In-home propagation
Final report

Ofcom

2216/IHP/R/2/1.4

30th June 2011

Table of Contents

1 SUMMARY

This document is the final report for the study "In-home propagation measurements" carried out for Ofcom by a consortium of Aegis Systems, the STFC Rutherford Appleton Laboratory, Signal Science Ltd and dB Spectrum Ltd.

The primary aim of the work was to determine the usefulness of a range of different frequencies between 500 MHz and 5 GHz for providing in-home wireless coverage. The study was intended to be application-neutral (e.g. the network might be used for audio/video distribution, or for computer networking), but used the 802.11n wireless networking standard as a representative example. The overall requirement was to be able to make statements of the form "*X% of UK homes would be adequately covered using a frequency of Y GHz*".

Given the rise in spectrum occupancy at 2.4 GHz, and the potential for congestion, it was felt that Ofcom should investigate whether the allocations at 5 GHz provide a practical alternative for in-home networking, or whether alternative allocations may be required in the future. Frequencies in the band 500–800 MHz are relevant to this study as they will form part of the 'Digital Dividend' of spectrum released in the transition to digital terrestrial TV broadcasting. These frequencies are also of interest in the context of the growth of interest in 'White Space' devices.

It was concluded that the most promising study approach would be to make measurements of (i) path loss using CW transmissions at the four frequencies and (ii) the throughput obtained using representative 802.11 equipment at 2.4 and 5.7 GHz. If a clear correlation could be demonstrated between the statistics of path loss and throughput at the higher frequencies, this would allow extrapolation of the results to the 500 MHz and 800 MHz cases. Ofcom were anxious that measurements made should use techniques broadly representative of available consumer technology.

Measurements were carried out, in the winter of 2010/2011, in a selection of houses broadly representative of the UK housing stock. The CW measurements were made simultaneously at four frequencies, with continuous logging of signals received at a 'hub location' from a transmitter carried around each room of interest. Measurements of throughput were made using standard 802.11n equipment. Significant problems were encountered in configuring 802.11n equipment to run at anything close to the rates expected; these problems were largely due to device driver issues, and were only resolved when an integrated solution in the form of a high-end laptop was adopted.

In the course of the measurements, the following observations were made:

- Coverage, defined either in terms of 'points where connection was possible' or 'throughput speed attained' was invariably better at 2.4 GHz than at 5.7 GHz.

- The statistics of CW signal variation within rooms were similar at all four frequencies.
- No statistical difference was observed between the coverage to points around the room perimeter and the coverage within the entire room.
- In general, there was no significant difference in coverage with interior doors open or closed, although one house did exhibit a slight excess loss at 5 GHz with doors closed.

The throughput measurements alone, when weighted according to the UK distribution of the house types in which they were made, allow the coverage of homes at 2.4 GHz and 5 GHz to be estimated.

The results of the throughput and path-loss measurements showed a clear correlation, broadly similar at both 802.11n frequencies. This allowed the results for the two higher frequencies to be extrapolated to 500 and 800 MHz.

Figure 1.1 shows the percentage coverage within UK homes that might be obtained using 2x2 MIMO techniques in a 20 MHz bandwidth at each of the four frequencies. These coverage figures assume that the same system configuration (including transmitted power) is used at all frequencies, and that adequate spectrum is available.

Figure 1.1: Predicted coverage, UK weighted average, MIMO extrapolated

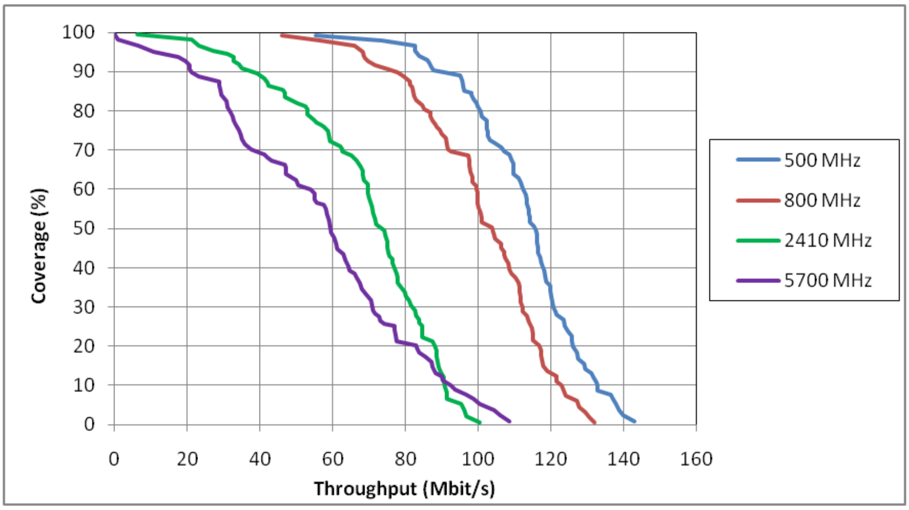

The concept of ‘coverage’ is not simple, and the discussion in Section 9, particularly Section 9.1, should be consulted in interpreting Figure 1.1.

2 INTRODUCTION

The Ofcom 'Mini-competition request' M/C 049 "*In Home Propagation Measurements*" sought to *"... determine the usefulness of a range of different frequencies for providing in-home coverage to UK homes* ..." Frequencies of interest included 500 MHz, 800 MHz, 2.4 GHz and 5 GHz.

Given the rise in spectrum occupancy at 2.4 GHz, and the potential for congestion, it was felt that Ofcom should investigate whether the allocations at 5 GHz provide a practical alternative for in-home networking, or whether alternative allocations may be required in the future. Frequencies in the band 500–800 MHz are relevant to this study as they will form part of the 'Digital Dividend' of spectrum released in the transition to digital terrestrial TV broadcasting. These frequencies are also of interest in the context of the growth of interest in 'White Space' devices.

The primary aim of the work was to determine the usefulness of a range of different frequencies between 500 MHz and 5 GHz for providing in-home wireless coverage. The overall requirement was to be able to make statements of the form "*X% of UK homes would be adequately covered using a frequency of Y GHz*".

The study was intended to be application-neutral (e.g. the network might be used for audio/video distribution, or for computer networking), but used the 802.11n wireless networking standard as a representative example. A brief overview of the 802.11 family of standards is given in Annex B.

3 PREVIOUS STUDIES AND MODELS

3.1 Empirical studies

Given the popularity and wide deployment of wireless LAN systems in domestic environments, there have been surprisingly few empirical studies that have considered coverage, throughput or propagation in residential buildings. The best characterised environments, by far, are university engineering departments and industrial laboratories [Mikas, 2008], [Cheung, 2002].

One exception is reported in [Cheung, 2002], where measurements were made at 2.4 GHz and 5 GHz in a townhouse in Oregon, with dimensions comparable to some of the houses examined in this study. It was found that line-of-sight paths showed a path-loss exponent slightly less than 2.0 at both frequencies, but that the difference in loss between the two frequencies was less than theoretically expected, perhaps due to higher levels of reflected power at the higher frequency[1]. For non-line-of-sight paths, the difference between the two frequencies was more stark, a best-fit exponent of 3.7 and 4.6 respectively.

[1] See Chapter 8 for a discussion of 'Room Gain'

3.2 Path-loss models (CW)

The measurements and modelling in this study are primarily concerned with determining the throughput achieved by representative wireless LAN equipment in a variety of domestic environments. The main output required by the study, i.e. coverage statistics, could be derived directly from such observations if information was only required for the 2.4 and 5.7 GHz bands.

The study, however, is required to provide data on the equivalent coverage that would be achieved at lower frequencies. If it can be shown that there is a simple relationship between path loss and throughput, and that this relationship is frequency-independent, it will be possible to extrapolate the measured coverage statistics to the lower frequencies. To do so will require an understanding of path loss at different frequencies, through modelling or measurement. Existing data on indoor path loss is summarised here.

3.2.1 Channel models of the 802.11n project

A series of six representative channel models for wireless LAN systems were developed in the course of the 802.11n standardisation process. While most of the development focussed on wideband representations for the MIMO channel, path-loss models were also developed [Perahia, 2008].

These models are all extremely simple, consisting as they do of an assumption of free-space propagation (d^2 law) up to a given ‘breakpoint’ distance, with a more rapid increase of loss ($d^{3.5}$) thereafter. In the two ‘residential’ models the breakpoint is set at 5 m. The assumed standard deviation of shadow fading (location variability) is 3 dB before the breakpoint and either 4 or 5 dB beyond it.

3.2.2 ITU-R Recommendation P.1238

ITU-R Recommendation P.1238 presents a rather incoherent set of information regarding indoor path loss. A simple empirical expression is given which includes a distance exponent and a floor penetration loss factor.

Empirical distance exponents are tabulated for a few environments (office, residential, commercial) and for different frequencies; the table is, however, only partly populated. In particular (and surprisingly) there are no entries covering the 2.4 GHz band, and the 5.2 GHz exponent (3.1) is given only for the ‘office’ environment.

The table for floor penetration loss factors is also sparsely populated, and for 5.2 GHz only includes a figure of 16 dB for a single floor penetration.

It is noted that shadow fading follows log-normal statistics with a 12 dB standard deviation at 5.2 GHz in office environments. It is interesting that this is much greater than the value assumed in the 802.11n models.

Overall, this Recommendation is too simplistic and contains too little empirical data to be of much practical use.

3.2.3 COST 231

The ‘Multi Wall Model’ developed in the COST 231 project sought to avoid the tendency of earlier models to overestimate path loss by assuming a constant multiplicative attenuation with each wall or floor crossing. In practice, when attenuation of the ‘direct’ path reaches a certain value, lower-loss indirect paths tend to become available. The COST231 model thus rolls-off partition-transition losses as illustrated in Figure 3.1.

