamic range is therefore 88 dB plus the required SNR to demodulate the signal.

<u>3.4GHz:</u>

For intra-home of link of 30 m, the wanted signal from the aggregating system suffers a propagation loss of:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{6}\right) \approx 131dB$$

Assuming $G_{TX} = G_{RX}$ = 6 dBi, a transmitter power of P_{TX} = 21 dBm (125 mW) would give:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_W = 21dBm + 6dB + 6dB - 131dB = -98dBm$$

For the fragments identified around 3.4 GHz, the primary users are Fixed Wireless Access (FWA) links and Police helicopter downlinks. FWA typically uses of order P_{TX} = 30 dBm (1 W) and antenna gains up to G_{TX} = 30 dB for point-to-point operation. The maximum power normally assigned is 80 dBm EIRP according to (Ofw 30, 2004). Consider an FWA link 30 m away from the aggregating receiver. At D = 30 m and f = 3.4 GHz interferer loss is:

$$L_I = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 83dB$$

So the received power from the interferer is:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_I = 80dBm + 6dB - 83dB = 3dBm$$

Therefore a dynamic range of 101 dB plus the required SNR to demodulate the signal is required.

The long range link is D = 1 km at this frequency, giving:

$$L_W = 20\log\left(\frac{4\pi f}{c}\right) + 10\log\left(D^{2.7}\right) \approx 124dB$$

Again G_{TX} = G_{RX} = 6 dBi. If P_{TX} = 1 W = 30 dBm, this gives a received signal level of:

$$P_{RX} = P_{TX} + G_{TX} + G_{RX} - L_{total} = 30dBm + 6dB + 6dB - 124dB = -82dBm$$

For the long range link, an estimate of the dynamic range is therefore 85 dB plus the required SNR to demodulate the signal.

C Costs of aggregation and non-aggregating systems

C.1 Introduction

To consider the costs of developing a spectrum aggregating device a number of systems are considered below. It should be noted that these are illustrative for this study and the costs and designs should not be used for any other purpose.

C.2 Satellite communications terminal

UHF SATCOM is used primarily by the military for two-way voice communications. The number of user is therefore limited to several hundred. However, similar principles can be applied to other SATCOM services which deliver commercial satellite voice communications, of which there are estimated to be tens of thousands of users in the European region alone.

C.2.1 Conventional Design

Figure C-1 is a more detailed RF architecture which shows the individual RF components of the two chains that lie within the RF, conversion, IF and AF stages.

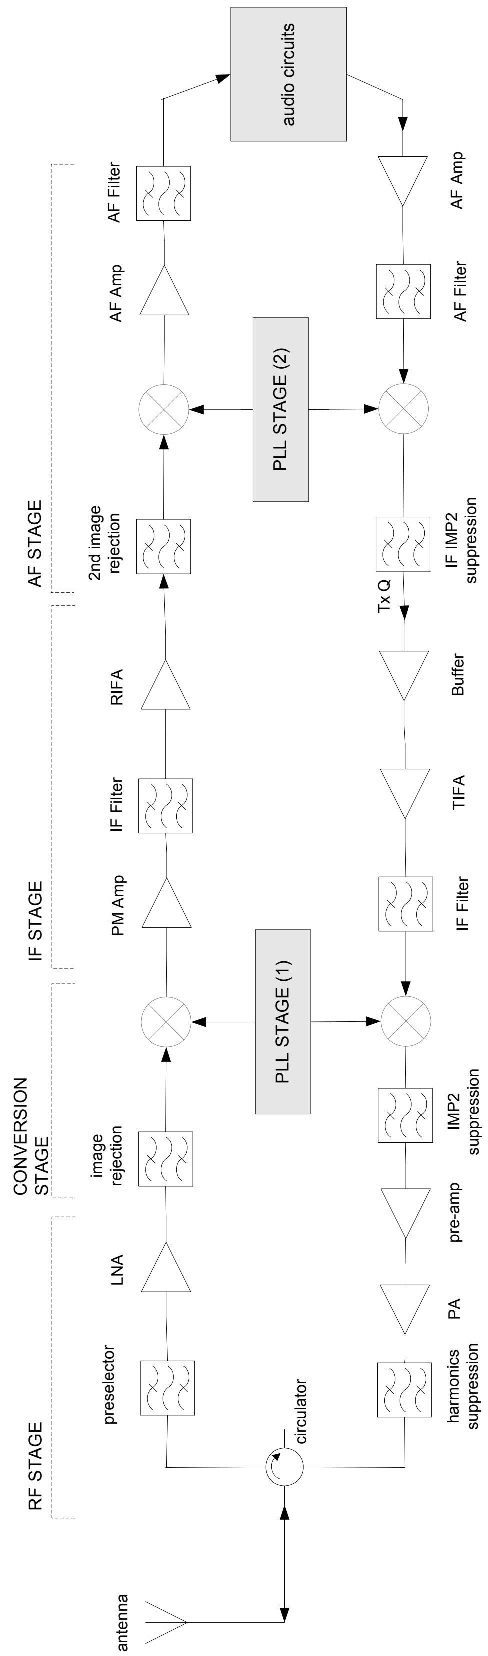

Figure C-1: RF Architecture for a conventional SATCOM terminal.

The PLL stage is considered separately and shown in the diagram below, Figure C-2. A dual synthesiser is used to select and generate the local oscillator (LO) frequencies, enabling the terminal to be tuned to the correct channel. The TCXO is a highly stable reference that determines the stability of the rest of the loop. LO amplifiers, on both RF and IF outputs, boost the signal to a level required by the relevant VCO. As for the VCOs themselves, the IF VCO can be narrowband as it does not need to change (the IF will remain the same regardless of RF tuning). The RF VCO, however, needs to be wideband so as to encompass the entire range of channels to which it needs to tune. A low pass filter (LPF) is placed within each phase locked loop to suppress frequency spurs. At the end of each loop output a splitter is used so that the LO can be shared between receive and transmit chains.

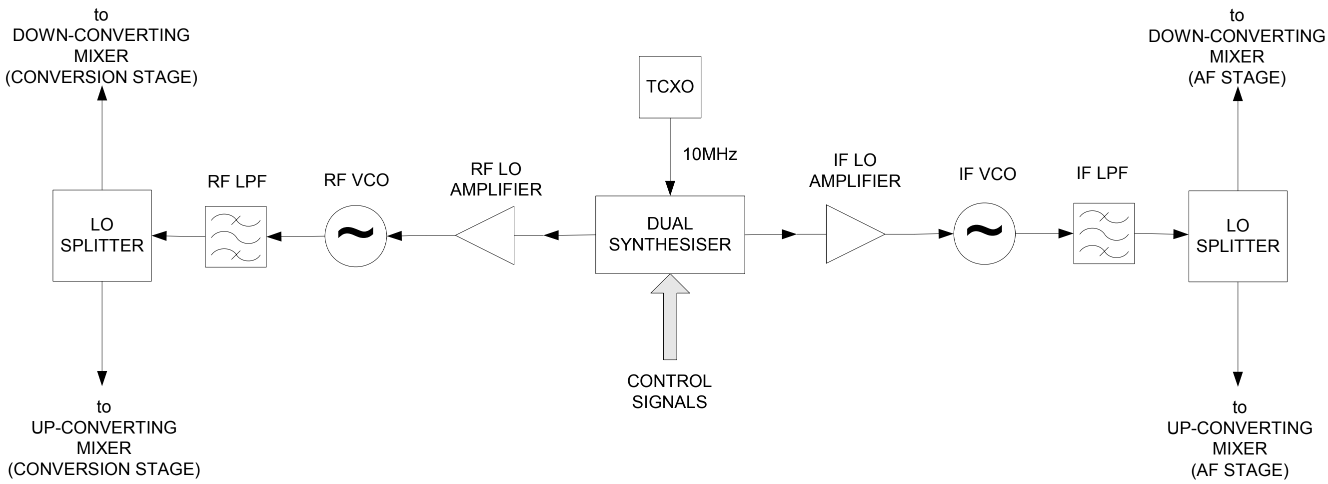

Figure C-2: Phase-Locked Loop Stage for a SATCOM terminal

Listed in Table C-1 are the prices for the analogue components contained within the RF, conversion, IF, AF and PLL stages of the system.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	Preselector	Kel-Com 2KCB10-255/T38.25-1.1	1	55.73	55.73
2	LNA	Agilent MGA-82563 + SGA-4186	1	2.36	2.36
		Sirenza SGA-4186	1	1.00	1.00
3	Image Reject Filters	Kel-Com 3KCB20-255/T38.25-1.1	2	55.73	111.46
4	Mixers	ADE-12MH	4	4.42	17.68
5	PM Amp	Watkins-Johnson AG604-86	1	1.14	1.14
6	IF Filters	70MHz SAW, 854651	2	48.72	97.44
7	RIFA & TIFA	Watkins-Johnson AG604-86	2	1.14	2.28
		Sirenza SGA-4186	1	1.00	1.00
8	AF Amplifiers	Differential Amplifier DRV1100P	2	5.97	11.94
9	AF Filters	IC Filter MF10CCN	2	3.26	6.52
10	IMP2 suppression filters	Kel-Com 3KCB20-305/T38.25-1.1	2	55.73	111.46
11	Tx IF Buffer	MAX436	1	5.50	5.50
12	Pre-amp	Watkins-Johnson AG604-86	1	1.14	1.14
		Sirenza SGA-4186	1	1.00	1.00
13	Power Amp (PA)	Polyfet PHM020	1	147.00	147.00
14	Harmonics supression filter	Kel-Com 4KCB20-305/T38.25-1.1	1	55.73	55.73
15	Circulator	Rennaisance 3A2BC 225-400MHz	1	1000.00	1000.00
16	Antenna	High gain UHF SATCOM antenna	1	400.00	400.00
				RF SUB-TOTAL =	**£2,030.38**
17	LO splitters	Mini-Circuits SCP models	2	10.45	20.90
18	Low pass filters	Lumped element custom design	2	0.50	1.00
19	RF VCO	UMC UMZ-362-A16	1	36.00	36.00
20	LO Amplifiers	Agilent MGA-82563	2	2.36	4.72
21	Dual Synthesiser	ADF4252 (Analog Devices)	1	3.10	3.10
22	TCXO	Vectron TC-350-CAF-106	1	200.90	200.90
23	IF VCO	Z-Comm SMV0135A	1	17.50	17.50
				PLL SUB TOTAL =	**£284.12**
				TOTAL =	**£2,314.50**

Table C-1: Non-aggregating SATCOM terminal cost breakdown

C.2.2 Spectrum-Aggregating Design

To turn the SATCOM terminal into a spectrum aggregating device, multiple receive and transmit chains are needed for each fragment. This could result in the component count escalating dramatically, save for the fact that the channel bandwidth is narrow at 25kHz so it is unlikely that it would be split into more than five fragments (of 5kHz each, perhaps).

As the conventional design stands, only the IF filters need to have their pass-band narrowed in order for the terminal to operate on a single fragment. Then it is a simple matter of adding more chains, one pair (transmit and receive) for each additional fragment.

However, the entire chain does not have to be replicated. It is feasible for all of the components in the RF stage, and some in the conversion stage, to be shared by all fragments, since these components are already wideband enough to encompass the entire tuneable band. This assumes that all usable fragments will lie somewhere within the band, which is a realistic assumption for SATCOM as using out-of-band fragments would mean changing the space element (i.e. launching new satellites). In the receiver chain the preselector, LNA and image rejection filter can be shared. In the transmit chain, the harmonics suppression filter can be shared. The IMP2 suppression filter could also be shared, but the pre-amp and PA cannot, and it is not worth combining transmit chains together to share the filter just have them split back out to separate amplifiers, then combined back to share the harmonics suppression filter. In the RF stage the circulator can also be shared, and so can the antenna.

The rest of the components in the RF, conversion and IF stages will need to be replicated, one replication per additional chain/fragment. However, the AF stage can be replaced entirely by analogue-to-digital and digital-to-analogue conversion, since combining of the fragments into a coherent message will require digital signal processing. This means that only one LO drive is required per chain, i.e. one dual synthesiser per transmit/receiver chain pair. Thus some saving can be made on AF and LO components.

A component level description of this system (Figure C-3) which indicates additional components required in red.

QinetiQ Proprietary

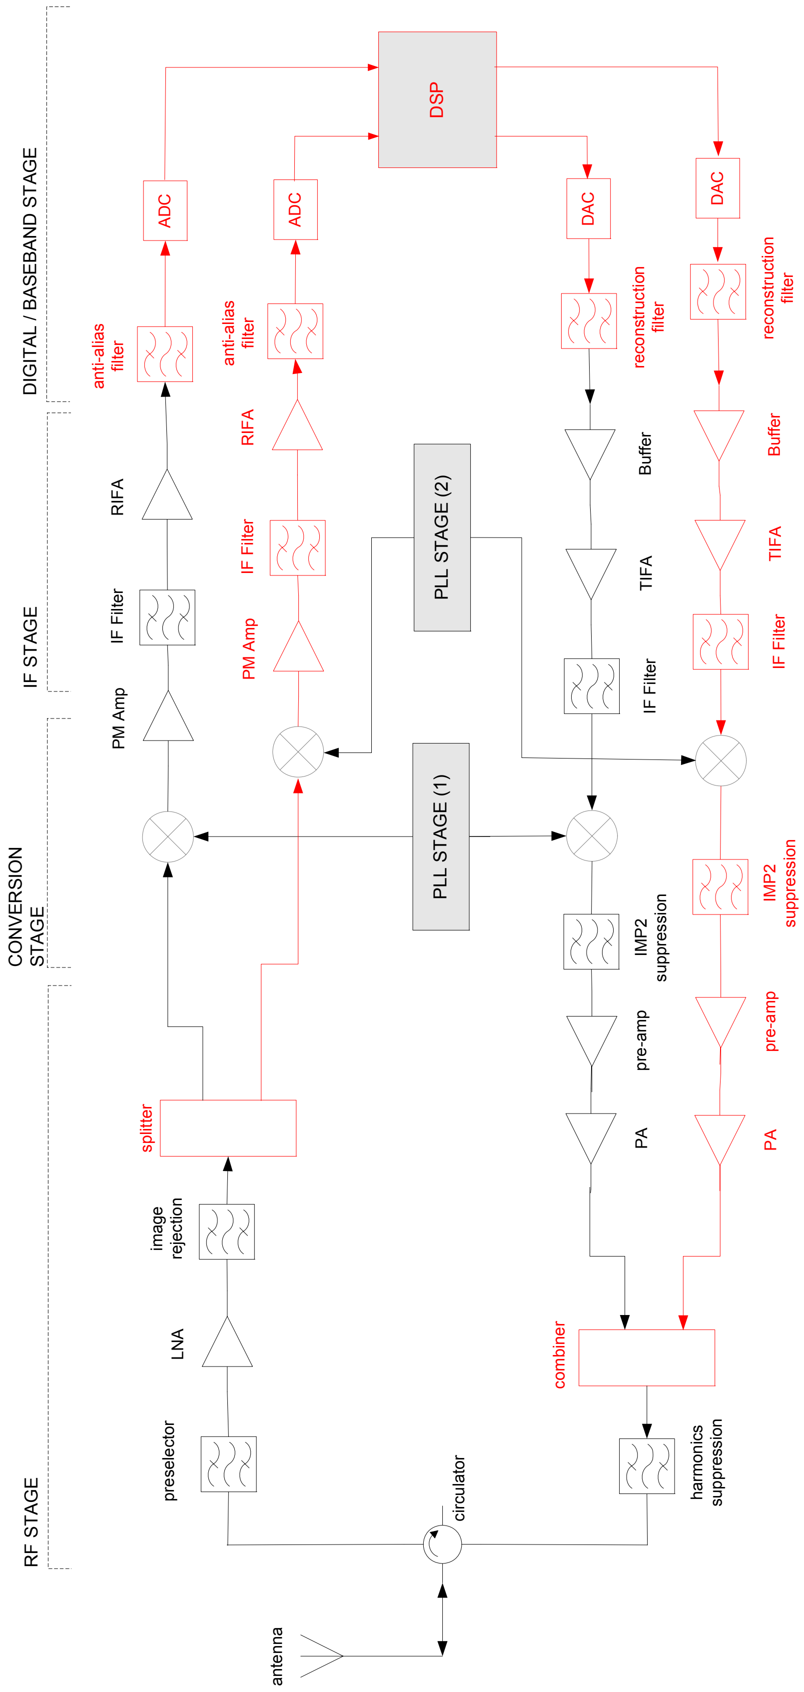

Figure C-3: RF Architecture for 2-fragment aggregating SATCOM terminal

The cost of these additional components is given in Table C-2 below. Compared to the cost of a non-aggregating SATCOM radio (£2,314.50), this additional cost is equivalent to a 73% increase in cost. The majority of this extra cost is the DSP/FPGA module required for processing the fragments. DSPs and FPGAs are still very expensive devices. Without the DSP module, the increase would only be 21%.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Mini-Circuits SBTC-2-10-75	1	3.49	3.49
2	2-way combiner	Mini-Circuits ZA2C5-500-15W	1	74.95	74.95
3	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
4	PM Amp	Watkins-Johnson AG604-86	1	1.14	1.14
5	IF Filters	70MHz low-loss SAW filter, p/n 854651	2	48.72	97.44
6	Anti-alias filters	IC Filter MF10CCN	2	3.26	6.52
7	RIFA	Watkins-Johnson AG604-86	1	1.14	1.14
		Sirenza SGA-4186	1	1.00	1.00
8	ADCs	Analog Devices AD6644	2	29.00	58.00
9	IMP2 suppression filters	Kel-Com Ceramic p/n 3KCB20-305/T38.25-1.1	1	55.73	55.73
10	Pre-amp	Watkins-Johnson AG604-86	1	1.14	1.14
		Sirenza SGA-4186	1	1.00	1.00
11	Power Amp (PA)	Polyfet PHM020	1	147.00	147.00
12	TIFA	Watkins-Johnson AG604-86	1	1.14	1.14
13	Tx IF Buffer	MAX436	1	5.50	5.50
14	Reconstruction filters	IC Filter MF10CCN	2	3.26	6.52
15	DACs	Analog Devices AD9772A	2	14.95	29.90
				RF SUB-TOTAL =	**£1,691.61**
				PLL SUB TOTAL =	**£0.00**
				TOTAL =	**£1,691.61**

Table C-2: Cost of additional components for 2-fragment spectrum aggregator

It is not a simple matter to say that each additional fragment will cost £1,691.61, for three reasons. First, for third and subsequent fragments, an additional dual synthesizer will be needed. Second, the two-way splitter and combiner components need replacing with n-way components, where n is the total number of fragments. As n increases, the cost of the splitter and combiner increase, but not necessarily linearly. This is because 2-, 3- and 4-way splitters are more common than larger n-way devices. The third reason is that a DSP or FPGA module is required for spectrum aggregation, but may be able to process several fragments before a larger or additional module is needed.

Table C-3, therefore, shows the additional cost (relative to the non-aggregating system) of a three-fragment system. This is 104% more expensive than the non-aggregating system. Similarly, Table C-4 shows the additional cost of a four-fragment aggregating system, which is 116% more expensive. Table C-5 shows the additional cost of a five-fragment aggregating system, which is 147% more expensive. Finally, the results for further costings up to ten fragments is shown in

Table 5-2. These results indicate that the cost increases on average by 29% per fragment, or £670. However we have not considered the need for additional DSP/FPGA processing at higher fragments, so this trend will not hold indefinitely. At the threshold number of fragments which requires an additional or better DSP, the trend will be interrupted by a step increase in cost.

Note there is nothing to stop a terminal from have different number of receiver fragments to the number of transmit fragments, so long as a complementary terminal is available at the other end of the link to receive/transmit the correct fragments.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	3-way splitter	Mini-Circuits JPS-3-1W	1	14.95	14.95
2	3-way combiner	Mini-Circuits ZB3CS-920-15W	1	114.95	114.95
3	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
4	Mixers	ADE-12MH	2	4.42	8.84
5	PM Amp	Watkins-Johnson AG604-86	2	1.14	2.28
6	IF Filters	70MHz low-loss SAW filter, p/n 854651	4	48.72	194.88
7	Anti-alias filters	IC Filter MF10CCN	3	3.26	9.78
8	RIFA	Watkins-Johnson AG604-86	2	1.14	2.28
		Sirenza SGA-4186	2	1.00	2.00
9	ADCs	Analog Devices AD6644	3	29.00	87.00
10	IMP2 suppression filters	Kel-Com C3KCB20-305/T38.25-1.1	2	55.73	111.46
11	Pre-amp	Watkins-Johnson AG604-86	2	1.14	2.28
		Sirenza SGA-4186	2	1.00	2.00
12	Power Amp (PA)	Polyfet PHM020	2	147.00	294.00
13	TIFA	Watkins-Johnson AG604-86	2	1.14	2.28
14	Tx IF Buffer	MAX436	2	5.50	11.00
15	Reconstruction filters	IC Filter MF10CCN	3	3.26	9.78
16	DACs	Analog Devices AD9772A	3	14.95	44.85
			RF SUB-TOTAL =		**£2,114.61**
17	LO splitters	Mini-Circuits SCP models	2	10.45	20.90
18	Low pass filters	Lumped element custom design	2	0.50	1.00
19	RF VCO	UMC UMZ-362-A16	1	36.00	36.00
20	LO Amplifiers	Agilent MGA-82563	2	2.36	4.72
21	Dual Synthesiser	ADF4252 (Analog Devices)	1	3.10	3.10
22	TCXO	Vectron TC-350-CAF-106	1	200.90	200.90
23	IF VCO	Z-Comm SMV0135A	1	17.50	17.50
			PLL SUB TOTAL =		**£284.12**
			TOTAL =		**£2,398.73**

Table C-3: Cost of additional components for 3-fragment spectrum aggregator

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	4-way splitter	Mini-Circuits PSC-4-1	1	37.95	37.95
2	4-way combiner	Mini-Circuits ZB4CS-960-12W	1	134.95	134.95
3	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
4	Mixers	ADE-12MH	4	4.42	17.68
5	PM Amp	Watkins-Johnson AG604-86	2	1.14	2.28
6	IF Filters	70MHz low-loss SAW filter, p/n 854651	4	48.72	194.88
7	Anti-alias filters	IC Filter MF10CCN	2	3.26	6.52
8	RIFA	Watkins-Johnson AG604-86	2	1.14	2.28
		Sirenza SGA-4186	2	1.00	2.00
9	ADCs	Analog Devices AD6644	2	29.00	58.00
10	IMP2 suppression filters	Kel-Com C3KCB20-305/T38.25-1.1	2	55.73	111.46
11	Pre-amp	Watkins-Johnson AG604-86	2	1.14	2.28
		Sirenza SGA-4186	2	1.00	2.00
12	Power Amp (PA)	Polyfet PHM020	2	147.00	294.00
13	TIFA	Watkins-Johnson AG604-86	2	1.14	2.28
14	Tx IF Buffer	MAX436	2	5.50	11.00
15	Reconstruction filters	IC Filter MF10CCN	2	3.26	6.52
16	DACs	Analog Devices AD9772A	2	14.95	29.90
				RF SUB-TOTAL =	**£2,115.98**
17	LO splitters	Mini-Circuits SCP models	4	10.45	41.80
18	Low pass filters	Lumped element custom design	4	0.50	2.00
19	RF VCO	UMC UMZ-362-A16	2	36.00	72.00
20	LO Amplifiers	Agilent MGA-82563	4	2.36	9.44
21	Dual Synthesiser	ADF4252 (Analog Devices)	2	3.10	6.20
22	TCXO	Vectron TC-350-CAF-106	2	200.90	401.80
23	IF VCO	Z-Comm SMV0135A	2	17.50	35.00
				PLL SUB TOTAL =	**£568.24**
				TOTAL =	**£2,684.22**

Table C-4: Cost of additional components for 4-fragment spectrum aggregator

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	5-way splitter	Mini-Circuits ZBSC-5-1	1	119.95	119.95
2	5-way combiner	Mini-Circuits ZBSC-5-1	1	119.95	119.95
3	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
4	Mixers	ADE-12MH	6	4.42	26.52
5	PM Amp	Watkins-Johnson AG604-86	3	1.14	3.42
6	IF Filters	70MHz low-loss SAW filter, p/n 854651	6	48.72	292.32
7	Anti-alias filters	IC Filter MF10CCN	3	3.26	9.78
8	RIFA	Watkins-Johnson AG604-86	3	1.14	3.42
		Sirenza SGA-4186	3	1.00	3.00
9	ADCs	Analog Devices AD6644	3	29.00	87.00
10	IMP2 suppression filters	Kel-Com C3KCB20-305/T38.25-1.1	3	55.73	167.19
11	Pre-amp	Watkins-Johnson AG604-86	3	1.14	3.42
		Sirenza SGA-4186	3	1.00	3.00
12	Power Amp (PA)	Polyfet PHM020	3	147.00	441.00
13	TIFA	Watkins-Johnson AG604-86	3	1.14	3.42
14	Tx IF Buffer	MAX436	3	5.50	16.50
15	Reconstruction filters	IC Filter MF10CCN	3	3.26	9.78
16	DACs	Analog Devices AD9772A	3	14.95	44.85
				RF SUB-TOTAL =	**£2,554.52**
17	LO splitters	Mini-Circuits SCP models	6	10.45	62.70
18	Low pass filters	Lumped element custom design	6	0.50	3.00
19	RF VCO	UMC UMZ-362-A16	3	36.00	108.00
20	LO Amplifiers	Agilent MGA-82563	6	2.36	14.16
21	Dual Synthesiser	ADF4252 (Analog Devices)	3	3.10	9.30
22	TCXO	Vectron TC-350-CAF-106	3	200.90	602.70
23	IF VCO	Z-Comm SMV0135A	3	17.50	52.50
				PLL SUB TOTAL =	**£852.36**
				TOTAL =	**£3,406.88**

Table C-5: Cost of additional components for 5-fragment spectrum aggregator

C.2.3 Additional Band Design

So far we have only considered adding fragments that lie within the original band of the SATCOM terminal (290-320MHz or 240-270MHz). To use fragments that are outside of this band, components in the RF stage need to be replaced with wider band parts or be duplicated to allow more than one band at a time. Duplication is the cheaper of the two, since wideband components are rare (and hence expensive), and this approach also allows for more flexibility (it may be possible to tune the bands, or at the very least have two bands that are not necessarily neighbouring each other in the spectrum).

Figure C-4 shows the detailed architecture of the modifications required for the RF stage, for a two-band aggregator with 2 fragments per band. Again, the additional/replacement components are highlighted in red. The cost of these components is shown in the following Table, and also in Table C-7 for a dual-band, five-fragment design.

