Final Report Annex

Annex to
A Study on the Efficient Dimensioning of Broadband Wireless Access Networks

Ofcom Ref: AY4463

Cotares Ltd.

Cambridge University Engineering Dept.

Cambridge Broadband Ltd.

© Crown Copyright 2004

All requests for reproduction should be addressed to Her Majesty's Stationery Office

C•O•T•A•R•E•S

Consulting and Research

Annexe Contents

1 Introduction and Summary

This document forms an Annex to the main report entitled “A Study on the Efficient Dimensioning of Broadband Wireless Access Networks”.

In Section 2 “Project Infrastructure Elements” the hardware and software building blocks of the project are described. The components of the Cambridge Broadband FWA equipment are explained in Section 2.1 and the choices of wireless Hotspot equipment are discussed in Section 2.2. The configuration of the core communications network is described in Section 2.3 and the way volunteer users of the system were handled is summarised in Section 2.4.

Section 3 on Cost Models shows how it may be possible to construct FWA systems which can economically backhaul last-hop service providers in rural areas.

Section 4 is a review of international issues relating to the bands around 3.5 GHz.

2 Project Infrastructure Elements

2.1 Cambridge Broadband FWA Equipment

The CBL FWA equipment has many features which make it a useful choice for this research exercise. Its efficient Media Access Control (MAC) layer carries data across the air with low overheads. It can operate at high power and its equalisation gives good non-line-of-sight support. These factors combine to give it excellent coverage, and 64QAM connections have been demonstrated at 27 km range.

In terms of networking it is based on Asynchronous Transfer Mode (ATM) from which it inherits a sophisticated capacity for priority scheduling of different connections across the air. The hardware supports multiple modulation levels for users within the same base station sector so that those with excellent signal to noise can benefit from 16QAM or 64QAM modulation while others use QPSK. It has a scaleable system architecture which allowed us to build a complicated radio network for this trial involving a base station with two jumps of radio backhaul.

2.1.1 System overview

A CBL VectaStar Base Station consists of one or more sectors, each served by an Access Point (AP) with up to 212 CPE’s (Customer Premises Equipment) per sector. The standard Access Point antennas are 90 degree sector antennas; hence four Access Points are required for 360 degree coverage, as shown in Figure 2-2. Each AP can support up to 1024 ATM Virtual Circuits, with Quality of Service (QoS) configured individually per Virtual Circuit. Each CPE can have one or more services, with each service carried over a separate ATM VC.

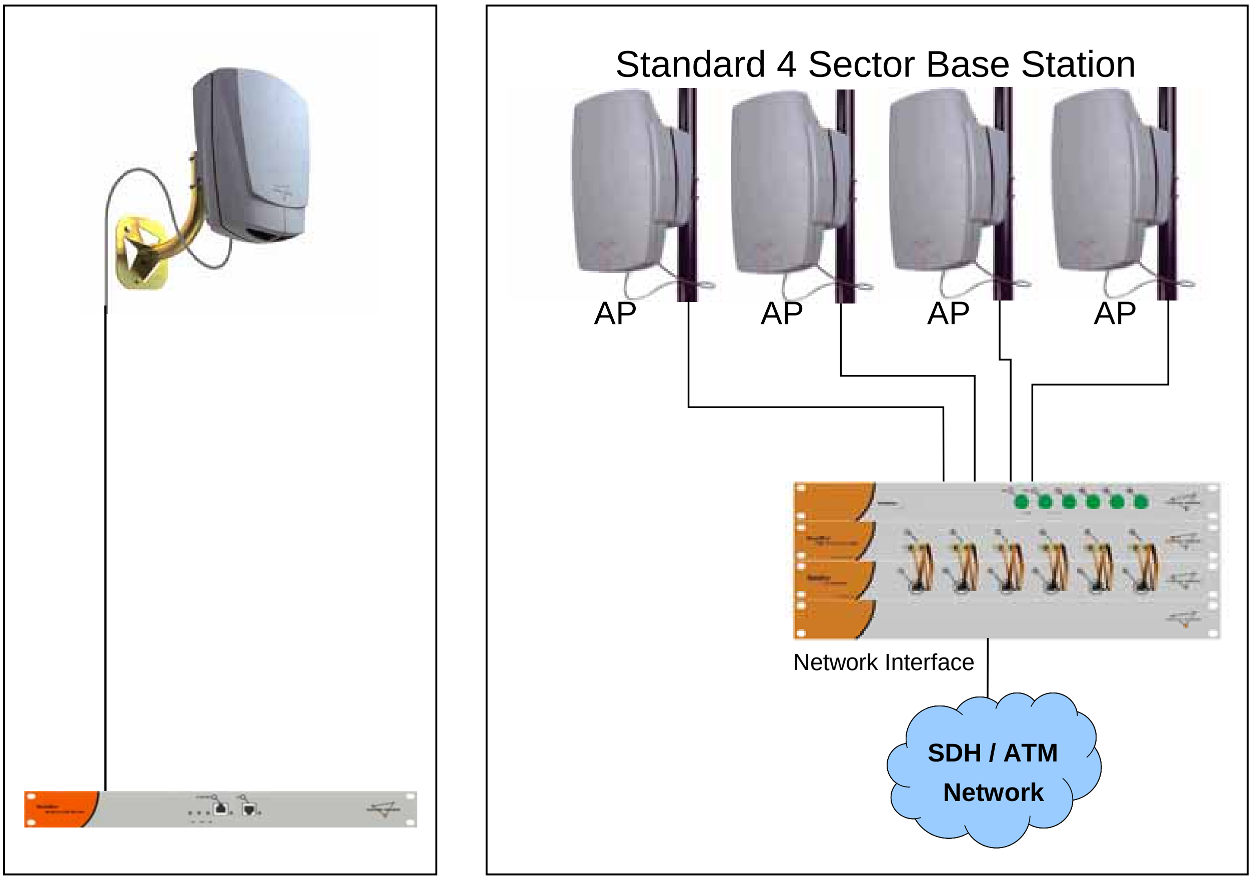

Figure 2-1: CBL Customer Premises Equipment (left) and base station equipment (right)

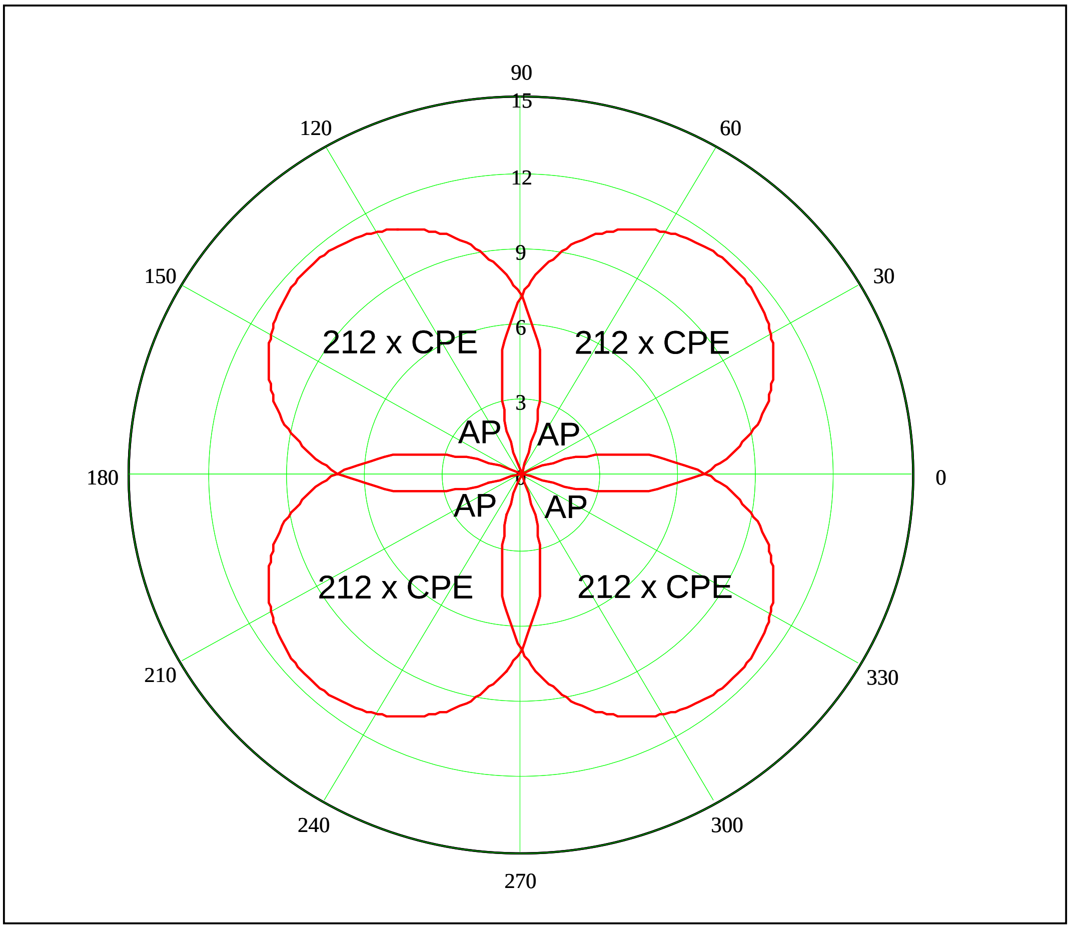

Figure 2-2: Four CBL AP's provide 360 degree coverage at a base station

The schematic on the left of Figure 2-1 shows the equipment used in this trial at customer premises. It consists of an ATM-based Out Door Unit (ODU) and an Interface and Control Unit (ICU), which is mounted indoors, to provide IP functionality and E1 (a synchronous protocol used by telecommunications operators). The ODU connects to the ICU over fibre and also requires a 48V power connection over several sets of twisted pairs with RJ-45 terminations.

On the right of Figure 2-1 is a schematic showing the elements of a base station. Four AP units are connected by fibre to an ATM multiplexer. This performs the role of a simplified ATM switch and is known as a MUX. The other rack elements are the PC which controls the base station, known as the APC, and power and cable distribution units.

Figure 2-3 and Figure 2-4 show the radiation patterns for AP and CPE antennas respectively. The fact that very little power is radiated on axis from the back of the APs allows back-to-back APs to operate on the same frequencies at base stations.

2.1.2 Supported services

Each CPE can support the following multiple service types:

- Ethernet bridge
- multiple routed IP connections
- multiple VLAN connections
- multiple E1 connections
- native ATM
- multiple VoIP

The choice of using Ethernet bridges was made for most connections both to volunteer subscribers and for the management LAN, all base station and backhaul kit was placed on the same Ethernet subnet.

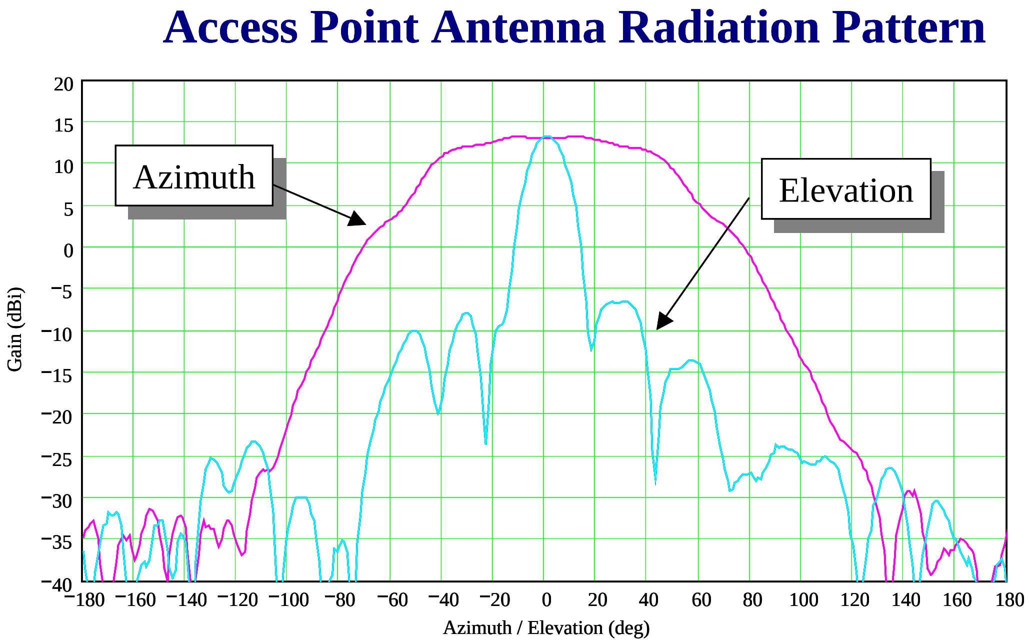

Figure 2-3: Access Point Radiation Pattern

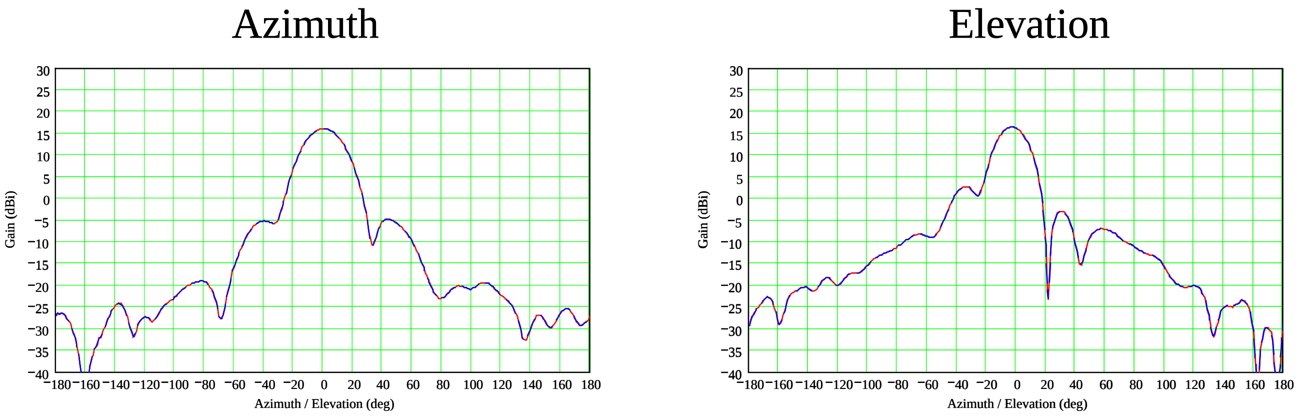

Figure 2-4: Customer Premises Equipment radiation pattern

Each service is carried over a unique over-the-air ATM virtual circuit which can be configured for:

- Modulation QPSK, 16 QAM or 64 QAM (uplink & downlink are independent)
- Bandwidth (peak rate limit and committed rate limit)
- Scheduler priority
- Constant Bit Rate (CBR) or non-CBR service
- CBR synchronisation clock scheme

Extensive use was made of the Quality of Service parameters when experimenting with the system configurations which illustrate the research targets of this project.

2.1.3 System flexibility

The flexibility of the CBL architecture was well illustrated by the chain of radio-backhauled base stations set up as part of the project deployment, see Figure 2-5. In this configuration a bundle of ATM virtual circuits representing all the user connections at the Bourn base station, plus all the management connections, is brought back across the Bourn to Lime Kiln Hill radio backhaul link. They continue through the MUX units at Lime Kiln Hill, where they are joined by a further collection of ATM virtual circuits representing subscriber and management connections at Lime Kiln Hill base station. The full collection of ATM virtual circuits are then brought back over the radio backhaul from Lime Kiln Hill to the network core in the Gates Building.