Figure 3.1: COST 231 roll-off to limit shielding due to multiple floors

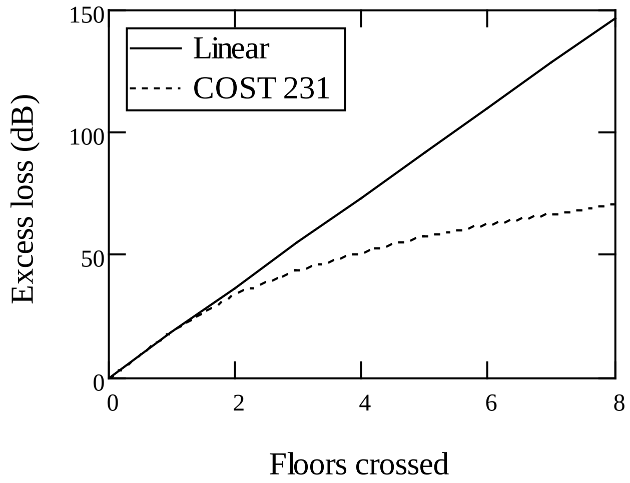

3.3 Wideband models

Prior to the measurement campaign, it was anticipated that the relationship between path loss and throughput might vary with frequency, perhaps on account of variations in the multipath channel and the degree of correlation in the MIMO channel matrix. In the event, a surprising degree of uniformity was found between frequencies (see section 7 below). This was convenient, as it supported the use of a straightforward extrapolation to derive throughput figures for the lower frequencies.

Had this not been the case, it would have been necessary to consider the characteristics of the wideband channel explicitly, and a brief review of relevant models is included here.

3.3.1 The Turin / Suzuki ‘Δ-K’ model

One of the first models for the wideband channel was that of Turin [Turin, 72], who suggested a Poisson arrival rate for the components, with a lognormal amplitude distribution.

The fit to the Poisson distribution is fair for insensitive receivers (more powerful multipath only is observed) but fails for more sensitive observations. In practice, the pattern in location of the scatterers in a building, or urban area, ensures a deviation from the standard Poisson model.

The discrepancies between the pure Poisson model, and experimental observation were noted by Turin, who tentatively suggested a remedy in the form of a two state model. Noting the poor fit of such a model to the observed 'clustering' of arrivals, he proposed, in a footnote, a modified Poisson arrival model (the Δ-K model) in which the arrival rate toggles between two values based on the occupancy, or otherwise, of the previous 'bin'.

Figure 3.2: The Turin / Suzuki 'Δ-K' model

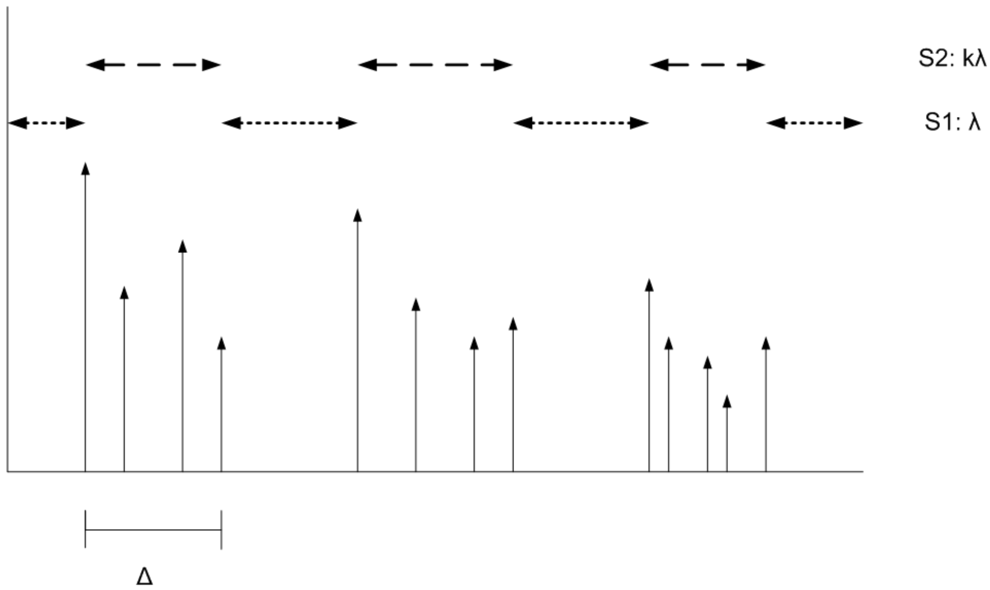

The proposal by Turin to use a two-state model better to represent component arrival statistics was developed by Suzuki [Suzuki 1977], also in the context of outdoor, urban, channel modelling. This model takes into account the clustering of path arrival times by toggling between two states, associated with different arrival rates. A transition to the 'higher' (or 'lower') state is triggered by a single arrival, after which the state is maintained for a period Δ seconds (see Figure 3.2 above).

3.3.2 The Salah & Valenzuela 'double exponential' model

An alternative approach, which also accounts for the clustering of arrivals, was proposed by Saleh & Valenzuela [Saleh, 87], in the light of a series of indoor measurements. In this model, separate arrival and decay rates are used for the clusters, and for the rays within a cluster. While the peaks, or initial ray-arrivals, of each cluster may be predicted with fair accuracy from ray-trace models, the detail of the subsequent rays in each cluster will be dependent on the detail of terminal clutter, and cannot readily be predicted.

Figure 3.3: The Saleh & Valenzuela model

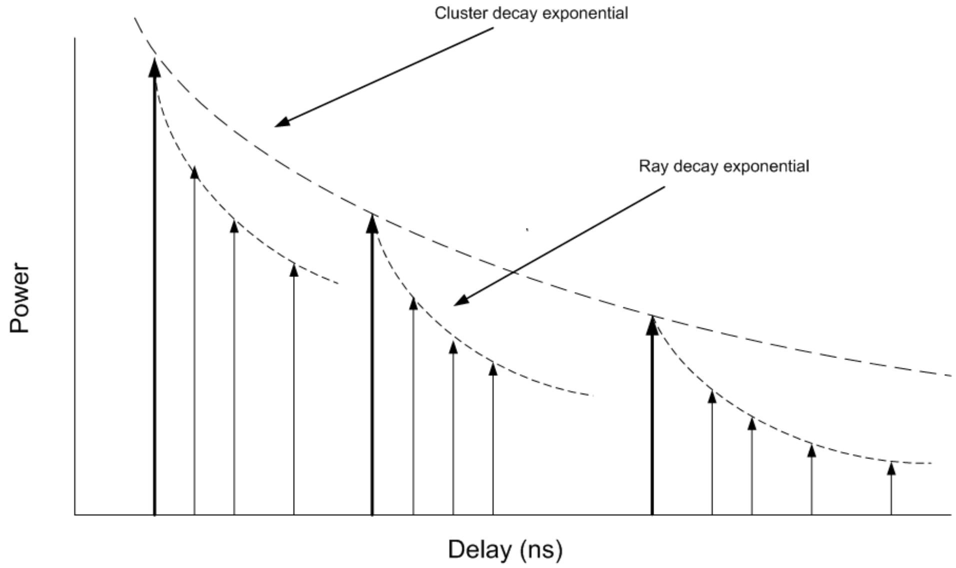

The approach of Saleh & Valenzuela employs a double-Poisson model in which cluster centres follow a Poisson distribution (i.e. have exponentially distributed inter-arrival times), and individual components within each cluster are also Poisson distributed. The two Poisson processes will generally have different arrival rates, as indicated in Figure 3.3 above. Based on some 200 impulse response in an office building, using a sounder with a temporal resolution of ~5 ns, estimated these arrival rates as Λ = 1/300 ns for the clusters and λ = 1/5 ns for the rays. It was also noted that, in half of the measurement locations, only one cluster was observed.

3.3.3 802.11 channel models

A set of six reference channel models was created from empirical data as part of the 802.11n development process, with values of delay spread between 0 and 150 ns. The impulse response models are based on the Saleh & Valenzuela approach with powers and angles of arrival specified for each tap of each cluster. The 'residential' model, for instance, has twelve taps in two clusters, giving an RMS delay spread of 15 ns.

For distances up to the breakpoint (5 m for the residential model) the channel is assumed to be Ricean, and Rayleigh beyond this.

For a MIMO channel, the taps are defined for all elements in the channel matrix, with a correlation dependant on the relative angles of arrival, wavelength and antenna element spacing.

4 BUILDING TYPES

The object of the study reported here was to provide information along the lines of "*50% of UK homes would be adequately covered at 5 GHz*". To make such assertions will require *"… databases of known distributions of homes across the UK …"*.

4.1 Building stock

A useful overview of the domestic building stock in England and Wales is provided within the data gathered in the census of 2001[2]. Data extracted from Table KS16 is summarised in Table 4.1.

Table 4.1: Building type distribution (2001 census)

Building type	Percentage of household spaces[3]
House or bungalow: detached	22.8
House or bungalow: semi-detached	31.6
House or bungalow: terraced	26.0
Flats: purpose-built	13.6
Flats: converted & commercial buildings	5.5
Mobile homes, etc.	0.4

No further information on building type, sizes, construction or age is captured by the census.

A much more detailed set of information is contained in the English Housing Survey (EHS), which samples a limited number of households in some detail, and extrapolates the results to derive statistics for England as whole. The latest survey is that dated 2008 [EHS, 2008], which incorporates data gathered between April 2007 and March 2009. The remit of the survey extends beyond physical stock, to include issues of security, utility services, stock condition repair costs and attitudes to housing.

[2] Available from http://www.statistics.gov.uk/census2001.

[3] 'A 'household space' is defined for census and other Government statistical purposes as a 'household's accommodation' (which may be within a shared dwelling).The term 'percentage of household spaces' does not imply any weighting by size.