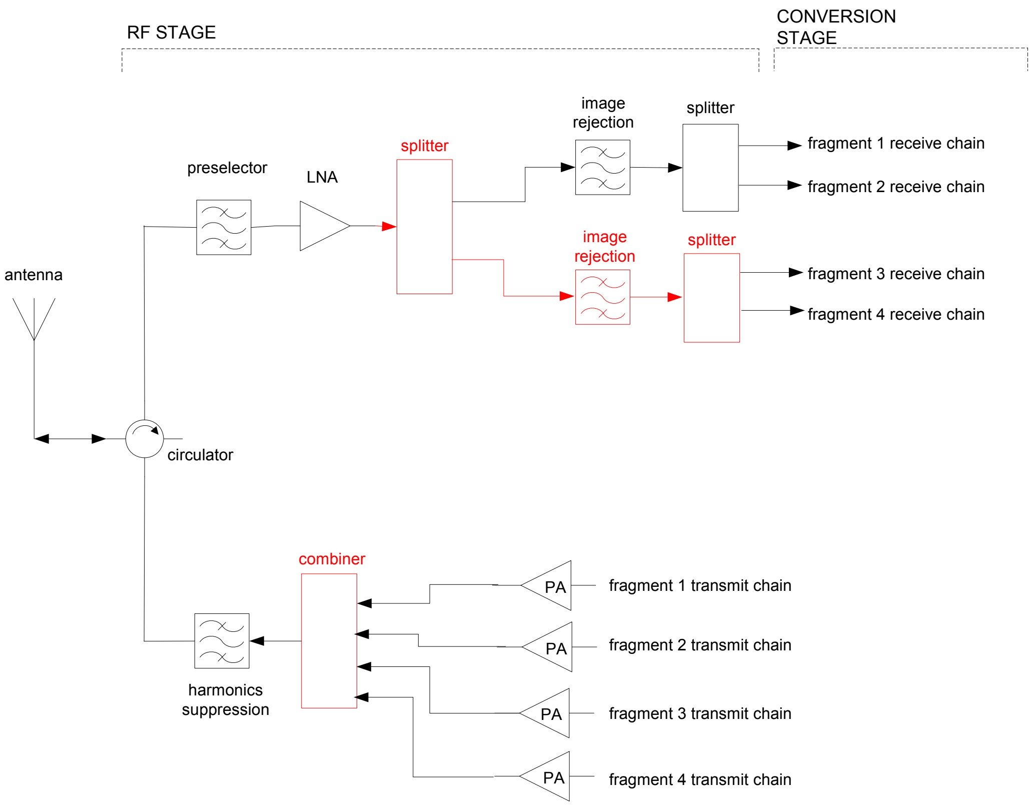

Figure C-4: RF stage architecture for a dual-band aggregating SATCOM terminal

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Mini-Circuits SBTC-2-10-75	2	3.49	6.98
2	Image Reject Filter	Kel-Com 3KCB20-255/T38.25-1.1	1	55.73	55.73
3	4-way combiner	Mini-Circuits ZB3CS-920-15W	1	114.95	114.95
		ADDITIONAL RF SUB-TOTAL =			**£177.66**

Table C-6: Cost of additional components for using four fragments across two bands

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Mini-Circuits SBTC-2-10-75	1	3.49	3.49
2	3-way splitter	Mini-Circuits JPS-3-1W	1	14.95	14.95
3	Image Reject Filter	Kel-Com 3KCB20-255/T38.25-1.1	1	55.73	55.73
4	5-way combiner	Mini-Circuits ZB3CS-920-15W	1	119.45	119.95
				ADDITIONAL RF SUB-TOTAL =	**£194.12**

Table C-7: Cost of additional components for using five fragments across two bands

This shows that dividing fragments across different bands does cost extra, but not very much. £177.66 is only 7.7% of the original non-aggregating design cost. The dual-band five-fragment design incurs a similar 8.4% increase. Compared to their 4- and 5-fragment single-band counterparts, the extra cost is only 3.5%. This is to be expected since the majority of component duplication is caused by adding a fragment to the system, requiring extra components in the conversion, IF and PLL stages, not the RF stage. Because most of this duplication occurs in the conversion, IF and PLL stages, the value of the carrier frequency – and hence the upper and lower limits of the bands – will have little impact on the cost of each fragment. Also, the size of a fragment will not affect the hardware cost, since final fragment size is determined by digital filtering within the DSP which can be changed via software.

Therefore the size of fragments and how they are distributed across multiple bands is not as significant as the total number of fragments and the total number of bands.

C.3 Private Business Radio (PBR) Transceiver

In this section we detail the design and costs of building an aggregator for VHF hand-portable or mobile PBR voice terminal with some limited data handling such as scrambling, 2/5 tone calling, and dual tone multi frequency (DTMF) and trunk networking capability.

C.3.1 Conventional Design

The outline design for a common hand-portable transceiver is given in Figure C-6. All PBR radios will constitute this basic design, with more advanced devices having additional hardware for various features such as trunk networking, or simply better quality parts for improved robustness and performance.

VHF is a low enough frequency for PBR radios to be made of discrete electronic components such as transistors, inductors, capacitors as well as some integrated circuits, as opposed to dedicated RF analogue components. Thus PBR design is split into a number of circuit blocks, categorised into three groups, as follows:

- Receiver Circuits:-
    - Antenna Switching Circuit

- RF Circuit
- 1st Mixer & IF Circuit
- 2nd IF & Demodulator Circuit
- AF Amplifier Circuit
- Squelch Circuit

- Transmitter Circuits:-
  - Microphone Amplifier
  - Modulation Circuit
  - Drive Amplifier Circuit
  - RF Power Amplifier Circuit
  - APC Circuit
- PLL Circuits:-
  - Reference Oscillator Circuit
  - Programmable Divider & Phase Detection Circuits
  - Charge Pump & Loop Filter Circuits
  - VCO Circuit

From a design point of view it is useful to split the design into four units:

- RF unit
- Main unit
- VCO board
- Logic unit

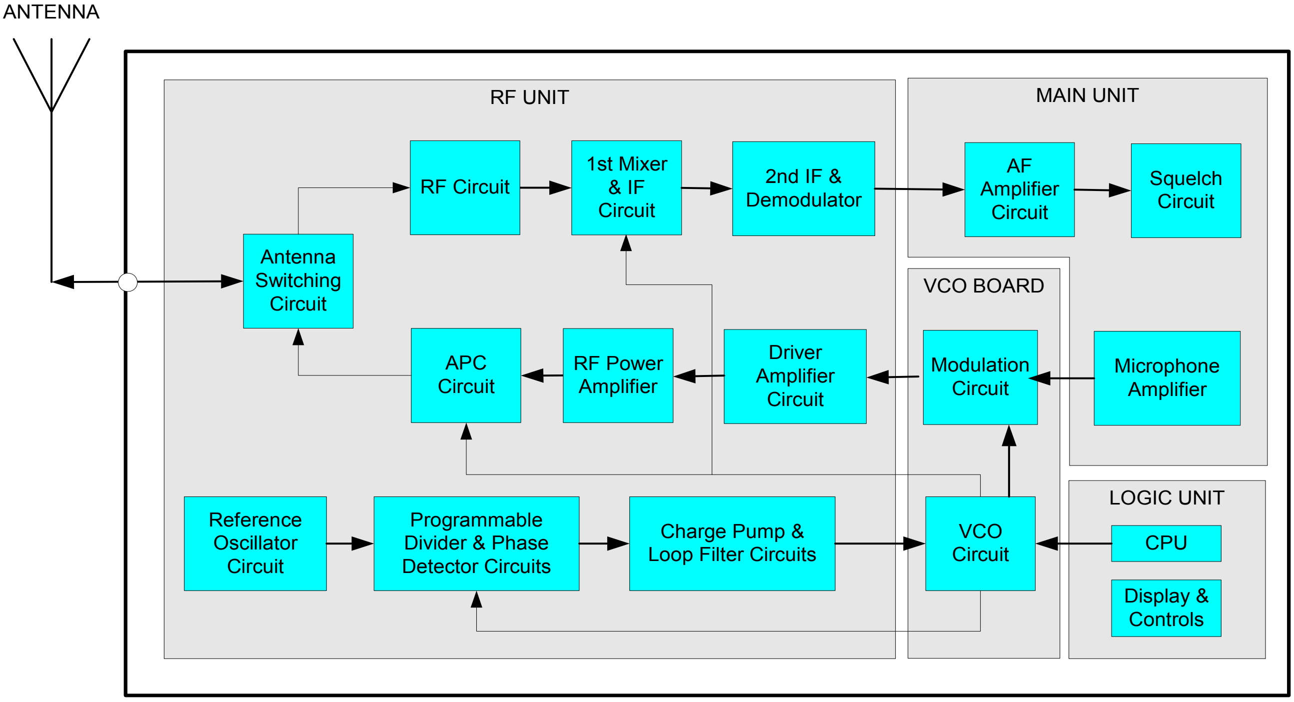

Figure C-5:: Outline design for PBR hand-portable

The antenna switching circuit functions as a low-pass filter while receiving, but upon detection of higher power transmit signals it goes into a high impedance state and blocks transmit signals from entering the delicate receiver circuits. This allows sharing of one antenna for both receive and transmit functions.

The RF contains amplifiers and filters to boost and clean up the weak and noisy received signal. The filters are tuneable and controlled by the PLL circuit, which allows the device to choose receive signals across the wide 146 – 174MHz band but at the same time offering good image rejection for the chosen frequency of interest.

The first mixer and IF circuits in the receive chain convert the receive RF into a lower IF, using PLL output from the VCO board. Crystal filters and transistors are used to filter and boost the IF signal, before it is passed onto the second mixer and IF circuit. This double superheterodyne approach improves the image rejection and stability of receiver gain. Typically, the second IF will be as low as a few hundred kHz; 455kHz is a common value. Such low frequencies mean that the second IF stage is usually implemented on a single integrated circuit with a minimum number of peripheral components, and its own local oscillator (LO) circuit is used (rather than a second separate set of PLL circuits).

A demodulator circuit (again, usually an IC) is used to convert the second IF into audio frequency (AF) signals that constitute the wanted message. CTCSS/DCS encoders, decoders, etc. are normally implemented at this point in PBR radios.

Finally, the AF amplifier circuits ensure the AF signal is strong enough to drive a speaker, and also control the volume, etc. The squelch circuit taps the AF signal from the AF amplifier circuit to control noise squelch, as well as cutting off the AF output to the speaker when there is no RF received signal.

On the transmit side, the microphone amplifier circuit feeds the modulation circuit on the VCO board, which converts AF to modulated RF. There is no need for double-stage up-conversion because of the relatively low RF values and partly because image rejection is not a critical performance characteristic in transmission circuits. Thus, the modulation circuits converts direct from AF to RF.

The output RF is fed into the driver amplifier circuit, which is equivalent to the pre-amplifier in the UHF SATCOM terminal. It boosts the RF signal to a level needed to drive the main RF power amplifier, which is normally implemented as a monolithic IC.

The automatic power control (APC) circuit protects the power amplifier from a mismatched output load and also selects ‘high’ or ‘low’ transmit power modes, a typical feature on most PBR devices. The APC does this by combining forward and rectified signals, the result of which should be a minimum when the circuit is matched. A detector monitors the sum, switching off the power amplifier if it exceeds a preset threshold.

The reference oscillator, programmable divider, phase detector, charge pump, loop filter and voltage controlled oscillator (VCO) circuits collectively make up the phase-locked loop (PLL) circuits. Together they provide a stable oscillation of the transmit frequency and the local oscillator (LO) source for driving the first mixer stage of the receive chain. ICs, amplifiers and crystals are the main components of the PLL circuits.

Figure C-6 is a more detailed RF architecture showing the functions of each circuit block. The individual components of each circuit are not shown as the circuit diagrams would become very complicated, beyond the level of detail necessary for this report. For the same reason the quantity of similar type components and their prices, rather than each individual component, are listed in the following tables, for common components such as resistors, capacitors, etc. The actual price of each component may vary dependant on its physical size, voltage rating, value, etc. so an average prices for each component type has been used.

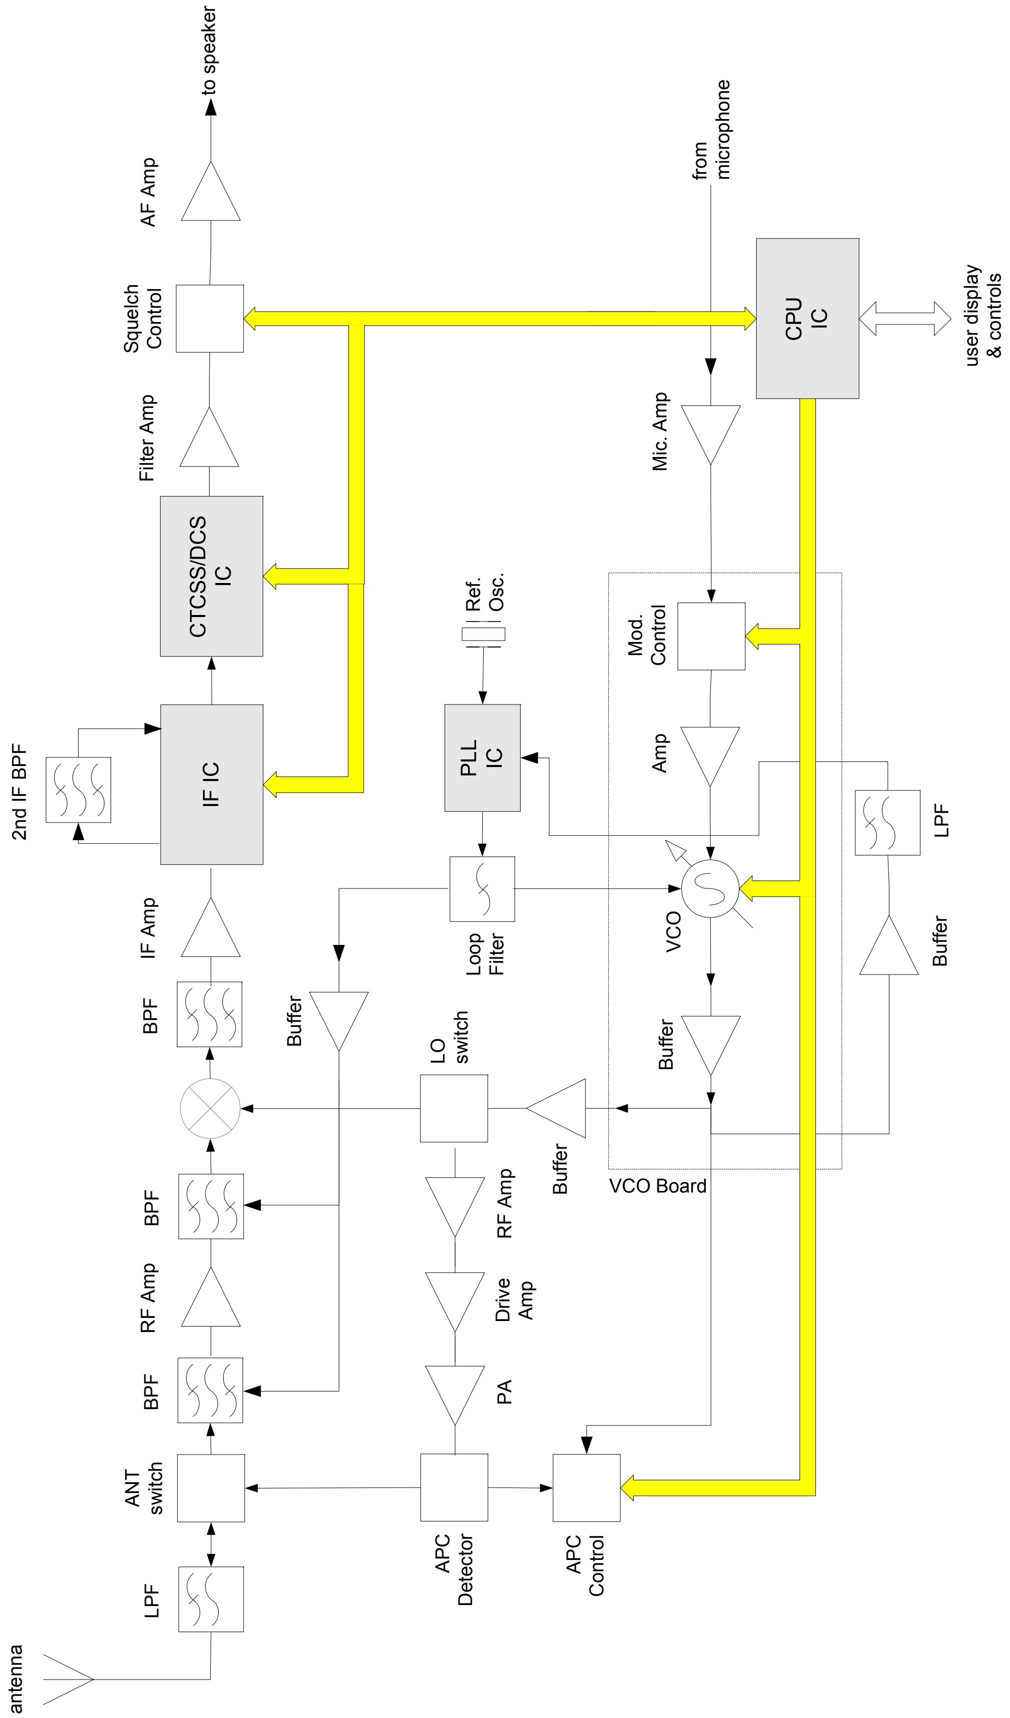

Figure C-6: Detailed architecture diagram of PBR design

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	VHF Antenna	Rubber duck	1	6.99	6.99
2	Low Pass Filter	Inductors	3	0.31	0.93
		Surface mount ceramic capacitors	8	0.23	1.84
3	Antenna Switch	Diode	3	0.20	0.60
		Surface mount ceramic capacitors	4	0.23	0.92
		ANTENNA SWITCHING CIRCUIT SUBTOTAL =			**£4.29**
RF CIRCUIT					
1	Band Pass Filter	Matching inductor pair	1	2.71	2.71
		Surface mouint ceramic capacitors	2	0.23	0.46
		Diode	1	0.20	0.20
2	RF Amplifier	VHF RF Amplifier FET	1	3.44	3.44
3	Band Pass Filter	Matching inductor pair	1	2.71	2.71
		Surface mouint ceramic capacitors	2	0.23	0.46
		Diode	1	0.20	0.20
		RF CIRCUIT SUBTOTAL =			**£10.18**
1st MIXER & IF CIRCUIT					
1	Mixer	ADE-12MH	1	4.42	4.42
2	Band Pass Filter	Crystal filters	2	0.69	1.38
3	IF Amplifier	RF transistor	1	0.12	0.12
		1st MIXER & IF CIRCUIT SUBTOTAL =			**£5.92**
2nd IF & DEMODULATOR					
1	IC1	FM IF mixer and processing chip	1	4.72	4.72
2	Crystals	Surface mount 21MHz	2	0.46	0.92
3	Resistors	Surface mount fixed-value	1	0.04	0.04
4	Capacitors	Surface mount tantalum	3	1.60	4.80
5	2nd IF Filter	Ceramic filter	1	0.29	0.29
		2nd IF & DEMODULATOR SUBTOTAL =			**£10.77**
MAIN UNIT					
1	IC1	Microcontroller	1	23.00	23.00
2	IC2	High precision voltage detector	1	1.61	1.61
3	IC3	8-bit shift register	1	0.17	0.17
4	IC4	Audio power amp	1	1.54	1.54
5	IC6	CTCSS encoder/decoder	1	14.00	14.00
6	IC7	Op-Amp	1	0.20	0.20
7	Transistors	Surface mount transistors	16	0.05	0.80
8	Diodes	Surface mount, small signal	5	0.20	1.00
9	Crystal	Surface mount 3.6MHz	1	0.34	0.34
10	Resistors	Surface mount fixed-value	87	0.04	3.48
11	Capacitors	Surface mount ceramic	89	0.23	20.47
12	Capacitors	Surface mount tantalum	23	1.60	36.80
		MAIN UNIT SUBTOTAL =			**£103.41**
LOGIC UNIT					

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	IC1	Central Processor Unit (CPU)	1	12.30	12.30
2	IC2	8K I2C serial EEPROM	1	0.34	0.34
3	IC3	High precision voltage detector	1	1.61	1.61
4	Transistors	Surface mount transistors	6	0.05	0.30
5	Diodes	Surface mount, small signal	2	0.20	0.40
6	Crystal	Ceramic crystal CPU clock	1	0.69	0.69
7	Resistors	Surface mount fixed-value	44	0.04	1.76
8	Capacitors	Surface mount ceramic	32	0.23	7.36
9	Capacitors	Surface mount tantalum	3	1.60	4.80
				LOGIC UNIT SUBTOTAL =	**£29.56**
VCO BOARD					
1	Transistors	Surface mount transistors	4	0.05	0.20
2	Diodes	Surface mount, small signal	3	0.20	0.60
		Varicap diode	1	0.11	0.11
3	Inductors	RF coils	3	1.17	3.51
4	Resistors	Surface mount fixed-value	16	0.04	0.64
5	Capacitors	Surface mount ceramic	16	0.23	3.68
				VCO BOARD SUBTOTAL =	**£8.74**
DRIVER AMPLIFIER CIRCUIT					
1	Transistors	Surface mount transistors (pre-drive & drive)	2	0.05	0.10
2	Inductors	surface mount	2	1.17	2.34
3	Resistors	Surface mount fixed-value	7	0.04	0.28
4	Capacitors	Surface mount ceramic	8	0.23	1.84
				DRIVER AMPLIFIER CIRCUIT SUBTOTAL =	**£4.56**
RF POWER AMP					
1	IC1	4W RF power module	1	38	38.00
2	Transistors	Surface mount transistors amplifier	1	0.12	0.12
	Diodes	Surface mount, small signal	2	0.20	0.40
4	Resistors	Surface mount fixed-value	5	0.04	0.20
5	Capacitors	Surface mount ceramic	8	0.23	1.84
				RF POWER AMP SUBTOTAL =	**£40.56**
AUTOMATIC POWER CONTROL					
1	APC detector circuit	Inductors	1	38	38.00
		Diodes	2	0.12	0.24
		Surface mount fixed-value resistors	4	0.04	0.16
		Surface mount ceramic capacitors	5	0.23	1.15
2	APC control	Transistors	3	0.12	0.36
		Surface mount ceramic capacitors	6	0.23	1.38
		Surface mount fixed-value resistors	4	0.04	0.16
				AUTOMATIC POWER CONTROL SUBTOTAL =	**£41.45**
PLL CIRCUITS					
1	Reference Oscillator	Transistor amplifier	1	0.12	0.12

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
		Crystal 12.8MHz	1	0.83	0.83
		Capacitor	1	0.23	0.23
2	Phase Detector	RF PLL Frequency Synthesizers (550MHz)	1	4.21	4.21
3	Charge Pump	Transistors	2	0.12	0.24
4	Loop Filter	resistors	1	0.40	0.40
		capacitors	1	0.23	0.23
		Buffer transistor	2	0.12	0.24
5	LO switch	Buffer transistor	1	0.12	0.12
		PIN diodes	2	0.20	0.40
				PLL CIRCUITS SUBTOTAL =	**£6.50**

PBR SYSTEM TOTAL COMPONENT COST = £265.94

Table C-8: Conventional PBR system costs

C.3.2 Spectrum-Aggregating Design

To turn a PBR hand-portable into a spectrum aggregating device, additional receive and transmit circuit blocks are required, just as additional components were required for a spectrum-aggregating UHF SATCOM terminal. The number of additional blocks will depend upon the number of fragments being aggregated.

To turn a PBR hand-portable into a spectrum aggregating device, additional receive and transmit circuit blocks are required, just as additional components were required for a spectrum-aggregating UHF SATCOM terminal. The number of additional blocks will depend upon the number of fragments being aggregated.

To aggregate fragments within the 146 – 174MHz band, some components and circuits can be shared across all fragments since they already operate over the entire band. They are: the antenna; low pass filter (LPF), antenna switch, and automatic power control (APC). The reference oscillator part of the PLL circuitry can also be shared, as long as its output is boosted. All these shared blocks reside in the RF unit.

On the receive path a single VCO could be used to share one LO source between multiple mixers, IF circuits and demodulators. Traditionally, each IF IC would require a different 2nd IF band pass ceramic filter, some of them away from common IF values (e.g. 455kHz), making them more expensive. But an IF of 21MHz can be easily digitised, so the 2nd IF and demodulator circuitry would be replaced altogether by digital signal processing (DSP). In DSP, the centre frequency of a filter can be changed easily and does not affect the component cost. It also has the additional benefit of being able to change the fragment centre frequency through software control only, no hardware changes or tuning required. To select any fragment within the 146 – 174MHz band means the 1st IF may get pushed up to 49MHz (21 + 28), but this is still low enough to be fully digitised by today's ADCs.

However, since there is direct up-conversion from AF to RF on the transmit path, it is not possible to share one LO source (the VCO board and PLL circuits) and then use different IF values to create different fragment centre frequencies. Therefore separate VCO boards and PLL circuits dedicated to each fragment are needed for transmission. And because a PBR system is half duplex, the use of LO switches allows the same VCO boards to be used for receive as well as transmit. This means there is not actually much saving to be had by sharing an LO source when receiving, but it would cost more due to faster ADCs and the need of an RF splitter. Common LO sharing is unlikely to be the most economically viable design for mass production.

Regardless of dedicated or shared LO sources, digitisation of the 21MHz IF will reduce the component count and make the PBR device more flexible. DSP is most likely required in any case to manage fragmentation, so it is sensible to optimise its use elsewhere in the system. Figure C-8 details the overall design changes required to change the standard PBR into a two-fragment aggregator system.