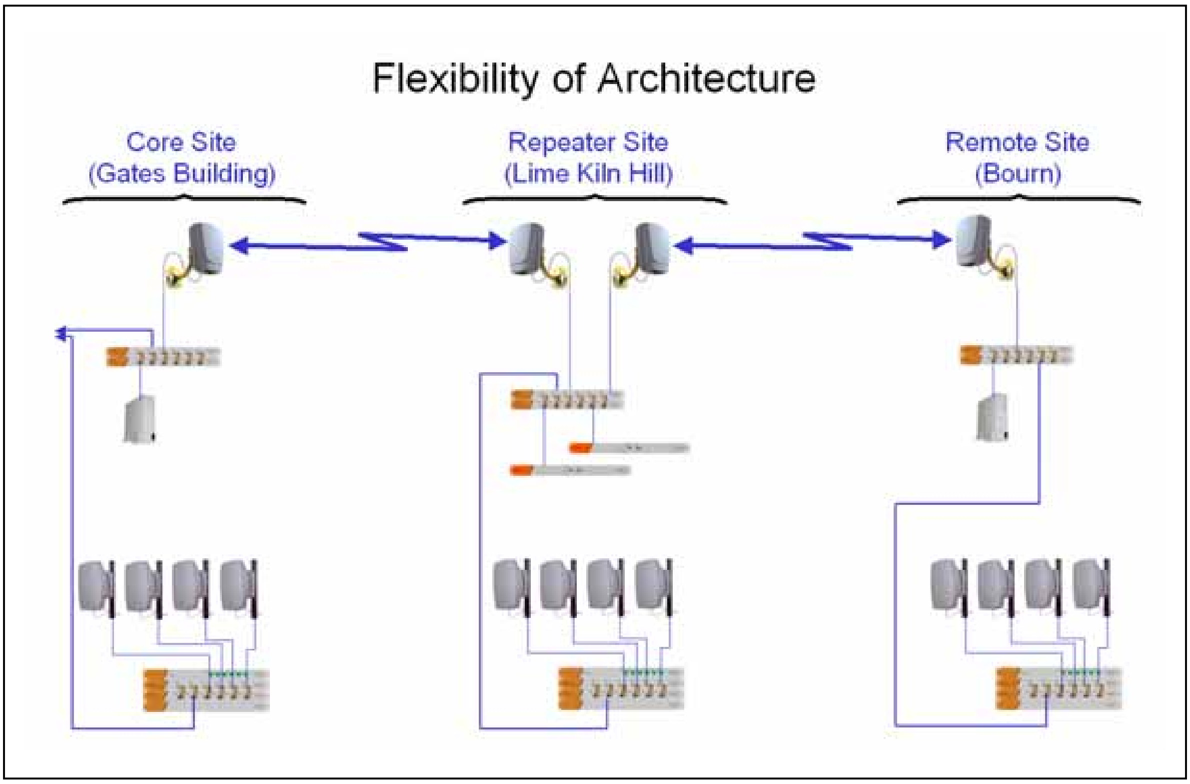

Figure 2-5: A part of the project deployment showing a chain of radio backhauls illustrates the flexibility of the CBL architecture

2.1.4 Non-line of sight propagation

There is no universal definition of non-line-of-sight (non-LOS) propagation. Any radio link where there is a partial obstruction between radios will suffer from excess path loss (meaning that the path loss is higher than free space loss). Excess path loss will be the single most important factor in reducing coverage for an FWA system. Any link can suffer from delay spread (reflections), if the reflections are very strong and from the ground then they can cause flat fading, meaning total loss of signal power, and no equalisation technique can overcome this. If the reflections are from nearby buildings then they can cause delay spread (frequency selective fading) and an equaliser (or OFDM) is used to overcome this. Figure 2-5, above, classifies the various factors affecting good radio coverage.

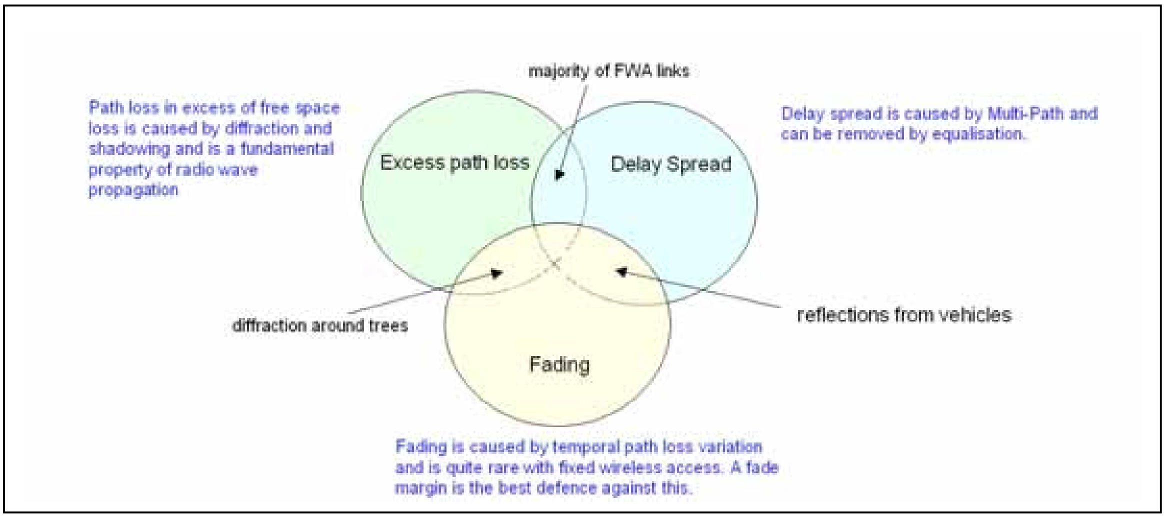

Figure 2-6: Factors affecting radio propagation

The CBL design incorporates a number of features to improve radio coverage. The high link budget means that 3.5 GHz FWA non-LOS coverage is power limited not delay spread limited and the higher the link budget, the better the coverage. Secondly, circularly polarised antennas reduce multi-path resulting in better non-LOS performance. The benefit of this is that reduced reflections allow a higher re-use of bandwidth in the base station frequency colouring plan, and so result in more efficient use of bandwidth. Finally, adaptive equalisation is used to improve the non-LOS coverage. A 128 symbol sequence used to sound the radio channel and a large punctured decision feedback equaliser is employed, which is the optimal solution for handling post cursor delay spread. The equaliser has a delay spread tolerance of 3.6 µS in a 14 MHz channel, rising to 28.8 µS in a 1.75 MHz channel.

2.1.5 Choice of modulation and frequency re-use

Higher order modulations offer a higher capacity for any given bandwidth, but require higher Signal to Noise Ratios (SNR). In the CBL system each CPE can be configured with its own modulation level, according to the SNR available. Figure 2-7, shows the constellation mappings for each of QPSK, 16QAM and 64QAM modulations. It is clear that the amount of space around the constellation points for QPSK will permit operation with a much lower SNR than the higher order modulations, which require progressively better SNR.

As an example of the channel capacity of the system, in a 2.5 MHz channel with a 2.0 MHz symbol rate the gross throughputs at different modulations would be 4Mbps for QPSK, 8 Mbps for 16QAM and 12 Mbps for 64QAM. This corresponds to rates of 2 bits per symbol for QPSK, 4 bits per symbol for 16QAM and 6 bits per symbol for 64QAM.

For an isolated base station, SNR is simply a function of transmit power and path loss. In LOS deployments path loss is simply a function of range and in non-LOS deployments other effects will dominate. In order to maximise capacity frequencies are re-used in real deployments but this means that self-interference can reduce the SNR. This is very relevant to the interests of this research project.

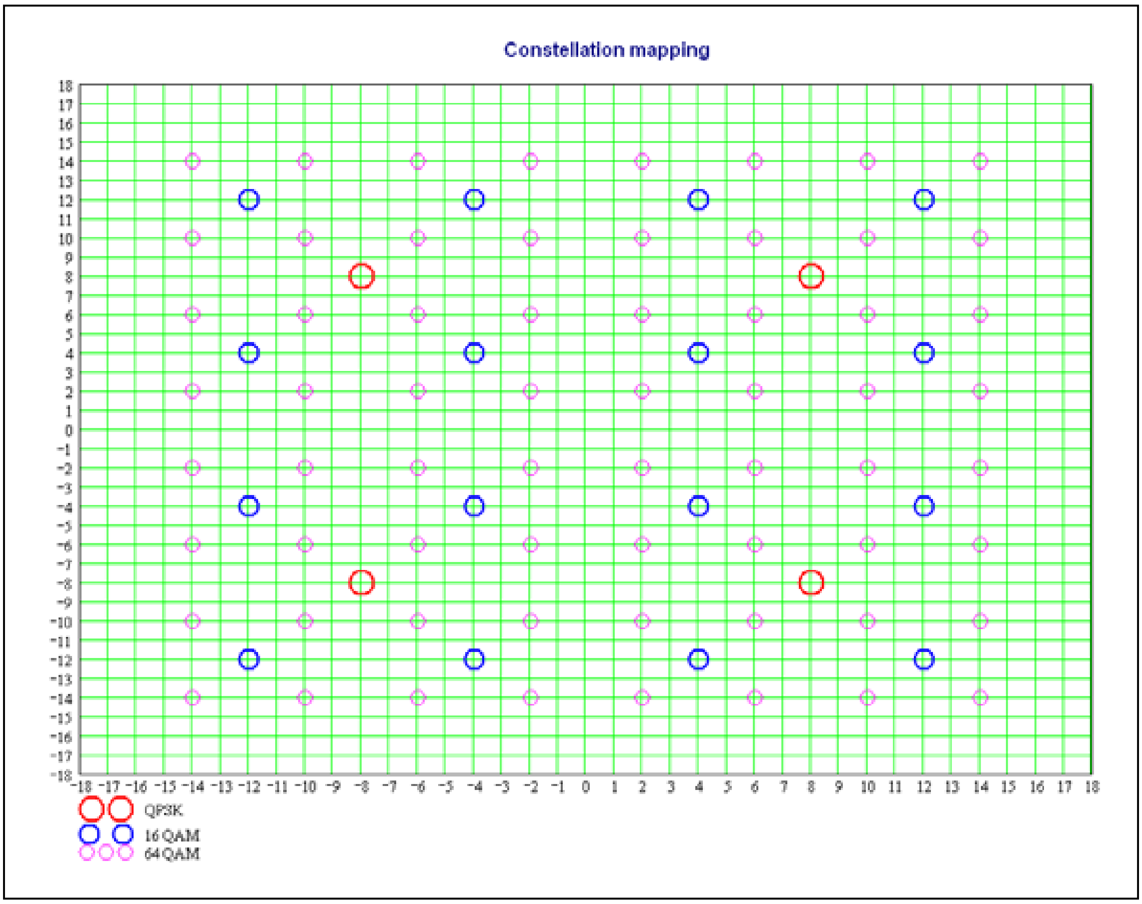

Figure 2-7: QPSK, 16QAM and 64QAM constellations all shown on the same diagram

As an example of a theoretical frequency re-use pattern employing 4 frequencies see Figure 2-8 below. In a real deployment many factors come into play which force ad hoc modification of such a pattern to accommodate reflections from buildings, diffraction by man-made and natural obstacles, and the long or short reach of some base station because of their height.

The key to good frequency planning is to optimise frequency re-use within the network to maximise average SNR. One of the benefits of a real trial system with multiple base stations is the ability to experiment with frequency re-use patterns and see the effect on overall capacity. Some examples are presented in the main report of how different frequency re-use patterns affect the efficient use of spectrum.

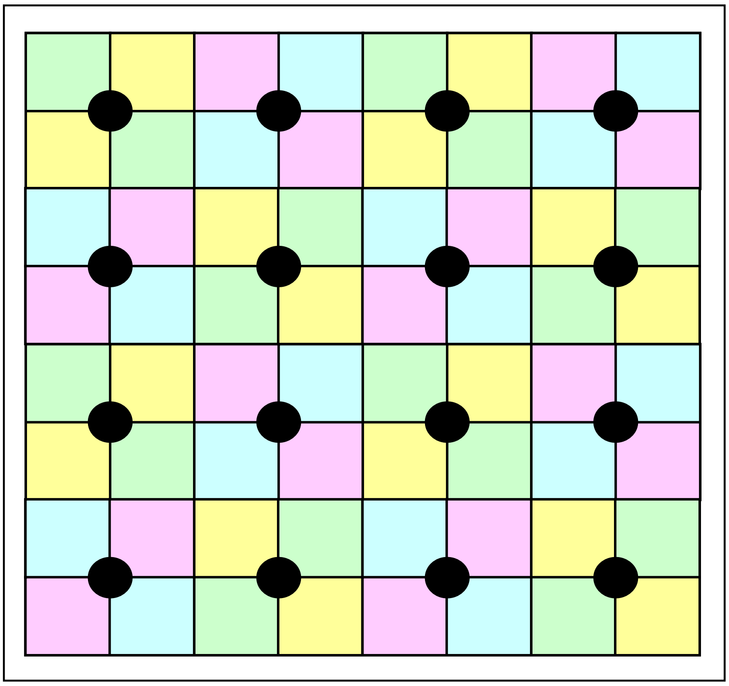

Figure 2-8: A theoretical frequency re-use pattern using four colours

2.2 Wireless Hotspot Equipment

A remarkable amount of recent investment and development effort has created a healthy market in domestic quality IEEE 802.11 wireless local area network (WLAN) equipment. The Wi-Fi Consortium is an industry group dedicated to promoting interoperability between different manufacturer's equipment. The term Wi-Fi has now become synonymous with IEEE 802.11 compliant WLANs.

The new generations of equipment are cheap, allow a high data rate, have good range and are fairly straightforward to connect to up-to-date home PC systems, with a range of hardware connection possibilities. Commercial WLAN "Hotspots" began to appear a few years ago in hotels, airports and business centres. Rural communities realised that Wi-Fi Hotspots could be an effective and cheap way to provide some form of broadband connectivity if IP backhaul could be made available. There are now many examples of self-help communities and small Hotspot businesses in the UK who are providing broadband to rural locations in this way.

Working with IEEE 802.11b, which has only 3 channels at 11 Mbps available, makes it difficult to install multiple levels of wireless backhaul around villages. With only a single leased line connection at some point in a community a number of cross-village WLAN connections must be made to prospective Hotspot locations. All the WLAN links can be subject to unintended interference from users with their own in-house WLAN systems so a high degree of co-operation is required to avoid problems.

Leased line connections to remote rural locations are expensive because they will often involve laying new cable. The tariffs for such lines have an "Out of Area" component which penalises remote areas. A FWA system capable of operating at high bandwidths over long distances, such as the one used in this trial, would appear to provide a cheaper backhaul solution for rural Hotspots. The cost of the relatively expensive, high-performance FWA system can be amortised by using cheap and hopefully self-installable WLAN equipment to provide the last hops to rural homes.

For this reason experimental Hotspot installations were included in this FWA trial to evaluate the following aspects of the technology:

- Co-sited installations of FWA CPEs and WLAN Hotspots
- Interoperability with high-performance FWA equipment
- Reliability, configurability and availability of suitable domestic-quality WLAN equipment
- Suitability for commercial operation
- Connection throughput and latency
- Radio interference problems

2.2.1 Choice and Configuration of Hotspot Equipment

For the Wi-Fi Hotspots, the emphasis is on lowering the cost per end-user, so systems where self-install of the Hotspot clients (though not the Hotspot access point) are of interest. This ruled out the use of specially mounted antennas, coax and power wiring at the client end. A solution might be to use medium gain indoor antennas placed in preferred locations within houses, ideally on windowsills, where reasonable signal strength could be expected without necessarily having a line of sight to the Wi-Fi base stations.

To quickly gauge the effectiveness of this model, some initial measurements were made using an 8dBi antenna mounted on the wall of a property. Several surrounding houses were visited, and good coverage was obtained in some windows of most of the surrounding houses. This was a quick check that something would work, and was carried out using a laptop with an Orinoco Silver PCMCIA card driving a 5dBi indoor range extender antenna. In fact this turned out to be one of the better receivers. The results of the original survey are shown in Figure 2-9.

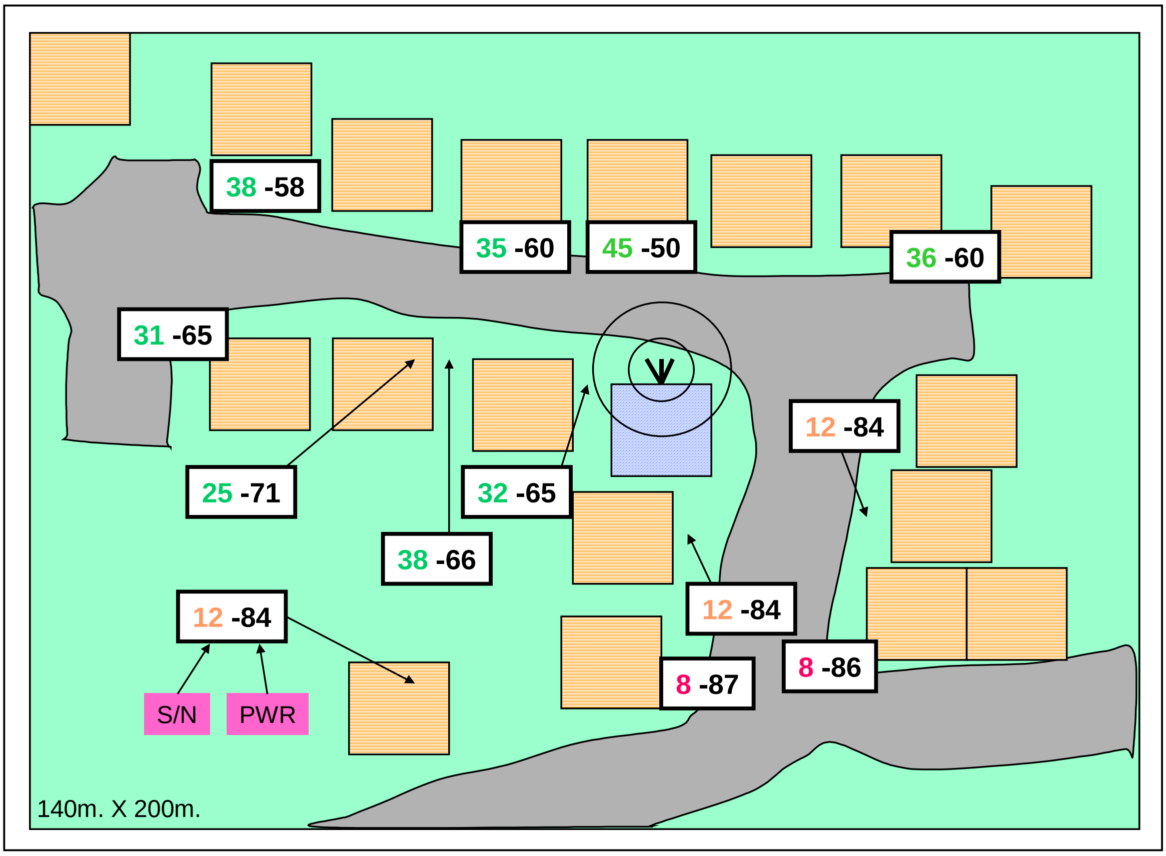

Figure 2-9: A survey of IEEE 802.11b reception in a typical residential street. The S/N figures are coloured green where reception is excellent, orange where it is adequate, and red where it will not work.