Table 4.2: Housing stock data from EHS 2008

Dwelling type	**Floor area (m^2)**	**Percentage of sample**
End terrace	*83*	*9.9*
Mid terrace	*82*	*18.7*
Small terraced house	*59*	*9.8*
Medium/large terraced house	*94*	*18.8*
All terrace	82	28.6
Semi-detached	92	26.0
Detached house	147	17.4
Bungalow	76	9.4
Converted flat	69	3.7
Purpose-built flat, low rise	56	13.4
Purpose-built flat, high rise	59	1.5

Note 1: terraced houses are broken down both by position and by size. These are not four separate categories.

Note 2: Floor areas are average across samples, and do not include conservatories (which have an average area of 10 m^2.)

Figure 4.1: Dwelling type statistics (from EHS, 2008)

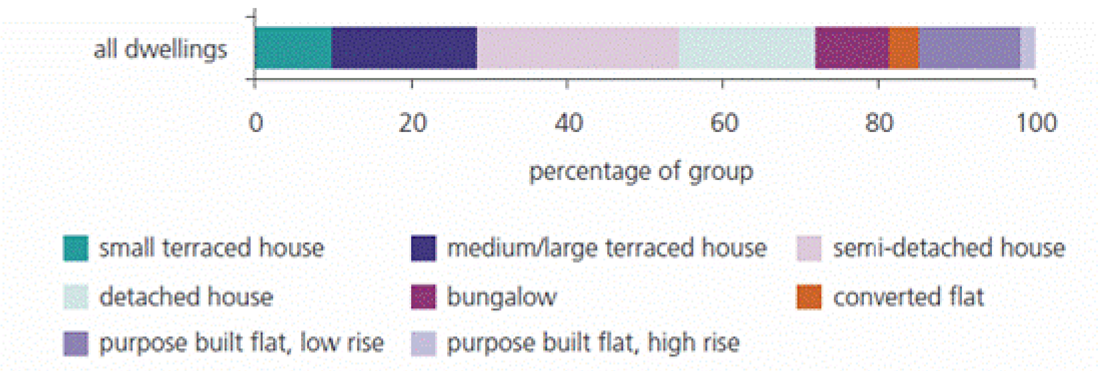

An attempt may be made to compare the data from the two sources; the main difficulty is that bungalows are a separate category in the EHS. If it is assumed that bungalows are distributed fairly evenly between the 'detached' and 'semi-detached' categories[4], this would give a reasonable correspondence between the two sets of data.

It should also be borne in mind that the census applies to England and Wales, while the EHS covers England only.

Table 4.3: Simplified housing stock data based on EHS 2008

Building type	Census (% of household spaces)	EHS (% of sample)
House or bungalow: detached	22.8	17.4
House or bungalow: semi-detached	31.6	26.0
House or bungalow: terraced	26.0	28.6
Flats: purpose-built	13.6	14.9
Flats: converted & commercial buildings	5.5	3.7
Mobile homes, etc.	0.4	n/a

4.2 Apportionment of measured coverage statistics

It has been possible to investigate only a limited sample of houses within this study. To derive statistics that are plausibly applicable to the UK housing stock as a whole, it will be necessary to weight the results from each category appropriately.

A correspondence is proposed in Table 4.4 between the EHS categories and the sample buildings measured in the study. Flats have not been measured within this work, as it was felt, on the basis of measurements in small houses, that little useful data would be gained as coverage would be achieved at all locations.

It can be seen from Table 4.4 that there is a reasonable correspondence between the floor areas of the houses measured in the study, and the representative building types with which they are associated, though this is less true for the unusually large bungalow and the smaller-than-average detached house.

[4] Anecdotally this seems unlikely to be the case; most bungalows appear to be detached. Furthermore, the inclusion of bungalows in the 'terraced' category seems surprising!

Table 4.4: Housing stock data from EHS 2008

Dwelling type	Average Floor area (m^2)	Measurement location	Measurement location floor area (m^2)	% weighting
Small terraced	59	MT	66	9.8
Medium/large terraced	94	VT	91	18.8
Semi-detached	92	SD & LSD	107 & 180	26.0
Detached	147	MD	83	17.4
Bungalow	76	MB	165	9.4
Flats	59	BF	14	18.6

5 CW MEASUREMENT METHOD

5.1 Introduction

The general approach is to explore the statistics of path loss within a building using a portable transmitter, and a fixed receiver located in one or more positions representative of typical access point locations.

Test and development licences were obtained for the following frequencies, all at a maximum of 200 mW eirp.

- 506.0 MHz
- 802.0 MHz
- 2410.0 MHz
- 5760.9 MHz

These frequencies were chosen to avoid interference with local TV transmissions in Sussex and Oxfordshire.

5.2 Transmitters

An arrangement has been adopted in which transmissions are made simultaneously on four frequencies from a single handheld antenna unit, fed from transmitters carried in a shoulder bag.

Figure 5.1: Transmitter unit

The transmitters themselves are powered by lightweight Lithium-ion polymer (LiPo) cells, which give over an hour of continuous operation.

Figure 5.2: Transmit antenna unit

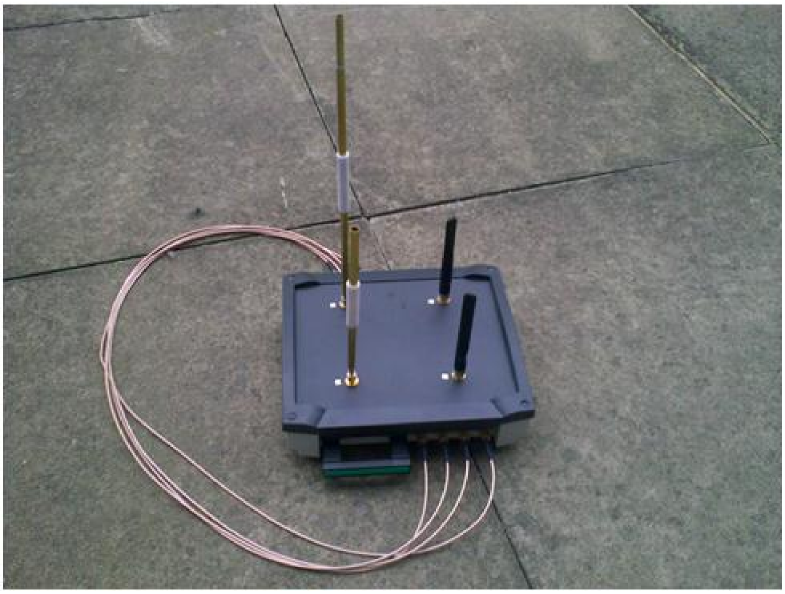

The antenna unit consists of four dipoles, two of which (for 2.4 GHz and 5 GHz) are commercial items while those for the lower frequencies were built at RAL.

It was essential, for reasons of measurement speed and comparability between frequencies, to measure simultaneously on the four frequencies. The disadvantage of this is that it is impossible to avoid interaction between the four antennas which, together with practical limitations on the size of the ground plane implies that the radiation patterns will not be uniform. In practice, it was found that the variation in the horizontal radiation pattern was averaged by the semi-random movement of the measurement personnel within the buildings.

5.3 Receivers

Signals from the transmitters are logged simultaneously by a system of four Rohde & Schwarz EB200 measuring receivers. As these have an upper frequency limit of 3 GHz, a down-converter has been constructed by RAL for the 5 GHz band measurements, and this is integrated with the receivers in a 19” rack unit.

Figure 5.3: Receiver rack with down-converter and Ethernet switch

The receivers are controlled over an Ethernet interface by a PC running Aegis Systems ‘Logger’ software.

5.4 Receive antenna system

The receive antenna system is constructed as a cluster of dipoles, identical to those used at the transmitter, but mounted on a tripod instead of being handheld.

The interaction between the four dipoles, and limited size of the ground plane give rise to variations in the azimuth pattern as noted in Section 5.2. In the case of the receive antennas, however, there is no movement of personnel to average these effects, and it is therefore possible that a radiation pattern null (or peak) might happen to be directed towards an area being explored by the transmitter.

Two approaches were investigated for the avoidance of such effects.

The more elegant approach made use of a discone antenna, designed to have a bandwidth sufficient to cover all the measurement frequencies. The output from this antenna was fed to the receivers using a wideband amplifier and splitter. Simple considerations of physical symmetry would ensure that the same horizontal pattern would be obtained at all frequencies.

Figure 5.4: Experimental discone antenna

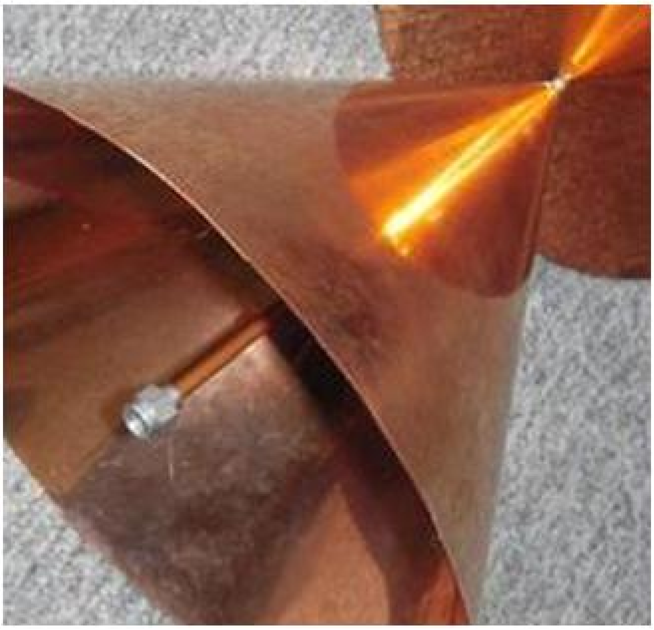

Unfortunately, time available within the project did not allow for the absolute calibration of this antenna, nor for the vertical radiation pattern at the different frequencies to be investigated.

An alternative arrangement was therefore used for the in-home measurements, in which the receive antenna cluster was mounted on a stepper motor, which was configured to rotate the antennas continuously during measurements, oscillating the cluster backwards and forwards through 180°.