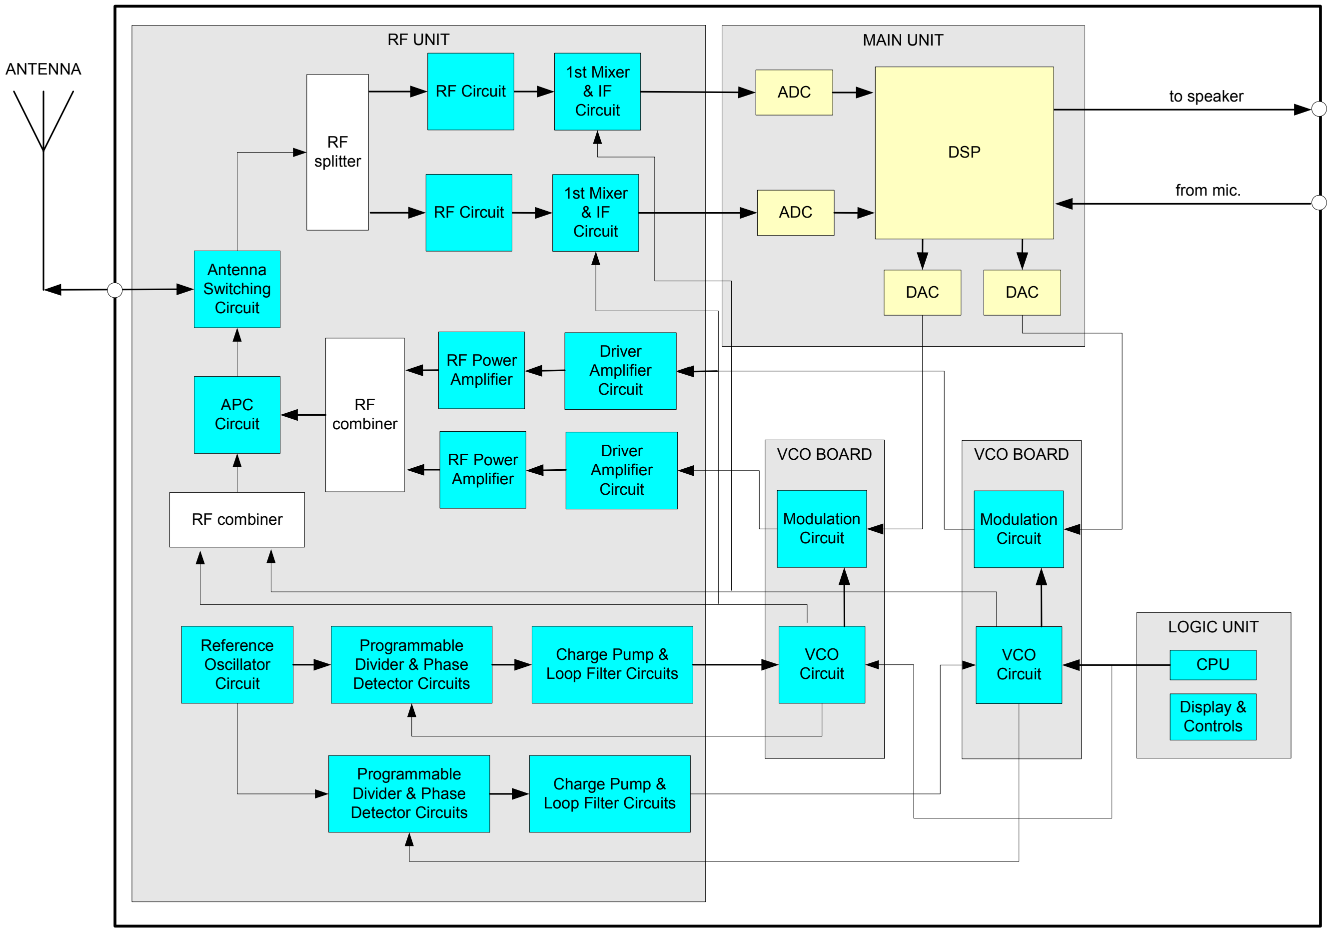

Figure C-7: Outline design for spectrum-aggregating PBR radio

Figure C-8 is a detailed diagram of the circuit blocks and components required for the two-fragment design. Parts required in addition to the conventional, non-aggregating system are highlighted in red. The cost of these parts is detailed in Table C-9.

As in the UHF SATCOM terminal example, adding DSP to the design constitutes the majority of the additional cost. Even the smallest DSP devices are likely to be able to cope with processing of more than two fragments, so the DSP cost does not need to be duplicated for three, four, five or more fragment designs (although eventually additional DSP would be needed as the number of fragments goes up). Also the type (and hence price) of RF analogue splitters and combiners will change depending on the number of fragments.

For 3, 4 and 5-fragment systems, then, the additional cost is shown in Table C-10, Table C-11, and Table C-12 respectively.

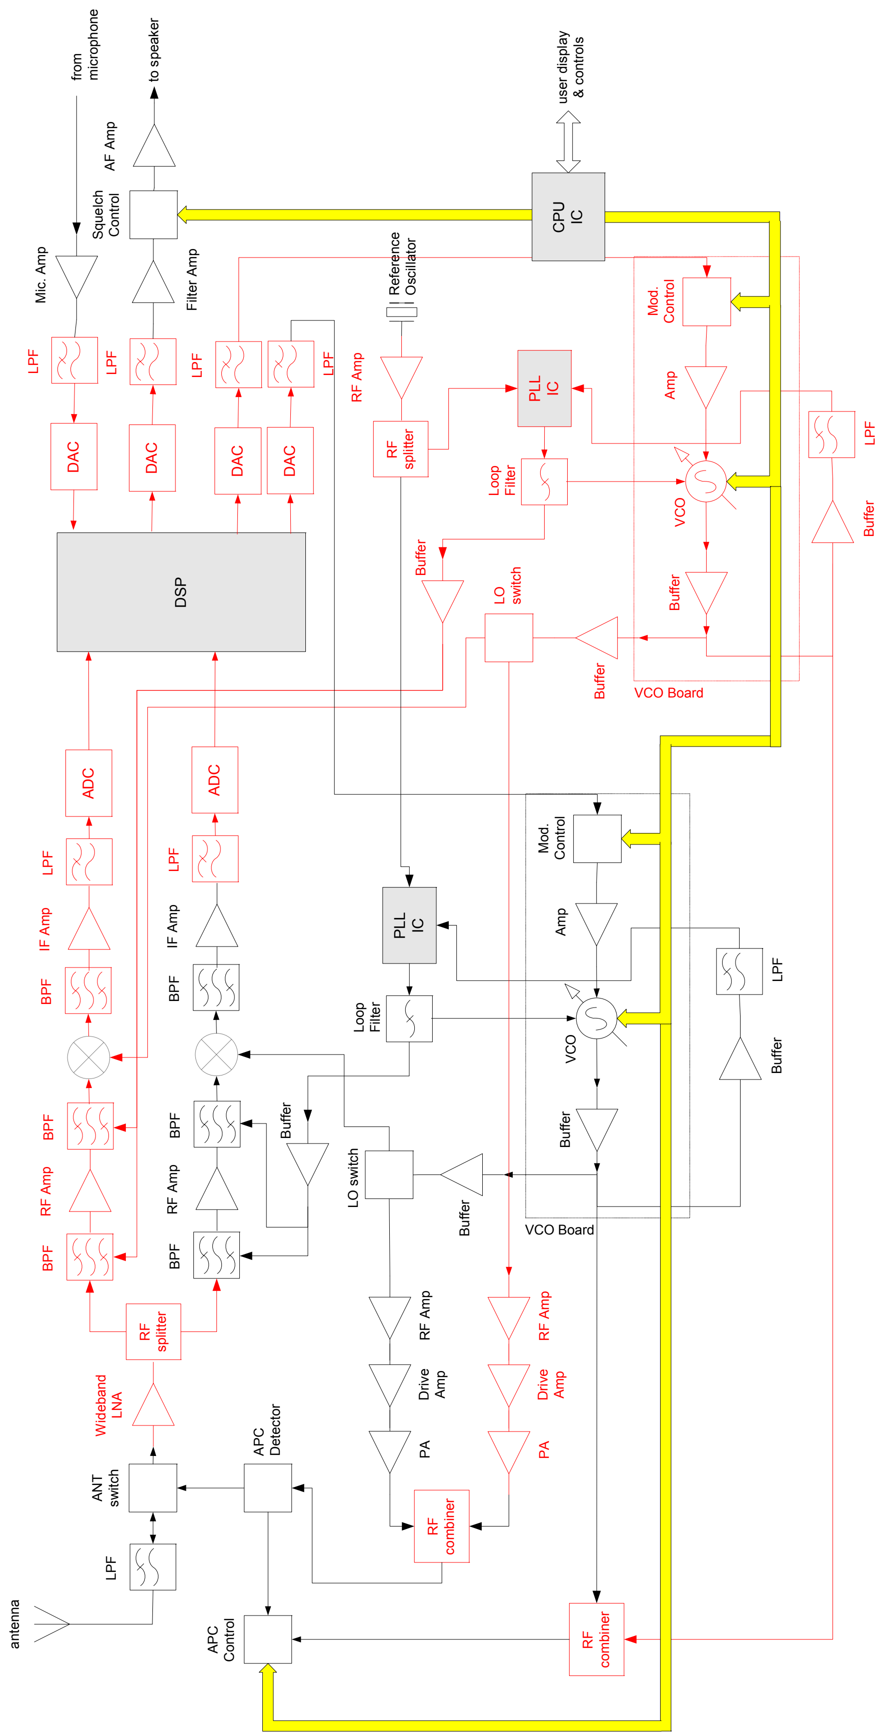

Figure C-8: System architecture for a two-fragment PBR device

IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
2-way splitter	Mini-Circuits SBTC-2-10-75	2	3.49	6.98
2-way combiner	Mini-Circuits ZA2C5-500-15W	2	74.95	149.90
Wideband LNA	Agilent MGA-82563	1	2.36	2.36
ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14
DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
ADC (microphone)	ADS7823EB	1	5.95	5.95
LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
DAC (speaker)	AD5320BRT	1	4.44	4.44
RF Circuit	See table in section 5.3.1	1	10.18	10.18
1st Mixer & IF Circuit	See table in section 5.3.1	1	5.92	5.92
LPF (IF anti-alias)	Kel-Com Ceramic type filter	1	55.73	55.73
ADC	Analog Devices AD6644	1	29.00	29.00
DAC (modulation control)	AD5320BRT	1	4.44	4.44
LPF (modulation control)	LTC1062CN8	1	5.64	5.64
VCO circuit board	See table in section 5.3.1	1	8.74	8.74
PLL circuits (less Ref. Osc.)	See table in section 5.3.1	1	5.84	5.84
Driver Amplifier circuit	See table in section 5.3.1	1	4.56	4.56
RF power amp circuit	See table in section 5.3.1	1	40.56	40.56
			ADDITIONAL TOTAL =	**£1,582.84**

Table C-9: Additional cost of a 2-fragment PBR design

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	3-way splitter	Mini-Circuits JPS-3-1W	2	14.95	29.90
2	3-way combiner	Mini-Circuits ZB3CS-920-15W	2	114.95	229.90
3	Wideband LNA	Agilent MGA-82563	1	2.36	2.36
4	ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
5	LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
6	RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14
7	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
8	ADC (microphone)	ADS7823EB	1	5.95	5.95
9	LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
10	DAC (speaker)	AD5320BRT	1	4.44	4.44
11	RF Circuit	See table in section 5.3.1	2	10.18	20.36
12	1st Mixer & IF Circuit	See table in section 5.3.1	2	5.92	11.84
13	LPF (IF anti-alias)	Kel-Com Ceramic type filter	2	55.73	111.46
14	ADC	Analog Devices AD6644	2	29.00	58.00
15	DAC (modulation control)	AD5320BRT	2	4.44	8.88
16	LPF (modulation control)	LTC1062CN8	2	5.64	11.28
17	VCO circuit board	See table in section 5.3.1	2	8.74	17.48
18	PLL circuits (less Ref. Osc.)	See table in section 5.3.1	2	5.84	11.68
19	Driver Amplifier circuit	See table in section 5.3.1	2	4.56	9.12
20	RF power amp circuit	See table in section 5.3.1	2	40.56	81.12
				ADDITIONAL TOTAL =	**£1,856.37**

Table C-10: Additional cost of a 3-fragment PBR design

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	4-way splitter	Mini-Circuits PSC-4-1	2	37.95	75.90
2	4-way combiner	Mini-Circuits ZB4CS-960-12W	2	134.95	269.90
3	Wideband LNA	Agilent MGA-82563	1	2.36	2.36
4	ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
5	LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
6	RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14
7	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
8	ADC (microphone)	ADS7823EB	1	5.95	5.95
9	LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
10	DAC (speaker)	AD5320BRT	1	4.44	4.44
11	RF Circuit	See table in section 5.3.1	3	10.18	30.54
12	1st Mixer & IF Circuit	See table in section 5.3.1	3	5.92	17.76
13	LPF (IF anti-alias)	Kel-Com Ceramic type filter	3	55.73	167.19
14	ADC	Analog Devices AD6644	3	29.00	87.00
15	DAC (modulation control)	AD5320BRT	3	4.44	13.32
16	LPF (modulation control)	LTC1062CN8	3	5.64	16.92
17	VCO circuit board	See table in section 5.3.1	3	8.74	26.22
18	PLL circuits (less Ref. Osc.)	See table in section 5.3.1	3	5.84	17.52
19	Driver Amplifier circuit	See table in section 5.3.1	3	4.56	13.68
20	RF power amp circuit	See table in section 5.3.1	3	40.56	121.68
				ADDITIONAL TOTAL =	**£2,112.98**

Table C-11: Additional cost of a 4-fragment PBR design

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	5-way splitter	Mini-Circuits ZBSC-5-1	2	119.95	239.90
2	5-way combiner	Mini-Circuits ZBSC-5-1	2	119.95	239.90
3	Wideband LNA	Agilent MGA-82563	1	2.36	2.36
4	ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
5	LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
6	RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14
7	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
8	ADC (microphone)	ADS7823EB	1	5.95	5.95
9	LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
10	DAC (speaker)	AD5320BRT	1	4.44	4.44
11	RF Circuit	See table in section 5.3.1	4	10.18	40.72
12	1st Mixer & IF Circuit	See table in section 5.3.1	4	5.92	23.68
13	LPF (IF anti-alias)	Kel-Com Ceramic type filter	4	55.73	222.92
14	ADC	Analog Devices AD6644	4	29.00	116.00
15	DAC (modulation control)	AD5320BRT	4	4.44	17.76
16	LPF (modulation control)	LTC1062CN8	4	5.64	22.56
17	VCO circuit board	See table in section 5.3.1	4	8.74	34.96
18	PLL circuits (less Ref. Osc.)	See table in section 5.3.1	4	5.84	23.36
19	Driver Amplifier circuit	See table in section 5.3.1	4	4.56	18.24
20	RF power amp circuit	See table in section 5.3.1	4	40.56	162.24
				ADDITIONAL TOTAL =	**£2,417.59**

Table C-12: Additional cost of a 5-fragment PBR design

C.3.3 Additional Band Design

All of the above PBR examples assume that the fragments lie within the PBR band of 146 – 174MHz. But what about using fragments outside of this band? Because the original PBR band is comparatively narrow anyway, it is easy to widen the spectrum-aggregating PBR design to exploit out-of-band fragments without too much cost or extra design effort. Components that are already shared: antenna, antenna switch, low pass preselection filter, APC circuitry, LNA, RF splitters and RF combiners, are available or can be made to operate over wider bandwidths than in the designs considered above.

Figure C-4 shows the basic block diagram of a system using one fragment close to (or inside) the PBR band, and a second fragment far enough away to warrant the use of dedicated RF analogue components and a second stage of conversion. In both cases half the number of fragments are in or close enough to the PBR band to enable the use of discrete components similar to those in a conventional PBR radio, while half the fragments are at frequencies sufficiently distanced from the PBR band for which dedicated analogue components are required (similar to the SATCOM terminal).

Note that these costs are the same regardless of the number of bands. For example the cost is the same for a device with two bands with three fragments in each, three bands with three, two and one fragment per band, four bands with one or two fragments in each band, or any other band/fragment combination with three near-PBR and three non-PBR fragments. That is to say, for the PBR multiple band case bands and fragments are interchangeable.

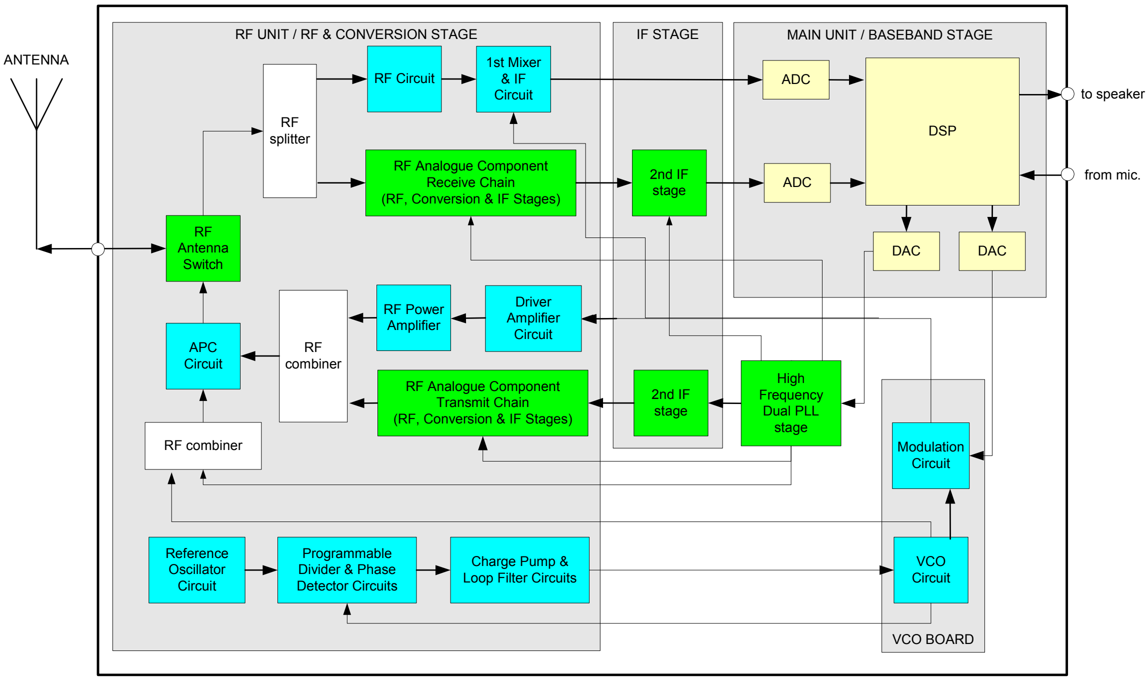

Figure C-9: RF stage architecture for a dual-band aggregating SATCOM terminal

Table C-13 shows an approximate additional cost of parts required to build a two band system. Table C-14 and Figure C-10 then show the additional cost for four fragments and six fragments respectively. In both cases half the number of fragments are in or close enough to the PBR band to enable the use of discrete components similar to those in a conventional PBR radio, while half the fragments are at frequencies sufficiently distanced from the PBR band for which dedicated analogue components are required (similar to the SATCOM terminal).

Note that these costs are the same regardless of the number of bands. For example the cost is the same for a device with two bands with three fragments in each, three bands with three, two and one fragment per band, four bands with one or two fragments in each band, or any other band/fragment combination with three near-PBR and three non-PBR fragments. That is to say, for the PBR multiple band case bands and fragments are interchangeable.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Mini-Circuits SBTC-2-10-75	2	3.49	6.98
2	2-way combiner	Mini-Circuits ZA2C5-500-15W	2	74.95	149.90
3	Wideband LNA	Agilent MGA-82563	1	2.36	2.36
4	ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
5	LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
6	RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14

7	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
8	ADC (microphone)	ADS7823EB	1	5.95	5.95
9	LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
10	DAC (speaker)	AD5320BRT	1	4.44	4.44
11	RF antenna switch	Mini-Circuits ZSDR-230	1	89.95	89.95
12	RF analogue Rx chain	based on SATCOM costing	1	141.82	141.82
13	RF analogue Tx chain	based on SATCOM costing	1	229.72	229.72
14	RF PLL stage	based on SATCOM costing	1	284.12	284.12
				ADDITIONAL TOTAL =	**£2,157.84**

Table C-13: Additional cost of two bands (one PBR, one non-PBR), one fragment per band

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	4-way splitter	Mini-Circuits PSC-4-1	2	37.95	75.90
2	4-way combiner	Mini-Circuits ZB4CS-960-12W	2	134.95	269.90
3	Wideband LNA	Agilent MGA-82563	1	2.36	2.36
4	ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
5	LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
6	RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14
7	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
8	ADC (microphone)	ADS7823EB	1	5.95	5.95
9	LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
10	DAC (speaker)	AD5320BRT	1	4.44	4.44
11	RF antenna switch	Mini-Circuits ZSDR-230	2	89.95	179.90
12	RF analogue Rx chain	based on SATCOM costing	2	141.82	283.64
13	RF analogue Tx chain	based on SATCOM costing	2	229.72	459.44
14	RF PLL stage	based on SATCOM costing	2	284.12	568.24
11	RF Circuit	see	1	10.18	10.18
12	1st Mixer & IF Circuit	see	1	5.92	5.92
13	LPF (IF anti-alias)	Kel-Com Ceramic type filter	1	55.73	55.73
14	ADC	Analog Devices AD6644	1	29.00	29.00
15	DAC (modulation control)	AD5320BRT	1	4.44	4.44
16	LPF (modulation control)	LTC1062CN8	1	5.64	5.64
17	VCO circuit board	see	1	8.74	8.74
18	PLL circuits (less Ref. Osc.)	see	1	5.84	5.84
19	Driver Amplifier circuit	see	1	4.56	4.56
20	RF power amp circuit	see	1	40.56	40.56
				ADDITIONAL TOTAL =	**£3,262.98**

Table C-14: Additional cost of four fragments, two PBR and two non-PBR

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	6-way splitter	Mini-Circuits JCPS-6-3	2	69.95	139.90
2	6-way combiner	Mini-Circuits ZFSC-6-110	2	109.95	219.90
3	Wideband LNA	Agilent MGA-82563	1	2.36	2.36
4	ADC (original channel fragment)	Analog Devices AD6644	1	29.00	29.00
5	LPF (original channel fragment)	IC Filter MF10CCN	1	3.26	3.26
6	RF Amplifier (Ref Osc. Booster)	Watkins-Johnson AG604-86	1	1.14	1.14
7	DSP or FPGA module	based on competitive HERON prices	1	1200.00	1200.00
8	ADC (microphone)	ADS7823EB	1	5.95	5.95
9	LPF (microphone + speaker)	Linear Tech. LTC1043CN Filter IC	2	4.60	9.20
10	DAC (speaker)	AD5320BRT	1	4.44	4.44
11	RF antenna switch	Mini-Circuits ZSDR-230	3	89.95	269.85
12	RF analogue Rx chain	based on SATCOM costing	3	141.82	425.46
13	RF analogue Tx chain	based on SATCOM costing	3	229.72	689.16
14	RF PLL stage	based on SATCOM costing	3	284.12	852.36
11	RF Circuit	see	2	10.18	20.36
12	1st Mixer & IF Circuit	see	2	5.92	11.84
13	LPF (IF anti-alias)	Kel-Com Ceramic type filter	2	55.73	111.46
14	ADC	Analog Devices AD6644	2	29.00	58.00
15	DAC (modulation control)	AD5320BRT	2	4.44	8.88
16	LPF (modulation control)	LTC1062CN8	2	5.64	11.28
17	VCO circuit board	see	2	8.74	17.48
18	PLL circuits (less Ref. Osc.)	see	2	5.84	11.68
19	Driver Amplifier circuit	see	2	4.56	9.12
20	RF power amp circuit	see	2	40.56	81.12
				ADDITIONAL TOTAL =	**£4,193.20**

Figure C-10: Additional cost of six fragments, three PBR and three non-PBR

From these costs it can be deduced that the distribution of bands will affect the cost, as will the number of fragments in a given band. But the width and distribution of fragments within a band will not affect the cost, nor will the number of bands.

For example, if an aggregating device only has bands distributed in or close to the PBR region, the number of bands will not really change the cost and the total additional cost will depend only upon the number of fragments. But introducing one or more bands far away from the PBR region will add significantly to the cost, and this cost will vary depending on the number of fragments in these far-out bands.

This is different to the situation for UHF SATCOM, where the total number of bands affected cost more than the distribution.

C.4 Wireless Local Area Network (WLAN)

WLAN uses RF to link two or more computers together in a small area (for example in the office or at home).

C.4.1 Conventional Design

The variation in prices and types of components makes is very difficult to estimate the one-off cost of a 2.4GHz WLAN system. In the table below, the average price for each component has been used, but the total cost could still vary greatly depending on the application intended.

OFDM				
IDENTITY	**DESCRIPTION**	**QTY**	**Item Price**	**Sub Total**
Antenna	Indoor use	1	20.00	20.00
Transceiver Chip	Based on Texas Instrument prices	1	3.60	3.60
DSP	Based on competitive HERON	1	1200.00	1200.00
		OFDM TOTAL =		**£1,223.60**

FHSS				
IDENTITY	**DESCRIPTION**	**QTY**	**Item Price**	**Sub Total**
Antenna	Indoor use	1	20.00	20.00
Transceiver Chip	Based on Texas Instrument prices	1	3.60	3.60
Microcontroller	Based on 16-bit devices	1	12.00	12.00
		FHSS TOTAL =		**£35.60**

Table C-15: Estimated cost of one-off WLAN devices (non-aggregating)

C.4.2 Spectrum-Aggregating Design

For a two-fragment aggregator, there are no additional parts and hence no additional cost if OFDM is already being used. Since the 802.11g standard can have up to 52 OFDM sub-carriers, it is highly unlikely that additional DSP will be needed to aggregate two fragments, or indeed much higher numbers of fragments.

For a FHSS design, there is additional cost in the form of a splitter/combiner, an extra transceiver chip and more capable MCU. These additional costs are given in Table C-16. Note that the original 16-bit microcontroller device cost has been subtracted as the more advanced MCU replaces this.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Mini-Circuits SCN-2-27	1	2.50	2.50
2	Transceiver Chip	Based on Texas Instrument prices	1	3.60	3.60
3	MCU	Based on 32-bit M68000 series	1	66.25	66.25
		Minus 16-bit microcontroller	1	-12.00	-12.00
			ADDITIONAL TOTAL =		**£60.35**

Table C-16: Estimated cost of 2-fragment aggregating FHSS WLAN system

For higher numbers of fragments, in OFDM system the cost will be the same and in FHSS systems the cost in given below in Table C-17 for 3- 4- and 5-fragment systems. Table 5-8 summarises the costs for up to ten fragments.

As DSP was the dominant expense in SATCOM and PBR aggregating systems, so the improved MCU is the initial and dominant expense in moving from a conventional to a spectrum-aggregating WLAN device. Once this initial increase is accounted for, however, the cost steadily increases as the number of fragments increases, roughly £5.70 per fragment, for up to seven fragments. Devices with eight or more fragments see another step increase, this time caused by the sudden increase in the cost of 8-way splitter/combiners with 2.4GHz performance.

3 FRAGMENTS					
ITEM	**IDENTITY**	**DESCRIPTION**	**QTY**	**Item Price**	**Sub Total**
1	3-way splitter	Mini-Circuits SCN-3-28	1	3.95	3.95
2	Transceiver Chip	Based on Texas Instrument prices	2	3.60	7.20
3	MCU	Based on 32-bit M68000 series	1	66.25	66.25
		Minus 16-bit microcontroller	1	-12.00	-12.00
		3-FRAGMENT ADDITIONAL TOTAL =			**£65.40**

4 FRAGMENTS					
ITEM	**IDENTITY**	**DESCRIPTION**	**QTY**	**Item Price**	**Sub Total**
1	4-way splitter	Mini-Circuits Sdb-4-25	1	9.95	9.95
2	Transceiver Chip	Based on Texas Instrument prices	3	3.60	10.80
3	MCU	Based on 32-bit M68000 series	1	66.25	66.25
		Minus 16-bit microcontroller	1	-12.00	-12.00
		4-FRAGMENT ADDITIONAL TOTAL =			**£75.00**

5 FRAGMENTS					
ITEM	**IDENTITY**	**DESCRIPTION**	**QTY**	**Item Price**	**Sub Total**
1	5-way splitter	Mini-Circuits SCN-2-27	2	2.50	5.00
		Mini-Circuits SCN-3-28	1	3.95	3.95
2	Transceiver Chip	Based on Texas Instrument prices	4	3.60	14.40
3	MCU	Based on 32-bit M68000 series	1	66.25	66.25
		Minus 16-bit microcontroller	1	-12.00	-12.00
		5-FRAGMENT ADDITIONAL TOTAL =			**£77.60**

Table C-17: Additional cost of 3-, 4- & 5-fragment FHSS WLAN designs

C.4.3 Additional Band Design

To aggregate fragments outside of the WLAN bands will incurr large costs for both OFDM and FHSS type systems. The cost will depend on the specific bands in which fragments reside, but in general a superhetrodyne approach using individual RF analogue components will be required. This means ADCs, DACs and possibly DSP – or more MCU processing power – will also be needed.