Properly mounted antennas were then installed at this location. A 12dBi antenna was mounted above chimney height, and an 8dBi antenna mounted at the side of the chimney below roof height, see Figure 2-10. Cable runs of 7 metres of LMR400 were used to each antenna, coming through into an attic with mains electricity and lighting where the CBL ICU (the CPE indoor unit) was also sited. This allowed experimentation with a range of different radios driving the two WLAN antennas. It is important to minimise antenna cable runs if maximum coverage is to be achieved, so the Hotspot installations should ideally either be made indoors as close as possible to the antennas, perhaps in a roof void, or utilise outdoor weatherproof RF equipment sited close to the antenna. Both of these options will increase the installation costs, which should be borne in mind when considering budgets.

A Dlink DWL-6000AP was used as the standard Wi-Fi Access Point (Wi-Fi AP). It offered reasonable radio performance and good configuration options, especially the ability to allow or disallow the routing of traffic between the wireless clients. The head-end equipment was also configured to isolate users from each other so inter-subscriber traffic still had to pass through the wireless trial firewall at the main router.

At a later stage a Dlink DI-624 router was also used which offered better reporting options, a local firewall, and the opportunity to attach a traffic generator through a wired connection. The status and error reports were configured to be reported by email every few hours. These reports allowed the discovery of a variant of the Nachi worm virus on one subscriber PC, and also gave early indications of poor links shown up by the high frequency of re-associations of the Wi-Fi clients with the router. However, this second router was much harder to configure to prevent inter-client traffic as it was not designed for this purpose.

Figure 2-10: Histon Wi-Fi hotspot co-installed with CBL Outdoor Unit

With hindsight, an indispensable option in the Wi-Fi AP would be a site survey feature that could be interrogated remotely to see what signal strengths or noise characteristics were being observed for each subscriber. A combination of the features found on the Wi-Fi AP and the Wi-Fi Router used here would have been ideal, but are rarely found in the same equipment.

The DWL-6000AP Wi-Fi AP was also used for the other Hotspots as the ability to isolate wireless users from each other and to allocate addresses from the head-end made their setup almost exactly the same as subscribers that attached directly to the CBL CPEs.

It is interesting to note that although the network was reasonably well protected from the Internet by the use of NAT in the main project router, any computers which could also dial up, were used while travelling, or could be prone to Trojans could easily bring viruses within the protected domain. This means it is very important to isolate users from each other. It has been reported that some last-hop operators using IEEE 802.11b have been prone to severe outbreaks of viruses within their systems and have had to insist that all subscribers run a prescribed anti-virus regime on their computers.

The Wi-Fi AP connects directly to the Ethernet port on the CBL ICU using a standard crossover cable. Figure 2-11 shows some of this configuration.

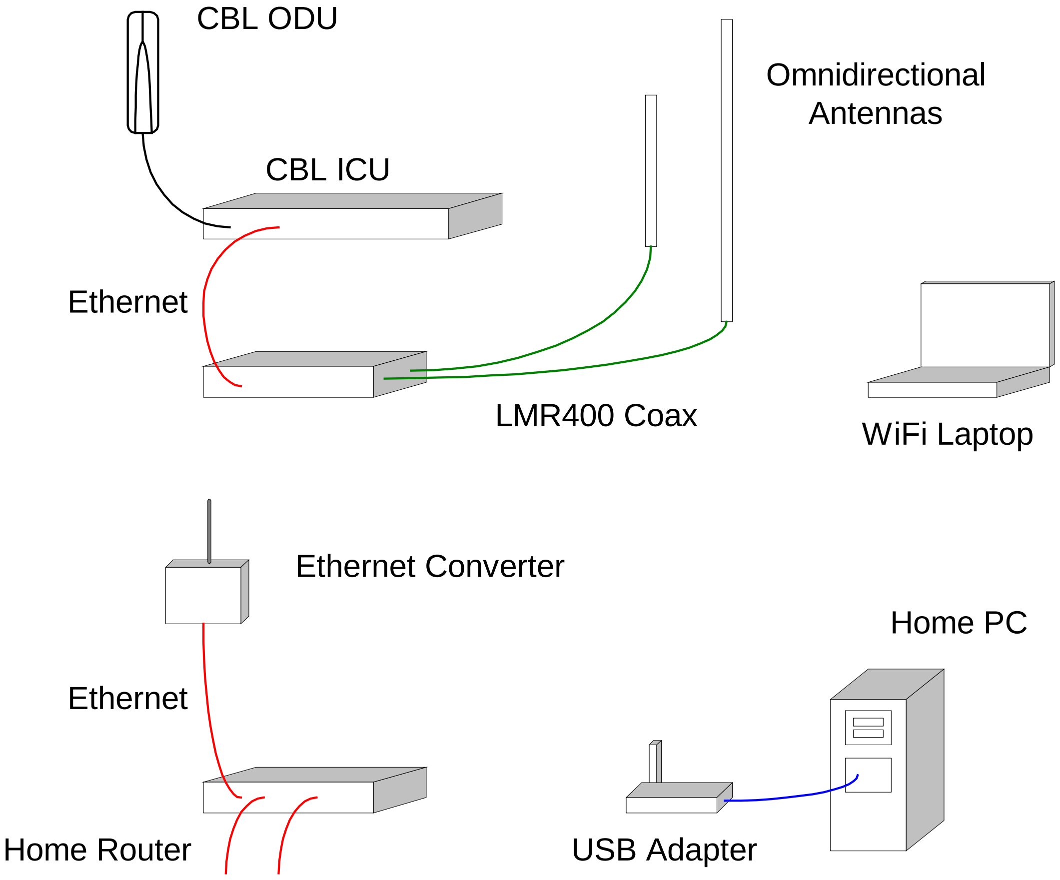

Figure 2-11: Wi-Fi Hotspot Equipment

Ethernet Converters (ECs) make a convenient choice for the Wi-Fi client equipment. These are self-contained units with a single 10/100Mbps Ethernet Port for connection to a subscriber’s PC. They are relatively easy to install since the Ethernet software is already installed on PCs with an Ethernet NIC. There is then just a small amount of configuration to be done to ensure that the IP protocol is enabled, and that the computer is set to obtain its IP, gateway and DNS server addresses via DHCP.

At the beginning of the project there were two ECs on the market, the Netgear ME101 and the Linksys WET11. On testing these it was found that the Netgear ME101 would connect in only a few preferred locations within nearby houses whereas the WET11 would connect in almost all rooms. This allowed us to make installation much more straightforward, as users could mount the WET11 in the room where their PC was kept, and so avoid having to run Ethernet cable around the house or install their own indoor Wi-Fi base station, which for most people, would rule out self-install.

For users whose PCs did not have Ethernet built-in USB-connected units were tried where available. This was problematic as it involved the use of Wi-Fi client drivers in whatever operating system was being used. This is the situation for Microsoft Windows:

- Windows 95 has an edition with USB support, but even after much effort the Linksys WUSB11 could not be made to work with it.
- Windows 98Se has USB support, and the USB device eventually worked reliably, but only after several reboots and blue screens, and the persistence of the subscriber.
- Windows XP Home has better support, with its own Wi-Fi client drivers built-in. The device installed easily, but was prone to locking on to other Wi-Fi networks and failing to return to the preferred network. It was eventually discovered that after settings of the network priorities and various other parameters, stability could finally be achieved.

It is very important to be able to prevent association with any network except the local Hotspot since if the Hotspot goes away and another network is found the Microsoft drivers by default will associate with the new network and not return to the Hotspot even though it has been given greater priority. They only reassess the priorities when a network connection is lost.

After some experience it was evident that the Ethernet Converters could be part of a self-install with the right instructions and a good signal but the WUSB11 would be unsuitable.

A laptop with a PCMCIA card was tried in one home on the Histon Hotspot but the coverage was limited to near the window. An Ethernet converter performed much better, working in all parts of the room. At the house of one subscriber to the Harston Hotspot, perhaps 50m from the 12dBi antenna, the signal coverage was sufficiently good that a laptop could be used with PCMCIA card for connection to the trial, but this was exceptional.

Most of the equipment trialled had Web and telnet interfaces through which it could be remotely configured. The Ethernet Converters even allowed firmware updates over the radio interface, which in a commercial deployment could potentially save many house visits.

The greatest advantage of the Ethernet Converters was that the latest release of the firmware added a site survey feature. They could then be interrogated remotely and the signal quality could be monitored for all the networks that they were seeing. This could be a very useful tool in detecting and characterising other interfering networks as well as diagnosing faults and obtaining early warning if a connection is becoming marginal.

2.2.2 Security and Encryption

Shared Key authentication was experimented with, but this was not supported on all of the project Wi-Fi equipment, and even when it was, it was not always compatible between different manufacturers. Instead, Open System authentication was used. A single WEP key was used for all networks and no unencrypted traffic was passed.

MAC address filtering was available, but for simplicity of administration this was not used. Commercial Hotspots use MAC addresses for access control in conjunction with a shared key (password), but for FWA, this may be too lax as a cloned MAC address could be used indefinitely.

A commercial system would be safer if based on a shared key authentication system with separate keys for each subscriber that could easily be revoked.

2.3 An Overview of the Core Communications Network

The IP switching and routing functionality at the network core is provided by a Cisco 3725 router. This is a highly configurable multiservice access router capable of switching 100,000 IP packets per second. The Cisco IOS software version 12.3 provides all the necessary IP bridging and routing functions, as well as NAT and a DHCP server. The router comes equipped with two fast ethernet ports as standard. To this was added two OC-3 155Mbps ATM modules to provide connections to the Cambridge Broadband PFWA equipment, and two 2Mbps serial interfaces for the Internet connection via PSINet. See Figure 2-12, below, for a schematic of the core of the radio network.

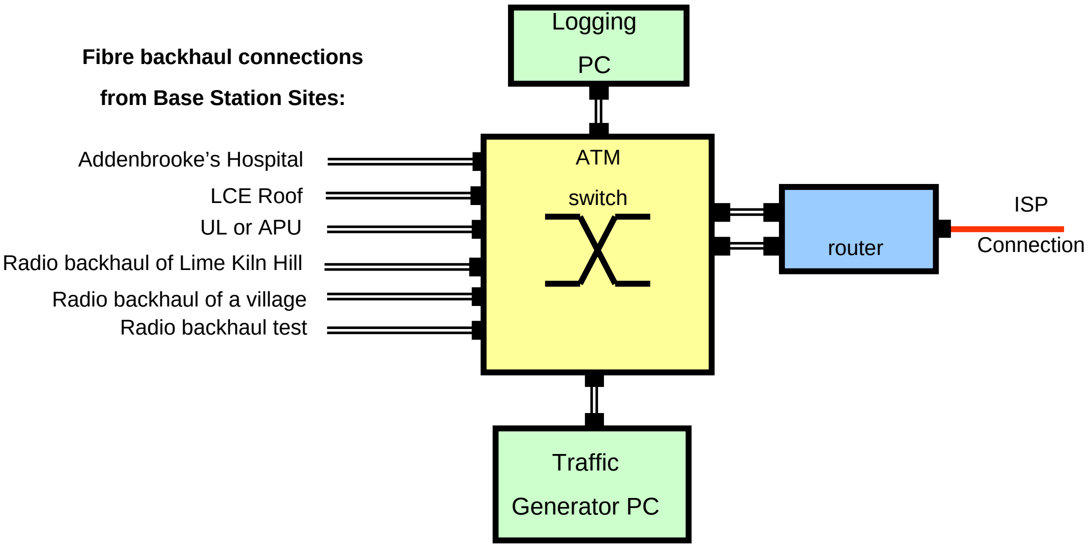

Figure 2-12: Schematic of the network core

2.3.1 IP Network Design

The overall design of the network envisaged provision of IP connectivity and an Internet service to around 50 mainly domestic subscribers, with the addition of up to 25 subscribers connected via Wi-Fi hot spots. All subscribers were offered the possibility of being able to connect essentially any number of their own machines to the Ethernet outlet provided by the wireless CPE.

In order to provide the maximum flexibility in the allocation of IP subnets to subscribers, it was decided to use an address scheme based on RFC1918 private address within the network, and rely on the NAT capabilities of the router to map internal addresses to globally-routable ones. This meant that the allocation of IP addresses internally wasn't constrained by the number of globally-routable addresses allocated to us by the ISP. Initially a block of 32 global addresses was applied for, and this has proven sufficient for the lifetime of the project.

The default allocation of IP address space to subscribers was based on a 29-bit wide subnet mask, giving the subscriber the ability to configure up to 5 independent IP entities for connection to the wireless network. Address allocation to individual hosts was managed by the DHCP server at the router. In most cases this meant that moving subscribers' PCs from an existing broadband connection to the wireless network was simply a matter of changing their ethernet connection and renewing their DHCP lease. If a subscriber needed more addresses, for example when a connection was required for Wi-Fi hot spot purposes, their configuration could be quickly changed at the router, meaning the end stations simply had to obtain a new DHCP address for the move to the new address range to be completed. At the end of the trial, three different subnetwork size allocations were in use, namely at 8, 16, and 32 address widths.

2.3.2 Bridging and Routing

The CBL kit provides a number of possible methods of configuring the connection between the CPE and the core network. The simplest and most flexible method of configuring subscribers is to opt for an Ethernet bridge connection.

In practice this involves creating a RFC 1483 bridge between the Ethernet interface on the ICU and an ATM virtual connection. This is configured through the VectaStar EMS software running on each individual base station. The unique ATM virtual connection specified at this stage is then used to create the link from the base station to the Cisco router. For base stations which are connected directly via fibre to the 8-port MUX, this simply involves making a switch table entry using the software controlling the 8-port MUX (which runs on the logger PC) to map the incoming ATM virtual connection from the base station to the Cisco port connection. For base stations backhauled via radio, each subscriber ATM virtual connection must be mapped through one or more radio links, refer to Figure 2-12.

Each subscriber ATM virtual connection is terminated at a unique ATM subinterface on the Cisco router, which in turn is allocated to a unique bridge group. Each subscriber’s connection can therefore be visualised as a distributed ethernet switch, with one interface at the subscriber’s location, the other at the network core. In order to provide IP routing, use was made of the Cisco IOS integrated bridging and routing capability (IRB), which permits routing of IP datagrams over bridged interfaces. Each of the bridge subinterfaces is allocated to a unique bridged virtual interface (BVI), and each of these is assigned a unique IP address and subnet mask. It is this address / netmask pair which determines the IP subnet for each bridge linking the subscribers’ Ethernet to the central router.

Where there is a requirement to have more than one individual ATM virtual connection connection assigned to the same IP subnet, this is easily achieved through allocating their ATM subinterfaces to the same bridge group. This was how each of the base station and radio backhaul APCs was networked, enabling them to interoperate exactly as if they had all been connected to an Ethernet LAN. This enabled a user of the VectaStar EMS running on any of the base stations to query and configure any of the other APCs, since each appeared on the same LAN.

The net result is that a simple star-topology network is arrived at with a router at its core, with each subscriber connected to a unique IP interface at the router. At the start of the network deployment, there was only one “uplink” or service connection at the core router, namely the Internet link to PSINet. In terms of routing configuration, all that was required was to specify a single static default route pointing towards the Internet connection. All IP routing within the radio network was implicit, and no routing protocols were required.

Subsequently, it was possible to connect directly to the Cambridge University Data Network (CUDN), via the spare fast Ethernet port on the router, and make a direct connection into the Computer Laboratory network. This was immediately beneficial to those of the users who have accounts on various University systems, as they were then able to connect directly to these without having to be routed out of the trial network via the 4Mbps Internet link. They then enjoyed a much faster Ethernet connection. Following discussion with counterparts in the Computer Laboratory, the IP routing for this connection was again handled using static routes.