Figure 5.5: Stepper motor arrangement for antenna oscillation

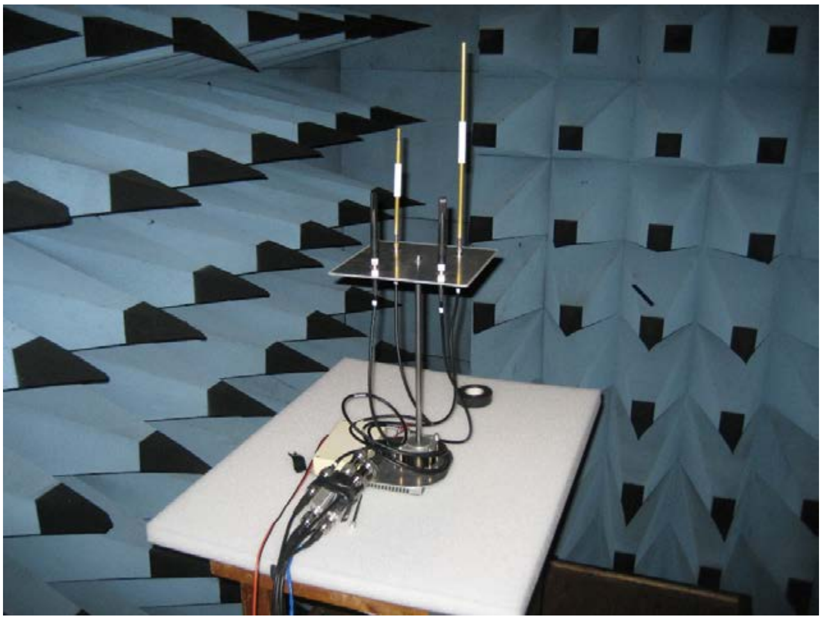

This arrangement was found to average the pattern satisfactorily, and to allow repeatable loss measurements to be made.

Annex C describes the calibration of the CW measuring equipment.

6 THROUGHPUT MEASUREMENT METHOD

In parallel with the CW measurements described above, measurements of throughput were obtained using commercially-available WLAN equipment.

Whereas the CW measurements were made on a continuous basis, with the aim of covering most of the floor area of the house, it was necessary to make the throughput measurements in a few, discrete, locations. This is partly for reasons of practicality (the measurement method adopted requires a significant period of time to make a single throughput measurement and partly for realism, as the 802.11n system would not normally be used for connections to mobile terminals (though this is perfectly possible, and may become the more common scenario in the future).

6.1 Hardware

The key hardware element in the throughput trials is a Cisco Aironet 1252AG dual band 802.11n access point, which allows MIMO operation at both 2.4 and 5 GHz. This device is used a reference in the interoperability testing work of the Wi-Fi alliance, and allows connection via a Gigabit Ethernet port.

Fig.6.1: Cisco 1252AG dual-band 802.11n

This unit is highly configurable, and allows channel-bonding and 3 x 2 MIMO operation.

It was initially intended to use an external USB transceiver as the client device for the throughput measurements. An apparently suitable device, the Ubiquiti SR-71-USB was obtained; this device supports 2 x 2 MIMO operation in both bands and has external ports which would allow the use of calibrated antennas, and the exploration of different antenna spacings.

Fig.6.2: Ubiquiti SR-71-USB

Testing revealed, however, that the maximum throughput rates obtained were very much lower than would be expected. A variety of other external USB ‘dongle’ transceivers were evaluated, but, in all cases, it was found that the throughput obtained was poor, or very variable.

It was not possible, in the time available, to diagnose the exact problems with these units, but it appears that there were incompatibilities between the firmware and device drivers of the transceivers and those of the laptop PCs with which they were used.

The final measurement arrangement therefore made use of an integrated solution, in the form of a Toshiba ‘Portege’ R700 15X laptop, which includes a dual band 802.11n transceiver. This provided reliable connectivity at the expected rates.

One disadvantage of this arrangement is that the specification and location of the transceiver antennas is beyond the control of the project team. The antenna elements used in the laptop are flat patches, size 27 x 15 mm, located behind the upper part of the LCD screen. The position of the two antenna elements is indicated in Figure 6.3, below. No information is available regarding the gain or radiation pattern of the elements.

Figure 6.3: Showing the location of 802.11n antenna on client laptop screen

The antenna patches are separated by 103 mm, which corresponds to 0.8λ at 2.4 GHz and 2.0λ at 5.7 GHz.

6.2 Method

The Aegis Systems 'Logger' software is a general purpose tool originally intended to provide an interface to measuring instruments via serial, Ethernet or GPIB connections. The software has been extended to allow measurements to be made of a wide variety of computer networking characteristics.

One option available in Logger is to use FTP to transfer large quantities of data across a network, recording the bit rate achieved in the process. Each test consists of ten repeats of a trial in which ten files are copied from the source to the destination computer. The files are not, however, written to the destination computer—the data is thrown away as soon as it arrives to keep the transfer rate as high as possible. The smallest files (ten of which are transferred in each of trials 1-10) are 1 kilobyte large, and the file size increases by a factor of ten after every tenth trial.

Two data sets are plotted on the graphs below:

- The 'user' throughput, which measures the data transfer rate based on the complete FTP transaction including setting up FTP connections, issuing file transfer requests and changing to the right folder on the remote machine. This data set is plotted in red and has a general 'staircase' appearance as the file size increases, indicating improving efficiency with increasing file size.

- The ‘file transfer’ throughput, which measures the rate at which file data is transferred (in chunks of 32 kB). This is closer to the ‘wire speed’ and surprisingly reliable given the approximate nature of PC clocks and the relatively short intervals involved. This data set is plotted in green and has, ideally, a constant value close to the maximum rate.

The measurements from transfers of large files are generally, but not always, the highest and the most reliable. In the example shown in Figure 6.4, a Mac Mini is used as the FTP server and is connected to the logging machine via the gigabit Ethernet port, a Gigabit Ethernet switch (Sitecom LN-120) and Cat-6 (Gigabit) patch cables. The size of files being transferred increases in decade s from 1 kB to 100 MB.

It can be seen that for the largest file sizes a throughput of around 300 Mbps is achieved, and this will be sufficient to determine the extent to which an 802.11n WLAN, configured as a 2x2 MIMO system in a 20 MHz channel, is constraining throughput. In earlier tests, it was found that the maximum wired throughput was only around 110–120 Mbps, which is comparable to the maximum expected from the 802.11n connection. The exact type of cable, switch and hard disk drives were found to be critical.

Figure 6.4: Throughput test over Gigabit Ethernet

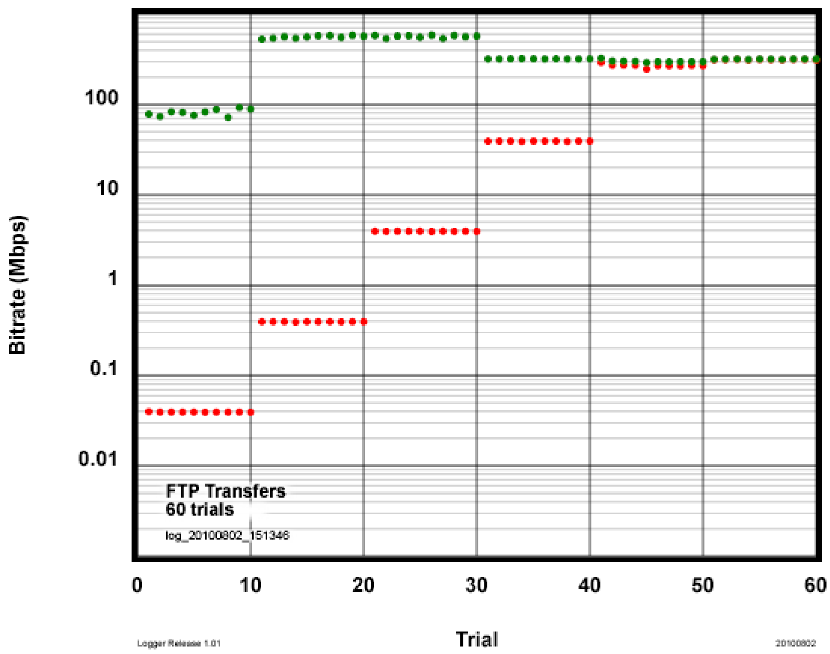

When the FTP approach was used over wireless connections it was found that the bit rates achieved were rather poor, even taking into account the expected overheads of the MAC layer.

Throughput measurements made using FTP are somewhat lengthy, owing to the overheads of the protocol and the elapsed time taken to transfer sufficient data to give reliable and repeatable statistics. Lower overheads should be associated with TCP, and the Windows version of the well-known Unix tool 'Test TCP' (TTCP), which is available from Microsoft as NTTCP was compiled and incorporated into the 'Logger' suite. In addition, the 'Netperf' utility has been incorporated in the Aegis software, although the operation of this widely-used benchmarking tool is more opaque than in the case of NTTTCP.

Both of these tools allowed throughput measurements to be made more rapidly. The results have, however, proved very variable, with widely different performance[5] being obtained from different client devices or configurations. Owing to the time constraints of the project, it was concluded that slower, but reliable and repeatable FTP measurement was to be preferred over the potentially faster, but unreliable and less well-documented alternatives.

Practical experience during the initial stages of the measurement campaign showed that reliable and repeatable results could still be obtained if the maximum size of file transferred was limited to 10 Megabytes, and that it was only necessary to make five trials at this size (i.e. transfer a total of 500 MB).

7 MEASUREMENT RESULTS

Details of all the measurement locations are given in Annex D, and a summary of all measurements in Annexes E and F. The raw and reduced measured data have been made available to Ofcom as XML files and Microsoft Excel work books respectively.