Thus the cost of aggregating outside the WLAN band will be similar in nature to that of the PBR examples, i.e. for WLAN additional band design, bands and fragments are interchangeable. Also, the distribution of bands will affect the cost, as will the number of fragments in a given band, but the width and distribution of fragments within a band will not affect the cost, nor will the number of bands.

C.5 Microwave Fixed Link

There are many fixed links operating at frequencies between 5 - 40GHz. They are generally used as large bandwidth "pipelines" between remote or hard-to-access areas, such as rural GSM base stations. Indeed, a major part of the fixed link market constitutes.

The outline design for a fixed microwave link transceiver station is given in Figure C-14. It consists of five main parts:

- Receiver Chain
- Transmit Chain
- Antenna and antenna sharing device
- Phase Locked Loop (PLL) Stage
- Analogue-to-Digital and Digital-to-Analogue Conversion

From a design point of view it is useful to split the design into five stages: RF, Conversion, IP, Baseband and the PLL.

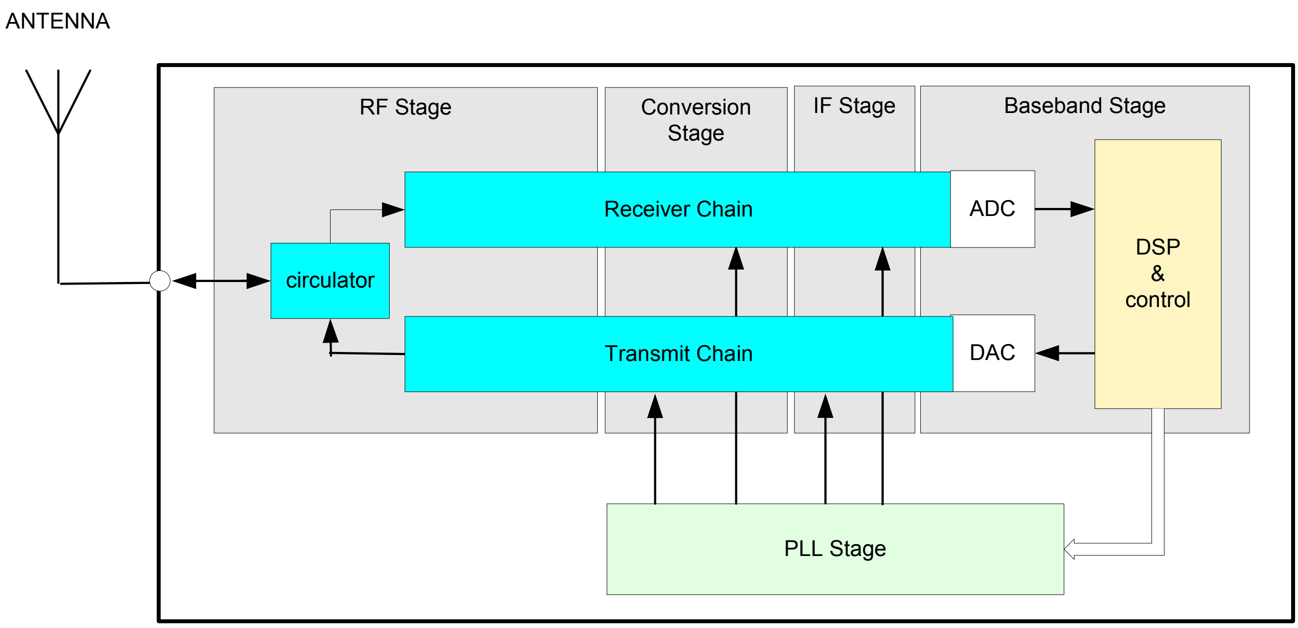

Figure C-11: Outline design for microwave fixed link transceiver

Some microwave fixed links may be acting as relays, in which case they will be using two antennas (pointing in different directions), one for receive and one for re-transmit. In this case, an antenna sharing device is not needed as either chain connects directly to an antenna. But in most cases, and in this example, one antenna will be shared by the chains.

The receiver chain contains the analogue components such as low noise amplifiers, and filters to boost the weak microwave signal received at the antenna. The chain also includes mixers, in the conversion stage, for down-converting the microwave frequency to an IF suitable for further analogue signal conditioning and final down-conversion to baseband. Baseband may be analogue, but more often it is converted to digital and some extensive DSP will be installed to control, monitor and modulate the data and the behaviour of the transceiver station.

Digital baseband intended for transmission is converted to analogue and fed into the transmit chain at a suitably low IF. After initial analogue conditioning (amplification, filtering), one or two mixers (single or dual stage conversion) will turn the IF into microwave frequencies, ready for amplification by a power amplifier (PA) and final filtering before entering the antenna. The high gain of the antenna significantly relaxes the power requirement of the PA.

Driving the mixers in the conversion stage are precise local oscillator (LO) frequencies, whose value combined with the value of the wanted signal frequency will determine the IF. LO frequencies are provided by the phased-lock loop (PLL) stage.

Because the mode of operation is frequency division duplex (FDD), a circulator is used to enable transmit signals go out through the antenna (and not back into the receive chain), while receive signals pass into the receive chain (but not the transmit chain), at the same time.

Figure C-12 is a more detailed RF architecture which shows the individual RF components of the architecture.

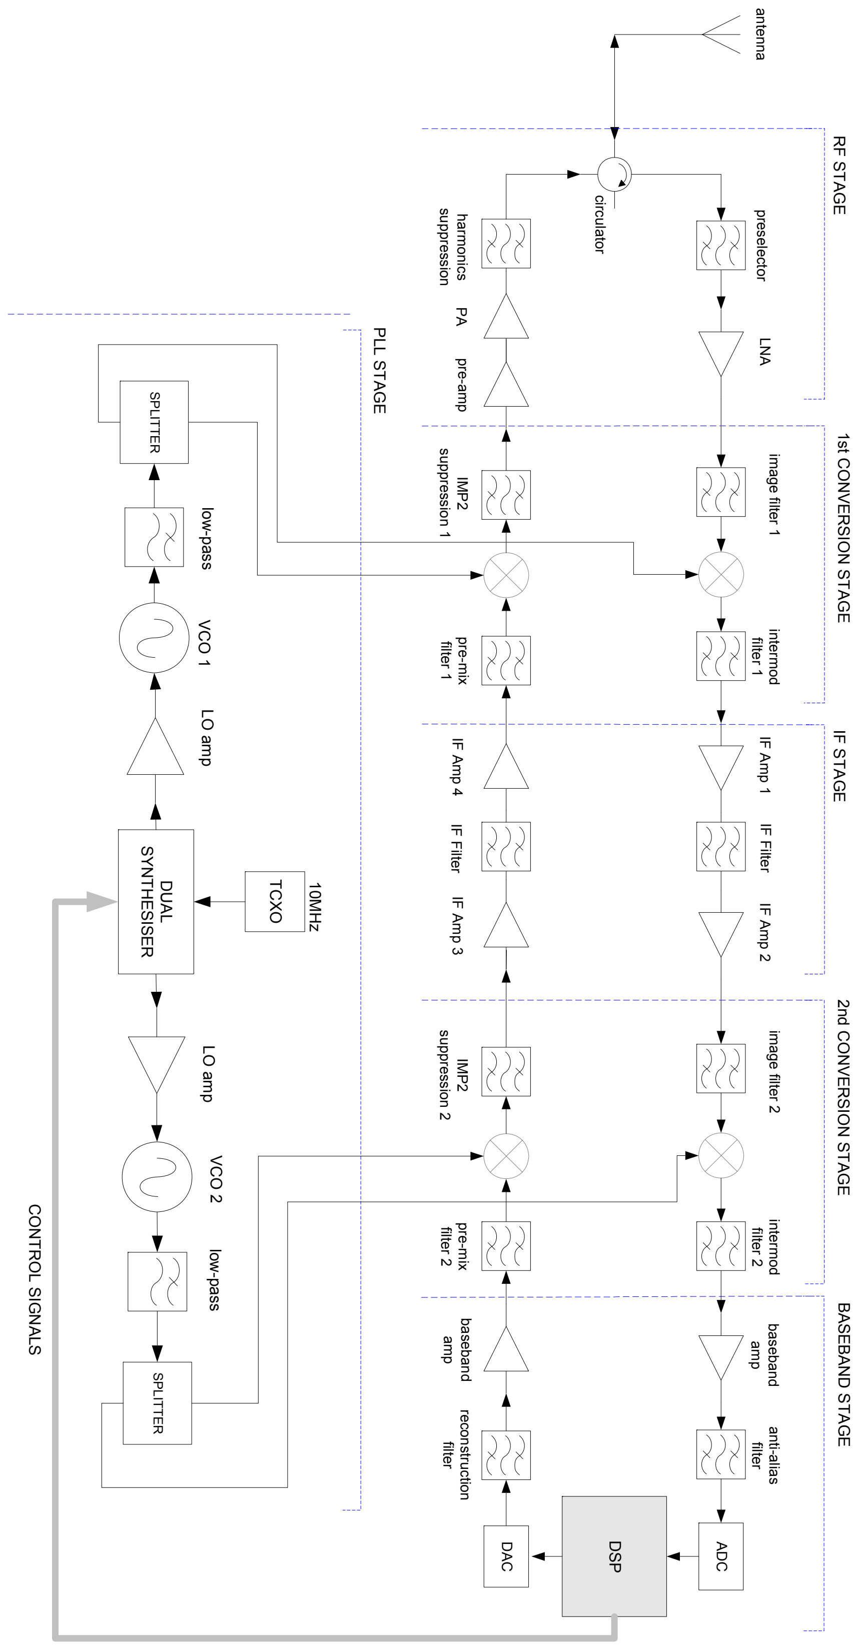

Figure C-12: RF Architecture for a conventional microwave fixed link transceiver station.

Listed in Table C-18 are the prices for the analogue components contained within the RF, conversion, IF, AF and PLL stages of the system.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	Antenna	Parabolic reflector dish	1	827.20	827.20
2	Circulator	Renaissance broadband coaxial	1	2000.00	2000.00
3	Preselector	Waveguide bandpass low Q	1	125.00	125.00
4	LNA	Lucix S180265L3201	1	500.00	500.00
5	Image Reject Filter 1	Waveguide bandpass high Q	1	350.00	350.00
6	Mixers (1st stage)	Mini-Circuits MCA1-12GL	2	13.95	27.90
7	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
8	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	4	1.52	6.08
9	IF Filters	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
10	Image filter 2	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
11	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	2	10.95	21.90
12	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	1	55.73	55.73
13	Baseband amplifier	Sirenza SGA-4186	2	1.00	2.00
14	Anti-alias filter	IC Filter MF10CCN	1	3.26	3.26
15	ADC	Analog Devices AD6644	1	29.00	29.00
16	DSP board	based on competitive HERON prices	1	1200.00	1200.00
17	DAC	Analog Devices AD9772A	1	14.95	14.95
18	Reconsctruction filter	IC Filter MF10CCN	1	3.26	3.26
19	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	1	55.73	55.73
20	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
21	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
22	IMP2 suppression 1	Waveguide bandpass high Q	1	350.00	350.00
23	Pre-amplifier	Mini-Circuits ZX60 series	1	64.95	64.95
24	Power Amplifier	Mini-Circuits ZHL-4240W	1	1495.00	1495.00
25	Harmonics suppression	Waveguide bandpass high Q, high power	1	500.00	500.00
26	LO splitters	Mini-circuits ZC range	2	149.95	299.90
27	Low pass filters	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
28	RF VCO	Micronetics MW500 series	1	79.95	79.95
29	LO Amplifiers	Sirenza SNA-100	2	4.72	9.44
30	Dual Synthesiser	ADF4252 (Analog Devices)	1	3.10	3.10
31	TCXO	Vectron TC-350-CAF-106	1	200.90	200.90
32	IF VCO	UMC UMZ-362-A16	1	36.00	36.00
				TOTAL =	**£8,420.85**

Table C-18: Non-aggregating microwave link transceiver station cost breakdown

C.5.1 Spectrum-Aggregating Design

To turn the microwave link transceiver station into a spectrum aggregating device, multiple receive and transmit chains are needed for each fragment. Figure C-14 is a two-fragment example. The conventional channel is already high-bandwidth, making it unlikely that any other components could be shared to give satisfactory performance (e.g. preselector, LNA or PA). Also, highly directional antennas tend to have relatively narrow bandwidths. This means that if fragments are too widely scattered, multiple antennas might be needed to cover all of them.

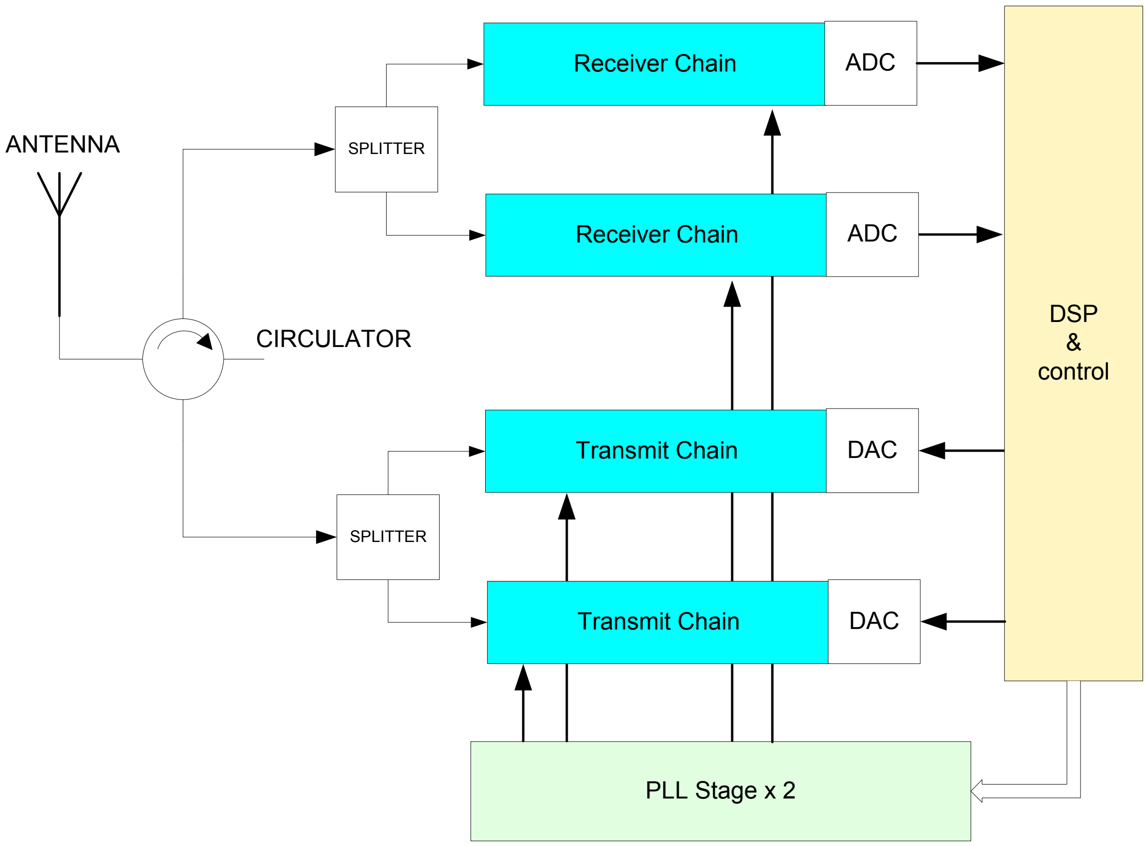

Figure C-13: System design overview of two-fragment aggregator

A more detailed architecture of the two-fragment example is shown in Figure C-14. It shows that, with just two fragments, the component count and complexity is starting to become significant. The breakdown in cost of the additional parts required on top of the conventional design is given in Table C-19

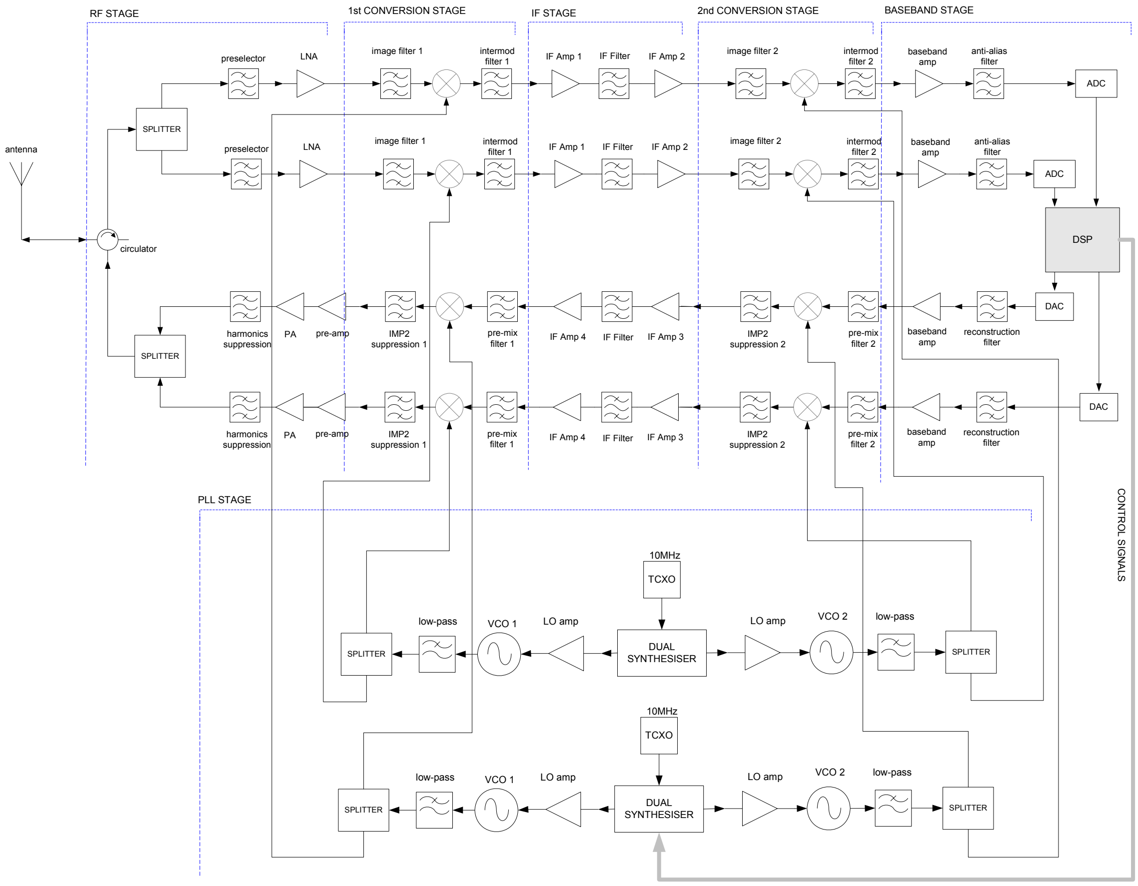

Figure C-14: RF Architecture for 2-fragment aggregating microwave station

The cost of these additional components is given in Table C-19 below. Compared to the cost of a non-aggregating conventional microwave transceiver station, this additional cost is equivalent to a 58% increase. Almost all of this extra cost is due to the duplication of components, since there is little scope for sharing common parts across chains/fragments.

Cost breakdowns for 3-fragment, 4-fragment and 5-fragment aggregating systems are given in the subsequent tables (Table C-20, Table C-21, and Table C-22,. The summary of costs for higher number fragment, from 6 to 10. In these estimates, for 6 or more fragments an additional antenna and circulator have been included to represent the likelihood that a single highly directional antenna will not be sufficient to cover more distributed fragments.

On average, a 58% increase cost per fragment pair is observed (approximately £4,884). This is to be expected due to the large duplication of components.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Microwave International	2	249.95	499.90
2	Preselector	Waveguide bandpass low Q	1	125.00	125.00
3	LNA	Lucix S180265L3201	1	500.00	500.00
4	Image Reject Filter 1	Waveguide bandpass high Q	1	350.00	350.00
5	Mixers (1st stage)	Mini-Circuits MCA1-12GL	2	13.95	27.90
6	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
7	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	4	1.52	6.08
8	IF Filters	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
9	Image filter 2	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
10	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	2	10.95	21.90
11	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	1	55.73	55.73
12	Baseband amplifier	Sirenza SGA-4186	2	1.00	2.00
13	Anti-alias filter	IC Filter MF10CCN	1	3.26	3.26
14	ADC	Analog Devices AD6644	1	29.00	29.00
15	DAC	Analog Devices AD9772A	1	14.95	14.95
16	Reconsctruction filter	IC Filter MF10CCN	1	3.26	3.26
17	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	1	55.73	55.73
18	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
19	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	1	19.95	19.95
20	IMP2 suppression 1	Waveguide bandpass high Q	1	350.00	350.00
21	Pre-amplifier	Mini-Circuits ZX60 series	1	64.95	64.95
22	Power Amplifier	Mini-Circuits ZHL-4240W	1	1495.00	1495.00
23	Harmonics suppression	Waveguide bandpass high Q, high power	1	500.00	500.00
24	LO splitters	Mini-circuits ZC range	2	149.95	299.90
25	Low pass filters	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
26	RF VCO	Micronetics MW500 series	1	79.95	79.95
27	LO Amplifiers	Sirenza SNA-100	2	4.72	9.44
28	Dual Synthesiser	ADF4252 (Analog Devices)	1	3.10	3.10
29	TCXO	Vectron TC-350-CAF-106	1	200.90	200.90
30	IF VCO	UMC UMZ-362-A16	1	36.00	36.00
				ADDITIONAL TOTAL =	**£4,893.55**

Table C-19: Breakdown of additional cost of 2-fragment microwave aggregator

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	3-way splitter	Microwave International	2	500	1000.00
2	Preselector	Waveguide bandpass low Q	2	125.00	250.00
3	LNA	Lucix S180265L3201	2	500.00	1000.00
4	Image Reject Filter 1	Waveguide bandpass high Q	2	350.00	700.00
5	Mixers (1st stage)	Mini-Circuits MCA1-12GL	4	13.95	55.80
6	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
7	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	8	1.52	12.16
8	IF Filters	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
9	Image filter 2	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
10	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	4	10.95	43.80
11	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	2	55.73	111.46
12	Baseband amplifier	Sirenza SGA-4186	4	1.00	4.00
13	Anti-alias filter	IC Filter MF10CCN	2	3.26	6.52
14	ADC	Analog Devices AD6644	2	29.00	58.00
15	DAC	Analog Devices AD9772A	2	14.95	29.90
16	Reconsctruction filter	IC Filter MF10CCN	2	3.26	6.52
17	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	2	55.73	111.46
18	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
19	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	2	19.95	39.90
20	IMP2 suppression 1	Waveguide bandpass high Q	2	350.00	700.00
21	Pre-amplifier	Mini-Circuits ZX60 series	2	64.95	129.90
22	Power Amplifier	Mini-Circuits ZHL-4240W	2	1495.00	2990.00
23	Harmonics suppression	Waveguide bandpass high Q, high power	2	500.00	1000.00
24	LO splitters	Mini-circuits ZC range	4	149.95	599.80
25	Low pass filters	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
26	RF VCO	Micronetics MW500 series	2	79.95	159.90
27	LO Amplifiers	Sirenza SNA-100	4	4.72	18.88
28	Dual Synthesiser	ADF4252 (Analog Devices)	2	3.10	6.20
29	TCXO	Vectron TC-350-CAF-106	2	200.90	401.80
30	IF VCO	UMC UMZ-362-A16	2	36.00	72.00
				ADDITIONAL TOTAL =	**£9,787.30**

Table C-20: Breakdown of additional cost of 3-fragment microwave aggregator

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	4-way splitter	Microwave International	2	750	1500.00
2	Preselector	Waveguide bandpass low Q	3	125.00	375.00
3	LNA	Lucix S180265L3201	3	500.00	1500.00
4	Image Reject Filter 1	Waveguide bandpass high Q	3	350.00	1050.00
5	Mixers (1st stage)	Mini-Circuits MCA1-12GL	6	13.95	83.70
6	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
7	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	12	1.52	18.24
8	IF Filters	Mini-Circuits coaxial VLF series 6700	6	19.95	119.70
9	Image filter 2	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
10	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	6	10.95	65.70
11	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	3	55.73	167.19
12	Baseband amplifier	Sirenza SGA-4186	6	1.00	6.00
13	Anti-alias filter	IC Filter MF10CCN	3	3.26	9.78
14	ADC	Analog Devices AD6644	3	29.00	87.00
15	DAC	Analog Devices AD9772A	3	14.95	44.85
16	Reconsctruction filter	IC Filter MF10CCN	3	3.26	9.78
17	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	3	55.73	167.19
18	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
19	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
20	IMP2 suppression 1	Waveguide bandpass high Q	3	350.00	1050.00
21	Pre-amplifier	Mini-Circuits ZX60 series	3	64.95	194.85
22	Power Amplifier	Mini-Circuits ZHL-4240W	3	1495.00	4485.00
23	Harmonics suppression	Waveguide bandpass high Q, high power	3	500.00	1500.00
24	LO splitters	Mini-circuits ZC range	6	149.95	899.70
25	Low pass filters	Mini-Circuits coaxial VLF series 6700	6	19.95	119.70
26	RF VCO	Micronetics MW500 series	3	79.95	239.85
27	LO Amplifiers	Sirenza SNA-100	6	4.72	28.32
28	Dual Synthesiser	ADF4252 (Analog Devices)	3	3.10	9.30
29	TCXO	Vectron TC-350-CAF-106	3	200.90	602.70
30	IF VCO	UMC UMZ-362-A16	3	36.00	108.00
				ADDITIONAL TOTAL =	**£14,680.95**

Table C-21: Breakdown of additional cost of 4-fragment microwave aggregator

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	5-way splitter	Microwave International	2	1250	2500.00
2	Preselector	Waveguide bandpass low Q	4	125.00	500.00
3	LNA	Lucix S180265L3201	4	500.00	2000.00
4	Image Reject Filter 1	Waveguide bandpass high Q	4	350.00	1400.00
5	Mixers (1st stage)	Mini-Circuits MCA1-12GL	8	13.95	111.60
6	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
7	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	16	1.52	24.32
8	IF Filters	Mini-Circuits coaxial VLF series 6700	8	19.95	159.60
9	Image filter 2	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
10	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	8	10.95	87.60
11	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	4	55.73	222.92
12	Baseband amplifier	Sirenza SGA-4186	8	1.00	8.00
13	Anti-alias filter	IC Filter MF10CCN	4	3.26	13.04
14	ADC	Analog Devices AD6644	4	29.00	116.00
15	DAC	Analog Devices AD9772A	4	14.95	59.80
16	Reconsctruction filter	IC Filter MF10CCN	4	3.26	13.04
17	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	4	55.73	222.92
18	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
19	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
20	IMP2 suppression 1	Waveguide bandpass high Q	4	350.00	1400.00
21	Pre-amplifier	Mini-Circuits ZX60 series	4	64.95	259.80
22	Power Amplifier	Mini-Circuits ZHL-4240W	4	1495.00	5980.00
23	Harmonics suppression	Waveguide bandpass high Q, high power	4	500.00	2000.00
24	LO splitters	Mini-circuits ZC range	8	149.95	1199.60
25	Low pass filters	Mini-Circuits coaxial VLF series 6700	8	19.95	159.60
26	RF VCO	Micronetics MW500 series	4	79.95	319.80
27	LO Amplifiers	Sirenza SNA-100	8	4.72	37.76
28	Dual Synthesiser	ADF4252 (Analog Devices)	4	3.10	12.40
29	TCXO	Vectron TC-350-CAF-106	4	200.90	803.60
30	IF VCO	UMC UMZ-362-A16	4	36.00	144.00
				ADDITIONAL TOTAL =	**£20,074.60**

Table C-22: Breakdown of additional cost of 5-fragment microwave aggregator

C.5.2 Additional Band Design

Figure C-15 shows an example of a dual-band four-fragment aggregator, with two fragments in each band.