2.3.3 Discussion

The IP network configuration as described above is both simple and easy to manage. All IP configuration is in one place, on the central router, and can be written out to a simple text file. One major benefit of this approach was the simplification of any fault diagnosis at the IP level.

The Cisco is capable of generating extensive logging information (via syslog), which was captured on the logger PC. Any issues arising from, for example, problems with DHCP allocation, were easily identified and mostly resolved quickly. Additionally, the DHCP server on the Cisco is able to write out its allocation tables to another host. This acts as a cache of the current state, so that if the router is rebooted, this state is not lost. This allowed us to keep a rough tally of how many individual subscribers’ hosts were in use, by simply inspecting the state file.

The flexibility of the network design can be illustrated by reference to some specific examples. One of the subscribers was CBL itself, who initially asked for a connection to be set up to facilitate their office move. They were initially provided with a default subscriber configuration as described above. However, the flexibility of the network design is illustrated by the fact that it was possible to give a CBL employee, at a different location, access to the same bridge (and therefore the same IP subnet) simply by mapping the ATM virtual connection from his ICU directly into the bridge configured in the ICU at CBL's premises. It can easily be seen how widely distributed subscribers could be brought into what would effectively be a VPN thorough this type of arrangement.

Another example derived from setting up the direct connection to the Computer Laboratory (CL) network. Initially the connection was set up in the default manner, with RFC1918 IP addresses allocated to the subscriber via DHCP. However, the user, a member of the CL staff, required access to restricted hosts within the CL network, which would have been disallowed from the DHCP addresses. This was overcome by enabling the user to use IP addresses allocated from the Computer Laboratory's address space on his radio broadband Ethernet connection, and using Cisco's Policy-based Routing (PBR) to ensure all traffic from these hosts was routed via CL, irrespective of destination IP address, thereby avoiding asymmetric routing. This was accomplished alongside the default subscriber configuration, so that two different IP subnet ranges were in use on the bridged Ethernet at the subscriber's location. This illustrates the flexibility of the bridging solution for subscriber connections: the radio part of the network is completely transparent to IP, which is handled exclusively at the router.

2.3.4 PC servers at the network Core

Two Dell PowerEdge 400SC PC servers each with a 2.4GHz Pentium IV processor and two 120GB hard disks were deployed at the network core for management purposes. These were each equipped with a 155Mbps OC-3 ATM card.

The first, called "logger", was primarily intended for use as the main repository for the logging information derived (via syslog) from the components of the radio network. Additionally, it served as syslog host for the Cisco router, and general purpose host for management and monitoring purposes. It also ran the software which controlled the 8-port MUX at the network core.

The second PC, "traffic", was used primarily as a traffic generation endpoint. Its ATM interface allowed us to set up specific ATM VCs to any location on the radio network, and to perform traffic throughput measurements across these point-to-point links. Additionally, its disk capacity was used to keep backup copies of the logging information collected on logger.

As well as their ATM interfaces, which were connected directly to the core 8-port MUX, the two PCs were connected to a fast Ethernet switch on a management LAN to which the core Cisco was also attached. The advantage of this arrangement was that if ever there was a problem with the ATM part of the network, it was always possible to log in to the main management PC (logger) via the Internet, and be able to check logs and test connectivity for diagnostic purposes.

Both PCs ran RedHat 9 Linux as their operating system. This was chosen principally because of previous experience with the linux-ATM software application, required to interface with the ATM hardware. Linux in general provides the platform for all the essential tools for monitoring and management that were required (syslog, ssh, perl, python, bb, mrtg, rrdtool, tftp server, http server).

For security reasons, access control measures were put in place on the APCs, essentially restricting access to hosts on the management LAN described above.

2.4 Volunteer Users

2.4.1 The Requirement for Volunteer Users

The design of the trial system included the idea of placing CPE's on real user's premises and giving these volunteer users access to the Internet via their FWA connection. There are several benefits from this strategy. Firstly, it would not have been possible to find so many sites for CPEs if the carrot of a fast broadband connection during the period of the trial had not been available. Secondly, the users generate a certain level of real traffic in the system, and it is interesting to see how people make use of what is a relatively fast domestic broadband connection. Finally, the presence of the users adds an urgency to fixing problems in the network and has a positive effect on limiting the downtimes.

2.4.2 Attracting and Managing Volunteer Users

It was noted that the number users required was quite large and the number of potential applicants for trial users would possibly run into hundreds. Some form of database approach was needed to deal with all the required information. It was decided to set up a Web database using the php and mysql database tools available from the host of cotares.com. This allowed the creation of an online application form for volunteer users to fill in. Apart from questions about their contact details they were asked specific questions in relation to the suitability of their premises, and of their computer systems, for use in this trial. It was necessary to know whether they had the authority to permit the mounting of antennas on their buildings. At the time of application applicants agreed to a set of Terms and Conditions. The terms included sections on the return of the equipment at the end of the trial, it included an obligation not to use the trial Internet connection for illegal activities such as breaching copyright, and it asked users not to make public any passwords or encryption keys. The latter condition was because Wi-Fi Hotspots were not intended to be freely available to non trial users.

Applicants to take part in the trial were placed in a succession of categories as they moved towards eventual installation and connection to the Internet. Many applicants were rejected at an early stage because they were too far out of range of the Cambridge-based trial. The remainder progressed through the following set of categories:

New Applicants An email acknowledging their application was sent at this stage with an indication of when things might progress.

Measurement Required If their location and legal circumstances looked acceptable they were placed in a queue to have the signal strength and quality at their site measured by the radio measurement vehicle. This, of course, required their likely base station to have been installed and working. It turned out that rather few applicants were rejected at this stage because most of the unlikely signal paths had been rejected at the New Applicants stage.

Installation Required It was quicker to make measurements than to make system installations so a queue of site installations was formed. It was found that it was usually not practical to make more than three installations per day, although four were done on some occasions. Of course on other occasions, when it was not possible to arrange convenient times with users only one or two installations could be done.

Registered At this point the user has a powered-up CPE on their house which is registered with a base station. They are then in a queue to be connected to the Internet.

Networked Once the system administrator has added appropriately configured ATM virtual circuits to the system and also made entries in the core IP router they can be classified as networked.

The final stage was to ensure that their PC was configured to DHCP an address from the project system and connect to the Internet. Many volunteer applicants were very used to configuring their home PCs and had little trouble with this. But with so many users a number of home PCs required considerable coercion. A small number of users were still using Windows 95, which does not have very flexible networking capabilities. Many people had Windows 98 which did not present many real problems. However some PCs were old enough that working device drivers for the Ethernet ports they actually had could not be found and device drivers for the new Ethernet cards supplied by the project did not work on the old PC hardware. This wasted a lot of time and designers of future trials are recommended not to get involved with such out-of-date hardware.

Once the users were up and running the online database was invaluable for looking up details of their installations. The database contained 53 fields for each user and included many aspects of their CPE equipment and networking setup. The database could be downloaded from the Web site in comma separated value (csv) format. This file was used to drive the equipment monitoring software since it contained the all the necessary details relating to the CPE and how to connect to it. The advantages of having one central up-to-date set of records which everyone in the trial Consortium could easily access are obvious.

2.4.3 Installations at volunteer subscribers' premises

The project used a team of highly experienced installers for placing CPEs on private residences. This made it easy to plan and execute the many and varied installations. The units were usually sited on chimney stacks, but could also be put on masts bracketed to the side of the house. The physical installation involves bringing down a fibre and power line from the outside unit to a small indoor unit which provides an Ethernet socket to the resident. Cables were run discreetly and the point of entry into the house was always carefully discussed with the volunteer subscriber. When a potential installation site was being checked for acceptable signal strength an installation plan was made so that the installers would be able to arrive with the correct bracketing and ladders. The installation details were always carefully checked with the volunteer subscriber before commencement. Details of importance included: method of attaching the bracketing to the chimney or gable end; cable runs; point of entry for cables into the loft or a room; and the situation of the indoor unit. A typical installation took about two hours, some were slightly quicker than this, and some took up to twice as long.

Figure 2-13, below, shows the installation elements. The subscriber units weighed about 7 kg and present a low wind load. The fibre and twisted pair low voltage power cable were both of outdoor quality. The Indoor Control Unit (ICU) is a 19 inch rack-mount unit which could be conveniently place on top of a typical PC box or, for example, underneath a printer. The ICU was powered by a small 48V power supply unit, and volunteer subscribers were asked to leave the equipment powered at all times so that they could be monitored continuously.

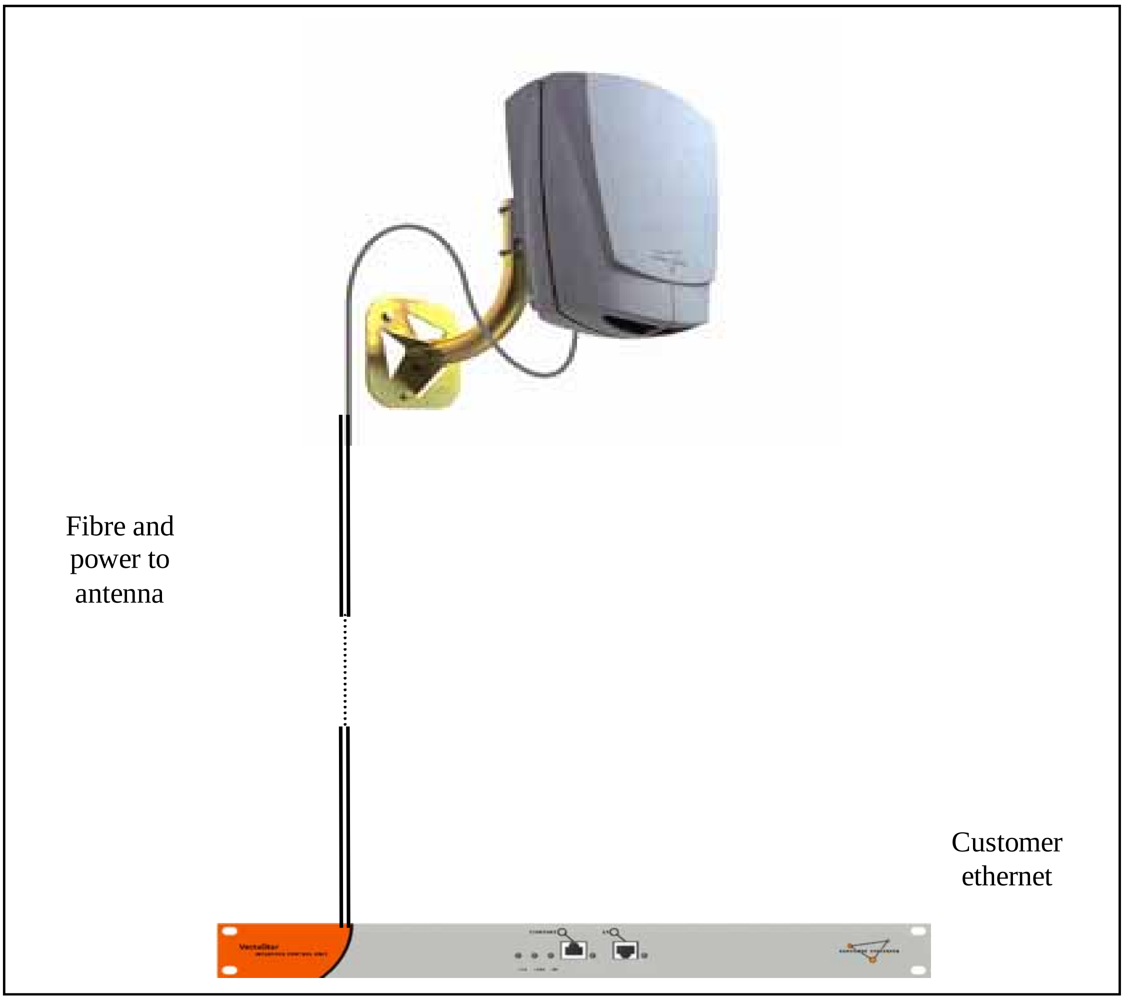

Figure 2-13: Installation at a volunteer subscriber's premises

2.4.4 A Survey of Volunteer Users

2.4.4.1 The design of the questionnaire

At the end of the radio trial Volunteer Users were asked to complete a simple survey form in order to get some feedback on what they had thought of the trial and whether, on the basis of their experience, they would actually use a radio broadband service if they had the option. The survey form is shown on the next page in Figure 2-14.

The Questionnaire asked responders to rate the Reliability, Service Quality and Usefulness of the service. Users were also asked what their opinion was on how much a radio broadband service might cost, both for “domestic quality” (contention ratio of 50:1) and “business quality” (contention ratio of 20:1) connections. Finally, they were asked if they would prefer a wireless or wired broadband service if they had the choice.

Cambridge Public Fixed Wireless Access Trial

User Feedback Questionnaire

Please complete the questionnaire below and return to us by Feb 18th. Replies can emailed to radio@cotares.com Please note that your reply will not be published or passed on and will only be used for compiling user feedback statistics.

Our Questions	Your Answers				
Your name:					
Please put an X in the relevant box					
	very bad	bad	fair	good	Very good
Reliability:					
Service quality:					
Usefulness:					
This is only a trial which will not result in a commercial service, but, if it were, how much do you think it would cost? How much would it cost per month for each of these domestic services, assuming a contention ratio* of 50?					
Please put an X in the relevant box					
A 4 Mbps service (which is what you had):	£0	£40	£80	£130	£200
A 1 Mbps service:	£0	£25	£35	£60	£100
A more typical 0.5 Mbps service:	£0	£15	£25	£40	£60
How much would it cost per month for each of these business services, assuming a contention ratio* of 20?					
Please put an X in the relevant box					
A 4 Mbps service (which is what you had):	£0	£80	£150	£250	£400
A 1 Mbps service:	£0	£60	£90	£150	£200
A more typical 0.5 Mbps service:	£0	£40	£60	£80	£120
If you had the choice of similar price and performance wired or wireless broadband which one would you go for?	wired	wireless	none		

*All ADSL lines and wireless access systems – regardless of service provider - are subject to contention, meaning that the network bandwidth available is shared between a number of subscribers. The ratio to which the available bandwidth is shared between users is called the "contention ratio". All UK providers, including BT, are subject to the same contention ratios. Services targeted at casual home users are frequently cheaper, but have a higher contention ratio. This means that the available network capacity is shared between a greater number of users. Business services have a lower contention ratio, which will provide a more consistent level of performance. A 50 to 1 contention ration means the bandwidth may be shared with up to 50 other subscribers, whereas a 20 to 1 contention ration means the bandwidth may be shared with no more than 20 other subscribers

Figure 2-14: The Radio Trial Questionnaire

	very bad	bad	fair	good	very good	other
Reliability	2.8	8.3	22.2	44.4	22.2	0.0
Service quality	0.0	0.0	13.9	41.7	44.4	0.0
Usefulness	0.0	0.0	5.9	14.7	79.4	0.0

Figure 2-15: Performance responses (%) for FWA volunteer users

	very bad	bad	fair	good	very good	other
Reliability	12.5	12.5	37.5	25.0	12.5	0.0
Service quality	0.0	0.0	12.5	62.5	25.0	0.0
Usefulness	0.0	0.0	12.5	37.5	50.0	0.0

Figure 2-16: Performance responses (%) for Wi-Fi volunteer users

	very bad	bad	fair	good	very good	other
Reliability	4.5	9.1	25.0	40.9	20.5	0.0
Service quality	0.0	0.0	13.6	45.5	40.9	0.0
Usefulness	0.0	0.0	7.1	19.0	73.8	0.0

Figure 2-17: Performance responses (%) for combined volunteer users

4 Mbps service cost	**£0**	**£40**	**£80**	**£130**	**£200**	**other**
	0.0	50.0	30.6	11.1	8.3	0.0
1 Mbps service cost	**£0**	**£25**	**£35**	**£60**	**£100**	**other**
	0.0	50.0	38.9	8.3	2.8	0.0
0.5 Mbps service cost	**£0**	**£15**	**£25**	**£40**	**£60**	**other**
	8.3	52.8	36.1	2.8	0.0	0.0

Figure 2-18: Cost responses (%) for FWA volunteer users – domestic quality service

4 Mbps service cost	**£0**	**£80**	**£150**	**£250**	**£400**	**other**
	5.6	47.2	25.0	11.1	8.3	2.8
1 Mbps service cost	**£0**	**£60**	**£90**	**£150**	**£200**	**other**
	0.0	75.0	16.7	5.6	0.0	2.8
0.5 Mbps service cost	**£0**	**£40**	**£60**	**£80**	**£120**	**other**
	2.8	86.1	8.3	0.0	0.0	2.8

Figure 2-19: Cost responses (%) for FWA volunteer users – business quality service