This section gives an overview of some of the findings from the measurements.

7.1 CW measurements

For the CW measurements, the transmitter units were carried around the bedroom or reception room in a regular pattern intended to explore the area of the room uniformly and without undue overlap. Some runs were also carried out in which only the perimeter of the room was explored, as this is the most likely location for fixed wireless device (printers, desktop PCs, multimedia devices, etc.), with the intention of determining whether there is any systematic difference in coverage with respect to the 'whole room' statistic.

[5] In some cases an actual throughput of around 4 Mbps was achieved over a link with a reported PHY layer speed of 300 Mbps!

In all the figures below, the following colour-coding is used:

500 MHz

800 MHz

2.4 GHz

5 GHz

Results have been processed to generate four types of plot:

- **Time-series** data, showing transmission loss against time for the four frequencies.
- **Cumulative loss distribution**, showing probability with which given values of loss are exceeded. If it is assumed that, to a first approximation, devices operating at all frequencies would have similar radiated powers, antenna gain and noise figure, these plots give an indication of the coverage that might be expected at each frequency.
- **Normalised time-series** and **normalised cumulative loss** distributions. As above, but with the 20 log10 (frequency) term removed to indicate, for each frequency, the impact of losses in excess of free space due to diffraction, penetration loss, multipath, etc.

The full set of measurement results has been made available as an Excel workbook. The following sections summarise the findings.

7.1.1 Variation of transmission loss with frequency

As the frequency of operation increases, the isotropic[6] path loss seen for a given path will tend to increase. This is partly due to the fact that the 'effective aperture' of an isotropic antenna becomes smaller as wavelength decreases and so less of the transmitted power flux is intercepted and, secondly, because waves of shorter wavelength will tend to suffer higher diffraction and absorption losses.

A measurement made in the 'Modern Bungalow' (main bedroom to living room) shows this trend clearly (Figure 7.1).

[6] The loss between antennas of 0 dBi gain.

Figure 7.1a: Path-loss statistics, bedroom to living room (Run 11)

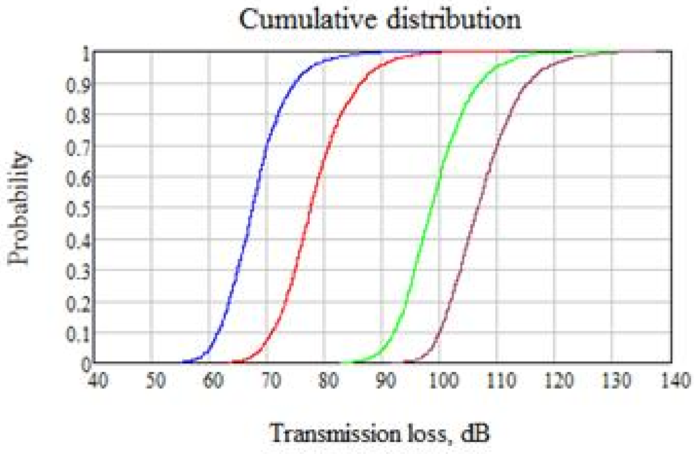

Figure 7.1b: Normalised path-loss statistics, bedroom to living room (Run 11)

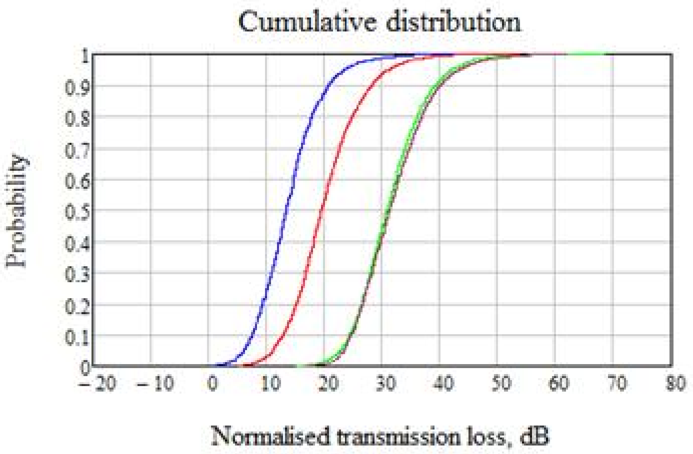

In the first figure, it can be seen that the median loss at 5 GHz is some 40 dB greater than that at 500 MHz, and that the loss increases monotonically with frequency. The second figure, in which the 20 log(f) dependence is removed, shows that the *excess* loss (diffraction and absorption) increases by around 7 dB between 500 MHz and 800 MHz, and by a further 11 dB to 2.4 GHz. There is, however, no increase in the excess loss between 2.4 GHz and 5 GHz.

7.1.2 Interdecile spread of measurements

At the start of this project it was expected that there might be significant differences in the multipath environment seen at the different frequencies; for example, it might be the case that coverage at the lower frequencies is provided by coupling along a direct path between terminals (through intervening partitions), while at higher frequencies propagation might be via multiple reflections.

Any such difference in propagation modes between frequencies would have implications for network performance, particularly in terms of the diversity gain and MIMO advantage that may be available.

As the measurements did not make use of channel-sounding techniques, it was not possible to assess the multipath environment directly, but difference between frequencies can be indirectly observed in the statistics of path loss gathered within each room. This variability has been characterised in terms of the interdecile spread of the path loss, and has been found to be surprisingly constant across the four frequencies, as evidenced by the shape of the four CDFs in Figure 7.1 above. The results are also very similar for the different houses.

The results for the 'Modern Detached' and 'Modern Terraced' are shown in Figures 7.2 to 7.4. In all cases the interdecile spread is around 15 dB.

Figure 7.2: Interdecile spread (MD)—hub in study

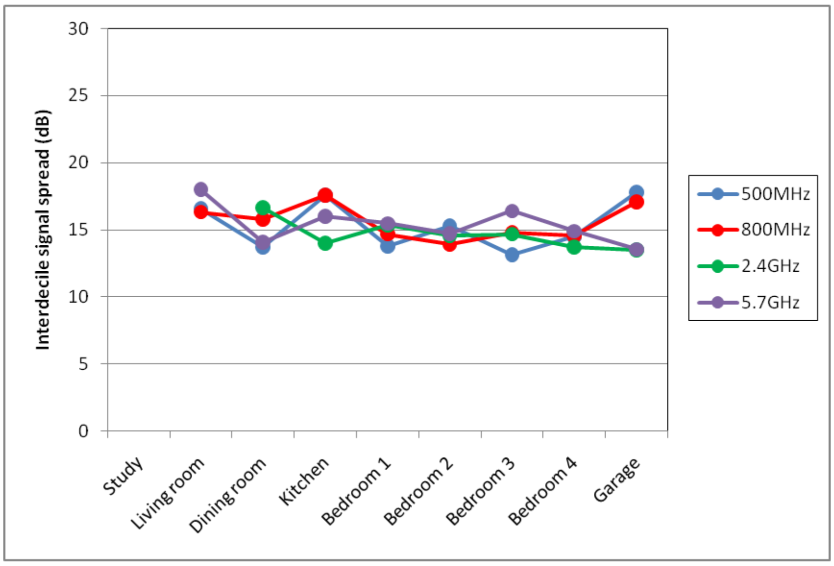

Figure 7.3: Interdecile spread (MD)—hub in hallway

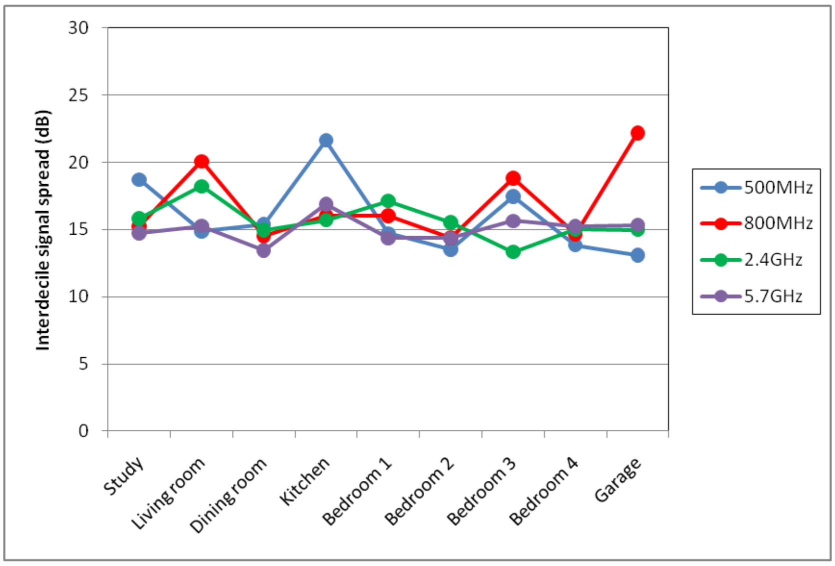

Figure 7.4: Interdecile spread (MT)

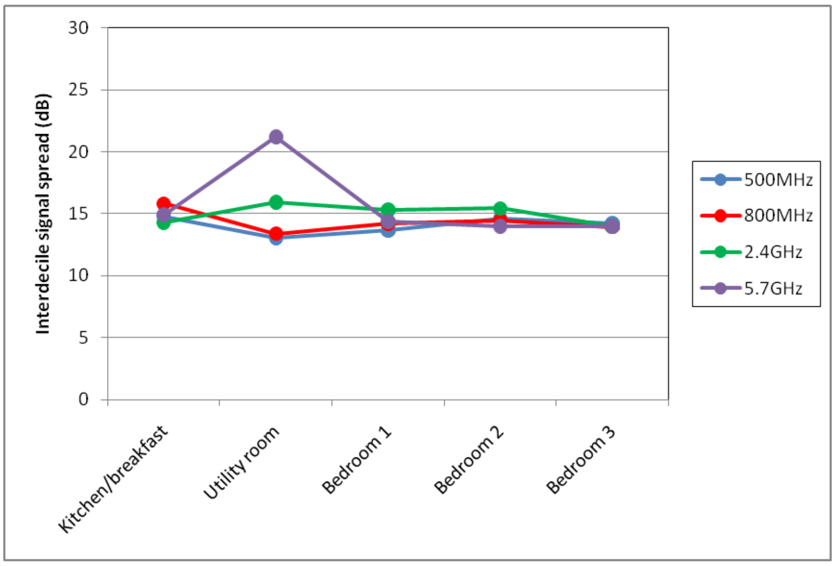

A little more variation is evident for the case of the Victorian Terrace in Figures 7.5a and 7.5 b below. The relatively high values for the kitchen and living room (measured from the hall) reflect the large percentage change in path length across these rooms, while measurements made in the back bedroom move from a line-of-sight to non-line-of-sight situation, a transition which affects the higher frequencies most strongly.