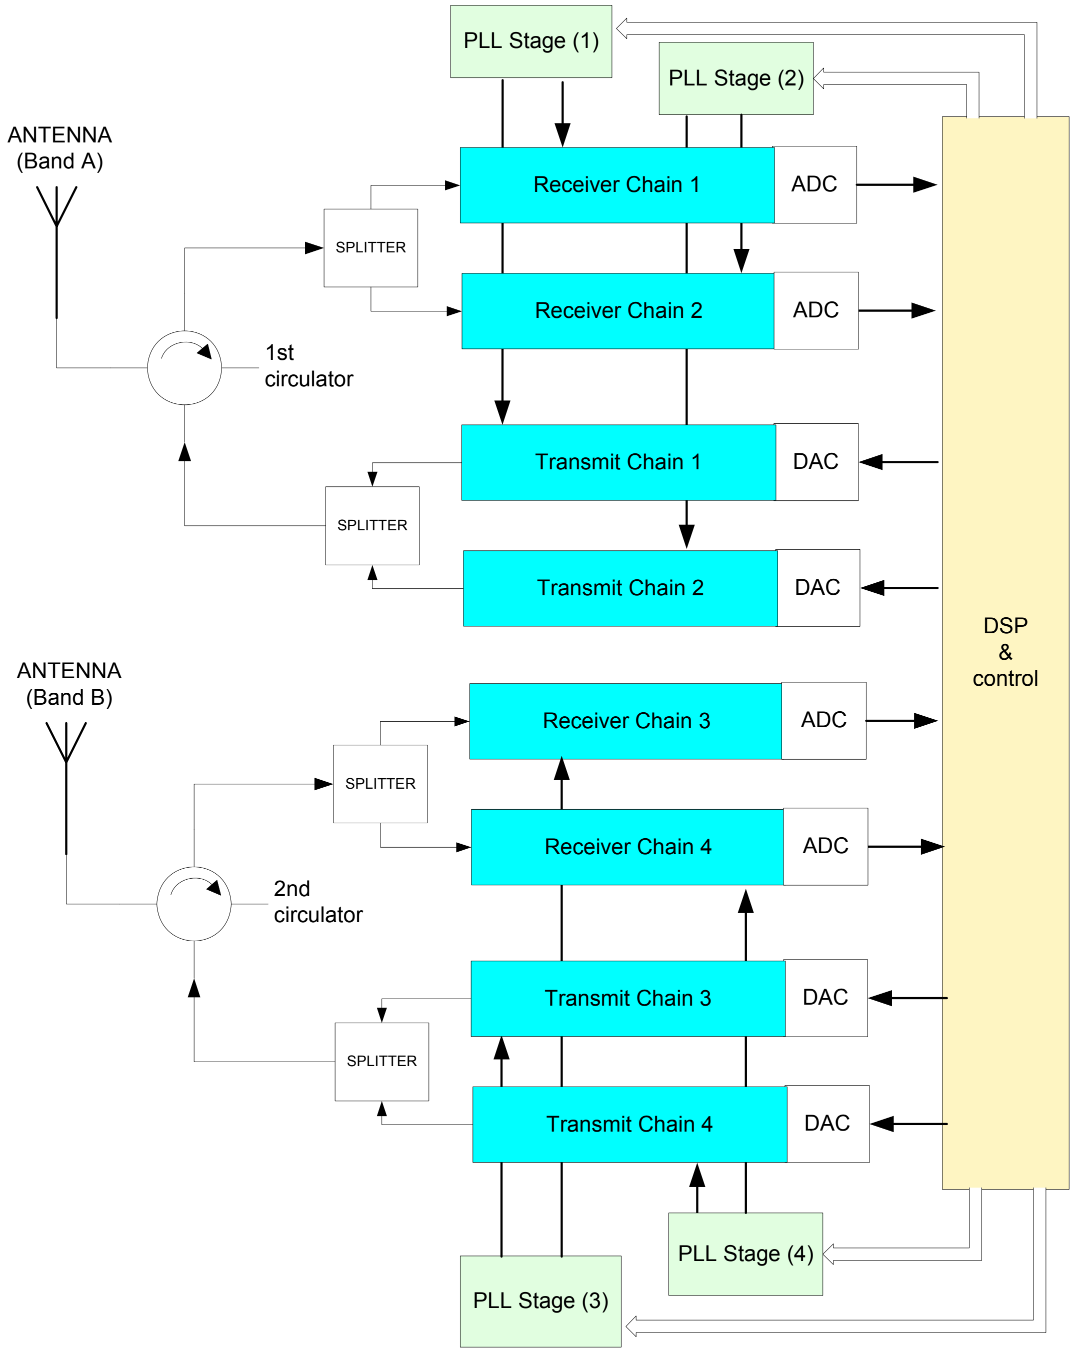

Figure C-15: System design for a dual-band microwave aggregator

The additional cost of a dual-band aggregator with five fragments, two in one band and three in the other, are shown below. We have already seen that the approximate cost for a fragment pair is around £5,000. These dual-band examples show that this cost per fragment is still the same, and that the cost of adding a new band is largely due to the additional antenna and circulator. However, even though these are expensive parts in a microwave system (accounting for a third of the cost of our conventional system example), it does not significantly add to the overall cost – only 10% at most. Therefore, the number of fragments is the dominant factor in cost, while the cost can be pushed up slightly further depending on the distribution of fragments in relation to bands, or in other words the fragment density within a band. For a given total number of fragments, if the fragment density is low, then this means there are more bands and thus higher cost than if all the fragments were concentrated within just one or two bands.

The cost for aggregating microwave fragments in different microwave bands will be similar to that of aggregating higher number of fragments, because multiple antennas (and therefore circulators) will be needed to cover the various bands. The breakdown of components and costs of a dual-band, four-fragment microwave aggregator are detailed in Table C-23.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Microwave International	4	249.95	999.80
2	Preselector	Waveguide bandpass low Q	3	125.00	375.00
3	LNA	Lucix S180265L3201	3	500.00	1500.00
4	Antenna	Parabolic reflector dish	1	827.20	827.20
5	Circulator	Renaissance broadband coaxial	1	2000.00	2000.00
6	Image Reject Filter 1	Waveguide bandpass high Q	3	350.00	1050.00
7	Mixers (1st stage)	Mini-Circuits MCA1-12GL	6	13.95	83.70
8	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
9	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	12	1.52	18.24
10	IF Filters	Mini-Circuits coaxial VLF series 6700	6	19.95	119.70
11	Image filter 2	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
12	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	6	10.95	65.70
13	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	3	55.73	167.19
14	Baseband amplifier	Sirenza SGA-4186	6	1.00	6.00
15	Anti-alias filter	IC Filter MF10CCN	3	3.26	9.78
16	ADC	Analog Devices AD6644	3	29.00	87.00
17	DAC	Analog Devices AD9772A	3	14.95	44.85
18	Reconsctruction filter	IC Filter MF10CCN	3	3.26	9.78
19	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	3	55.73	167.19
20	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
21	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	3	19.95	59.85
22	IMP2 suppression 1	Waveguide bandpass high Q	3	350.00	1050.00
23	Pre-amplifier	Mini-Circuits ZX60 series	3	64.95	194.85
24	Power Amplifier	Mini-Circuits ZHL-4240W	3	1495.00	4485.00
25	Harmonics suppression	Waveguide bandpass high Q, high power	3	500.00	1500.00
26	LO splitters	Mini-circuits ZC range	6	149.95	899.70
27	Low pass filters	Mini-Circuits coaxial VLF series 6700	6	19.95	119.70
28	RF VCO	Micronetics MW500 series	3	79.95	239.85
29	LO Amplifiers	Sirenza SNA-100	6	4.72	28.32
30	Dual Synthesiser	ADF4252 (Analog Devices)	3	3.10	9.30
31	TCXO	Vectron TC-350-CAF-106	3	200.90	602.70
32	IF VCO	UMC UMZ-362-A16	3	36.00	108.00
				ADDITIONAL TOTAL =	**£17,007.95**

Table C-23: Breakdown cost of a dual-band, four-fragment microwave aggregator

The additional cost of a dual-band aggregator with five fragments, two in one band and three in the other, are shown in Table C-24. We have already seen that the approximate cost for a fragment pair is around £5,000. These dual-band examples show that this cost per fragment is still the same, and that the cost of adding a new band is largely due to the additional antenna and circulator. However, even though these are expensive parts in a microwave system (accounting for a third of the cost of our conventional system example), it does not significantly add to the overall cost – only 10% at most. Therefore, the number of fragments is the dominant factor in cost, while the cost can be pushed up slightly further depending on the distribution of fragments in relation to bands, or in other words the fragment density within a band. For a given total number of fragments, if the fragment density is low, then this means there are more bands and thus higher cost than if all the fragments were concentrated within just one or two bands.

ITEM	IDENTITY	DESCRIPTION	QTY	Item Price	Sub Total
1	2-way splitter	Microwave International	2	249.95	499.90
2	Preselector	Waveguide bandpass low Q	4	125.00	500.00
3	LNA	Lucix S180265L3201	4	500.00	2000.00
4	3-way splitter	Microwave International	2	500	1000.00
5	Antenna	Parabolic reflector dish	1	827.20	827.20
6	Circulator	Renaissance broadband coaxial	1	2000.00	2000.00
7	Image Reject Filter 1	Waveguide bandpass high Q	4	350.00	1400.00
8	Mixers (1st stage)	Mini-Circuits MCA1-12GL	8	13.95	111.60
9	Intermod Filter 1	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
10	IF Amplifiers (1-4)	Mini-Circuits ERA-1SM	16	1.52	24.32
11	IF Filters	Mini-Circuits coaxial VLF series 6700	8	19.95	159.60
12	Image filter 2	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
13	Mixers (2nd stage)	Mini-Circuits MCA1-80MH	8	10.95	87.60
14	Intermod Filter 2	Kel-Com Ceramic 3KCB20 series	4	55.73	222.92
15	Baseband amplifier	Sirenza SGA-4186	8	1.00	8.00
16	Anti-alias filter	IC Filter MF10CCN	4	3.26	13.04
17	ADC	Analog Devices AD6644	4	29.00	116.00
18	DAC	Analog Devices AD9772A	4	14.95	59.80
19	Reconsctruction filter	IC Filter MF10CCN	4	3.26	13.04
20	Pre-mix filter 2	Kel-Com Ceramic 3KCB20 series	4	55.73	222.92
21	IMP2 suppression 2	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
22	Pre-mix filter 1	Mini-Circuits coaxial VLF series 6700	4	19.95	79.80
23	IMP2 suppression 1	Waveguide bandpass high Q	4	350.00	1400.00
24	Pre-amplifier	Mini-Circuits ZX60 series	4	64.95	259.80
25	Power Amplifier	Mini-Circuits ZHL-4240W	4	1495.00	5980.00
26	Harmonics suppression	Waveguide bandpass high Q, high power	4	500.00	2000.00
27	LO splitters	Mini-circuits ZC range	8	149.95	1199.60
28	Low pass filters	Mini-Circuits coaxial VLF series 6700	8	19.95	159.60
29	RF VCO	Micronetics MW500 series	4	79.95	319.80
30	LO Amplifiers	Sirenza SNA-100	8	4.72	37.76
31	Dual Synthesiser	ADF4252 (Analog Devices)	4	3.10	12.40
32	TCXO	Vectron TC-350-CAF-106	4	200.90	803.60

33	IF VCO	UMC UMZ-362-A16	4	36.00	144.00
				ADDITIONAL TOTAL =	**£21,901.70**

Table C-24: Breakdown cost of a dual-band, five-fragment microwave aggregator

D Target cost calculation tables and NPV analysis

D.1 Satcom Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£6,664.38
+	Standard Dealer Margin		30%	£ -	£1,537.93
=	Cost to Retailer				£5,126.45
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£5,111.45
+	Distribution Cost/Mark-up		15%	£ -	£664.49
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£4,429.95
+	Profit Margin		8%	£ -	£268.48
+	Warranty Cost		2%	£ -	£67.12
+	Corporate Allocations		10%	£ -	£335.60
+	Business Unit Selling, General & Admin		12%	£ -	£402.72
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£3,356.03
+	Overhead		45%	£ -	£1,041.53
=	Direct Target Cost (Labour & Material)				£2,314.50

Table D-1: Target Cost Calculation for the Standard Satcom Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£11,514.19
+	Standard Dealer Margin		30%	£ -	£2,657.12
=	Cost to Retailer				£8,857.07
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£8,842.07
+	Distribution Cost/Mark-up		15%	£ -	£1,151.10
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£7,673.97
+	Profit Margin		8%	£ -	£465.09
+	Warranty Cost		2%	£ -	£116.27
+	Corporate Allocations		10%	£ -	£581.36
+	Business Unit Selling, General & Admin		12%	£ -	£697.63
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£5,813.62
+	Overhead		45%	£ -	£1,804.23
=	Direct Target Cost (Labour & Material)				£4,009.39

Table D-2: Target Cost Calculation for the Single Band 2-Fragment Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£16,302.11
+	Standard Dealer Margin		30%	£ -	£3,762.02
=	Cost to Retailer				£12,540.08
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£12,525.08
+	Distribution Cost/Mark-up		15%	£ -	£1,631.49
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£10,876.59
+	Profit Margin		8%	£ -	£659.19
+	Warranty Cost		2%	£ -	£164.80
+	Corporate Allocations		10%	£ -	£823.98
+	Business Unit Selling, General & Admin		12%	£ -	£988.78
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£8,239.84
+	Overhead		45%	£ -	£2,557.19
=	Direct Target Cost (Labour & Material)				£5,682.65

Table D-3: Target Cost Calculation for the Dual Band 4-Fragment Terminal

D.2 Private Business Radio (PBR)

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£802.57
+	Standard Dealer Margin		30%	£ -	£185.21
=	Cost to Retailer				£617.36
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£602.36
+	Distribution Cost/Mark-up		15%	£ -	£76.35
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£509.01
+	Profit Margin		8%	£ -	£30.85
+	Warranty Cost		2%	£ -	£7.71
+	Corporate Allocations		10%	£ -	£38.56
+	Business Unit Selling, General & Administrative		12%	£ -	£46.27
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£385.61
+	Overhead		45%	£ -	£119.67
=	Direct Target Cost (Labour & Material)				£265.94

Table D-4: Target Cost Calculation for the Conventional PBR Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£5,331.75
+	Standard Dealer Margin		30%	£ -	£1,230.40
=	Cost to Retailer				£4,101.35
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£4,086.35
+	Distribution Cost/Mark-up		15%	£ -	£530.78
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£3,538.56
+	Profit Margin		8%	£ -	£214.46
+	Warranty Cost		2%	£ -	£53.61
+	Corporate Allocations		10%	£ -	£268.07
+	Business Unit Selling, General & Administrative		12%	£ -	£321.69
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£2,680.73
+	Overhead		45%	£ -	£831.95
=	Direct Target Cost (Labour & Material)				£1,848.78

Table D-5: Target Cost Calculation for the 2-Fragment PBR Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£10,139.36
+	Standard Dealer Margin		30%	£ -	£2,339.85
=	Cost to Retailer				£7,799.51
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£7,784.51
+	Distribution Cost/Mark-up		15%	£ -	£1,013.15
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£6,754.35
+	Profit Margin		8%	£ -	£409.35
+	Warranty Cost		2%	£ -	£102.34
+	Corporate Allocations		10%	£ -	£511.69
+	Business Unit Selling, General & Administrative		12%	£ -	£614.03
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£5,116.93
+	Overhead		45%	£ -	£1,588.01
=	Direct Target Cost (Labour & Material)				£3,528.92

Table D-6: Target Cost Calculation for the Multi-Band 4-Fragment PBR Terminal

D.3 Wireless Local Area Network (WLAN)

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£143.47
+	Standard Dealer Margin		30%	£ -	£33.11
=	Cost to Retailer				£110.36
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£95.36
+	Distribution Cost/Mark-up		15%	£ -	£10.22
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£68.14
+	Profit Margin		8%	£ -	£4.13
+	Warranty Cost		2%	£ -	£1.03
+	Corporate Allocations		10%	£ -	£5.16
+	Business Unit Selling, General & Administrative		12%	£ -	£6.19
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£51.62
+	Overhead		45%	£ -	£16.02
=	Direct Target Cost (Labour & Material)				£35.60

Table D-7: Target Cost Calculation for the Conventional WLAN Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£316.15
+	Standard Dealer Margin		30%	£ -	£72.96
=	Cost to Retailer				£243.20
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£228.20
+	Distribution Cost/Mark-up		15%	£ -	£27.55
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£183.65
+	Profit Margin		8%	£ -	£11.13
+	Warranty Cost		2%	£ -	£2.78
+	Corporate Allocations		10%	£ -	£13.91
+	Business Unit Selling, General & Administrative		12%	£ -	£16.70
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£139.13
+	Overhead		45%	£ -	£43.18
=	Direct Target Cost (Labour & Material)				£95.95

Table D-8: Target Cost Calculation for the 2-Fragment WLAN Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£800.31
+	Standard Dealer Margin		30%	£ -	£184.69
=	Cost to Retailer				£615.62
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£600.62
+	Distribution Cost/Mark-up		15%	£ -	£76.12
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£507.50
+	Profit Margin		8%	£ -	£30.76
+	Warranty Cost		2%	£ -	£7.69
+	Corporate Allocations		10%	£ -	£38.45
+	Business Unit Selling, General & Administrative		12%	£ -	£46.14
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£384.47
+	Overhead		45%	£ -	£119.32
=	Direct Target Cost (Labour & Material)				£265.15

Table D-9: Target Cost Calculation for the 10-Fragment WLAN Terminal

D.4 Microwave Fixed Link

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£24,137.27
+	Standard Dealer Margin		30%	£ -	£5,570.14
=	Cost to Retailer				£18,567.13
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£18,552.13
+	Distribution Cost/Mark-up		15%	£ -	£2,417.63
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£16,117.51
+	Profit Margin		8%	£ -	£976.82
+	Warranty Cost		2%	£ -	£244.20
+	Corporate Allocations		10%	£ -	£1,221.02
+	Business Unit Selling, General & Administrative		12%	£ -	£1,465.23
	Non-Recurring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£12,210.23
+	Overhead		45%	£ -	£3,789.38
=	Direct Target Cost (Labour & Material)				£8,420.85

Table D-10: Target Cost Calculation for the Conventional Fixed Link Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£38,139.82
+	Standard Dealer Margin		30%	£ -	£8,801.50
=	Cost to Retailer				£29,338.33
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£29,323.33
+	Distribution Cost/Mark-up		15%	£ -	£3,822.56
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£25,483.76
+	Profit Margin		8%	£ -	£1,544.47
+	Warranty Cost		2%	£ -	£386.12
+	Corporate Allocations		10%	£ -	£1,930.59
+	Business Unit Selling, General & Administrative		12%	£ -	£2,316.71
	Non-Recuring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£19,305.88
+	Overhead		45%	£ -	£5,991.48
=	Direct Target Cost (Labour & Material)				£13,314.40

Table D-11: Target Cost Calculation for the 2-Fragment Fixed Link Terminal

Sign	Price/Cost Element	Estimate	% Factor	Per Unit Factor	Amount
	Manufacturers Suggested Retail Price				£72,804.33
+	Standard Dealer Margin		30%	£ -	£16,801.00
=	Cost to Retailer				£56,003.33
+	Shipping/Distribution Cost to Retailer		0%	£15.00	£15.00
=	Selling Price to Retailer				£55,988.33
+	Distribution Cost/Mark-up		15%	£ -	£7,300.61
+	Shipping/Logistics Cost to Distributor Centre		0%	£17.00	£17.00
=	Manufacturers Selling Price				£48,670.72
+	Profit Margin		8%	£ -	£2,949.74
+	Warranty Cost		2%	£ -	£737.44
+	Corporate Allocations		10%	£ -	£3,687.18
+	Business Unit Selling, General & Administrative		12%	£ -	£4,424.61
	Non-Recuring Development Cost	1200000			
	Estimated Production Volume	20000			
+	Allocated Non-Recurring Development Cost			£6.00	£6.00
=	Business Unit Target Cost				£36,871.76
+	Overhead		45%	£ -	£11,442.96
=	Direct Target Cost (Labour & Material)				£25,428.80

Table D-12: Target Cost Calculation for the Dual Band 4-Fragment Fixed Link Terminal

D.5 NPV Analysis – Satcom Terminal

Year	Discount factor 10%	Units Built @ £2,315	Units Sold @ £4,430	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£46,300,000	£0	-£46,300,000	-£46,300,000	-£46,300,000
1	0.91		4000	£0	£17,720,000	£17,720,000	£16,109,091	-£30,190,909
2	0.83		4000	£0	£17,720,000	£17,720,000	£14,644,628	-£15,546,281
3	0.75		4000	£0	£17,720,000	£17,720,000	£13,313,298	-£2,232,983
4	0.68		4000	£0	£17,720,000	£17,720,000	£12,102,998	£9,870,016
5	0.62		4000	£0	£17,720,000	£17,720,000	£11,002,726	£20,872,742

Table D-13: NPV Analysis for the Satcom Standard Terminal

Year	Discount factor 10%	Units Built @ £4,009	Units Sold @ £7,674	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£80,180,000	£0	-£80,180,000	-£80,180,000	-£80,180,000
1	0.91		3600	£0	£27,626,400	£27,626,400	£25,114,909	-£55,065,091
2	0.83		3600	£0	£27,626,400	£27,626,400	£22,831,736	-£32,233,355
3	0.75		3600	£0	£27,626,400	£27,626,400	£20,756,123	-£11,477,232
4	0.68		3600	£0	£27,626,400	£27,626,400	£18,869,203	£7,391,971
5	0.62		3600	£0	£27,626,400	£27,626,400	£17,153,821	£24,545,792

Table D-14: NPV Analysis for the Satcom Single Band 2-Fragment Terminal

Year	Discount factor 10%	Units Built @ £5,683	Units Sold @ £10,877	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£113,660,000	£0	-£113,660,000	-£113,660,000	-£113,660,000
1	0.91		3200	£0	£34,806,400	£34,806,400	£31,642,182	-£83,995,455
2	0.83		3200	£0	£34,806,400	£34,806,400	£28,765,620	-£57,027,686
3	0.75		3200	£0	£34,806,400	£34,806,400	£26,150,563	-£32,511,533
4	0.68		3200	£0	£34,806,400	£34,806,400	£23,773,240	-£10,224,121
5	0.62		3200	£0	£34,806,400	£34,806,400	£21,612,036	£10,037,163

Table D-15: NPV analysis for the Satcom Dual Band 4-Fragment Terminal

D.6 NPV Analysis – PBR Terminal

Year	Discount factor 10%	Units Built @ £266	Units Sold @ £509	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£5,318,800	£0	-£5,318,800	-£5,318,800	-£5,320,000
1	0.91		8000	£0	£4,072,073	£4,072,073	£3,701,885	-£3,469,091
2	0.83		3000	£0	£1,527,027	£1,527,027	£1,262,006	-£1,786,446
3	0.75		3000	£0	£1,527,027	£1,527,027	£1,147,278	-£256,769
4	0.68		3000	£0	£1,527,027	£1,527,027	£1,042,980	£1,133,846
5	0.62		3000	£0	£1,527,027	£1,527,027	£948,164	£2,398,042

Table D-16: NPV analysis for the PBR Standard Terminal

Year	Discount factor 10%	Units Built @ £1,849	Units Sold @ £3,539	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	5000	0	-£9,243,900	£0	-£9,243,900	-£9,243,900	-£36,980,000
1	0.91		500	£0	£1,769,282	£1,769,282	£1,608,439	-£33,119,273
2	0.83		500	£0	£1,769,282	£1,769,282	£1,462,217	-£29,609,521
3	0.75		500	£0	£1,769,282	£1,769,282	£1,329,288	-£26,418,837
4	0.68		500	£0	£1,769,282	£1,769,282	£1,208,444	-£23,518,215
5	0.62		500	£0	£1,769,282	£1,769,282	£1,098,585	-£20,881,287

Table D-17: NPV analysis for the PBR 2-Fragment Terminal

Year	Discount factor 10%	Units Built @ £3,529	Units Sold @ £6,754	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£70,578,400	£0	-£70,578,400	-£70,578,400	-£70,580,000
1	0.91		2000	£0	£13,508,706	£13,508,706	£12,280,642	-£68,124,000
2	0.83		2000	£0	£13,508,706	£13,508,706	£11,164,220	-£65,891,273
3	0.75		2000	£0	£13,508,706	£13,508,706	£10,149,291	-£63,861,521
4	0.68		2000	£0	£13,508,706	£13,508,706	£9,226,628	-£62,016,292
5	0.62		2000	£0	£13,508,706	£13,508,706	£8,387,843	-£60,338,810

Table D-18: NPV analysis for the PBR Multi-band 4-Fragment Terminal

QinetiQ Proprietary

D.7 NPV Analysis – WLAN (FHSS) Terminal

Year	Discount factor 10%	Units Built @ £38	Units Sold @ £68	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£720,000	£0	-£720,000	-£720,000	-£720,000
1	0.91		8000	£0	£544,000	£544,000	£494,545	-£472,727
2	0.83		3000	£0	£204,000	£204,000	£168,595	-£247,934
3	0.75		3000	£0	£204,000	£204,000	£153,268	-£43,576
4	0.68		3000	£0	£204,000	£204,000	£139,335	£142,203
5	0.62	1500	3000	-£54,000	£204,000	£150,000	£93,138	£311,094