4 Mbps service cost	**£0**	**£40**	**£80**	**£130**	**£200**	**other**
	16.7	83.3	0.0	0.0	0.0	0.0
1 Mbps service cost	**£0**	**£25**	**£35**	**£60**	**£100**	**other**
	0.0	100.0	0.0	0.0	0.0	0.0
0.5 Mbps service cost	**£0**	**£15**	**£25**	**£40**	**£60**	**other**
	0.0	100.0	0.0	0.0	0.0	0.0

Figure 2-20: Cost responses (%) for Wi-Fi volunteer users – domestic quality service

4 Mbps service cost	**£0**	**£80**	**£150**	**£250**	**£400**	**other**
	16.7	66.7	16.7	0.0	0.0	0.0
1 Mbps service cost	**£0**	**£60**	**£90**	**£150**	**£200**	**other**
	0.0	100.0	0.0	0.0	0.0	0.0
0.5 Mbps service cost	**£0**	**£40**	**£60**	**£80**	**£120**	**other**
	0.0	100.0	0.0	0.0	0.0	0.0

Figure 2-21: Cost responses (%) for Wi-Fi volunteer users – business quality service

4 Mbps service cost	**£0**	**£40**	**£80**	**£130**	**£200**	**other**
	2.4	54.8	26.2	9.5	7.1	0.0
1 Mbps service cost	**£0**	**£25**	**£35**	**£60**	**£100**	**other**
	0.0	58.1	32.6	7.0	2.3	0.0
0.5 Mbps service cost	**£0**	**£15**	**£25**	**£40**	**£60**	**other**
	7.0	60.5	30.2	2.3	0.0	0.0

Figure 2-22: Cost responses (%) for combined volunteer users – domestic quality service

4 Mbps service cost	**£0**	**£80**	**£150**	**£250**	**£400**	**other**
	7.1	50.0	23.8	9.5	7.1	2.4
1 Mbps service cost	**£0**	**£60**	**£90**	**£150**	**£200**	**other**
	0.0	78.6	14.3	4.8	0.0	2.4
0.5 Mbps service cost	**£0**	**£40**	**£60**	**£80**	**£120**	**other**
	2.4	88.1	7.1	0.0	0.0	2.4

Figure 2-23: Cost responses (%) for combined volunteer users – business quality service

Preference	**wired**	**wireless**	**none**	**both**	**no response**
	50.0	41.7	0.0	5.6	2.8

Figure 2-24: Preferences (%) expressed by FWA volunteer users

Preference	**wired**	**wireless**	**none**	**both**	**no response**
	37.5	62.5	0.0	0.0	0.0

Figure 2-25: Preferences (%) expressed by Wi-Fi volunteer users

Preference	**wired**	**wireless**	**none**	**both**	**no response**
	47.7	45.5	0.0	4.5	2.3

Figure 2-26: Preferences (%) expressed by combined volunteer users

2.4.4.2 Analysis of the questionnaire responses

The responses can be broken down into those for FWA volunteer users, Hotspot users, and combined wireless users.

In Figure 2-15, Figure 2-16 and Figure 2-17 we see the Performance responses for FWA users, Wi-Fi users and Combined users respectively. In terms of reliability the FWA volunteers scored “good” more often than any others, while the Wi-Fi users scored “fair” more often. This is perhaps not surprising, the Wi-Fi connections took some time to settle down and with some Wi-Fi APs persistently crashing there were periods of poor service. Some FWA users also experienced very poor service because of either interference problems or poor signal reception.

The Service Quality category also scores higher for FWA volunteers, and no-one thought it was “bad” in either case. FWA users found the service more Useful than the Wi-Fi users, perhaps because it was more reliably available. Most people found the service Usefulness to be “very good”. Certainly when a 4 Mbps Internet connection is working well it can be very useful. Overall the responses are “good” to “very good” in each category.

The cost responses are a little difficult to draw conclusions from because it became clear that some responders interpreted the questions as “how much would you pay” rather than “how much do you think you would be charged”. The former question was not asked in the survey as it was feared that most responses would be £0.

In Figure 2-18 it can be seen that the highest number of responses are for the values that are rather lower than might be paid for a cable or ADSL service of domestic quality, although a few people were under the impression that it might cost a lot more. For a business quality connections the results shown in Figure 2-19 have the most responses for rather low rental values compared to traditional cable prices, it is possible that the volunteer users do not have any experience in this category, or that they do not understand the advantages of a lower contention ratio service.

Looking at the responses on monthly rental for Wi-Fi customers in Figure 2-20 and Figure 2-21 they are firmly of the opinion that it should be delivered at quite low prices compared to current cable and ADSL monthly rental costs. The combined cost figures shown in Figure 2-22 and Figure 2-23 again suggest that people are hoping that radio broadband prices will be lower than current wired prices.

The responses concerning the preference for wired or wireless broadband connections were surprising, see Figure 2-24, Figure 2-25 and Figure 2-26. The Wi-Fi users showed a very high preference for a wireless service. Although their service quality had not been good they could evidently see that if these problems could be removed the Hotspot approach has a very low impact on their home and property. It requires no invasive cable laying in the garden or digging up of the pavements, and no holes drilled through walls. Nearly 42% of the FWA users also expressed a preference for wireless services though the FWA systems used were certainly a lot more invasive than the Wi-Fi endpoints.

3 Cost Models

A range of different business models can be envisaged based on fixed wireless access. In cities and larger towns there may be a sufficient number of profitable customers to warrant a rollout. Customers here would include: office blocks; factories; schools and colleges and perhaps residential blocks. Of all the possible business models the most topical one would be provision of broadband in rural locations.

Some cost models are examined here for fixed wireless broadband rollouts. They are not intended to be viewed as business plans, the intention is merely to compare the relative sizes of various setup and running expenses. The analysis suggests that using a high specification FWA chiefly as a backhaul system it is possible to deliver broadband connectivity to rural locations cheaper than with current leased-line alternatives.

3.1 Rural Wireless Broadband

In this scenario the margins are low and a self-install or simple-install solution at the customer premises is necessary. The high-performance FWA system used in this trial is currently too expensive to serve individual rural domestic customers. A more appropriate use for it would be to backhaul a last hop technology. Current potential last hop technologies are Wi-Fi (IEEE 802.11a, b and g) and perhaps mini-DSLAMs. Other last hop technologies may appear in future, perhaps from the developing IEEE 802.16a fixed wireless access standard, see Section 4.6.

The first business model analysed is the backhaul of last hop systems in rural villages and small towns. This commercial separation of backhauling last hop systems from that of the business of last hop systems themselves might be advantageous since there are already a large number of small last hop operators. A high performance radio backhaul system based on fixed wireless could significantly undercut the current connection costs, which are based on leased lines, of these operators.

Providing a service covering large rural areas, which have a relatively low population, places the emphasis on being able control the costs of backhaul of the large-scale high performance system. An interesting possibility, used in the current project, is to backhaul FWA base stations in chains. Each base station would serve all the communities within range and would be backhauled by the previous base station in the chain. Line of sight connections can be made over several tens of km in the bands around 3.5 GHz. By adding higher gain antennas where appropriate high order modulation levels can be used widely in the network, particularly in the radio backhaul chain, giving optimal data throughput.

Figure 3-1 shows a regional hub with one backhaul chain of base stations extending into the countryside. In practice a number of backhaul chains would converge on regional Points of Presence (POPs). A POP in a city or large town would connect to a conventional 4 sector base station. The FDD system used in the trial requires CPE-type units to connect with AP-type units, so along the length of the backhaul chain CPE and AP-types would be alternated. In Figure 3-1 the village labelled A has a repeater station on it consisting of two CPE-type units connected back-to-back, just as was done on the Lime Kiln Hill base station in this trial. All ATM virtual circuits can be run continuously from end to end by this arrangement. The repeater A then backhauls a 4 sector base station, labelled B. Here two of the AP units are in the backhaul chain while the other two are providing connections to CPEs in surrounding villages which could be between 5 and 30 km away. Base station C continues the chain deeper into the countryside, it has two back-to-back CPEs in the backhaul chain and again a pair of APs provide connections to local villages.

The data bandwidth budgets in such a chain of base stations need consideration. The data bottleneck is most acute on the last jump back to the POP. Working within the UK allocation of 2 x 20 MHz of FWA bandwidth it should be possible to operate with one pair of 10 MHz bands operating for the uplink/downlink in the backhaul chain. Another pair of 10 MHz bands could then be used for distribution to either side of the backhaul chain. Assuming that there are good lines of sight and using high-gain antennas for the backhaul path the data bandwidth available upstream and downstream will be around 40 Mbps. This is more than enough for the POP bandwidth that would be required. A refinement to this chain topology would be to bifurcate the chain in the outer parts of the network and use 5 MHz bands for links on the very edges of the network. This would allow the network to spread out evenly towards the perimeter of coverage from the POP.

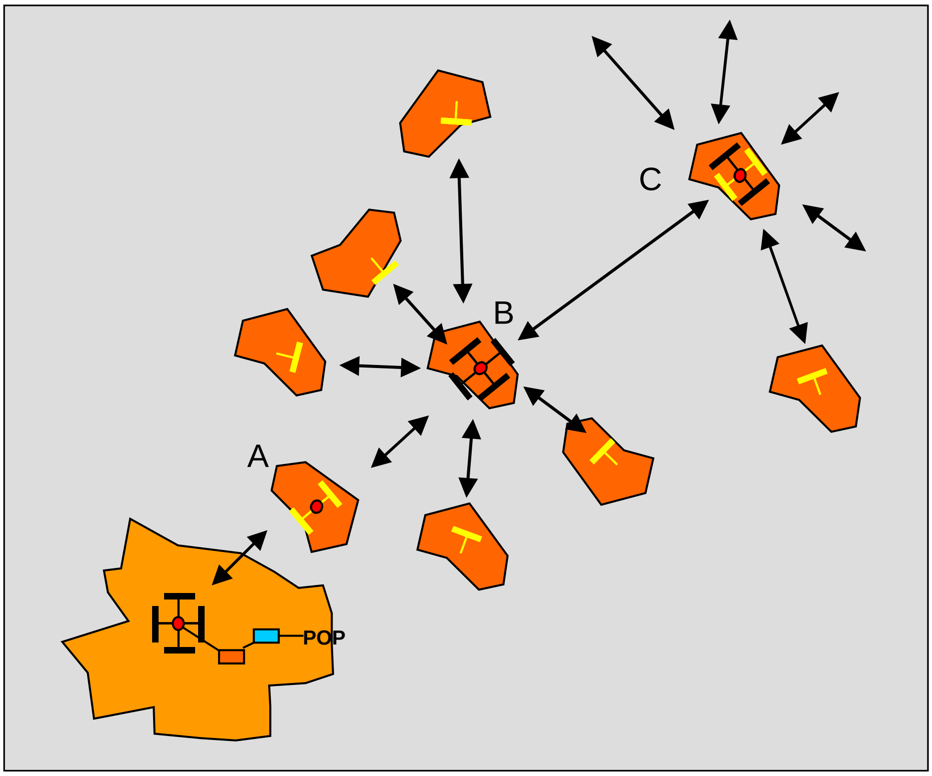

Figure 3-1: Rural base stations forming a backhaul chain and serving local distribution last hop systems in communities

3.2 Scaling and Balancing the Network

How many communities could be served with a chain of base stations, and how long would the chain become? A contention ratio of 50:1 on connections which average 0.5 Mbps will be assumed, which matches typical ADSL and cable offerings. How many customers are needed in a locality to justify the expense of installing a CPE system to provide local backhaul, and what customer take-up can be assumed? Wi-Fi Hotspot operators report that there is a reasonable expectation of take up of 5% in rural locations. Assuming each CPE gives service to communities of around 1000 households this would give the last hop operator about 50 customers, and it will seen in the Cost Model that this would allow the FWA backhaul operator to charge a profitable rental for the local backhaul. So assuming delivery of on average 0.5 Mbps to a community of 50 users with an average 50:1 contention across the network then 68 such communities could be served from a 34 Mbps POP. It would be reasonable to assume that each rural base station served 6 communities. This would mean that each chain would only have two rural base stations in it. The maximum throughput required in the link to the hub base station would average 8.5 Mbps, which is well within the capability of the system used in this trial.

A more cost-effective connectivity solution would place more base stations in each chain, and if a 155 Mbps POP was used it would be possible to operate with an average throughput of the full 40 Mbps along each chain. Now 80 communities could be served in each base station chain and it might be expected that the chain would be between 10 and 15 base stations long. In practice the chains would be bifurcated further out from the hub in order to cover the area more evenly. In total 320 communities containing a total of around 16000 last hop customers could be serviced from a single hub base station at a POP. This idea presents the opportunity to leap-frog leased-line connections over distances of hundreds of km using a long reach FWA system configured for backhaul.

With so many wireless jumps being proposed in this network latency issues should be considered. The equipment used in this trial added about 1.25 ms of one-way latency for each wireless jump. Over an average 12 jumps this would amount to 15 ms. This is relatively small compared to typical Internet latencies. More importantly it is has a low variability. It is large fluctuations in latency which cause problems with the Internet Protocol. These issues are examined in more detail in Section 12 of the main report.

3.3 Using a Self-Reconfiguring Mesh for Reliability

It is possible to set the link budgets on individual wireless hops to satisfy a particular reliability constraint, eg 99.999% uptime over a typical annual cycle of weather and foliage related path loss causes. Connecting 10 or 15 of these links serially would make the constraints much more restrictive. An interesting solution to this would be to bring both ends of the base station chains back to a hub base station at a POP. The chain could either take a sweep around the countryside and return to the same hub base station, or it could link right through to another hub system.

The chains of base stations would then form a mesh which could be rapidly re-configured if there was an outage on any single backhaul link in a chain. An automatic software system which could achieve this would be constantly monitoring all the connectivity paths in a chain of base stations from both ends. When a broken link was detected groups of pre-allocated ATM virtual connections running down the chain could be connected to CPEs which had become disconnected. The backhaul would then be made in the other direction in the chain until either radio propagation conditions changed or a repair could be made. The additional reliability brought about by the mesh topology would mean that 20 or 30 base stations could be backhauled by both ends of the chain. During an outage the system might become temporarily over-booked, but it would continue to work. High priority customers need not be affected.

If the service outage was caused by equipment failure on a link then it could be expected that the other links would still be operating to specification. On the other hand if weather conditions had caused a deep fade on a link it might also affect neighbouring backhaul links. Clearly some research is required to find how wide-ranging weather-related fades might be for 3.5 GHz wireless systems, if they extend over multiple 10 km long links then outages will occur in some parts of the chains.

Only snow caused a serious outage of the backhaul links in the trial system, see Section 11.1.2 of the main report. How this actually occurred is not known. It is possible that snow packed itself into the air ventilation gap behind the front cowling of the antenna and that an improved design could mitigate the problem. The FWA system software used in the trial was not configured to automatically change modulation level according to propagation conditions. This meant that the connection was broken when the error rate became too high for the demanded modulation level. It would be a requirement, especially in the backhaul chain, to adjust the modulation level as necessary to keep some sort of continuous connectivity in adverse conditions.

3.4 Advantages of Wireless Data Bandwidth

A high-specification data delivery system as outlined above could compete head-on with cabled systems, and is capable of outperforming them in a number of ways.

The perception gained in the trial system was that delivering higher peak-rate services with higher contention ratios would be very attractive to the average customer. This is because small blocks of data like Web pages or emails are likely to be delivered at the highest possible rate. The contention ratio of 50:1 on 0.5 Mbps services is a constraint arising from the difficulties of reliably delivering higher bandwidths to all customers over ADSL. There is no such difficulty in this type of radio network, the high-level network links could all peak at 40 Mbps bi-directionally, and a last hop system based on, e.g. IEEE 802.11a, should be able to deliver at least several Mbps on the final link, and up to more than 30 Mbps.

Secondly FWA systems can be operated to give symmetric service. As the use of multimedia continues to expand in home computer systems more people will be operating as sources of large amounts of data as well as sinks. They will be increasingly attracted to broadband systems which allow them to upload pictures, music and video to friends and family at rates much faster than the current asymmetric cable offerings of 128 Kbps or 256 Kbps.

3.5 Cost Elements for Rural Wireless Broadband

Before elaborating a cost model the various elements are examined. This will begin to show where the bulk of expenses really lie. In terms of the hardware, software and running costs of the wireless infrastructure vendors pricing models may vary. Typically a complete base station setup will be offered as a unit cost. The licensing arrangement may mean that having a higher number of CPEs attached to a base station may attract higher license costs. There will be arrangements for extended warranty after year one as well as software updates and support.

A model based roughly on the costs of CBL wireless kit used in this trial will give some representative figures for comparison.

3.5.1 Capital Expenditure on Wireless Equipment

Fixed Wireless Access is a young industry and few systems have reached optimum design and production levels. The result is that prices are continually falling for wireless broadband equipment as engineering improvements drive down the production costs and production volumes rise. By the end of this project, in March 2004, it should be possible to buy 4 sector FWA base stations of the type used in the radio trial for around $40K. A base station unit which was part of a backhaul chain and which had two APs for connection to communities on either side of the chain might cost around $45K.