Figure 7.5a: Interdecile signal spread (VT)—hub in hallway

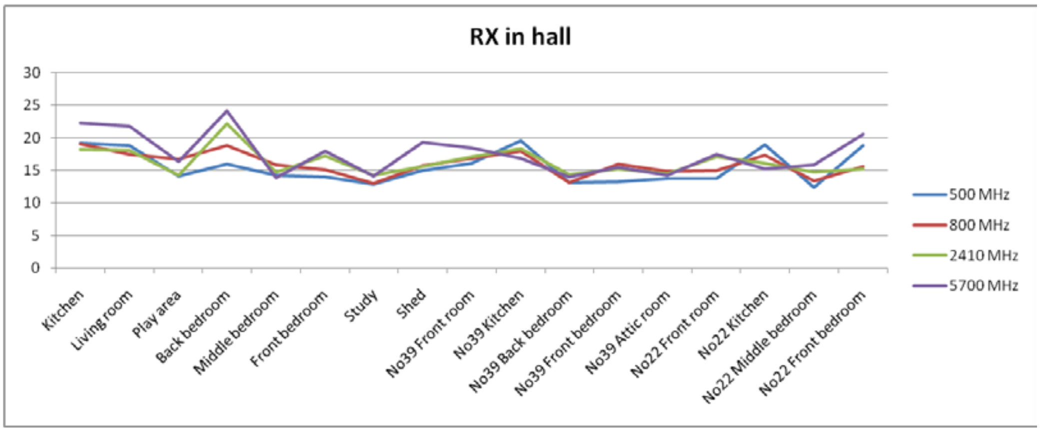

Figure 7.5b: Interdecile signal spread (VT)—hub in study

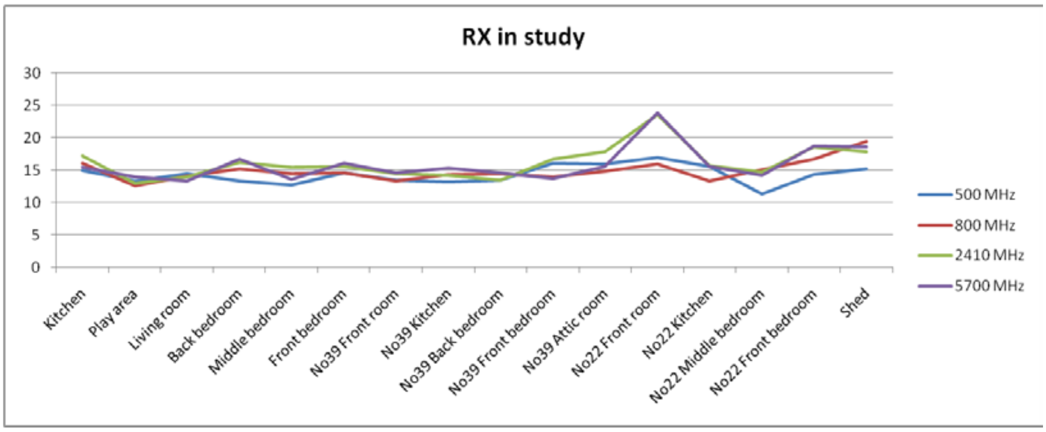

In general there is only a very slight, though monotonic, trend with frequency.

The apparently anomalous higher-frequency results for the front room at No.22 (diagonally across the road, see Fig.D.2, with receiver in VT study) have no obvious explanation. The individual measurements for this case are shown below.

Figure 7.6a: CDF of path loss, No 22 front room, RX in study

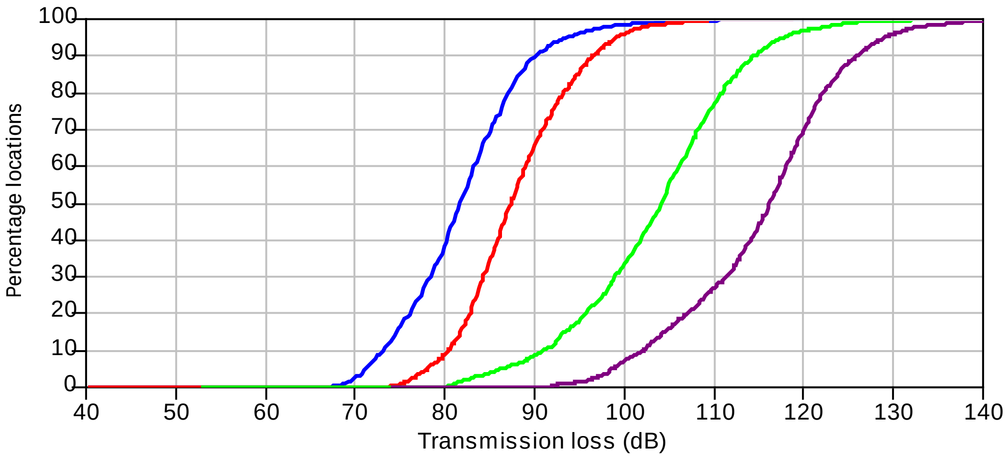

Figure 7.6b: Time-series of path loss, No 22 front room, RX in study

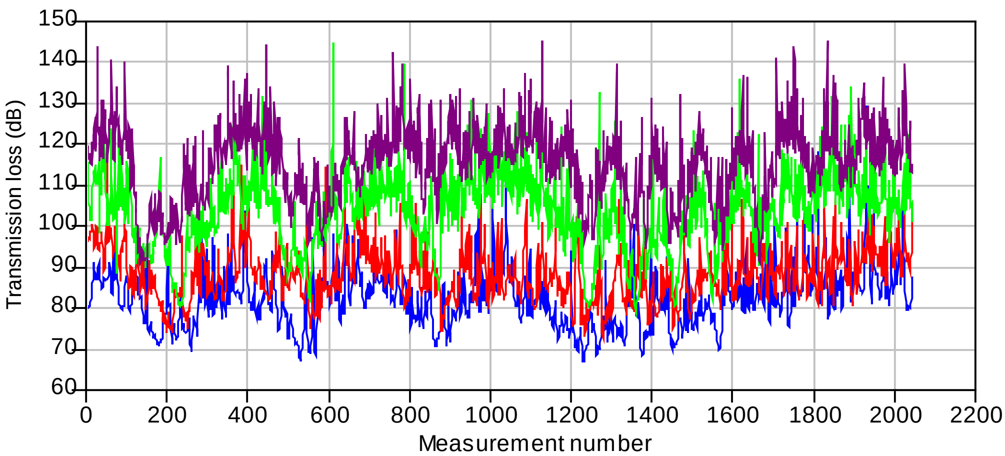

Although it appears that the measurements in the front room may be moving between quasi-LOS and diffracted regimes, the same would be expected to apply in the front bedroom, but in that case no anomaly is apparent (Fig. 7.7).

Figure 7.7: Time-series of path loss, No 22 front bedroom, RX in study

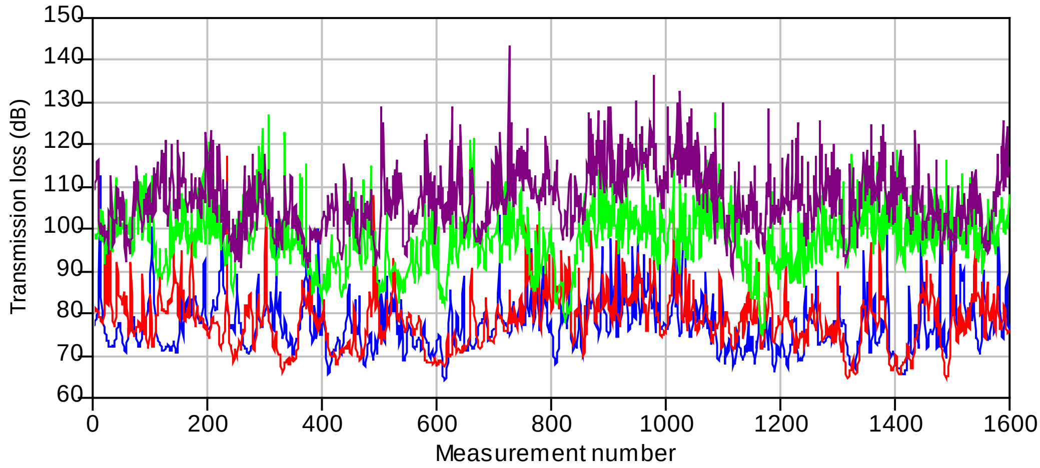

7.1.3 'Perimeter' versus 'overall' coverage

It has been suggested that access points and client devices will, statistically, tend to cluster around the edges of rooms, being located on desks, bookcases or other furniture (the use of laptops on sofas or chairs in the middle of rooms is a counter-example). It was, therefore, felt worthwhile to investigate whether there is any significant trend in the statistics of propagation to the room perimeter when compared to those for the room as a whole.

The statistics for the two cases are compared in Figure 7.8 and 7.9 for measurements made in the 'Modern Bungalow'. It can be seen that any divergence is marginal; similar results were obtained in other buildings.