Table D-19: NPV analysis for the WLAN Standard Terminal

Year	Discount factor 10%	Units Built @ £96	Units Sold @ £184	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£1,920,000	£0	-£1,920,000	-£1,920,000	-£1,920,000
1	0.91		3000	£0	£552,000	£552,000	£501,818	-£1,418,182
2	0.83		3000	£0	£552,000	£552,000	£456,198	-£961,983
3	0.75		3000	£0	£552,000	£552,000	£414,726	-£547,258
4	0.68		3000	£0	£552,000	£552,000	£377,023	-£170,234
5	0.62		3000	£0	£552,000	£552,000	£342,749	£172,514

Table D-20: NPV analysis for the WLAN 2-Fragment Terminal

Year	Discount factor 10%	Units Built @ £265	Units Sold @ £508	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£5,300,000	£0	-£5,300,000	-£5,300,000	-£5,300,000
1	0.91		1200	£0	£609,600	£609,600	£554,182	-£4,745,818
2	0.83		1200	£0	£609,600	£609,600	£503,802	-£4,242,017
3	0.75		1200	£0	£609,600	£609,600	£458,002	-£3,784,015
4	0.68		1200	£0	£609,600	£609,600	£416,365	-£3,367,650
5	0.62		1200	£0	£609,600	£609,600	£378,514	-£2,989,136

Table D-21: NPV analysis for the WLAN 10-Fragment Terminal

D.8 NPV Analysis – Fixed Link Terminal

Year	Discount factor 10%	Units Built @ £8,421	Units Sold @ £16,118	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£168,420,000	£0	-£168,420,000	-£168,420,000	-£168,420,000
1	0.91		8000	£0	£128,944,000	£128,944,000	£117,221,818	-£109,809,091
2	0.83		3000	£0	£48,354,000	£48,354,000	£39,961,983	-£56,526,446
3	0.75		3000	£0	£48,354,000	£48,354,000	£36,329,076	-£8,087,678
4	0.68		3000	£0	£48,354,000	£48,354,000	£33,026,433	£35,947,565
5	0.62	1500	3000	-£12,631,500	£48,354,000	£35,722,500	£22,180,862	£75,979,605

Table D-22: NPV analysis for the Fixed Link Standard Terminal

Year	Discount factor 10%	Units Built @ £13,314	Units Sold @ £25,484	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£266,280,000	£0	-£266,280,000	-£266,280,000	-£266,280,000
1	0.91		3600	£0	£91,742,400	£91,742,400	£83,402,182	-£182,877,818
2	0.83		3600	£0	£91,742,400	£91,742,400	£75,820,165	-£107,057,653
3	0.75		3600	£0	£91,742,400	£91,742,400	£68,927,423	-£38,130,230
4	0.68		3600	£0	£91,742,400	£91,742,400	£62,661,294	£24,531,064
5	0.62		3600	£0	£91,742,400	£91,742,400	£56,964,812	£81,495,876

Table D-23: NPV analysis for the Fixed Link 2-Fragment Terminal

Year	Discount factor 10%	Units Built @ £25,429	Units Sold @ £48,671	Cash Flow			Net Present Value	NPV Payback & Total
				Expenses	Revenues	Balance		
0	1.00	20000	0	-£508,580,000	£0	-£508,580,000	-£508,580,000	-£508,580,000
1	0.91		3000	£0	£146,013,000	£146,013,000	£132,739,091	-£375,840,909
2	0.83		3000	£0	£146,013,000	£146,013,000	£120,671,901	-£255,169,008
3	0.75		3000	£0	£146,013,000	£146,013,000	£109,701,728	-£145,467,280
4	0.68		3000	£0	£146,013,000	£146,013,000	£99,728,844	-£45,738,437
5	0.62		3000	£0	£146,013,000	£146,013,000	£90,662,585	£44,924,149

Table D-24: NPV analysis for the Fixed Link Dual Band 4-Fragment Terminal

E Conference Papers

Resource Trading for Spectrum Aggregation and Management

Dr David L. Robinson, QinetiQ Malvern UK, DRobinson@QinetiQ.com
Dr Anil K Shukla, QinetiQ Malvern UK, AShukla@QinetiQ.com
John Burns, Aegis Systems Ltd UK, John.Burns@aegis-systems.co.uk

***Abstract*—QinetiQ has developed a solution to manage a system providing services to customers in a dynamic way. This solution, Sky Dust^{TM}, is based on an economic model. Radio spectrum is a good example of a valuable but limited resource that needs to be managed as efficiently as possible. There are numerous ways of applying Sky Dust to do this, including trading spectral bands or managing the use of a single band. By trading spectrum in this way, different parts of the spectrum can be shared and aggregated dynamically to utilize the spectrum as efficiently and fairly as possible. The results so far are encouraging and show an improvement in information throughput of around 200% in a bandwidth limited network. A working demonstrator has been constructed at QinetiQ's resource trading laboratory in Malvern, UK.**

INTRODUCTION

The fragmentation of spectrum - due to the adoption of more spectrally efficient technologies, of old techniques not using frequencies efficiently, (*eg* analogue TV), of allocations not being used and also for historical reasons of allocating spectrum - presents considerable challenges for efficient management of the radio spectrum.

The tradition of the 'command and control' approach of spectrum licensing is seen as economically inefficient and has been linked with restricting technical innovation [1]. This has led regulators (such as Ofcom) to look at new measures to enable market forces to eliminate the weaknesses of the command and control approach and improve efficiencies.

The fragmentation of spectrum not only represents changes in user requirements and services but is also the product of the regulatory environment. Making effective use of multiple and simultaneously available spectral fragments, vacated or otherwise unused / underutilised, for new or alternative services represents an evolutionary development in the use and management of the radio spectrum.

Spectrum aggregation is the collective term for this. The identification of resource usage both practically and as expected from a knowledge of the licensed allocations might well show some differences. This represents a starting block on which to establish whether there are usable fragments available for use, through trading or otherwise, by others.

The general consensus is that there are gains to be had by using spectrum more efficiently and, for the case of fragmentation, some advantage will occur through aggregating these fragments. This bold concept raises many issues dealing with the human perceptions of ownership and of technical solutions to meet these challenges.

There are also significant challenges in identifying how spectrum can be identified, allocated and de-allocated ('traded') in an effective and legitimate way that meets with wide acceptance from regulators, managers and service providers. Even the local Scout troop might have a view when they need radio communications for their Open Day.

There is a need therefore, to look at both market mechanisms and technical solutions to enable 'virtual aggregations' to provide wider bandwidth services. Such aggregations are perceived to be both dynamic and regional in extent.

If it is assumed that fragments of spectrum can be utilized (e.g. using spread spectrum such as frequency hoping spread spectrum, direct sequence spread spectrum or hybrid spread spectrum) in an IP packet type concept to transmit data. The challenge then becomes one of how to ensure that the fragments of spectrum are best exploited to ensure "optimum" use. The method being investigated is via the application of a trading model where a central server is the trading engine [2] and the spectrum segments and the data to be transmitted are traded. A trading model, Sky DustTM, has been developed by QinetiQ to optimize the resources used within a network. This model is to be applied in a hierarchical manner to the spectrum where trading is carried out at many levels. In this paper we describe the technique being developed and the results of simulations to show the improvements achievable when trading takes place. The simulation referred to in this paper is based on trading network capacity for 50 users, and this demonstrates the potential benefits of trading resource and some potential indication of the benefits when trading is achieved in a fragmented spectrum.

THE SCENARIO

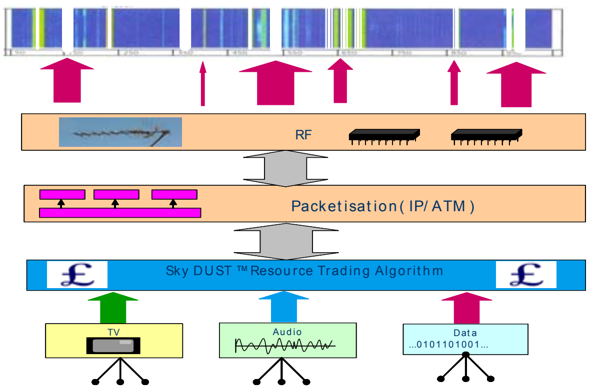

Figure 1 – The spectrum aggregation and resource trading approach

The trading model developed must be scalable for it to be applied ideally to the trading of spectrum and aggregations of spectrum. A number of techniques are being considered to make it possible to aggregate and share spectrum at a number of technology levels. The above approach illustrated in figure 1 involves aggregating spectrum across a number of frequencies using RF techniques. This is illustrated at the top of figure 1. This aggregation results in a number of bandwidth channels which can be packetised. This in turn makes it possible to share users and services across the channels by treating the channels as networks. Finally, the resource trading algorithm is used to trade users demanding different services across the packetised channels, and then to trade their demand within the channels. So in summary, the approach is to take fragments of radio spectrum, aggregate them, packetise them and then trade them across users and services. This approach will be developed under a one-year contract commissioned by Ofcom (UK Office of Communications), leading to modeling work based upon a scenario and a spectrum trading demonstrator.

SPECTRUM FRAGMENTS

The approach taken in identifying available spectrum fragments involved a comprehensive analysis of the current use of spectrum in the range 100 MHz to 5 GHz. This was based on a number of sources, including the UK national frequency allocation table, Ofcom's published channel plans for services such as PMR and fixed links, other Ofcom documents, in particular those relating to the spectrum framework review and dialogue with Ofcom's Business Radio Licensing team. Data were provided by Ofcom on current PMR assignments at four locations in the UK that were intended to represent different intensities of spectrum usage. In each case all frequencies in use within 100 km of the designated location were considered to be unavailable; other frequencies were assumed to be available for aggregation purposes.

The four locations chosen were:

v) Central London– dense urban environment (lowest availability of unused spectrum)

vi) Newcastle – urban environment

vii) Brough (North Yorkshire) – small town environment with nearby urban areas

viii) Ullapool (Scottish Highlands) – remote rural area

The following spectrum has been identified as potentially available on a national basis:

- 862 – 863 MHz
- 1375 – 1389 MHz
- 1399 – 1400 MHz
- 2290 – 2302 MHz
- 3440 – 3442 MHz
- 3475 – 3480 MHz

The following guard bands have been identified and may be useable subject to appropriate interference mitigation measures:

- 915 – 917 MHz (GSM cellular)
- 1350 – 1350.5 MHz (fixed links)
- 3600 – 3605 MHz (FWA)
- 3641 – 3650 MHz (FWA)
- 3875 – 3925 MHz (fixed links)
- 3961 – 3970 MHz (fixed links)
- 4195 – 4200 MHz (fixed links)

Within the existing PMR band for each of the four locations studied it was found that a total of 2.65MHz was found to be unassigned in London, 11.275MHz in Newcastle, and 4.925MHz in Brough. There is therefore a significant amount of fragmented spectrum available for use once an approach to making the best use of this spectrum has been determined.

Spectrum Trading and Sky Dust

The Sky DustTM Dynamic Resource Manager is illustrated in overview in Fig. 2 below. The system being managed is composed of two services (e.g. TV, internet) that can be delivered to a number of customers using a shared resource (ie the fragmented spectrum).

The Sky DustTM agents control customer demand for service. These agents are in turn managed by the Sky Dust trader, which keeps track of the overall use of the shared resource. The need of the customer is taken into account along with the type of service requested. Also use of resource is shared fairly between the customers over time.

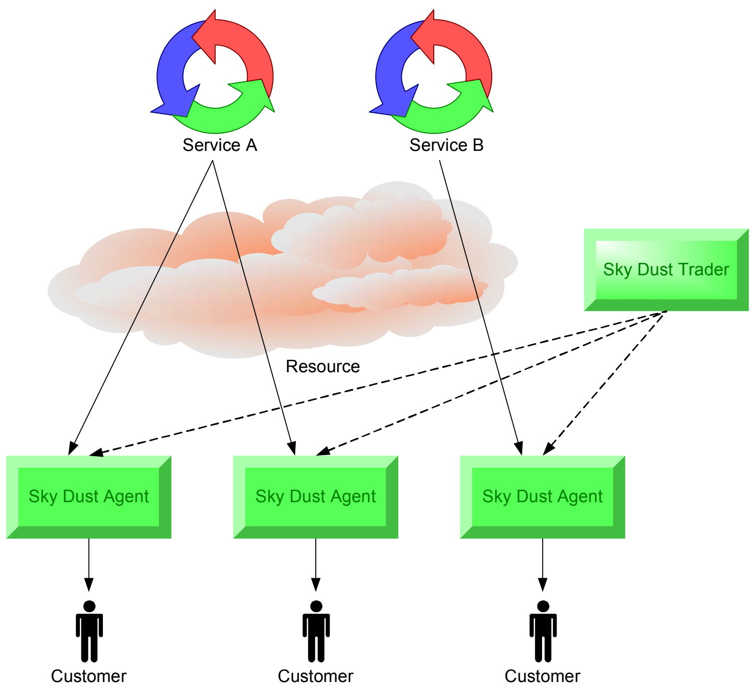

Figure 2 - Overview of the Sky Dust resource manager.

The Sky DustTM trader acts as an overarching co-coordinator of the agents and is a software system that the agents communicate with to request some of the shared resource (e.g. fragments of spectrum). The trader does not need to maintain detailed information about how the shared resource works or its current state. There is no need to take measurements from the shared resource either. All the trader needs to know is the capacity of the resource and then it keeps track of the resource demanded by the agents.

Resource Trading applies a market based approach to sharing a resource, which could be spectrum, across a community of customers by trading on behalf of customers who pay for service. Therefore the approach is to adapt this solution in a number of different ways to manage spectrum through aggregation and sharing. This will involve a hierarchical approach, starting with low level RF aggregation techniques. A packet service approach will then be built on this to provide a number of channels of fixed bandwidth. These channels will then be traded using Sky DustTM to share these channels across a community of users. To further develop this approach, consideration will be given to the possibility of assigning spectrum to operators by trading spectrum between operators. These operators would then use their assignments to provide services to communities of users. To do this, the operators would trade the use of their assignment of spectrum across their community of users as described above.

This hierarchical approach to trading could be extended to trade between areas to allocate spectrum within areas. Under the control of a super trader, it is a natural step to then trade pan geographical regions.

Scalable Implementations of Resource Trading

Resource trading can be applied in the following ways as defined in the table below, along with example applications. An explanation of how each of these implementations would work is then given.

A key idea, which makes many of these implementations possible, is that the resource trading algorithm can be applied to determine the price for a resource, by running the algorithm without allocating resource to the users. This can then be used by traders to bid for their resource allocation. This opens up a whole range of different ways resource trading can be scaled which are described below.

Implementation	Example application
Local trading within an allocated resource	Network bandwidth management in a wireless LAN
Trading across alternative resources	A wireless LAN where each user can join one of many allocations of bandwidth and resource trading domains
Peer to peer across many resources to establish a path to a service	Allocation of bandwidth across a network of networks. A price has to be paid for transit traffic
Peer to peer to trade for a single allocation of resource	Frequency assignment across a geographical region, area by area
Peer to peer to trade for a multiple allocation of resource	Spectrum assignment across geographical regions. Regions are allocated a part of the spectrum which can in turn be traded within a region
Hierarchically to trade for an allocation of shared resource	To share network bandwidth between different communities of users in a shared network, or on a shared single link. For example, different departments may be allocated different amounts of shared bandwidth
Hierarchically to trade across multiple resources for an allocation of single resource	To allocate a frequency to an operator from an allocation of spectrum in an area
Hierarchically to trade across multiple resources for an allocation of multiple resources	To trade for spectrum within a region's allocation of spectrum, to allocate spectrum to an area

In all of these cases the algorithm can be implemented on a central trader which trades with local traders, or it could be implemented by trading between traders. These implementations can be combined as needed to trade resource allocations and resources as required. So hierarchical trading can be combined with peer to peer trading etc. Each of the potential implementations is briefly described below.

Local trading within an allocated resource

This is the basic resource trading algorithm as described in detail. This is the most basic way of applying the algorithm for managing demand in a single resource such as network bandwidth.

Trading across alternative resources

This is one further step to extend the use of the algorithm where there is more than one resource a user could joint and use. For example there could be a number of alternative networks that a user could choose to join. This is a simple extension of the way the algorithm is used. When a user wants to use resource, then the user requests the current price for resource, for each alternative resource from a trader. The trader then responds with the current price associated with each resource and the user then joins the auction for the resource that has the lowest current price. The trading algorithm is then operated in the normal way.

There can be more than one trader offering to sell resource or a number of resources for a given price. So a user can chose the least expensive source of resource across traders as well as across available resources offered by a single trader.

Peer to peer across many resources to establish a path to a service

To trade for resource across many resources the algorithm can be applied to bid for resource along an end to end path. For example, if a user needs to access a remote server outside the user's local network, then the user will need bandwidth from the local network and bandwidth from all the networks along the way to the remote server. This means that networks can have a need to support transit traffic, which is traffic that their users do not source or sink, but does require bandwidth. Each network would then apply the algorithm, to determine a price for the request for transit bandwidth and apply a tariff as the traffic is not locally owned. Therefore the user that bids to buy the bandwidth across many networks will have to pay for local bandwidth, and add the charges for bandwidth across all the other networks involved.

So the algorithm would be modified as follows. A user would trade with its local trader to gain a quote for local resource. The user would also have to obtain a quote for all other resources required along the peer to peer path. A tariff factor would be applied by each remote trader to the price they quote because they don't own this users demand. The users would then apply the algorithm by adding all the quotes, to decide how much resource to bid for. So the price P for resource used by the user in the trading algorithm would be:

$$P= PL+T1*P1+T2*P2+\ldots.Tn+Pn$$

Where PL is the price quoted for local resource, Tn is the tariff applied by resource n for transit demand and Pn is the price quoted by resource n. The user's trader could construct this price on the user's behalf, rather than have the user trade directly with all the traders along the path.

Just as the resource required by a user could be bandwidth, some of the resource could be server capacity, for example. So using the resource trading algorithm, all resources involved in providing a service to a user can be traded in this way, peer to peer, trader to trader, end to end.

Peer to peer to trade for a single allocation of resource

Here a trader trades for an allocation of resource by trading for resource with all the traders that surround it. In this way, traders are allocated resource area by area. This is done as follows.

The resources available to be allocated are ordered by size and therefore value to a trader. These resources are marked as unavailable in a local table if another trader claims them as a result of an auction. Provided the first resource is available, the trader takes the size of the resource and uses this to run the algorithm without allocating resource to it users, and by doing so determines a price for that resource. The trader then bids with all of the surrounding traders using this price. If the trader bids the highest price then it claims the resource and begins to trade it amongst its users. Otherwise it marks the resource as unavailable within its local table and repeats the process for the next resource in the list. This is repeated until the trader wins an auction and claims a resource, or finds that there is no unallocated resource left to bid for. In this case the trader would have to wait for a future bidding round.

The relative amount of money allocated to each trader, which in turn is shared amongst its users, determines how successful they will be in bidding for resource because it affects the price bid. This can be used to control this process to ensure resources are allocated to traders as required. For example those traders that did not win a resource allocation could be paid compensation so that in a future bidding round they are more likely to win an allocation. The amount of money allocated to each trader could also be related to how much real money the traders have paid for a license to bid for resource.

Another surrounding trader can call an auction at any time with the trader to bid for any resource in the list. The trader responds by bidding against the other trader if it has not already claimed a resource to use.

Peer to peer to trade for a multiple allocation of resource

In this case the trader can bid for more than one resource. Again the resources available are ordered by size in a local table and the trader makes a bid for the first available resource, using the algorithm to determine the price for the resource according to the resource size. All traders that lose mark the resource as unavailable in their local table. Once the trader has claimed a resource by winning a bid, the trader continues to bid for more resource in the same way. However, in further bids, the trader adds all the resource previously won to the size of the resource it is bidding for when determining its price. Therefore the bid price will drop as the trader wins more and more resource. This means that as the trader's requirement for resource is satisfied it will be less determined to win further resource when bidding with other traders. In this way, some traders may win one resource to trade and some traders may win many. Some traders could win none at all and would have to wait for a future bidding round.

Once again, the relative amount of money allocated to each trader determines how successful they will be in bidding for resource because it affects the price bid. This can be used to control this process to ensure resources are allocated to traders as required, particularly to those traders that lost out all together in previous bidding rounds. The amount of money allocated to each trader could also be related to how much real money the traders have paid for a license to bid for resource.

Any other surrounding trader can action an auction for any available resource at any time. As previously described, the availability of resource is indicated in a table and, after each auction, all traders that lose in the bidding mark the resource bid for as unavailable in each of their local tables.

Hierarchically to trade for an allocation of shared resource

The resource trading algorithm determines a price for resource as previous explained. This can be used to determine an allocation of part of a resource shared with another trader. For example, there could be two departments sharing the same campus network LAN, whereby each department has a different amount of money allocated to its users, different user needs and therefore different demand for resource. Each department could therefore share a proportion of the bandwidth available and trade that share amongst its users. So a third of the bandwidth could be allocated to department A to trade across its users. The remaining bandwidth would then be traded within department B.

The share of resource would be allocated as follows. Each trader would allocate all of the resource available to the algorithm and trade amongst its users to arrive at a price for resource if all the resource where available. This price reflects demand for resource within the traders market. These users would not be allowed to buy resource at this stage as the algorithm is being used to determine a price only, without selling resource. The resource allocated to each trader would then be divided up according to these prices. So if there were two traders a and b, and Pa is the price for resource determined for trader a, and Pb is for b, then:

Resource allocation for a, Ra = Pa/(Pa+Pb)*Rt, where Rt is the total resource.

Once the share of resource have been allocated, resource trading takes place as normal, using these shares to allocated resource to the users within each trading domain.

As for peer to peer, the relative amount of money allocated to each trader determines how successful they will be in bidding for a share of resource because it affects the price bid, and this can be used to control this process to ensure resources are allocated to traders as required. The amount of money allocated to each trader could be related to how much real money the traders have paid for a license to bid for resource.

Hierarchically to trade across multiple resources for an allocation of single resource

Here there are a number of resources of different sizes available for allocation to traders. The resources are ordered according to size, the largest resource being considered the most valuable and the smallest the least. All traders bid for each resource, starting with the largest resource and working towards the smallest. In bidding for the first resource each trader runs the algorithm using the size of the resource, ie the amount of bandwidth, to determine the price that the trader could sell resource to its users for. This price, as before, represents the demand for resource within the traders' domain. This price is then simply used as the bid to claim the resource. The trader that bids the highest price, and therefore has the greatest need to win the bidding, is allocated the resource. The winner can then use the resource to trade amongst its users using the algorithm and drops out of this auction process. The next resource available is then auctioned across the remaining traders in the same way. This process continues until all the available resource has been allocated to traders to trade, or all traders have the resource they need to trade. Traders that failed to win a resource allocation would have to wait for the next bidding round.

Once again, the relative amount of money allocated to each trader determines how successful they will be in bidding for resource because it affects the price bid, and this can be used to control this process to ensure resources are allocated to traders as required, and to compensate traders that failed to win resource. The amount of money allocated to each trader could be related to how much real money the traders have paid for a license to bid for resource.

Hierarchically to trade across multiple resources for an allocation of multiple resources

In this scenario, the aim is to allocate resource to traders, such that each trader can claim more that one allocation of resource, and then share these allocations across its users, by running more than one resource trading algorithm. This done by ordering the resources to be allocated as previously described. Traders bid for the first resource, as described above, using the algorithm to determine the price for resource each trader will bid. The resource is then allocated to the highest bidder. The highest bidder then continues to bid for more allocations of resource by joining the bidding for the next resource in order. This time though, the bidder that won the first allocation determines its price for resource, using the algorithm, by adding the size of the previously won resource to that of the resource currently being auctioned. This means that the traders demand for more resource is likely to be less and therefore the price bid is likely to be less, as the trader already has a resource allocation. So, using the algorithm and summing resource already won with the resource being bid for, bid prices are determined for traders that already have resource. The trader that bids the highest is then allocated the resource in this auction, which may or may not be a trader that already has resource. All traders then join the next auction in the sequence and the process continues until all resources have been allocated. Some traders may win one resource to trade and some traders may win many and some none at all.

Once again the success of a trader in bidding will be determined by how much money is allocated to the trader relative to the other traders.

The resulting spectrum trading hierarchy

By applying these techniques of hierarchical and peer to peer trading the concept illustrated in figure 3 is being developed as part of this ongoing work. Here a "super trader" makes spectrum available to regional traders through hierarchical trading and regional trader do the same for area traders. For some parts of the spectrum the area traders may trade with each other, peer to peer. Similarly operators trade hierarchically with area trader for there assignments of spectrum and again may also trade peer to peer. Operators then use their assignments of spectrum to provide services to domestic and mobile users whilst applying the Sky Dust solution to make the most of the bandwidth available.

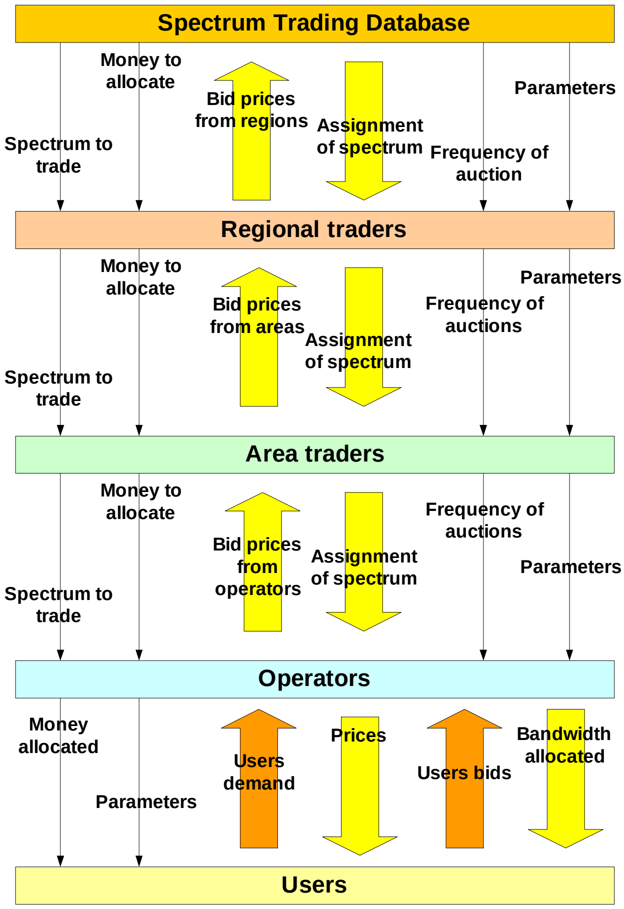

Figure 3 - The spectrum trading hierarchical and peer to peer concept.

RESULTS FROM SKY DUSTTM

In the Matlab simulation below the bandwidth in a computer network supporting 50 users is traded. In the simulation the network was loaded beyond the limit of its capacity by a factor of two and was therefore heavily congested. The simulation was run with and without trading. A message was deemed successful if it was delivered according to its quality of service requirements. In each case the simulation was run for 1000 seconds.

The results so far are encouraging and show an improvement in information throughput of around 200% in a bandwidth limited and congested network. These results are illustrated in Fig. 4.