The cost of CPE equipment has continued to fall, the cheapest type may now cost less than $1000, which is a convenient all-in-one unit with a direct Ethernet output from the antenna.

A useful network building block which was not used in this trial is a wireless ATM repeater station. This consists of two AP-type radio units or two CPE-type radio units connected by an ATM link. All ATM virtual circuits and their Qualities of Service would be preserved across such a connection. This could prove to be expensive if it was necessary to build it in the same way as a base station, which requires a PC controller and an ATM MUX. However, in this simple application it will soon be possible to simply plug two radio units back-to-back with fibre and rely on internal control of the connections. Such a two-way relay might cost about $10K.

3.5.2 Software Licences

In the system under consideration the base station license fee varies according to the maximum number of CPEs permitted to register with the base station. In the backhaul infrastructure model under consideration here no more than 10 registered units would be required.

3.5.3 Extended Warranty, updates, support

An operator building a large business on a particular vendor’s equipment would want to negotiate terms on extended warranty, software updates and general support. Typically there will be an initial 12 month warranty and in the case of CBL this can be extended after year one. Similarly, software updates and bug fixes are chargable at a rate based on the original base station license fee per annum, and the support desk costs a percentage of the list price of the network.

3.5.4 Sites and towers

This is often the major expense of a wireless networking system. It is somewhat mitigated for the new generations of FWA equipment because they are small and light. This means they both take up less vertical height on a tower and require must less substantial structures to support them. Most of the trial base station installations were easily achieved using existing metal handrails on top of buildings and towers. However, the annual rental costs on prime sites are increasing, and the costs of siting and installing a new mast can be high. See Table 3-1 for the example pricing elements used in the model.

3.5.5 Installation

Apart from tower rental or mast construction there are a number of other installation costs. The actual radio unit installation and cabling can often be done in one day by a small team, but there will be a requirement for a metered electricity supply and also an environmental enclosure. Such an enclosure will require both thermostatically controlled heating and cooling to prevent condensation, the effects of low temperatures and overheating.

Installation of CPEs in this Cost Model would be interesting because of the diversity of possible locations. They could be mounted at new sites with wooden or light metal poles; on prominent buildings in communities such as churches, halls or even tall houses; or on existing masts.

3.5.6 Maintenance, System Administration and other Staffing

In the scenario of 3.2, above, there would be over 600 radio units connected to the local hub. It is certain that some level of hardware failure either of radio units, fibre connections or power supply systems will be encountered. There will also be software maintenance issues, particularly of the 80 or so linux PCs which control all the base stations. For a large regional or national rollout there will need to be a team developing software systems for the entire infrastructure, as well as local offices for maintenance technicians. A national team would also be required to market bandwidth, and deal with revenue collection. However, this particular model does not supply end-users, it provides local backhaul to last hop operators, so marketing and revenue collection costs will be minimal.

3.5.7 POP connection

An STM-1 (155 Mbps) POP connection is likely to cost around £1M per annum. A 34 Mbps connection would cost around £260K per annum

3.5.8 Spectrum Access

There has been a recent auction of FWA bandwidth near 3.5GHz to help calibrate this cost. Further FWA allocations are possible, and moreover spectrum trading arrangements will now become possible. Typical regional licences, eg East Anglia, were auctioned for £0.5M, and amount raised for the whole auction was over £7M.

3.6 Rural Cost Model 1

This cost model looks at a scenario based on a relatively small-scale base station system fanning out from a hub base station which is near a POP. The costs of all the elements of the system are shown in Table 3-1. Pricing is in terms of the capital outlay, year 1 costs, and recurring costs thereafter. The system is designed to serve provide local backhaul to last hop operators serving at least 3200 households spread across 68 localities.

3.6.1 Assumptions of Cost Model 1

The availability of appropriate masts will always be a problem in this kind of rollout. New masts represent a heavy upfront capital penalty. With the figures used in the model it can be seen that there is a break-even over a five-year period and thereafter mast rental costs more. Base station AP units will require between one and two metres of vertical space on a mast and ideally the APs would be sited as far apart as possible. Directional units operating on the same frequency should be mounted strictly back-to-back to avoid interference.

CPEs should be much cheaper to site, and three different possibilities are included in Table 3-1. If a new mast is required to site a CPE giving service to a locality only a fairly lightweight structure is required. It could be a telegraph pole or a light metal pole. It is likely that villages and small towns will have suitable sites on existing structures. Either there will already be a mast of some type available or any high building may be suitable. For high buildings belonging to members of the public a suitable arrangement could be made, perhaps involving a cheap or free broadband service.

Since it is not possible to predict the mix of cheap and expensive sites in a rollout equal numbers of each option have been included in these rollout models.

Although savings in the cost of capital communications equipment will contribute to increased profits the major outlays will be in the POP connection and the company staffing. The possibility of being undercut by drops in wired broadband delivery costs to rural communities should not be a problem. The market price of a POP connection, representing wholesale bandwidth, will always be favourably related to the comparison cost of connecting up a particular rural community by wire. Prices for wired connections to rural locations are priced according to distance from particular POP positions. There is a standard charge for areas close to a POP which is then increased by a price per km for more distant end-points. Since this wireless broadband distribution system dispenses with the high costs of wired extensions to the network it will remain the cheaper option even if data bandwidth prices drop considerably.

The company staffing level assumed is modest, it would have to be carefully controlled. This is a backhaul network with relatively few customers in the form of community last-hop operators, some of which may be national in scope. This means billing and accounts should be a low overhead. There is clearly a requirement for a small maintenance team to deal with hardware and software failures. The number of people required will have to be inline with the MTBF of each piece of kit.

Assumptions need to be made concerning customer numbers, the take-up ratio, and the size of community that a CPE would serve. Indications from existing last-hop companies suggest that a market penetration of 5% would be a reasonable and conservative figure to aim at in rural communities. With the current prices for wholesale data bandwidth it is reasonable to imagine charging the customers of the last-hop companies at least £25 per month in the UK, in line with current wired broadband offerings. Last-hop operators will have to take a view on whether it is possible to charge a higher rate early in rollout when there is no other competition. This strategy could, of course, also cause people to decide once and for all that broadband is not worth it. The last-hop companies will need some margin so it might be reasonable to plan on gaining an average income per last-hop customer of around £16 per month. With a 5% take-up from a community of 1000 households this would give an income from a CPE location of £9600 per year. A more detailed Cost Model is required in order to find out how much it would cost a wireless broadband operator to deliver this service, and how much profit it could make, but this figure compares well to the cost of obtaining a dedicated wired broadband connection in a rural location.

3.6.2 Analysis of Cost Model 1

In Cost Model 1 data bandwidth will be delivered to a group of communities which are within a few tens of km of a POP location. The deployed infrastructure for the model corresponds with Figure 3-1, where the final link in the backhaul chain is the base station labelled C. This means there would be two base stations in each backhaul chain and any reasonable number of these chains could be added to the hub. This model to be examined will have 4 chains of base stations. A repeater station, labelled A in Figure 3-1, adds distance value to the delivered data bandwidth, taking the service into areas not wired for broadband. See Table 3-2 for the capital and recurring costs of the model where figures are in terms of thousands of UK pounds.

Table 3-2 shows that the capital costs are dominated by the costs of new masts, and if equipment can be sited on existing masts or high buildings significant capital savings can be made. The recurring costs over a few years quickly outweigh the capital costs however. The most significant contributions here are the cost of the central POP connection and the costs associated with staffing the company.

3.6.2.1 Potential Annual Income A

Table 3-3 contains a ten-year cashflow forecast for a business based on Cost Model 1 above. For the example of Potential Annual Income A a rollout of the infrastructure over three years has been assumed. The eight base stations installed support 68 CPEs situated in rural communities. In this model the CPEs can be thought of as mini-POPs delivering an average of 0.5 Mbps of uncontested bandwidth into communities of around 1000 households. Distribution from the mini-POPs will be done by last-hop operators. A 5% customer take-up of the last-hop operation has been assumed leading to 3200 domestic subscribers being connected to the system via the last hop operators by the end of year 3.

This does not look like a particularly good business, it does not break even until about year 6, although thereafter a steady £180K per annum profit is made. The maximum borrowing requirement peaks in year 3 at around £1.2M. The finance costs are not included in this model but in this case they will be significant, taking several hundred thousand pounds off the 10 year profits. The costs of radio licensing are not included either. In this case they will also be significant. The area covered will be a small fraction, perhaps one fifth of the areas covered in the recent UK 3.5GHz FWA auction. A pro-rata figure for radio licensing would then be about £100K.

Why does this look like a poor business model? The wireless distribution kit is being used in an inefficient way here. The final stage of backhaul is only carrying an average of 8.5 Mbps of data whereas it is specified for around 40 Mbps. Also the main infrastructure elements of the distribution system – the POP base station and the relay stations are supporting relatively few rural base stations.

3.6.2.2 Potential Annual Income B

Perhaps a higher customer take-up could make this small-scale model more attractive. In Table 3-4 it is assumed that the customer penetration level increases after year 3 and has doubled to 10% by year 10. The recurring costs increase because more POP bandwidth has to be rented to cover the increased demand, but the business looks sustainable. The finance costs and radio licensing costs should be absorbed by the end of year 6 and by year 10 an annual profit of over £500K is being made.

3.6.3 Conclusions on Cost Model 1

This model carries the penalties of running a small-scale operation. An interesting global figure which measures the efficiency of this business is the average cost per annum over the ten year period of delivering bandwidth to a CPE. It depends on the bandwidth being supplied to the POP so is different for the two different income scenarios. For Potential Annual Income A the figure is £7500 per annum, and this represents £12.50 to supply each end user per month. For Potential Annual Income B the bandwidh cost to the CPE is higher at £9000 per annum. But now there are twice as many customers by the end of the ten year period and the average cost of supplying the end users falls to £7.50 per month.

3.7 Cost Model 2

In this cost model the backhaul chain will be fully utilised to get the most efficient use out of the core infrastructure.

3.7.1 Assumptions of Cost Model 2

This cost model aims to improve the overall efficiency of the business by using much longer chains of base stations. As was discussed the reliability of the system can be dramatically improved if the chains either return to the original hub location or link right through to the hub of another system. In Table 3-5 an analysis is presented of costs for a system with four chains of 12 base stations being backhauled to a hub location. As before an ATM relay station moves the first base station in the chain well out of town. A 155 Mbps uncontended connection is now required at the POP, and this accounts for two-thirds or the recurring costs. In fact the higher this ratio is the more efficiently data bandwidth is being transported. The system supports 320 CPEs serving 16,000 customers via last hop operators representing a 5% coverage of a total rural population of 320,000 households.

There are about 25M households in the UK of which about 15% can be considered as being in rural locations. This represents 3,750,000 households. The planned coverage area of Cost Model 2 therefore takes in about 10% of the total rural population: it is a model for regional rural coverage.

Once again where cost variations occur for masts an equal mix of options has been chosen.

3.7.2 Analysis of Cost Model 2

It is not surprising to discover in Table 3-5 that the capital costs of masts and the recurrent cost of data bandwidth dominate the figures. If a regional licence had been acquired for say, £0.5M, then that would have to add that into capital costs.

The cashflow for Cost Model 2 looks more attractive than that of Cost Model 1 even only assuming a 5% end-user market penetration. Maximum borrowings peak in years 2 and 3 at about £4M but are only required until break even occurs early in year 4. Realistically, the cost of borrowings should be added in plus the cost of a radio license, which would be around £0.5M for this regional size operation. This would mean break even would occur by the end of year 4. After that the profitability looks excellent, and the business makes about £1.5M per year profit.

If the average cost per community of the wireless broadband connection is calculated it is evident why the profitability is so good. Now it only costs £5500 per community or £9 per end user per month, and this is for only 5% market penetration. However, the backhaul path is saturated by the end of year 3. There will be 80 communities each requiring an average of 0.5Mbps connected along each chain and the peak backhaul bandwidth is around 40 Mbps. In fact some room for growth could be engineered by partitioning the bandwidth a little differently in different parts of the chain. Near the hub the maximum 14 MHz bandwidth for backhaul could be used and only 5 MHz reserved for distribution. This would allow growth of the business by a further 25% without any additional capital outlay.

This rollout could be repeated across the other rural parts of the UK. Some regions may not have enough communities of 1000 households at a sufficient density for Cost Model 2 to operate efficiently. Other rural regions have quite a high density of villages and towns.

Linearly scaling Cost Model 2 up to include the entire rural population means 160,000 end-point subscribers of last-hop operators would be served via 3,200 CPEs. Some regions would clearly be less cost-effective to service but running a system 10 times as large would bring efficiency benefits as well. The figures show a profit of £99M over 10 years and a steady-state annual profit of £14.6M.

If the arguments of Section 3.4 above concerning the relative merits of wireless broadband over wired offerings are to be believed it may be that this type of model can compete head-on with wired services in suburban and urban locations. Retaining the modest figure of 5% market penetration the figures for national coverage would predict the cashflow of the lower sub-table of Table 3-6. There would be 1.25M households serviced using 25,000 CPEs supported by 3,750 base stations. Total capital costs would be £217M and the accumulated profit after 10 years would be £777M with £114M profit being generated per annum in the steady state. A license fee of £7.5M does not really make much difference to the initial maximum borrowing requirement of just over £300M, assuming a fast 3-year rollout.

3.8 Conclusions of the Cost Models

It appears that it would be possible to use a highly-specified FWA system, like the one used in this wireless broadband trial, to provide cost-effective delivery of data bandwidth to rural locations. Long backhaul chains of base stations help amortise the costs of hub equipment and distribution near the hub. The result is a system which can deliver rural data bandwidth more cheaply than a wired cost model which charges an Out of Area per km rate for rural locations.

A highly-specified wireless broadband system may be able to offer better apparent performance for the same price as traditional wired offerings. In particular it can deliver much higher peak data bandwidths to every customer. This may be expressed as delivering a higher bandwidth with a higher contention ratio than wired services. This and other advantages may allow wireless broadband to compete in the suburban and urban market as well as the rural one.

Using backhaul based on the same equipment as for FWA distribution has the further advantage of reducing the technical complexity of the complete network. Dealing with a mixture of technologies and products would involve the overhead of additional staff training and complicate the system administration.

	Capital (UK £)	Recurring costs (UK £)	
		Year 1	**After year 1**
Basic Unit Prices			
Ethernet CPE	600	0	51
High-gain Antenna for SU	500	0	0
High-gain Antenna for AP	700	0	0
External Rack Enclosure	1500	0	0
Internal Rack Enclosure	500	0	0
Compound Unit Prices			
4-way ATM base station	25000	0	2275
2-way ATM base station in chain	27000	0	1000
AP or SU two-way relay	7000	0	595
Site and Install Prices			
Base Station on existing mast	2000	7000	7000
Base station on new Mast	30000	1000	1000
Cheap existing CPE mast	200	300	300
Cheap new CPE mast	3000	100	100
Network Connection			
34 Mbps from POP	45000	259200	259200
155 Mbps from POP	100000	1000000	1000000
Installed Prices			
Installed 4-way POP ATM base station (Existing Mast, 34 Mbps POP connection)	73500	266200	268475
Installed 4-way POP ATM base station (Existing Mast, 155 Mbps POP connection)	128500	1007000	1009275
Installed 4-way POP ATM base station (New Mast, 34 Mbps POP connection)	101500	260200	262475
Installed 4-way POP ATM base station (New Mast, 155 Mbps POP connection)	156500	1002500	1004775

Installed chained base station station (existing mast)	30500	7000	8000
Installed chained base station station (new mast)	58500	1000	2000
Installed ATM relay (Existing Mast)	9000	5000	5595
Installed ATM relay (New Mast)	37000	1000	1595
Installed Ethernet CPE (Existing Mast)	3100	1000	1000
Installed Ethernet CPE (Cheap site)	1400	200	200
Installed Ethernet CPE (Cheap New Mast)	4100	200	200
Installed Ethernet CPE (customer site)	1200	0	51
Staffing and Office Space			
Small regional model manager, secretary, two support staff	15000	120000	120000
Large regional model manager, secretary, four support staff	25000	200000	200000

Table 3-1: Pricing elements for Cost Model

Cost Model 1
Short chain of Base Stations
backhauling Hotspot operators (UK £K)

	Capital	Year 1	Recurring after year 1
Hub 4-way AP Base Station existing mast and 34 Mbps POP	74	266	268
2 ATM relay stations, existing masts	18	10	11
2 ATM relay stations, new masts	74	2	3
4 2-way AP relay stations, existing masts	122	28	32
4 2-way AP relay stations, new masts	234	4	6
22 Ethernet CPEs, existing masts	68	22	22
23 Ethernet CPEs, existing cheap mast sites	32	5	5
23 Ethernet CPEs, new cheap masts	94	5	5
Staffing and office space	15	120	120
Totals	**731**	**461**	**472**