Figure 7.8a: Path loss statistics, living room perimeter

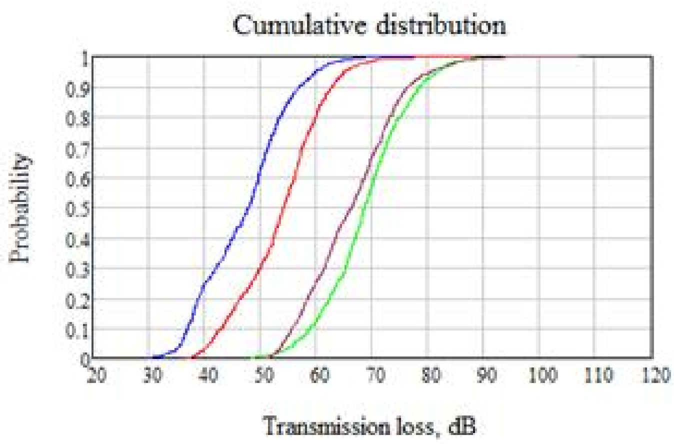

Figure 7.8b: Path loss statistics, living room

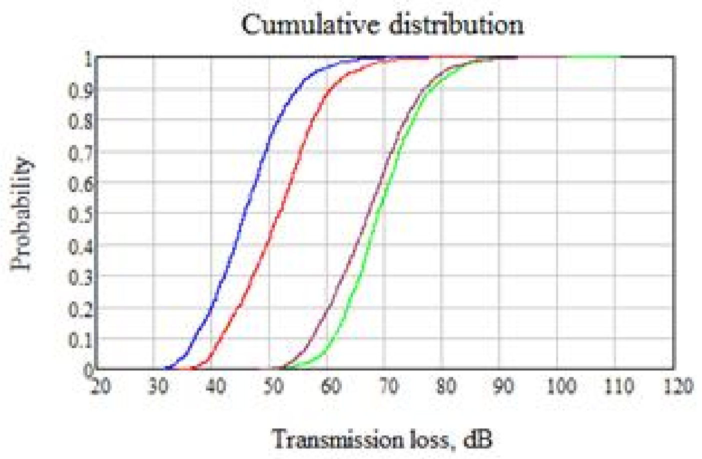

Figure 7.9a: Path loss statistics, main bedroom perimeter

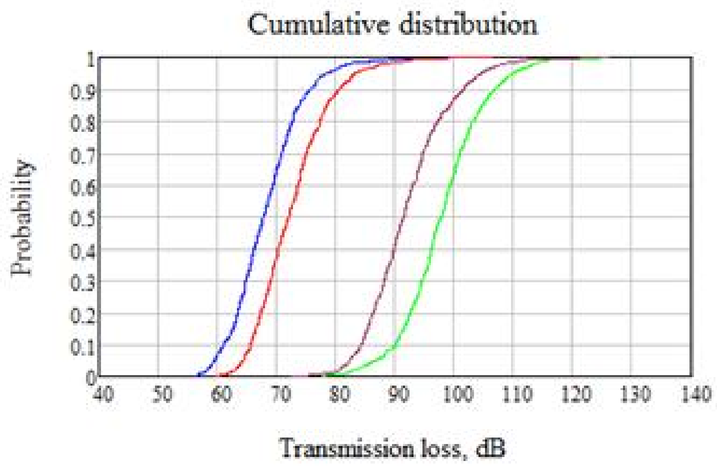

Figure 7.9b: Path loss statistics, main bedroom

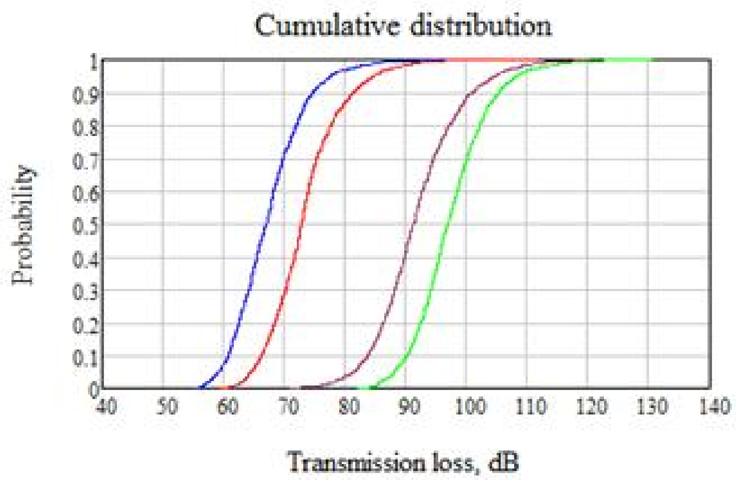

7.1.4 Repeatability of statistical results

Given unavoidably-uncontrolled aspects of the experimental method, in which the exact path taken by the transmitter will vary between runs, and in which the location and movement of other personnel is not strictly controlled, it was necessary to investigate the repeatability of the measurements. Figure 7.10 shows the statistics achieved from two separate measurement runs with the transmitter carried over similar, but non-identical, paths in the main bedroom ('Modern Bungalow').

Figure 7.10a: Path loss statistics, main bedroom

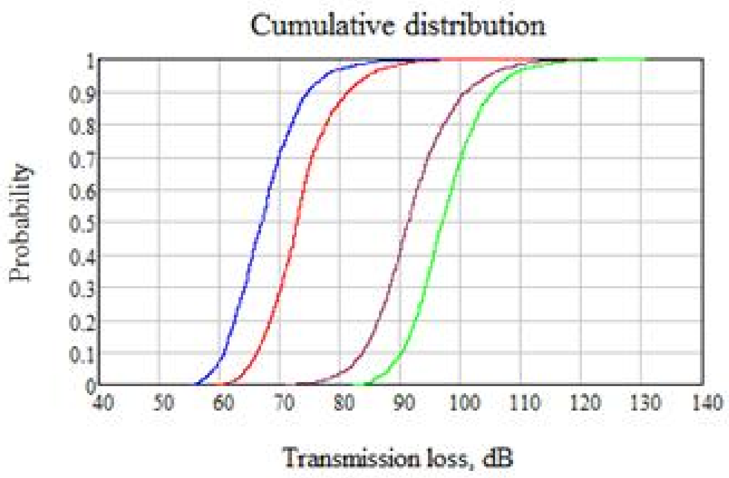

Figure 7.10b: Path loss statistics, main bedroom

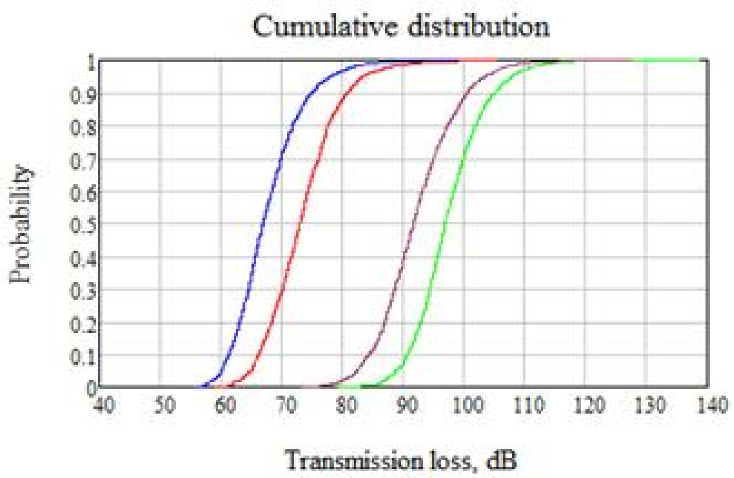

The path losses at median and decile points for the two runs are within 1 dB of each other, confirming that the sampling is adequate to ensure repeatability.

7.1.5 Impact of interior doors

In all the houses investigated, a number of measurements were made to examine the effect on path loss of the interior doors.

The results from the Victorian Terrace (VT), for a variety of rooms, and both hub positions, are given in Tables 7.1a and 7.1b, with cases of apparent 'door gain' highlighted.

Table 7.1a: Door attenuation—hub in VT hallway

RX in hallway	500 MHz 5	800 MHz 5	2410 MHz	5700 MHz
Living room	0.25	-0.25	1.3	1
Middle bedroom	-0.2	0.2	-0.2	0.1
Front bedroom	0.6	0.15	-0.2	-1
Play area	-1.3	-1.2	-0.2	-0.8
Back bedroom	-0.3	-1.4	0.4	0.5
Kitchen	1.5	0.2	1	4.2
Shed	-0.7	-1.4	0.2	0.4

Table 7.1b: Door attenuation—hub in VT study

RX in attic	500 MHz 5	800 MHz 5	2410 MHz	5700 MHz
Living room	0	-0.5	-0.5	-1.1
Middle bedroom	0.6	0.3	0.9	0.05
Front bedroom	-0.2	0.2	-0.2	0.1
Play area	1.25	-1.25	0.8	1.5
Back bedroom	0.25	-0.2	-0.1	0.3
Kitchen	-0.8	-0.2	-0.7	-0.1
Shed	0.9	-0.3	1	0

There appears to be no evidence of a systematic impact of path loss by the doors (which are of solid wood[7], except for the kitchen door, which has glass upper panes), and no obvious variation with frequency.

The term 'door attenuation' is, in any case, misleading, as there are likely to be several significant paths between the terminals, which may or may not pass through the doors.

Similar measurements were made in other houses with the results summarised in Tables 7.1 to 7.3.

[7] Pine, with frames ~3.3 cm thick and panels ~1.3 cm thick.

Table 7.2: Door attenuation—modern bungalow

Frequency	500 MHz	800 MHz	2.4 GHz	5.7 GHz
Mean	0.1	0.0	0.7	2.0
Std. Dev.	0.9	1.8	0.7	1.8

The doors in the bungalow are of solid pine construction and around 4 cm thick. This was the only house in which a clear trend for increased loss at the higher frequencies was seen, and where the loss at 5.7 GHz was significant.