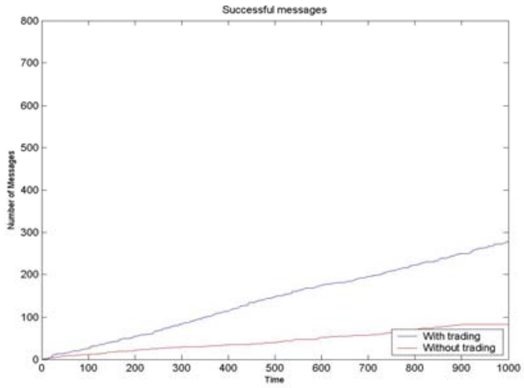

Figure 4 - Graph of successful messages in a congested network of 50 users, with and without trading.

SUMMARY AND CONCLUSIONS

Spectrum aggregation may be increasingly important as the spectrum becomes fragmented through market mechanisms and therefore there is a need to be able to aggregate and share spectrum to make the best use of it. A promising approach to make this possible is to trade the bandwidth with the services that need to use it. This paper describes the Sky DustTM approach that was developed to trade network bandwidth, which may also be applied to spectrum trading. The intention is to develop a spectrum trading solution by applying this approach in a hierarchical manor.

The potential benefits of trading are demonstrated via the simulation of sky dust in network bandwidth trading scenario. The results showed that a 200% improvement could be obtained if the resource was traded. It is hoped that similar benefits will be realized in our trading approach to spectrum aggregation.

REFERENCES

[1] Up the Revolution', Roger Dettmer, The IEE Review, May, 2005, *pp* 44-47.
[2] Pompermaier, A Market-Based Approaches to resource Allocation 2003
[3] Gibbens, R.J and Kelly, F.P Resource Pricing and the Evolution of Congestion Control. 1999

ACKNOWLEDGMENTS

The following are acknowledged for their support for this paper: Dr Alan C Smith, QinetiQ, who has kindly acted as a reviewer; John Burns, Aegis and Phillipa Marks, Indenpen Consultants, for their ongoing support in this study.

F Virtual aggregation solution

F.1 Introduction

If it is assumed that technically a system can be developed to exploit fragments, then the question arises on how the spectrum aggregation system be exploited fully to generate the maximum value. This may be achieved by a virtual aggregation solution in which the aggregating device supports many services (and hence users) to generate the most income. In this scenario, however, the access to the fragmented spectrum (which may be geographically variable) must be managed and it is proposed that a trading approach be adopted to maximise the benefits.

In this section we discuss the virtual aggregation of spectrum composed of spectrum fragments aggregated using RF techniques to provide a useful amount of bandwidth, which are then packetised so that they can be treated as a commodity that offers bandwidth. Therefore the emphasis here is the trading of spectrum aggregates to achieve virtual aggregation and hence maximise its value or use.

The overall spectrum aggregation concept is illustrated below in Figure F-1 and the methods of enabling the aggregation and trading of spectrum in a hierarchical and peer to peer trading system are examined below. The idea is that this trading is to be effected by a computer system, whereby the trading is internal to the system and does not involve the exchange of real money nor real spectrum licenses nor contracts. The money used in this trading system is virtual money and could be thought of as tokens.

In this trading system, spectrum is traded as a commodity that can provide bandwidth. By doing this, spectrum is shared across users and services according to the pattern of demand for bandwidth. In this way none contiguous pieces of spectrum are virtually aggregated according to where the demand for bandwidth is, by sharing the bandwidth associated with the pieces of spectrum across all the required services as if the bandwidth is a continuous block. Hence the bandwidth offered by the pieces of spectrum is virtually aggregated.

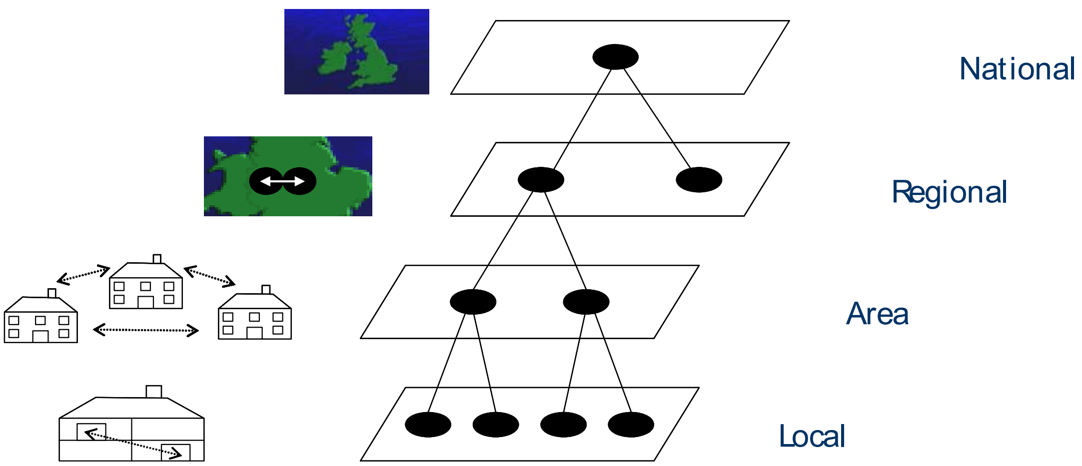

Figure F-1: Illustration of the hierarchical and peer to peer trading concept

In this trading hierarchy, spectrum aggregates (fragments) are traded starting at national level down to local level for assignment in the home. Peer to peer trading is permissible at any level in this hierarchy, particularly at local level, as this encourages the dynamic sharing of aggregates. It is assumed that this trading is supported by a network of computers that carry out the trading such that the trading traffic does not consume any of the bandwidth offered by the spectrum aggregates.

To facilitate this trading a resource trading algorithm could be used (for example the QinetiQ Sky DustTM Dynamic Resource Manager). This resource trading approach is key to enabling the virtual aggregation outlined here through making it possible to trade spectrum and trade aggregates of spectrum.

In this section an overview is given of resource trading is given and a proposal is made on how a resource trading algorithm can be used achieve the virtual aggregation previously described. The ways that resource trading can be scaled and applied to spectrum trading are explored. This leads to a description of the proposed aggregation and trading architecture. Having explored the potential architectures, likely application scenarios, where the proposed solution could be applied are outlined. The impact of the hidden node problem is also discussed and the practicalities of implementing the proposed trading solution are summarised.

F.2 Overview of the Resource Trading and Aggregation Approach

An example of a resource trading algorithm available is The Sky DustTM Dynamic Resource Manager. This is illustrated in Figure F-2 and used to outline the virtual aggregation concepts further. In this example, the virtual system being managed is composed of two services (e.g. TV, internet) that can be delivered to a number of customers using a shared resource (ie the fragmented spectrum).

The Sky DustTM agents control customer demand for service. These agents are in turn managed by the Sky DustTM trader, which keeps track of the overall use of the shared resource. The need of the customer is taken into account along with the type of service requested. Also use of resource is shared fairly between the customers over time.

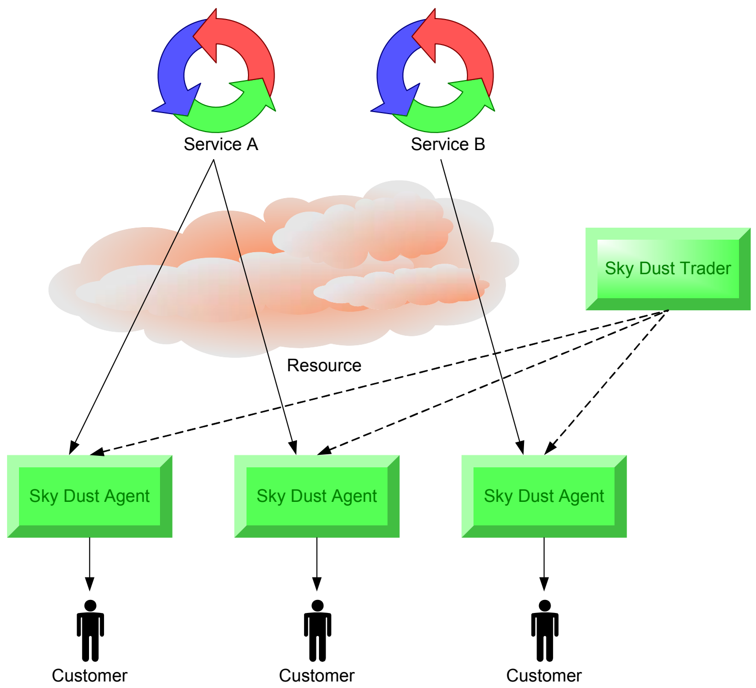

Figure F-2: The Sky Dust™ Dynamic Resource Manager

The Sky DustTM trader acts as an overarching co-coordinator of the agents and is a software system that the agents communicate with to request some of the shared resource (e.g. fragments of spectrum). The trader does not need to maintain detailed information about how the shared resource works or its current state. There is no need to take measurements from the shared resource either. All the trader needs to know is the capacity of the resource and then it keeps track of the resource demanded by the agents.

Resource Trading applies a market based approach to sharing a resource, in this case spectrum, across a community of customers by trading on behalf of customers who "pay" for service using virtual money. Each agent knows how much virtual money its customer is allocated and uses this to trade as part of the trading mechanism. Therefore the approach is to adapt this solution in a number of different ways to manage spectrum through aggregation and sharing. This involves a hierarchical approach, starting with low level RF aggregation techniques. A packet service approach is then built on this to provide a number of channels of fixed bandwidth. These channels are then traded using Sky DustTM to share these channels across a community of users. This is then further developed to assigning spectrum to operators by trading spectrum between operators. These operators would then use their assignments to provide services to communities of users. They would do this by trading the use of their assignments of spectrum across their community of users as described above. This trading could take place every minute, for example.

This hierarchical approach to trading could then be extended to trade between areas to allocate spectrum within areas. It is a natural step to then trade pan geographical regions, from regions down to areas and from areas down to operators. Once it is possible to trade assignments using this resource trading technique, the hierarchical and peer to peer trading concept initially introduced here becomes possible.

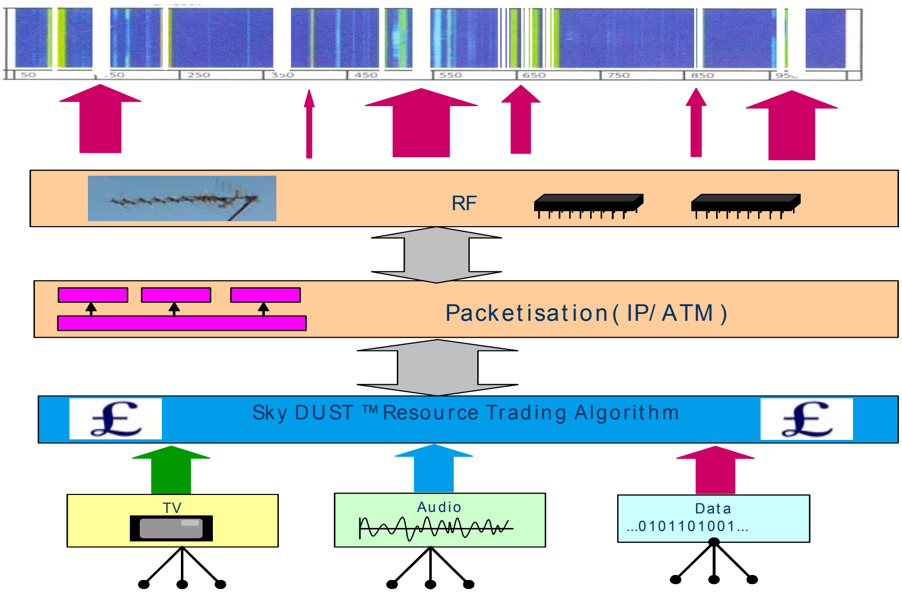

Figure F-3: Components of the proposed Aggregation and Trading Approach

In summary the approach involves aggregating spectrum across a number of frequencies using RF techniques. This is illustrated at the top of Figure F-3. This aggregation results in a number of bandwidth channels which can be packetised. This in turn makes it possible to share users and services across the channels by treating the channels as networks. Finally, the resource trading algorithm is used to trade users demanding different services across the packetised channels, and then to trade their demand within the channels. The approach is to take fragments of radio spectrum, aggregate them, packetise them and then trade them across users and services. The details of how this could be done now follow leading to a proposed trading architecture and scenarios that could be used for experimental emulation.

F.3 RF Aggregation and Packetisation Techniques

Listed below are the techniques and protocols initially identified for aggregating and packetising spectral fragments to generate a useful amount of bandwidth from pieces of fragments which would otherwise be of little use. These RF techniques are explored in the main report. The reason to include an overview here is to highlight that there are techniques to aggregate spectrum at the RF level and that there are MAC access protocols which can be used to deliver packetised services, which in turn can be resource traded in line with the trading model described above.

- Multi carrier parallel channels
- Multi carrier spread spectrum
- Frequency hopping spread spectrum

- Direct sequence spread spectrum
- Hybrid spread spectrum
- **Wi**reless **Fi**delity (Wi-Fi) MAC access protocol (IEEE 802.11)
- **W**orldwide **I**nteroperability for **M**icrowave **Acc**ess (WiMAX) MAC access protocol (IEEE 802.16)

Spread spectrum seems intuitively a suitable approach to aggregating the bandwidth available in spectral fragments. A simple approach would be to modulate several carriers by transmitting bits in parallel, assuming that the fragments were identical in capacity. Another approach might be to bit hop the bits across the frequency hops, thereby simplifying the design of the radios as only one modulator would be needed. Direct sequence would not work because all the frequencies generated by the spread spectrum sequence would be modulated by the data bits. Although this is a useful technique for transmitting data beneath the noise and for allowing multiple access through code separation, but would not directly increase the bandwidth available by aggregating fragments.

Another possibility is to make use of the code separation offered by direct sequence to create parallel channels of equal capacity spread across the spectrum to enable data bits to be transmitted in parallel. The code would have to be carefully engineered to fit within the available spectrum fragments. An advantage of this technique occurs where some fragments have a lower capacity than others. If modulating the whole sequence at the data rate defined by the larger capacity fragments, the only effect in the lower capacity areas of the spectrum is that the noise floor would be raised. Another advantage is that there could be as many carriers as there are orthogonal codes which would define the degree of parallelism possible and therefore the degree to which the data rate, or available aggregated bandwidth, could be multiplied. This approach could also be applied to direct sequence spread spectrum as a different way of spreading data bits across orthogonal channels.

By using direct sequence or frequency hopping, channels of different capacities could be created in accordance with the capacity of the spectral fragments chosen for aggregation. Channels of equal capacity could then be aggregated by spreading data bits across them in parallel to build up data channels of equal capacity. Several small channels generated through these spread spectrum techniques could be aggregated in this way to form the equivalent to a much larger channel. So large channels could be made available alongside aggregated smaller channels and treated as an equal source of bandwidth.

In the first instance the WiFi and WIMAX MAC access protocols appear suitable to use the aggregated fragments to provide network bandwidth. WiFi is based on the IEEE 802.11 standard which has a MAC access layer suited to local area networks, and is very similar to the Ethernet protocol. WiMAX is based on the IEEE 802.16 standard which has a MAC access layer best suited to metropolitan area networks.

These protocols enable the bandwidth offered by aggregates to be shared in a network of users. Also, by treating the bandwidth offered as a network, the bandwidth can be used to support a wide range of different services including video streaming, email, web, file transfer, audio streaming and voice services. It is therefore an important step to realise that in the spirit of sharing spectrum and making the most of it, one of the key ideas that supports this aim is the use of MAC access protocols to use the spectrum to provide packetised network bandwidth.

F.4 Resource trading architectures

Resource trading was originally developed for managing the bandwidth in a computer network by trading demand for network bandwidth between users. A representation of the resource trading algorithm is shown below in Figure F-4.

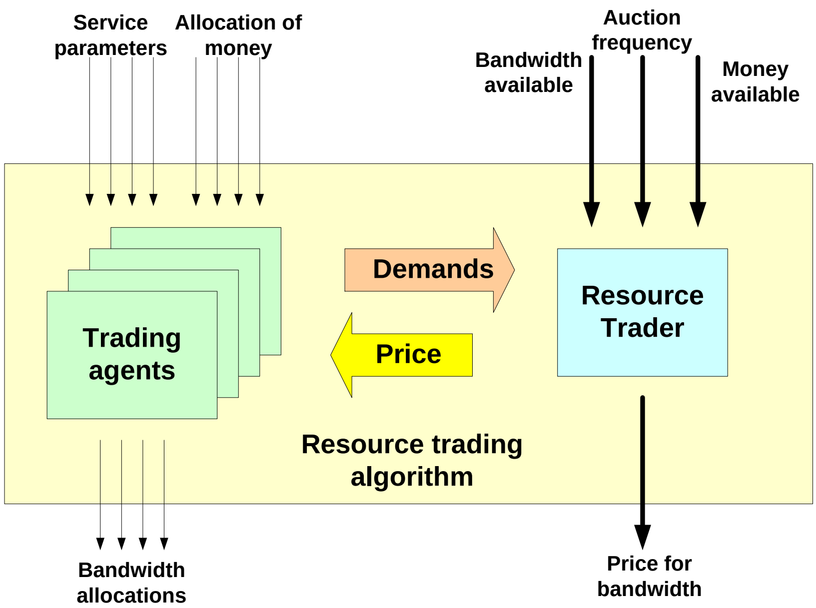

Figure F-4: The resource trading algorithm

There are a number of resource trading agents shown to the left of the diagram which are configured with service parameters and an allocation of virtual money for each agent. The service parameters represent the value or network information or service, such as web pages and emails, for each individual user. To the right of the diagram is the resource trader. This is configured with the total amount of virtual money available within the internal trading market to be divided amongst the trading agents, the frequency of the auctions and the total amount of bandwidth available to trade. Each trading agent acts on behalf of the user when the user actions an applications to transfer information. The agent will then trade with the trading agent to "buy" network bandwidth along side the other agents by taking part in an auction. The trading agents declare the bandwidth they need and the trader responds with a "price". This influences the bandwidth the agents are prepared to buy so they respond with a lower, for example, amount of bandwidth they would like. This negotiation continues as part of an auction process until the total demand for bandwidth matches the bandwidth available.

Normally this algorithm would be used to generate assignments of bandwidth for each individual user. However, the algorithm naturally generates a "price" for bandwidth as part of the auction process and reflects the demand for bandwidth in the network. This "price" can be used as a bid price to "buy" bandwidth and would be internal to the automated trading system considered here. It therefore becomes possible to use the algorithm to declare an internal "price" for some bandwidth, based on the amount of bandwidth offered for assignment and the demand for bandwidth in the network. As spectrum can be used to provide network bandwidth, the algorithm can be used to bid for assignments of spectrum. The realisation that the algorithm can be used to generate bid prices instead of assignments of bandwidth leads to a wide range of different ways that resource trading can be scaled up in complex computer network and spectrum trading scenarios. The different ways of applying resource trading are explained below along with example applications.

F.5 Scalable Implementations of Resource Trading

Resource trading can be applied in the following ways as defined in the table below, along with example applications. An explanation of how each of these implementations would work is also given.

A key idea, which makes many of these implementations possible, is that the resource trading algorithm can be applied to determine the price for a resource, by running the algorithm without allocating resource to the users. This can then be used by traders to bid for their resource allocation. This opens up a whole range of different ways resource trading can be scaled which are described below (Table F-1).

Implementation	Example application
Local trading within an allocated resource	Network bandwidth management in a wireless LAN
Trading across alternative resources	A wireless LAN where each user can join one of many allocations of bandwidth and resource trading domains
Peer to peer across many resources to establish a path to a service	Allocation of bandwidth across a network of networks. A price has to be paid for transit traffic
Peer to peer to trade for a single allocation of resource	Frequency assignment across a geographical region, area by area
Peer to peer to trade for a multiple allocation of resource	Spectrum assignment across geographical regions. Regions are allocated a part of the spectrum which can in turn be traded within a region
Hierarchically to trade for an allocation of shared resource	To share network bandwidth between different communities of users in a shared network, or on a shared single link. For example, different departments may be allocated different amounts of shared bandwidth
Hierarchically to trade across multiple resources for an allocation of single resource	To allocate a frequency to an operator from an allocation of spectrum in an area
Hierarchically to trade across multiple resources for an allocation of multiple resources	To trade for spectrum within a region's allocation of spectrum, to allocate spectrum to an area

Table F-1: Scaleable Resource Trading Examples

In all of these cases the algorithm can be implemented on a central trader which trades with local traders, or it could be implemented by trading between traders. These implementations can be combined as needed to trade resource allocations and resources as required. So hierarchical trading can be combined with peer to peer trading etc.

F.5.1 Local Trading within an Allocated Resource

This is the basic resource trading algorithm as described in detail. This is the most basic way of applying the algorithm for managing demand in a single resource such as network bandwidth.

F.5.2 Trading Across Alternative Resources

This is a further step to extend the use of the algorithm where there is more than one resource a user can auction for and use. For example there could be a number of alternative networks that a user could choose to join. This is a simple extension of the way the algorithm is used. When a user wants to use a resource, then the user requests the current price for the each alternative resource available from a trader. The trader then responds with the current price associated with each resource and the user then joins the auction for the resource that has the lowest current price. The trading algorithm is then operated in the normal way.

There can be more than one trader offering to sell resource or a number of resources for a given price. So a user can choose the least expensive source of resource across traders as well as across available resources offered by a single trader.

F.5.3 Peer to Peer Across Many Resources to Establish a Path to a Service

To trade for resource across many resources the algorithm can be applied to bid for resource along an end to end path. For example, if a users needs to access a remote server outside the user's local network, then the user will need bandwidth from the local network and bandwidth from all the networks along the way to the remote server. This means that networks can have a need to support transit traffic, which is traffic that their users do not source or sink, but does require bandwidth. Each network would then apply the algorithm, to determine a price for the request for transit bandwidth and apply a tariff as the traffic is not locally owned. Therefore the user that bids to buy the bandwidth across many networks will have to pay for local bandwidth, and add the charges for bandwidth across all the other networks involved.

The algorithm would be modified as follows. A user would trade with its local trader to gain a quote for local resource. The user would also have to obtain a quote for all other resources required along the peer to peer path. A tariff factor would be applied by each remote trader to the price they quote because they don't own this users demand. The users would then apply the algorithm by adding all the quotes, to decide how much resource to bid for. So the price P for resource used by the user in the trading algorithm would be:

$P = Pl + T1*P1 + T2*P2 + \ldots Tn + Pn$

Where Pl is the price quoted for local resource, Tn is the tariff applied by resource n for transit demand and Pn is the price quoted by resource n. The user's trader could construct this price on the user's behalf, rather than have the user trade directly with all the traders along the path.

Just as the resource required by a user could be bandwidth, some of the resource could be server capacity, for example. So using the resource trading algorithm, all resources involved in providing a service to a user can be traded in this way, peer to peer, trader to trader, end to end.

F.5.4 Peer to Peer to Trade for a Single Allocation of Resource

Here a trader trades for an allocation of resource by trading for resource with all the traders that surround it. In this way, traders are allocated resource area by area. This is done as follows.

The resources available to be allocated are ordered by size and therefore value to a trader. These resources are marked as unavailable in a local table if another trader claims them as a result of an auction. Provided the first resource is available, the trader takes the size of the resource and uses this to run the algorithm without allocating resource to it users, and by doing so determines a price for that resource. The trader then bids with all of the surrounding traders using this price. If the trader bids the highest price then it claims the resource and begins to trade it amongst its users. Otherwise it marks the resource as unavailable within its local table and repeats the process for the next resource in the list. This is repeated until the trader wins an auction and claims a resource, or finds that there is no unallocated resource left to bid for. In this case the trader would have to wait for a future bidding round.

The relative amount of money allocated to each trader, which in turn is shared amongst its users, determines how successful they will be in bidding for resource because it affects the price bid. This can be used to control this process to ensure resources are allocated to traders as required. For example those traders that did not win a resource allocation could be paid compensation so that in a future bidding round they are more likely to win an allocation. The amount of money allocated to each trader could also be related to how much real money the traders have paid for a license to bid for resource.

Another surrounding trader can call an auction at any time with the trader to bid for any resource in the list. The trader responds by bidding against the other trader if it has not already claimed a resource to use.

F.5.5 Peer to Peer to Trade for a Multiple Allocation of Resource

In this case the trader can bid for more than one resource. Again the resources available are ordered by size in a local table and the trader makes a bid for the first available resource, using the algorithm to determine the price for the resource according to the resource size. All traders that lose mark the resource as unavailable in their local table. Once the trader has claimed a resource by winning a bid, the trader continues to bid for more resource in the same way. However, in further bids, the trader adds all the resource previously won to the size of the resource it is bidding for when determining its price. Therefore the bid price will drop as the trader wins more and more resource. This means that as the trader's requirement for resource is satisfied it will be less determined to win further resource when bidding with other traders. In this way, some traders may win one resource to trade and some traders may win many. Some traders could win none at all and would have to wait for a future bidding round.

Once again, the relative amount of money allocated to each trader determines how successful they will be in bidding for resource because it affects the price bid. This can be used to control this process to ensure resources are allocated to traders as required, particularly to those traders that lost out altogether in previous bidding rounds. The amount of money allocated to each trader could also be related to how much real money the traders have paid for a license to bid for resource.

Any other surrounding trader can action an auction for any available resource at any time. As previously described, the availability of resource is indicated in a table and after each auction all traders that lose in the bidding, mark the resource bid as unavailable in each of their local tables.

F.5.6 Hierarchically to Trade for an Allocation of Shared Resource

The resource trading algorithm determines a price for resource as previous explained. This can be used to determine an allocation of part of a resource shared with another trader. For example, there could be two departments sharing the same campus network LAN, whereby each department has a different amount of money allocated to its users, different user needs and therefore different demand for resource. Each department could therefore share a proportion of the bandwidth available and trade that share amongst its users. So a third of the bandwidth could be allocated to department A to trade across its users. The remaining bandwidth would then be traded within department B.

The share of resource would be allocated as follows. Each trader would allocate all of the resource available to the algorithm and trade amongst its users to arrive at a price for resource if all the resource were available. This price reflects demand for resource within the traders market. These users would not be allowed to buy resource at this stage as the algorithm is being used to determine a price only, without selling resource. The resource allocated to each trader would then be divided up according to these prices. So if there were two traders a and b, and Pa is the price for resource determined for trader a, and Pb is for b, then:

Resource allocation for a, Ra = Pa/(Pa+Pb)*Rt, where Rt is the total resource.

Once the share of resource have been allocated, resource trading takes place as normal, using these shares to allocate resources to the users within each trading domain.

As for peer to peer, the relative amount of money allocated to each trader determines how successful they will be in bidding for a share of resource because it affects the price bid, and this can be used to control this process to ensure resources are allocated to traders as required. The amount of money allocated to each trader could be related to how much real money the traders have paid for a license to bid for resource.