Table 3-2: Capital and recurring costs for Cost Model 1

Potential Annual Income A, from Cost Model 1 (£K)

68 communities with 50 subscribers

Cashflow	**Capital**	**All Recurring**	**Income**	**Profit**	**Accumulated Profit**
Total Year 1	244	154	218	-180	-180
Total Year 2	487	465	435	-517	-697
Total Year 3	0	472	653	180	-517
Total Year 4	0	472	653	180	-337
Total Year 5	0	472	653	180	-156
Total Year 6	0	472	653	180	24
Total Year 7	0	472	653	180	205
Total Year 8	0	472	653	180	385
Total Year 9	0	472	653	180	565
Total Year 10	0	472	653	180	746
Total over 10 years	**731**	**4,398**	**5,875**	**746**	

Table 3-3: Cost Model 1 – Cashflow for Potential Annual Income A

Potential Annual Income B, from Cost Model 1 (£K)

68 communities with 50 subscribers
rising to 100 by year 10

Cashflow	Capital	All Recurring	Income	Profit	Accumulated Profit
Total Year 1	244	154	218	-180	-180
Total Year 2	487	465	435	-517	-697
Total Year 3	0	472	653	180	-517
Total Year 4	0	509	746	237	-280
Total Year 5	0	546	839	293	13
Total Year 6	0	584	933	349	362
Total Year 7	0	621	1,026	405	767
Total Year 8	0	658	1,119	462	1,228
Total Year 9	0	695	1,212	518	1,746
Total Year 10	0	732	1,306	574	2,320
Total over 10 years	**731**	**5,435**	**8,486**	**2,320**	

Table 3-4: Cost Model 1 – Cashflow for Potential Annual Income B

Cost Model 2
Long chain of Base Stations
backhauling Hotspot operators (UK £)

	Capital	Year 1	Recurring after year 1
Hub 4-way AP Base Station existing mast and 155 Mbps POP	129	1,003	1,009
2 ATM relay stations, existing masts	18	10	11
2 ATM relay stations, new masts	74	2	4
24 2-way AP relay stations, existing masts	732	168	192
24 2-way AP relay stations, new masts	888	24	38
106 Ethernet CPEs, existing cheap masts	329	106	106
107 Ethernet CPEs, existing cheap mast sites	150	21	21
107 Ethernet CPEs, new cheap masts	439	21	21
Staffing and office space	25	200	200
Totals	**2,783**	**1,555**	**1,604**

Table 3-5: Capital and recurring costs for Cost Model 2

Potential Annual income, from Cost Model 2 (£K)

320 communities each with 50 subscribers
(16000 subscribers)

Cashflow	Capital	All Recurring	Income	Profit	Accumulated Profit
Total Year 1	928	518	1,024	-422	-422
Total Year 2	1,855	1,571	2,048	-1,378	-1,800
Total Year 3	0	1,604	3,072	1,468	-332
Total Year 4	0	1,604	3,072	1,468	1,136
Total Year 5	0	1,604	3,072	1,468	2,605
Total Year 6	0	1,604	3,072	1,468	4,073
Total Year 7	0	1,604	3,072	1,468	5,542
Total Year 8	0	1,604	3,072	1,468	7,010
Total Year 9	0	1,604	3,072	1,468	8,479
Total Year 10	0	1,604	3,072	1,468	9,947
Total over 10 years	**2,783**	**14,918**	**27,648**	**9,947**	

10 year Total, all rural (160,000 subscribers)	**27,826**	**149,182**	**276,480**	**99,472**	

10 year Total, all households (1,250,000 subscribers)	**217,391**	**1,165,482**	**2,160,000**	**777,127**	

Table 3-6: Cost Model 2 - Cashflow

4 International Usage Comparisons

4.1 International Coordination

The European Radiocommunications Committee (ERC) has developed a Common European Frequency Allocation Table which members of CEPT (European Conference of Postal and Telecommunications Administrations) are working towards. It recommends 100 MHz duplex spacing at 3.4 GHz for Fixed Wireless Access. CEPT/ERC recommendation 14-03 E of 1997 gives guidelines on harmonised radio frequency channel arrangements and block allocations for low and medium capacity systems in the band 3400 MHz to 3600 MHz. Over 100 countries have adopted channel allocations in line with these recommendations but the details of channel centre frequencies and bandwidths vary. Some countries have generous allocations but restrict the number of licenses an operator may have, while in others operators are able to gain access to multiple contiguous FWA bands. To give an indication of the current situation Table 4-1, below, shows some typical FWA spectrum allocations conforming to the ERC recommendation.

COUNTRY	PARTITIONING (MHz)	SPECTRUM ALLOCATED
Australia	3425 - 3442.5 / 3475 - 3492.5 3442.5 – 3475 / 3542.5 – 3575	2 x 17.5 MHz and 2 x 32.5MHz
Austria	3450 - 3500 coupled to 3550 – 3600	25 MHz duplex max per operator
Denmark	3410 – 3590	2 licenses of 2 x 26.5 MHz 1 license of 27 MHz
France	3465 - 3495 and 3565 – 3595	2 x 15 MHz per licensee, 2 National Licensees
Finland	3410 – 3590	~2 x 20MHz to each operator
Germany	3410 – 3580	5 duplex blocks x 14 MHz each
Ireland	3410 – 3435 paired with 3510 - 3535 3475 – 3500 paired with 3575 – 3600	2 x 25 MHz each
Netherlands	3500 – 3580	80MHz TDD only
Portugal	3600 – 3800	3 licences 2 x 28 MHz
Spain	3400 – 3600	2 x 20 MHz per license
Sweden	3400 – 3600	3 licenses 2 x 28 MHz
UK	3480 – 3500 and 3580 – 3600	2 x 20 MHz

Table 4-1: Representative FWA spectrum allocations at 3.4 - 3.6 GHz

4.2 Potential FWA Services

It is appropriate to consider what uses this FWA bandwidth could be put to, and how much bandwidth is required for different service types. The bandwidth requirements of each service will be estimated, and the benefits of efficient co-existence of these services in a single FWA infrastructure will be discussed.

4.2.1 GSM Backhaul

In countries without sufficient wired infrastructure point-to-point radio backhaul of GSM base stations is a cost-effective alternative to laying new cables. The backhauls require N x 64 kbps synchronous fractional E1 specification links of up to 2 Mbps per GSM base station but are often used at much lower data bandwidths. If multiple backhauls to hub locations are required then point-to-multipoint radio systems become more economically viable. FWA systems supporting the E1 standard which can backhaul several 10's of GSM base stations per FWA base station are being installed in large numbers in China and Malaysia. In the UK it may be economical to backhaul 3G base stations using FWA bandwidth. Some 30,000 3G base station sites will be required to cover the whole UK and many will be in remote uncabled areas.

The problem with these E1 streams from the point of view of efficiently trading data bandwidth with other services is that they use data bandwidth continuously even when there is no real voice or data traffic being carried. This means that if they are carried they deplete the bandwidth available for mixed uses.

A typical usage level might be to backhaul between 10 and 20 GSM base stations to each FWA base station. This represents a maximum data rate which could range between 5 and 40 Mbps over four sectors. Because both the FWA base station and the GSM base stations are likely to be at good elevations it would be reasonable to expect to operate at 64QAM for these backhauls, this would require between 1.6 and 6.6 Msymbols per second to carry the data, spread over all the base station sectors.

4.2.2 Public Wireless Broadband

Currently, high specification FWA systems with ranges in excess of 5 km which fully support QoS and can operate at 3.4 GHz require rooftop installs of moderately expensive directional antennas. Using this grade of equipment it is not currently commercially viable to consider a model where mass-market wireless broadband is directly connected to consumers' homes. Such high-end systems, which support QoS and can make use of wide bandwidths, are more suited to backhauling less expensive last-hop solutions. In the future a family of cheap FWA wireless systems based on, for example, the IEEE 802.16a standard may allow direct domestic connections as well as supporting a wide range of other services with QoS controls, see Section 4.6. But if data bandwidths truly comparable with wired broadband are required to be delivered it will still be likely that high gain near line of sight antennas will be required to cover longer distance backhaul links, while omni antennas could be used for covering Hotspots over the last few hundred metres.

It is seen in the Cost Model of section 3 of this Annex that regional solutions for backhauling last-hop technologies can operate within the UK PFWA bandwidth allocation of 2 x 20 MHz, including backhaul to regional POP locations. This model assumed a 5% take-up, any higher customer level would require additional backhaul spectrum since the in-band backhaul links would saturate. No more PFWA spectrum would be necessary, however, since over long distances and large areas frequencies can be efficiently re-used, it was assumed that 2 x 5 MHz would be enough for connections to FWA CPEs. If it were a requirement to operate in urban locations base stations would have to be more densely deployed. The resultant interference problems could mean that frequency colouring patterns which demanded at least twice as much bandwidth would be required.

In this report's Cost Model the final connections to CPEs would not be saturated since, on average, only around 0.5 Mbps is required to be delivered to each community unit of 1000 households. As was pointed out this does have the advantage that higher bandwidths at higher contention ration could be offered, improving the apparent speed of the connection. Because of this unused bandwidth at the periphery it would be possible to exploit the network to a much greater extent by offering a range of other services. The backhaul would have to be adjusted to cope with the increased load as more services were introduced. This would have the effect of moving backhaul at other frequencies deeper into the network as demand grew. Optimal use of the network can be made by arranging that the backhaul from last-hop operators just fills the in-band backhaul links of the FWA network. In the Cost Model 14 MHz of bandwidth was needed in the in-band backhaul to accommodate all the end-user traffic.

4.2.3 Commercial Wireless Broadband

The service classes in this category are business-class connections to companies. In rural areas FWA could replace expensive leased-line connections. This class of traffic requires higher bandwidth and a higher priority than domestic connections. Small companies and rural local authority sites may typically want between 0.5 Mbps and 2 Mbps with a contention ratio of 20:1 or 10:1. The business-class connections will place more load on the final links to CPEs in the network. Each business-class connection will use a similar amount of bandwidth to between 10 and 40 domestic connections to the CPE and through the backhaul links. On average perhaps 10 such business-class connections might be expected from each CPE. If, on average, this placed a data bandwidth load of an extra 1 Mbps on each CPE connection it would triple the backhaul requirements to around 42 MHz of spectrum bandwidth if the idea of long chains of base stations was retained. However, since these business connections command higher charges shorter backhaul chains of base stations could then be tolerated. These shorter chains would still be more efficient than separately backhauling each FWA base station but would require independent backhaul. A new Cost Model would then be appropriate.

4.2.4 Broadband to Schools and Colleges

Schools and Colleges may need between 2 Mbps and 8 Mbps. These are high bandwidths for a wireless system, but on the other hand the times of day and days of week of peak use are complimentary to those of domestic broadband customers. Several such connections could be supported by each base station. It would not be possible to place many such streams into a long and heavily used in-band backhaul chain, but it may be possible early in a rollout. As demand for all classes of services grows these connections would have to be handed off to other backhaul routes.

4.2.5 Local VPNs

There will be point-to point connections between, for example, different sites in a locality of the same company or local authority. There was an example in the current trial where a local connection was set up between two sites of Cambridge Water. Such local VPNs, which do not involve backhaul, could easily be accommodated since an in-band backhaul model has more spare bandwidth at CPEs that in the backhaul chain. Typical requirements might be between 0.5 and 10 Mbps. Some of the traffic may be variable bit rate, but some may be constant bit rate, for example, security camera video links.

4.2.6 Long Distance VPNs

These will be harder to accommodate since they make demands on scarce backhaul bandwidth. The requirement may still be in the range 0.5 – 10 Mbps. Like links to schools and colleges the business could be taken on early in rollout and handed off to other backhaul methods at a later stage.

4.3 Allocating Bandwidth for Efficient Usage

The previous section discussed a variety of potential uses for FWA. In this section it is noted that the size and use made of bandwidth allocations can affect the efficiency with which bandwidth can be exploited. Different countries have different approaches on whether they allow multiple air link technologies to be used.

Some aspects of these problems have been considered by a report made to the Electronic Communication Committee (ECC) within the European Conference of Postal and Telecommunications Administrations (CEPT), (The Analysis of the coexistence of FWA cells in the 3.4 – 3.8 GHz Band ECC report 33).

4.3.1 The efficiency advantages of QoS-capable systems

The bandwidth allocations seen in different countries appear to have been arrived at by consideration of the competing and established uses of bandwidth at these frequencies. Bandwidth allocations are not expressed in terms of the mix of services that could be efficiently delivered in a commercially viable way. In Section 13.1 of the main report it was shown experimentally that a mixture of service types, each with their own specified QoS, can co-exist in an efficient manner. When high priority services are not carrying their full quota of traffic the spare bandwidth can be used by lower priority services. This means there will be a minimum bandwidth below which multiple services cannot be supported as in a narrow channel they will not be able to efficiently trade data bandwidth in this way. If this were to result in different services having to be carried in separately allocated bands then again no sharing of unused data bandwidth is possible and radio spectrum is being used at less than optimum efficiency.

Another consideration is the maximum bandwidth which FWA technologies are capable of exploiting. If the maximum bandwidth available to an operator is less than the maximum bandwidth that a chosen technology can use then that technology is penalised under that licensing regime since expensive wide-band components are not being fully utilised. An operator would want to make maximum use of their investment in base stations and CPEs by operating it over the maximum bandwidth of which it is capable.

FWA technology that does not support priority-based QoS will not be able to make any efficiency gains by running mixed services over a single point-to multipoint system. An operator using such an FWA system would not be able to offer cheap contended data bandwidth services on top of expensive guaranteed services. Using equipment of this specification the only way of selling a range of services of different priorities would be by the spectrally inefficient strategy of one operator or a number of different operators installing multiple FWA systems, each carrying a different service type.

4.3.2 FDD and TDD systems

There are two main duplexing approaches used in FWA networks. These alternative approaches do not co-exist very easily either in the same or adjacent bands. This has implications for bandwidth allocation and some countries have favoured one or the other approaches.

A Frequency Division Duplex (FDD) system, where transmissions to and from a base station are made on separate frequencies, was used in this trial. A Time Division Multiple Access (TDMA) protocol was used to synchronise upstream and downstream transmissions. FDD FWA systems have been operated at 50 MHz duplex spacing but the trend is towards 100 MHz spacing in the 3.4 GHz band. Time Division Duplex (TDD) FWA systems, which schedule upstream and downstream transmissions on the same frequency, are also available, though they have usually been more favoured for use in unlicensed spectrum.

4.3.2.1 FDD advantages and disadvantages

The splitting of upstream and downstream data into two different and comparatively well-separated channels brings a number of efficiency advantages. The upstream and downstream data can then be transmitted simultaneously, and because of the adequate frequency separation between upstream and downstream channels it is possible to effectively filter the receivers from the interference effects of the co-sited transmitters. Where multiple operators work in contiguous spectrum allocations it can be arranged that the upstream and downstream transmissions by different operators will not interfere. This allows several operators to share the same mast – a very important financial consideration.

FDD does have some disadvantages. A duplexer is required to allow both the transmitter and receiver to access the antenna, but then it could be argued that a TDD system will have to use more expensive filters. More importantly FDD systems are generally configured to have the same bandwidth upstream and downstream. This means that if the network load is not symmetric down to very short timescales there will be times when the bandwidth cannot be utilised, with a consequent efficiency loss.

4.3.2.2 TDD advantages and disadvantages

TDD systems have been used effectively in unlicensed bands where frequency hopping can work well. The systems can hop away from interference in the same channel or adjacent channels. It would not really be practical to have a frequency hopping duplex system. For carrying data traffic services which are asymmetric, in the sense that there is more downstream traffic to carry than upstream traffic, TDD may make more efficient use of the spectrum. In fact business connections are generally fairly symmetric on average and long term asymmetric traffic really only occurs to the current domestic broadband sector. The trends towards the generation and sharing of multimedia material in homes, the use of peer-to-peer networks, and the use of Voice over IP and real-time two-way video links will all increase demand for more symmetric services.

There is also an inherent inefficiency in TDD systems which arises from the issue of turnaround time. This turnaround time has two components, the smaller is the actual time it takes to reconfigure from being a transmitter to being a receiver, and vice versa. The larger component is the gap which must be left between transmitting and receiving as a result of propagation delays. This arises because radio units would do best to transmit early in order for transmissions to arrive on-time at the receiver, but they cannot do this because at the same time they have to listen until later to allow for the time of flight of radio waves from the other transmitter. In practice an upper bound on the range of TDD systems is set in order to minimise the turnaround time.

Using TDD in licensed bands presents band-planning, efficiency and regulatory problems. It is not possible to co-site TDD sector antenna on the same mast if they are using nearby bands because of the massive interference effect of having a co-sited transmitter and receiver operating on nearby channels. Band planning therefore becomes a complicated task. Worse still, operators working in contiguous spectrum allocations will not be able to co-site the transceiver equipment as it is not physically possible to make filters steep enough to prevent adjacent channel interference. Regulatory authorities which allow TDD in licensed bands must appreciate that extension, in a fair way, of FWA spectrum by adding contiguous allocations would be difficult because only the incumbent TDD operator could band-plan to make efficient use of the new channels.

4.4 Data Bandwidth Requirements of FWA Services

It is possible to estimate how much bandwidth might be required to support a particular range of services for the local distribution between base stations and their CPEs. This then dictates the level of backhaul bandwidth needed to support each base station. In architectures which group base stations together and use in-band backhaul the mix of FWA bandwidth and backhaul bandwidth will depend on the architecture and on the profitability of the various services being supported.

4.4.1 Bandwidth for Local Distribution from FWA Base Stations

In the minimal service model for rural domestic customers of Section 3.7.2 of this Annex 5 MHz wide channels were used to make connections between CPEs and 14 MHz channels were used to make connections to backhaul a chain of base stations. It may be possible to work at 64QAM modulation for many CPE connections since base stations and CPEs could be carefully sited for good lines of sight. But to avoid being over-ambitious 16QAM modulation can be assumed, giving about 14 Mbps maximum bi-directional throughput between a base station and a CPE.

Service Type	**Average CPE Data Bandwidth**	**Contention**	**Priority Level / Rate Limiting**
GSM/3G Backhaul	0.5 – 2.5 Mbps	1:1	High / Peak Rate
Public Wireless Broadband	0.5 Mbps	50:1	Low / Committed Rate
Commercial Wireless Broadband	0.5 – 2.0 Mbps	20:1	Medium / Committed Rate
Schools/Colleges	8 Mbps	10:1	Medium / Committed Rate
Local VPNs	0.5 – 10 Mbps	10:1	Medium / Committed Rate
Long distance VPNs	0.5 – 10 Mbps	10:1	Medium / Committed Rate

Table 4-2: The range of services which FWA could support

Table 4-2 above summarises some typical service types which could be accommodated by FWA. For operators specialising in one particular service type some estimates can be made of what spectrum allocation would satisfy their requirement. For GSM and 3G backhaul it may be possible to work with 2 x 10 MHz of spectrum, giving two 5 MHz channels for frequency colouring base stations. In urban deployments where inter-base station interference would become a problem it would be considerably easier to band-plan with twice this allocation.

If the advantages of using QoS are taken into account then spectrum can be used more efficiently and the allocation of spectrum required is less than the sum of the bandwidths needed for each service individually. All the service types marked as being prioritised at a committed rate in Table 4-2 above are candidates for being able to trade spare bandwidth in this way.

For the connection between FWA base stations and their CPEs an allocation of a minimum of 2 x 20 MHz may be sufficient to cater for a pessimistic 5% penetration level for domestic and commercial wireless broadband. Services requiring higher data rates might be accommodated at first but are generally not suited to general-purpose FWA systems and would have to be migrated to other carriers at a later stage. Since many cost savings, including mast sites and radio infrastructure, could be made by supporting 2G and 3G backhaul over the same system a more appropriate bandwidth for local distribution from FWA base stations capable of supporting a range of services would be 2 x 28 MHz, in line with many international allocations.

4.4.2 Backhaul Bandwidth

If FWA base stations are to be sited in locations remote from sources of wide bandwidth Internet access then substantial wireless backhaul bandwidth will be required. The Cost Model of Section 3.7.2 of this Annex demonstrated that for providing rural broadband it may not be cost-effective to backhaul small-scale base station systems to Internet Points of Presence. Backhauling larger groups or chains of FWA base stations will help amortise the costs of core infrastructure. The other cost efficiency of this architecture is that the same equipment is being used for backhaul as for local distribution. This means both uses can share the same switching infrastructure and be controlled by the same high-level management system.

However, for premium service types such as commercial broadband connections and 2G or 3G backhaul this in-band backhaul model will not be as appropriate for a number of reasons. Firstly, with local distribution bandwidth requirements as high as 2 x 28 MHz the backhaul bandwidths for large groups of base stations will be too large to be run in-band, even allowing for the effects of QoS control. Secondly, the higher profit margin for these services will make it economical to backhaul smaller groups of base stations directly back to Points of Presence using point-to-point links at higher frequencies.

Very few countries have allocated sufficient spectrum to support the model of a variety of services being carried using an in-band backhaul architecture. Notable exceptions would be Denmark, where single operators may use contiguous allocations; the Netherlands, where a large bandwidth is available (but uniquely in Europe it is only for use by TDD systems); and Sweden and Portugal where generous bandwidth allocations exist. Elsewhere the relatively small FWA allocations at 3.5 GHz would have to be used exclusively for local base station to CPE links unless the range of services offered was very restricted.

4.5 International FWA Usage

4.5.1 The General Situation

To date broadband FWA rollouts around the world have not been particularly commercially successful. There were early small-scale rollouts of broadband FWA in the USA in both unlicensed bands (2.4GHz and 5.8GHz) and the licensed multichannel multipoint distribution service (MMDS) bands, but for a number of reasons by October 2001 many of these were on hold. Early BFWA equipment was difficult and expensive to install and was not standards-based, so there was no competition from multiple vendor sources. Early rollouts often competed directly with wired broadband alternatives where the adoption of standards has created strong competition and the costs of installation and operation have been continuously falling. The IEEE 802.16 standard for FWA is now under development, of particular interest is the IEEE 802.16a standard which addresses lower frequency bands. If successful it is hoped that it could create growth in the FWA industry analogous to the way in that adoption of the IEEE 802.11 standard allowed a boom in wireless local area networks, see Section 4.6.

In general large operators would prefer to use licensed spectrum but some have been deterred by overly expensive and restrictive licences in many countries. Some early licenses were granted with an enforced deployment schedule and operators found that they had to roll out in areas with no demand. The business case for rollout in urban areas has not been successful and the Telecoms downturn dried up investment at a critical time. There are still very few customers in most countries with the exception of the US and Korea. Korea had 40,000 wireless broadband users in 2002.

The German experience was a particularly bad one. Germany held beauty contests for the 3.5GHz and 26GHz bands in 1999, when 860 licenses were awarded to 12 operators. Many later dropped out or became insolvent citing regulatory problems and the high initial price of European licenses.

4.5.2 Improvements in Regulation and PFWA Equipment

After a slow start regulatory and technical improvements are making FWA deployments more economically viable.

The international alignment of 3.4-3.6 GHz FWA allocations means equipment vendors can sell into a larger marketplace. A trend towards adoption of less rigid licensing constraints has meant that operators can react appropriately to market forces and expand their business into new geographical areas or new services at a rate depending on demand.

Only a few years ago most FWA equipment would only work with a good line of sight between base stations and CPEs. This gave coverage problems to operators who could often only give a reliable service to a half of the potential customers of a base station. Non line of sight systems which usually use either adaptive equalization or coding techniques such as Orthogonal Frequency Division Multiplexing (OFDM) improve the economic case allowing much higher coverage ratios.

4.5.3 Usage Trends for PFWA

Spectrum for Public Fixed Wireless Access has only been available in most countries for a short time. In countries which have auctioned bandwidth relatively few operators have started services yet, and many are still considering their usage of the bands. Some usage trends are beginning to emerge however.

The use of PFWA for pure residential connections is not generally considered to be a viable business at present. The emergence of PFWA standards, multi-vendor sourcing of radio kit and cheap commercial-grade last-hop systems may change this. A healthy business-oriented or backhaul-oriented service model is currently emerging in some countries.

4.5.4 International PFWA Licensing Conditions

Although FWA spectrum around 3.5 GHz is becoming more aligned internationally the detailed licensing conditions vary widely from country to country.

The UK 3.5 GHz FWA bands were auctioned in May 2003 and followed a technology neutral approach. License fees were staged into three five-year payments, allowing operators to withdraw after each five-year period. The licenses were not restrictive on services and usage but allowed for provision of FWA services to end users; wholesale provision to ISPs and Telcos and backhaul of an operator’s own FWA traffic including that resulting from wholesale traffic. Other countries are also following a similar “light touch” approach.

4.5.5 International PFWA Usage Examples

4.5.5.1 Geographical scope of licenses

National licenses have been granted in some countries, for example: Ireland, Greece, Denmark, Netherlands, Spain and Luxembourg. A system of regional licenses has been adopted in others, for example: UK, Germany and Australia. A further variation is that some countries, such as France, have both regional and national licences.

Regional licenses allow smaller operators to enter the market. But some of the largest operators would not consider buying a patchwork of regional licenses and would prefer one national license. This is often for reasons of uniformity in their advertising and service offerings.

4.5.5.2 Minimum rollout conditions

When many early FWA licenses were allocated the so-called ".com boom" was still gaining momentum and there appeared to be no shortage of investment for telecommunications. In this economic environment it is not surprising that regulators thought it reasonable to impose minimum rollout conditions on license holders. The thinking behind this is one the one hand to try and ensure that license holders do not sit on unused licenses, and on the other hand to force operators to provide services to geographical areas which otherwise might prove uneconomic.

The German experience has made this approach unpopular, but other countries also had rollout specifications. In Finland the licenses were granted for six years with a rollout requirement in year one, and in Ireland some basic services were required of operators.

4.5.5.3 Technology neutral regulations

Many recent allocations have followed a technology neutral approach, for example: UK, Denmark, Ireland and Sweden. However, the Netherlands has allocated bandwidth in a way which favours TDD-based technology.

More research needs to be done in this area in order to develop better guidelines for mixing technologies in the 3.4 to 3.8 GHz bands. The co-ordination of FDD, TDD and Mesh networks should addressed.

4.5.5.4 Operator selection criteria

Some countries pre-selected license applicants on the basis of a set of selection criteria designed to try and ensure that a well-run and competent FWA service would result. For example in the Netherlands the financial status of applicant companies was taken into consideration and applicants had to show that they possessed an appropriate level of radio expertise.

4.5.5.5 Efficient spectrum usage

The Netherlands and France explicitly mentioned that spectrum should be efficiently used. Of course it is quite possible for operators to present business plans which appear to make very efficient use of spectrum, but it would be difficult to enforce.

4.5.5.6 License allocation methods

A variety of license allocation methods have been used in different countries. Allocation by auction is quite common and was done in, for example: Austria, Greece, the Netherlands and the UK. So-called beauty contests were held in other countries, such as: Belgium, Denmark, France and Ireland. Licenses were allocated on the basis of how well applicants were judged to be able to efficiently and competently operate a useful service.

4.5.5.7 License payment schedules

A number of countries adopted a system of annual fees, for example: Finland, Belgium, Denmark and Ireland. In the UK the licenses are periodically renewed at an interval of five years. Luxembourg has a more complex way of raising revenue from the spectrum involving an initial fee, an annual fee, a percentage of turnover and a spectrum-related fee. Portugal has taken a different approach and takes an annual fee, a license fee, a per base station fee, and a spectrum block width related fee.

4.5.5.8 Length of license

Once again there is considerable variation around the world. The Spanish licenses are allocated for 20 years, the Swedish for 10 years, and in the UK they are renewable after each five year period for up to 15 years. The Finish licenses were for only six years.

4.5.5.9 Pairing of bands

Bands are usually allocated in pairs with a 50 MHz or 100 MHz spacing to accommodate the operators of FDD equipment. But in Australia it was possible to buy the unpaired bands independently, while in the Netherlands no such pairing was on offer.

4.5.5.10 Services offered

The domestic wireless broadband market has not been considered to be a profitable market until recently. However, in the last year there have been a number of announcements of small-scale FWA service provision to domestic customers, for example in Stuttgart, Sydney and some Spanish cities. Iceland has FWA rollouts specifically targeting residential customers, and in China whole apartment blocks are being connected to broadband via FWA systems.

More commonly FWA rollouts have first addressed the commercial sector. For example the Swedish operator Telia is currently rolling out a wireless broadband service to business customers. In Scandinavia generally most residential areas and apartment blocks are well covered by cabled broadband with a high percentage of take-up and FWA is targeted towards SMEs and larger corporations.

In countries without extensive cabled broadband FWA systems are being used for GSM cellular telephony backhaul, particularly in China and Malaysia. Another variant of this idea is the backhaul of ADSL mini-DSLAMs in Finland.

The South Korean government has a strategy to ensure that all its inhabitants will have access to broadband. This means there is a requirement to make broadband available in all new buildings. When the need cannot be met with cable solutions they have turned to FWA systems. This why the South Koreans have one of the largest installed bases of FWA broadband in the world.

In Japan FWA systems have met with little interest for a number of local reasons. ADSL and cable broadband is widely available and people are used to very high levels of service. Wireless systems are perceived to bring with them problems of service outages which would not be tolerated. In addition there is a local problem with the negotiation of building rights for siting radio equipment which makes gaining permission to install equipment very long-winded and expensive.

4.5.5.11 The situation in the United States of America

There is currently no allocation in the bands around 3.5 GHz in the use for fixed wireless access. Early deployments of FWA were made in unlicensed bands at around 2.4 GHz and at around 5.2 GHz and 5.7 GHz. Licensed operators use a large allocation of around 200 MHz around 2.1 GHz and in the range 2.5 to 2.7 GHz known as the Multichannel Multipoint Distribution System (MMDS) bands. Two-way transmissions have been permitted here since 1998 and about 20 metropolitan areas have some coverage although many rollouts were suspended a few years ago.

There are now plans to auction FWA spectrum in what is known as the extended C-Band. This is a 50 MHz band in the range 3.65 – 3.70 GHz.

4.6 International FWA Standardisation – IEEE 802.16

With a large number of proprietary FWA systems available, each with different attractions and niche applications, it is difficult for operators to put together an integrated fixed wireless access system capable of delivering a range of services to a range of customer types. A better situation for operators would be to have a range of multi-vendor systems conforming to a standard capable of supporting all the service types required. The IEEE 802.16 standard attempts to address this broad set of difficult problems. The original intention of the IEEE 802.16 Broadband Wireless Access Working Group was to promote a standard for FWA systems operating in the bands from 10 GHz to 66 GHz. Subsequently, the IEEE 802.16a standard has been drafted to extend the scope to include the bands in the 2 to 11 GHz range.

In summary, the standard covers the Media Access Control (MAC) and physical (PHY) layers. The single carrier PHYs support both FDD and TDD deployments, while the OFDM PHYs are designed for TDD. The MAC is designed to carry ATM, Ethernet and IP protocols at up to very high bit rates. It is compatible with ATM QoS as well as non-ATM services such as Multiprotocol Label Switching (MPLS) and VoIP. The uplink is Time Division Multiple Access (TDMA) with each CPE being assigned a slot to use with their individually assigned modulation level and error correction method. The full standard, which includes authentication, security, and many other features is very extensive.

An industry group known as the WiMAX Forum has been set up to promote interoperability of vendor equipment based on the IEEE 802.16 standard. WiMAX thus stands in relation to IEEE 802.16 in the same way that Wi-Fi does to IEEE 802.11. The WiMAX Forum is promoting a particular subset of the full IEEE 802.16 standard. It is to be hoped that this is a comprehensive enough specification to cover the very wide range of features required for all levels of FWA deployment. However, it seems unlikely that FWA equipment conforming to the minimum WiMAX standards can replace the functionality of existing well-developed high-specification systems. It could be argued that today FWA equipment should be thought of more as a piece of general networking kit than as a radio system and it is the scope of the professional networking features that will be the top consideration for operators.

Wi-Fi has been particularly successful and cheap interoperable wireless LAN equipment is now taken for granted. If WiMAX is similarly successful it will create great opportunities for vendors, operators and customers alike.

4.7 Conclusions on International Usage

Fixed Wireless Access systems appear to have had a false start when early technology was rolled out during a period with easy investment opportunities towards the end of the so-called ".com boom". When investment dried up the businesses of many of these early deployers failed and gave FWA a bad reputation.

Although a large number of countries now offer some FWA bands in the region of 3.5 GHz the assignments vary from the meagre to the more generous and the regulatory frameworks are very diverse. Small bandwidth allocations can only support a limited number of services in an inefficient way, whereas larger allocations could be used to efficiently multiplex a mix of services by QoS-capable FWA systems.

An effort to internationally align bands for fixed wireless access at around 3.5 GHz and the appearance of more competent FWA equipment has generated renewed interest and a new generation of small rollouts are now appearing all over the world. They are currently offering services mainly aimed at business customers.

The equipment vendor marketplace now holds a wide diversity of proprietary FWA solutions each with their own specialities. This means that operators must construct hybrid systems involving perhaps backhaul from one vendor and local distribution from another. Very local "Hotspot" wireless distribution equipment using the IEEE 802.11 standards present a number of problems to commercial operators. These unlicensed bands mean there will be interference problems, and the networking functionality of the equipment is not designed for commercial broadband distribution. It is hoped that the IEEE 802.16 initiative will give rise to a coherent range of FWA equipment covering all these niche markets. A uniform set of standards could allow multi-vendor sourcing of equipment and permit operators to offer a more sophisticated set of services to all levels of customers.