F.5.7 Hierarchically to Trade Across Multiple Resources for an Allocation of Single Resource

Here there are a number of resources of different sizes available for allocation to traders. The resources are ordered according to size, the largest resource being considered the most valuable and the smallest the least. All traders bid for each resource, starting with the largest resource and working towards the smallest. In bidding for the first resource each trader runs the algorithm using the size of the resource, i.e. the amount of bandwidth, to determine the price that the trader could sell resource to its users for. This price, as before, represents the demand for resource within the traders' domain. This price is then simply used as the bid to claim the resource. The trader that bids the highest price, and therefore has the greatest need to win the bidding, is allocated the resource. The winner can then use the resource to trade amongst its users using the algorithm and drops out of this auction process. The next resource available is then auctioned across the remaining traders in the same way. This process continues until all the available resources have been allocated to traders to trade, or all traders have the resource they need to trade. Traders that failed to win a resource allocation would have to wait for the next bidding round.

Once again, the relative amount of money allocated to each trader determines how successful they will be in bidding for resource because it affects the price bid, and this can be used to control this process to ensure resources are allocated to traders as required, and to compensate traders that failed to win resource. The amount of money allocated to each trader could be related to how much real money the traders have paid for a license to bid for resource.

F.5.8 Hierarchically to Trade Across Multiple Resources for an Allocation of Multiple Resources

In this scenario, the aim is to allocate resource to traders, such that each trader can claim more that one allocation of resource, and then share these allocations across its users, by running more than one resource trading algorithm. This done by ordering the resources to be allocated as previously described. Traders bid for the first resource, as described above, using the algorithm to determine the price for resource each trader will bid. The resource is then allocated to the highest bidder. The highest bidder then continues to bid for more allocations of resource by joining the bidding for the next resource in order. This time though, the bidder that won the first allocation determines its price for resource, using the algorithm, by adding the size of the previously won resource to that of the resource currently being auctioned. This means that the traders demand for more resource is likely to be less and therefore the price bid is likely to be less, as the trader already has a resource allocation. Using the algorithm and summing resource already won with the resource being bid for, bid prices are determined for traders that already have resource. The trader that bids the highest is then allocated the resource in this auction, which may or may not be a trader that already has resource. All traders then join the next auction in the sequence and the process continues until all resources have been allocated. Some traders may win one resource to trade and some traders may win many and some none at all.

Once again the success of a trader in bidding will be determined by how much money is allocated to the trader relative to the other traders.

F.6 Proposed Aggregation Trading Architecture

Drawing on all the different ways that resource trading can be applied as previously described, it becomes possible to architect a spectrum aggregation trading system by combining hierarchical trading with peer to peer trading. The first step to do this is to introduce a trading hierarchy where traders bid for their assignment of spectrum aggregations. This trading hierarchy is further developed by introducing peer to peer to increase the flexibility of the trading system. The trading hierarchy followed by the complete trading system is now explained along with the tables that would be needed to support the trading process.

F.6.1 The Spectrum Aggregation Trading Hierarchy

An example of trading for spectrum hierarchically is illustrated in Figure F-5. It is assumed that all of the traders here are part of a networked computer system, which is interconnected as part of a wired intranet.

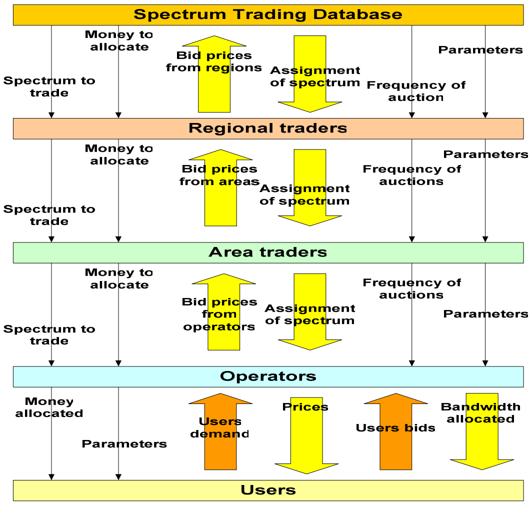

Figure F-5: Spectrum Trading Hierarchy

At the top of the hierarchy there is a spectrum trading database holding records of all the spectrum aggregates that are available for trade. This database could be a distributed database rather than a centralised database to ensure that the database is resilient. All of the traders in the hierarchy would hold a similar database of spectrum aggregates that they have assigned to them to trade. Each trader would periodically offer a spectrum aggregate for auction amongst the traders below and the highest bidder would win the aggregate as an assignment. Bid prices are generated for a particular spectrum aggregate from the bottom of the hierarchy using the resource trading algorithm. So the spectrum trading database offers an aggregated to the regional traders who currently cannot bid because they do not have bid prices. To obtain a bid price the regional trader offer the spectrum aggregate to the area traders and they in turn offer it to the operators. The operators then run the resource trading algorithm to determine the demand for the bandwidth the spectrum aggregate can offer. This then leads to a bid price per operator for the aggregate and each operator users this bid price to bid to their area trader. Each area trader then adds up all of the operator bids and uses the total as the bid to the regional trader. Each regional trader then adds up the area bids and uses the total to bid to the super trader. The highest regional bid wins the aggregate for the region. The winning region then assigns the aggregate to the highest area bid. The winning area then assigns the aggregate to the highest operator bid. The winning operator then uses the spectrum aggregate to provide services to its users.

Spectrum aggregates are put up for auction from the top down. Consequently, by running the resource trading algorithm between the users and the operators in response, bid prices are generated from the bottom up. Assignments of spectrum are then propagated down the hierarchy to the highest bidders. In this way, spectrum aggregates are assigned according to where the demand for spectrum is.

To support this process, the spectrum trading database also passes down the hierarchy all the resource trading parameters needed by the algorithm such as the frequency an aggregate should be auctioned, the assignment of money to each operator and the resource trading parameters which the algorithm needs to run. Operators would not be allowed to change their allocation of money as this would disrupt the trading market. Similarly operators would not be allowed to change the frequency of auction. However, operators would be allowed to tailor the trading parameters to reflect their own ideas of service priority and information value according to their business requirements. Therefore those parameters that are passed down from the spectrum trading database can be thought of as the default parameters that will enable in general optimal trades to make the most of the available bandwidth assigned to the operator's networks.

In this case the trading hierarchy has been considered to be composed of regional traders, area traders and operators and is therefore three trading layers deep for assignment. However, the hierarchy can be extended or contracted as required to include more layers or fewer layers.

F.6.2 The Complete Trading System

By introducing peer to peer trading to the hierarchy the complete trading system illustrated below in figure 4.3 becomes possible. In this case an aggregate can be peer to peer traded at any level in the hierarchy using the bid prices generated from the bottom up. Where an aggregate is to be peer to peer, instead of being assigned to the highest bidder at that level, it is assigned to all bidders and then they must in turn exchange bid prices with their local peers to decide which peer wins the assignment. For example, if an aggregate is to be traded between areas, peer to peer, instead of the regional trader assigning the aggregate to the area with the highest bid, the regional trader effectively assigns it to all areas. Each area then has to consider who its peers are and bid against them. The area that bids the highest price in a peer group wins the assignment. Winning areas can then hierarchically assign the aggregate to their subordinate operators.

The benefit of peer to peer trading is that it allows multiple assignments of aggregates in an area or region taking into account the interference between peers. It also allows peers to share an assignment in a dynamic way.

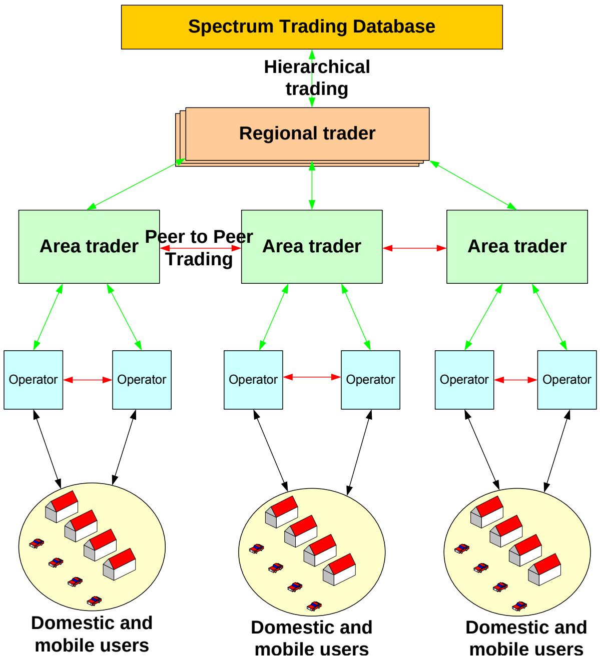

Figure F-6: The complete trading system

The complete trading system is therefore as illustrated above in Figure F-6 where a spectrum trading database makes a spectrum aggregates available for hierarchical trading to the regional traders. The regional traders assign the aggregates to the area traders for peer to peer trading. The area traders then trade their assignments hierarchically down to the operators. The operators may then have to trade their assignments peer to peer. Having won a peer to peer bid between operators for an assignment, they then use their assignment to provide a service to their mobile and domestic users using the resource trading algorithm. Once again, as the algorithm generates a price for bandwidth, the mobile and domestic users can chose to use the operator who offers the lowest price for bandwidth at any instance. In this way user demand is shared between operators.

F.6.3 Tables to Support the Spectrum Aggregation Trading Process

The tables needed to support this trading are illustrated below in

Table F-2. These tables would be implemented at every level in the hierarchy and would keep track of the assignment of spectrum aggregates and control the trading process.

Master Table of Spectrum Available to Trade

Spectrum to Trade														
Freq. and Aggregations			Assignment by Level				Trading Method by Level			Trading Frequency by Level			Bids	
Freq.	Agg.	BW	Assign.	Region	Area	Operator	Region	Area	Operator	Region	Area	Operator	Bid	Status

Supporting Tables

Service List	
Freq.	Service

Restrictions			
Freq	Region	Area	Operator

Peers				
Region	Area	Operator	Peer	Bid Received?

Subordinates				
Region	Area	Operator	Subord.	Bid Received?

Configuration Tables

Operators						
Region	Area	Operator	Money	License	Start	End

Service Trading Parameters				
Region	Area	Operator	Service	Trading parameters

Table F-2: Tables to support the trading of spectrum aggregates

These tables are organised as follows;

- Master Table
  - Spectrum to Trade – Maintains a list of the aggregates available for trading, where, how and how often they are to be traded, and provides support for bidding
- Supporting Tables
  - Service List – Maps service to aggregates to ensure assignments are appropriate
  - Restrictions – Contains a list of where aggregates should not be assigned
  - Peers – Defines the peers with which to trade at each level of the trading hierarchy
  - Subordinates – Defines the subordinates with which to trade at all level in the hierarchy
- Configuration Tables

- o Operators – Contains a list of all the operators in the system along with their money allocations and license details
- o Service Trading Parameters – For each operator, contains all the parameters needed by the resource trading algorithm

The master table contains records of all the frequencies and aggregations available to trade, along with the amount of bandwidth the fragments would offer. For example the Frequency and Aggregations field could be “400MHz”, “24” and 500KHz” to mean frequency 400MHz is part of aggregation 24 and this aggregation offers 500Khz of bandwidth. The spectrum trading database would hold all of the spectrum aggregates available whilst the regional, area and operators master tables would contain only those aggregates assigned to them, or those currently offered to them as part of an auction. To the right of the table is a bid status set of fields which would be used to keep track of assignments that are “Won”, “Lost”, ”Bidding” or “Assigned”, for example. If the aggregate had been won the Status field would be “Won” and the entry would be kept, if the assignment was lost it would be “lost” and the whole entry would subsequently be removed. If the trader was bidding then the Status field would be “Bidding” and the Bid field would contain the bid offered. Finally, once an operator has won an assignment to use, the field would be “Assigned”. The bid offered would be the sum of the bids offered by the subordinates, for example the bid offered by an area trader would be the sum of the bids offered by the operators.

The Assignment by Level fields to the left of the table determines where assignment should be static or dynamic. The first field would be “Dynamic” if assignment was to be dynamic at all level using the trading process. The first field would be “Static” if assignment was to be static and the Region, Area and Operator fields would contain the data that defines at what level in the hierarchy the assignment is static. For example the fields could be “Static”, “Worcestershire”, “Malvern”, ”NA” to mean the aggregate is to be statically assigned to Malvern but can be dynamically trading in the Malvern area. The Trading Method by Level fields determines where peer to peer trading is to take place in the hierarchy. For example if the Region, Area, Operator fields were “Hierarchical”, “Peer to Peer” and “Hierarchical” then the aggregate should be peer to peer traded at the area level and hierarchically traded above and below. Finally, the master table has as set of field called Trading Frequency by level. These fields simple contain data on how often a spectrum aggregate is to be traded. If the Region, Area and operator field were “Monthly”, “Daily” and “Hourly”, then the aggregate should be traded every month at the regional level, every day at the area level and every hour at the operator level.

The tables which support the master table are Service Lists, Restrictions, Peers and Subordinates. The service list table contains a list of frequencies and the services that they are appropriate for. This is to prevent a trade taking place that results in an aggregation of frequencies suited to wide area coverage being assigned to a local area home network, which the equipment could not support and would be a waste of wide area spectrum. The Restrictions table simple allows certain frequencies to be barred from particular Regions, Areas and Operators and contains a list of frequencies associated with where they are not to be assigned. The peers and subordinates tables contain lists of who the peers and subordinates are for a particular trader in the hierarchy. For example an area trader in Lancaster may have entries in the Subordinates table “Lancashire”, “Lancaster”, “Operator A” meaning Operator A is a subordinate to the Lancaster area trader. An entry in the Peers table “Lancashire”, “Lancaster”, “Operator A”, “Operator B” would indicate that Operator B is a peer of Operator A in Lancaster and they should peer to peer trade aggregates that must be traded peer to peer.

Finally, the configuration tables consist of a table of Operators and a table of Trading Parameters. The Operators table contains information on who and where the operators are in the hierarchy, their allocation of virtual money and their license details. The Trading Parameters table contains all the trading parameters needed by the resource trading algorithm to trade.

F.7 Scenario for Spectrum Aggregation

F.7.1 Proposed Trading Scenario for Emulation

A domestic trading scenario is proposed for emulation and this is illustrated in Figure F-7below. In this scenario, a number of users have their own home networks which are supported by an operator. The operator provides access to the internet through a wireless access network to these home networks.

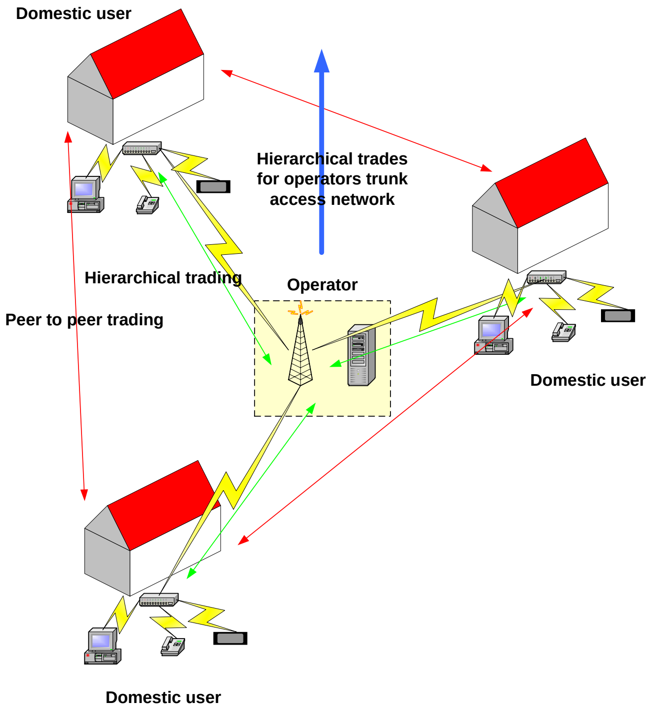

Figure F-7: Domestic Trading Scenario

Each home would have a home network access unit, which manages the spectrum aggregates assigned within the house to the home networks and controls access to the operator's network through the operator's spectrum. Some computing devices, digital TVs and phones may be directly connected to the access unit. Trading for aggregate assignment would take place hierarchically through the operator and peer to peer with the neighbours between the access units. The operator in turn would have to bid for an assignment of spectrum for the wireless access network, to provide connectivity to the home network access units.

F.7.2 Experimental Scenario

Below illustrates the overarching experimental scenario proposed.

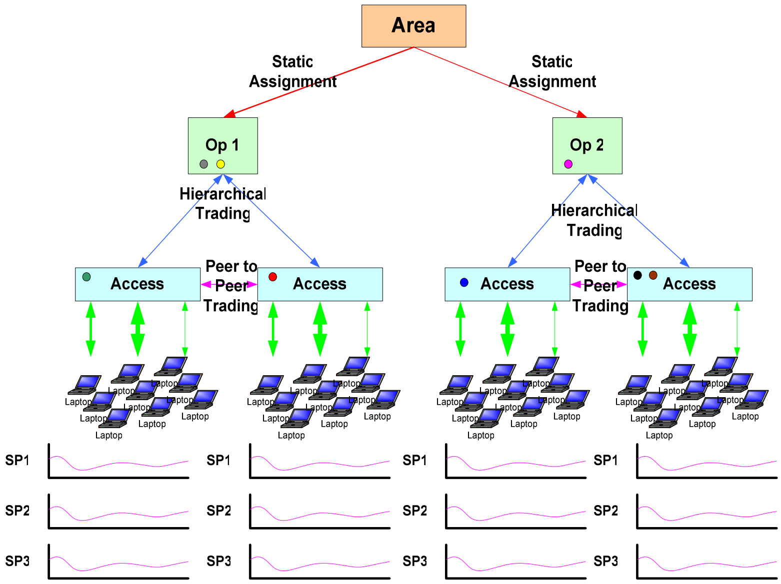

Figure F-8: Experimental Scenario

Here four home access units providing connectivity to an operator's network. Each access unit trades for spectrum aggregates to provide a number of wireless networks in the home, shown here as green arrows. The arrows have different thicknesses to indicate differing amounts of bandwidth. A number of domestic network devices trade for bandwidth across these wireless networks and within these wireless networks using Sky DustTM. Peer to peer trading for one of the aggregates takes place between two of the access units for each operator. The associated wireless network will therefore move between these homes as the demand for bandwidth changes within the homes. Finally the operators will make available for hierarchical auction further aggregates from time to time. This will mean that additional networks will become available to the homes according where the demand is. The coloured dots in the diagram are to indicate where spectrum fragments are at any time. A default aggregate should be statically assigned to each home, to enable trading to take place for assignments in the first instance when trading for assignments has yet to start. This assignment could represent the basic amount of bandwidth each home is initially guaranteed. Alternative approaches to managing this assignment need to be considered further.

F.8 The Impact of Hidden Nodes

The diagram shown in Figure F-9 is intended to represent the hidden node problem. To explain the problem the case of a wireless local area network (LAN) has been chosen as an example.

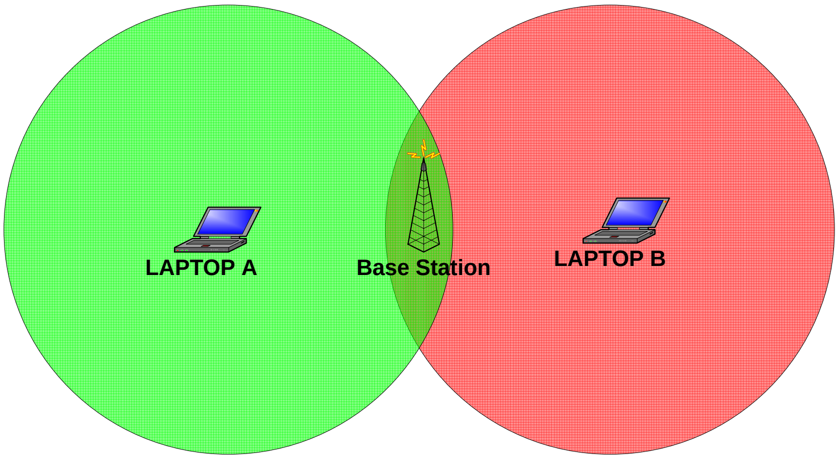

Figure F-9: Illustration of the Hidden Node Problem in a wireless LAN

Here there are two laptops who communicate with a base station using wireless LAN. The coloured circles represent the area of coverage for each laptop's wireless LAN. It can be seen that both laptops can reach the base station. However laptop A cannot see laptop B. This means that as both laptops try to send a packet to the base station at the same time, neither laptop will detect a collision and will continue to transmit the packet. The base station will then receive a corrupted packet. This is the hidden node problem and clearly the medium access control (MAC) layer in the wireless LAN is not working because laptops A and B cannot see each other, even though they can see the base station. The result is that the wireless LAN operates inefficiently with a high error rate in the MAC layer that the higher layer protocols must compensate for.

In the spectrum aggregation trading system this effect would be reduced in general as far as possible, through determining where spectrum is assigned taking into account the anticipated propagation distances. However, on a packet by packet basis, due to local interference and propagation effects, the hidden node problem can be expected to occur and the MAC layer protocols and higher layer protocols will still have to compensate for the loss of packets caused. This is an example of a reason to map services to spectrum aggregates to ensure that the aggregates assigned are suitable for the service they are to be used for. This is also a reason to use the tables to ensure that the aggregates assigned in a particular area and environment can offer the coverage needed. In the case of wireless LAN and the potential hidden node problem, aggregates would be assigned that were able to offer the required area of coverage so that this problem is generally less likely to arise. The area to be covered by a wireless LAN would have to be agreed and defined for the trading tables to be set up. Therefore, provided that the operators of wireless LANs stick to the agreed rules for the area of coverage, the hidden node problem would have a reduced impact, as specified in a service level agreement (SLA), in the spectrum aggregate trading scenario. If the rules were ignored by the operators then, in this example, the impact would be an increase in the rate of occurrence of the hidden node problem, causing a higher rate of packet loss in the MAC layer and a reduced performance for the operators' wireless LANs.

By using the trading tables to control where and what aggregates are assigned by the trading mechanism the impact of the hidden node problem can be generally reduced, provided that operators obey the rules which could be mandated in a contract or agreement.

F.9 Practicalities of implementation

To enable the implementation of a resource trading system a radio of appropriate complexity is needed to support the resource trading algorithms and the RF spectrum aggregation techniques. Encouragingly, it happens that there is a move towards the cognitive radio concept in which radios have in built intelligence and RF complexity. The diagram below in Figure F-10 is intended to show the relevance of cognitive radio to the components of the aggregation and trading approach.

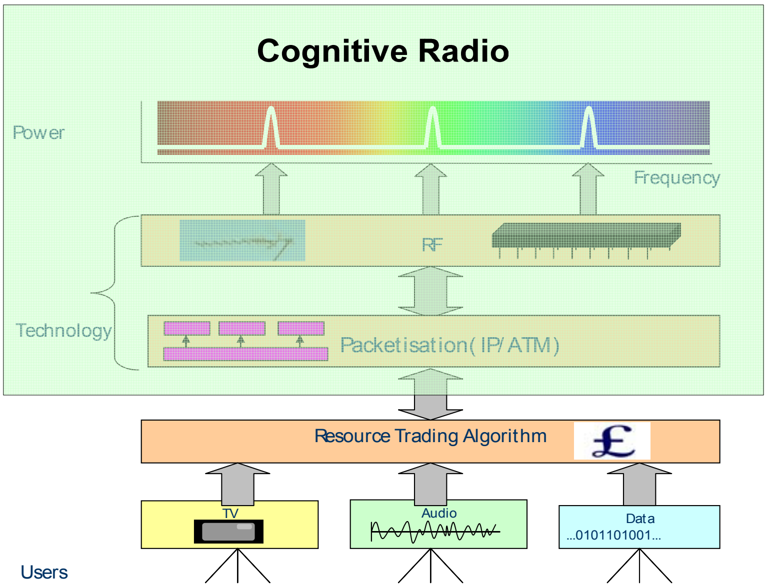

Figure F-10: The relevance to Cognitive Radio

As a generalisation, cognitive radio deals with the MAC layer of spectrum assignment and applies techniques such a game theory to decide which frequencies to use when forwarding a packet, for example. By doing this the cognitive radio can provide the means to aggregate fragments of spectrum and make the aggregation available to the higher layers as a useful piece of bandwidth. Resource trading can then be run at the higher layers to deal with resource assignment, and it is compatible with the cognitive radio concept of policy control as the trading parameters can be managed according to policies. Resource trading is a natural fit within the cognitive radio concept and therefore as the cognitive radio area develops the sophistication and capability of the radios available, the potential to implement the trading approach is facilitated.

In addition to the need for more intelligent radios that are capable of running the resource management and spectrum aggregate assignment algorithms, everywhere this is applied would have to have a spectrum access manager. For example, this would be the home spectrum access unit. It has been noted that there is a current progressive deployment of Digi boxes in the home due to the role out of digital TV. A device like this, which has already become acceptable to the domestic user, has the basic capabilities to act as a home spectrum access unit because they have the basic components of a computer which would be capable of running the assignment algorithms. An enhanced version of the Digi box could be a means to provide broadband internet access to wireless home networks as well as providing access to digital TV. This would also make it possible to trade bandwidth and assignment for the digital TV service to the home, in support of the increasing interactive nature of Digital TV services.

Finally, almost all domestic devices which would use spectrum for networking and communications within the home have a basic computing capability, which would be potentially able to run the resource trading algorithm. For example, wireless phone handsets, wireless LAN computers and peripherals.

F.10 Exploitation

A band manager may use this trading approach to manage the spectrum and fragments of spectrum, which the band manager has purchased, to generate the most revenue by selling it to users and operators. Service providers and operators may use these techniques to generate the most business in providing services to customers using the spectrum they have bought from band managers or have licensed to them.

The most likely and tangible initial application appears to be the domestic wireless broadband access market where this trading approach could be introduced under license by an broadband ISP to sell more bandwidth to domestic users in support of new and novel services. A limiting factor may be the availability of equipment that is capable of supporting RF spectrum aggregation and trading. However, it is noted that wireless LAN cards for laptops and personal computers currently support 4 alternate operating frequencies and 4 alternate codes. It is therefore possible to introduce initially a trading solution in the domestic wireless broadband access scenario, by trading these frequencies and codes. A solution like this could be treated as a reference application, which demonstrates the potential and the practicalities of the approach with a view to developing it further by introducing the aggregation of fragments.

F.11 Conclusions and Recommendations

Trading and aggregation is possible provided that the aggregates can be treated as independent commodities. Trading tables should be used to bound and control real time dynamic trading to ensure that assignments are appropriate and sensible, and that the system behaves in a stable way.

The consideration given to the hidden node problem shows that by using the trading tables to control where and what spectrum aggregates are assigned by the trading mechanism, the hidden node problem can be engineered out.

To fully ascertain the benefits of a virtual aggregation trading system, it was concluded that greater knowledge was required on single service aggregation. In particular how a single service system may be developed in hardware and the comparative cost difference between an aggregated solution and a conventionally designed system.

Initial distribution list

External
Ahmad Atefi, Ofcom
John Burns, Aegis
Phillipa Marks, Indepen

QinetiQ
Information Resources
Alan Taylor
Anil Shukla
Dr David Robinson
Aaron Stockton
Edward Burbidge
Liz Baker