Table 7.3: Door attenuation—modern detached

Frequency	500 MHz	800 MHz	2.4 GHz	5.7 GHz
Mean	0.4	0.2	0.6	0.3
Std. Dev.	0.9	0.4	0.9	1.2

The doors in this house (MD) were composite, and the downstairs room doors had vertical glass panes. For some individual measurements, particularly at 5.7 GHz, the loss was lower with the doors closed than with them open. Any difference appears to be small and not obviously frequency-dependant.

Table 7.4: Door attenuation—modern terraced

Frequency	500 MHz	800 MHz	2.4 GHz	5.7 GHz
Mean	−0.3	−0.3	0.6	0.5
Std. Dev.	0.7	1.0	0.2	1.5

The doors in the modern terraced (MT) were composite and again the difference between the doors closed and doors open cases is small. The mean “attenuation” is less than the variability between rooms.

It is concluded that the state of the internal doors generally has no significant impact on coverage at any frequency, although solid wood doors may increase path loss at 5 GHz in some cases.

7.1.6 Variability due to movement of people

A number of measurements were made in the ‘Modern Bungalow’ with the transmitter stationary in the bedroom, to assess the impact of movement by people within the building.

Figure 7.11 shows a time series from a typical run with the transmitter being carried, and this can be compared with Figure 7.12, in which the transmitter is stationary, and in which all personnel remained stationary for around a third of the duration of the measurement.

Figure 7.11: Time series of path loss, bedroom

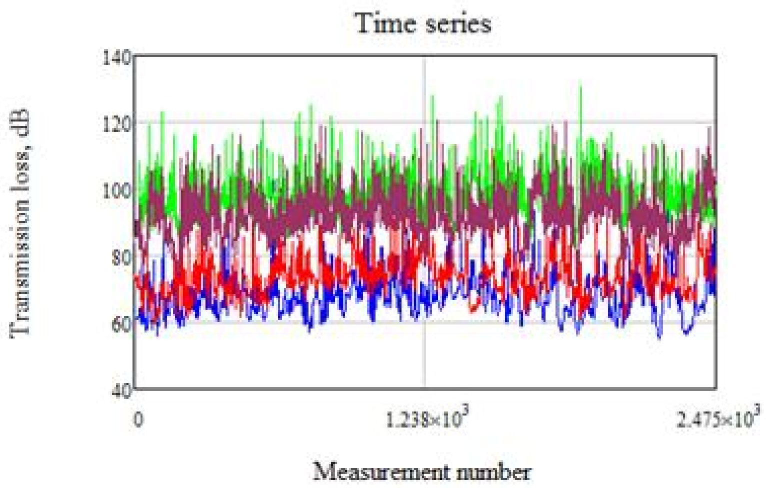

Figure 7.12: Time series of path loss, TX static in bedroom

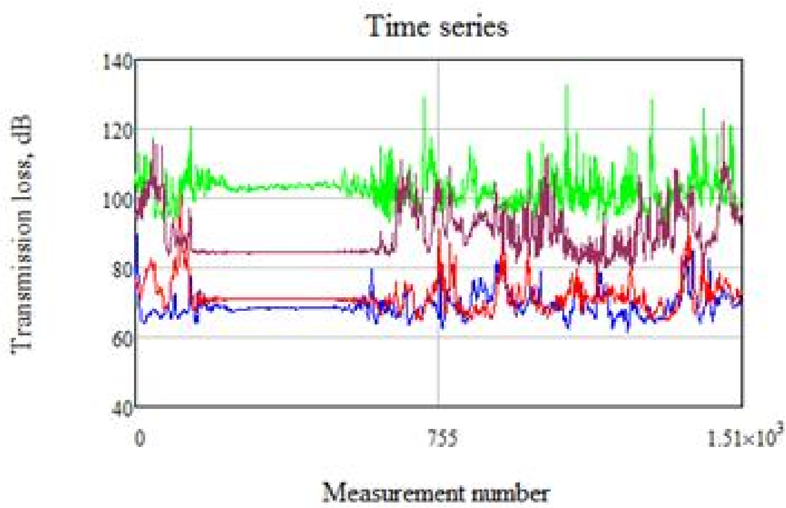

It can be seen that the fading due to multipath is, as might be expected, absent with the transmitter stationary and no movement by personnel. With movement by people, fading and enhancement can be seen to occur due both to multipath effects and from absorption and diffraction losses. This fading is, however, less intense than that seen when the transmitter is also mobile.

7.2 Throughput measurements

Throughput measurements were made by placing the 802.11n-enabled laptop at different locations (subjectively judged to be 'typical') within rooms previously characterised by CW measurements. Measurements were made at each location with the Cisco access point configured[8] for operation in each of the four modes (2.4 GHz SISO, 2.4 GHz MIMO, 5.7 GHz SISO and 5.7 GHz MINO).

The 'representative throughput' in all the results described in this report is taken to be the average of the throughput achieved during the last five trials, i.e. 50 transfers of a 1 MB file (see section 6.2). Where no connectivity was possible, the throughput is recorded as 0 Mbps.

The throughput results from one house (VT) are shown in Figure 7.13 and 7.14.

Figure 7.13: Throughput (VT)—AP in hall

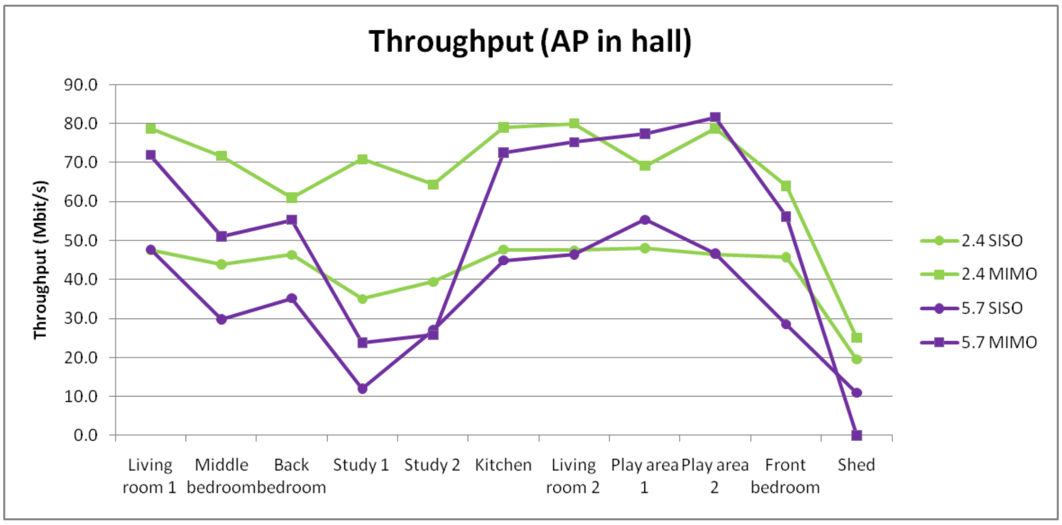

Figure 7.14: Throughput (VT)—AP in study

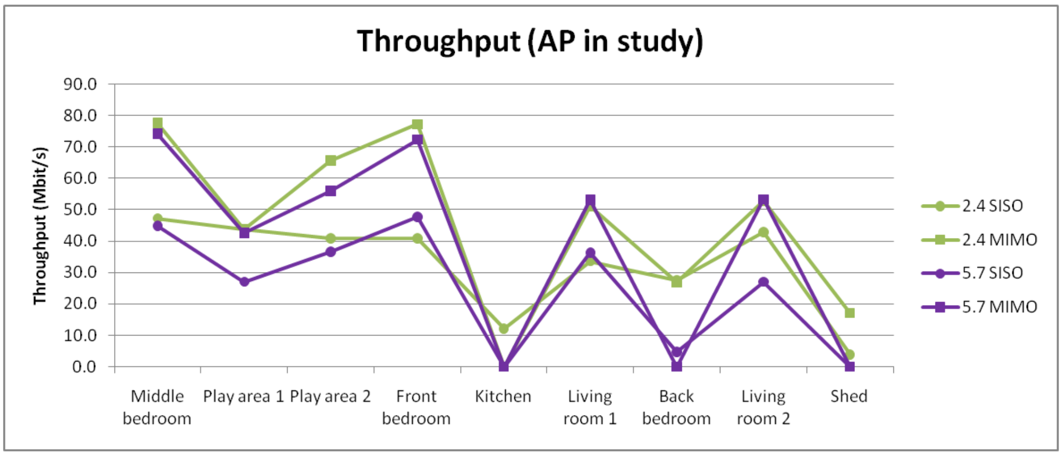

These figures illustrate several characteristics that were found to be generally true in all locations tested.

- As expected, MIMO offers a significant gain in throughput relative to SISO, typically by a factor between 1.5 and 1.8.

[8] In general, the AP would be configured for a particular mode, and then all rooms in the house would be measured in that mode, before moving to the next configuration.

- In most houses tested there were one or more locations where no association was possible between the terminals when using 5.7 GHz. In such cases, switching to 2.4 GHz always allowed operation, albeit at low speed.
- Somewhat obviously, the selection of a more central location for the access point (e.g. a hallway) ensured more uniform coverage, and better average throughput within a building.

One limitation of the measurement method relates to the relatively small number of locations it was possible to sample, owing to the length of time required to carry out the necessary FTP file transfers in all four modes tested. Thus, although each room was sampled at thousands of points for the CW measurements, only one or two 'representative' locations could be assessed in the throughput measurements.

Sets of small-scale measurements were therefore made in two of the houses. In the Victorian Terraced house, a number of locations some 10 cm apart in the living room were sampled, with the AP in the study—one of the longest paths possible. The results are indicated in Figure 7.15.

Figure 7.15: Variability of throughput at 5.7 GHz

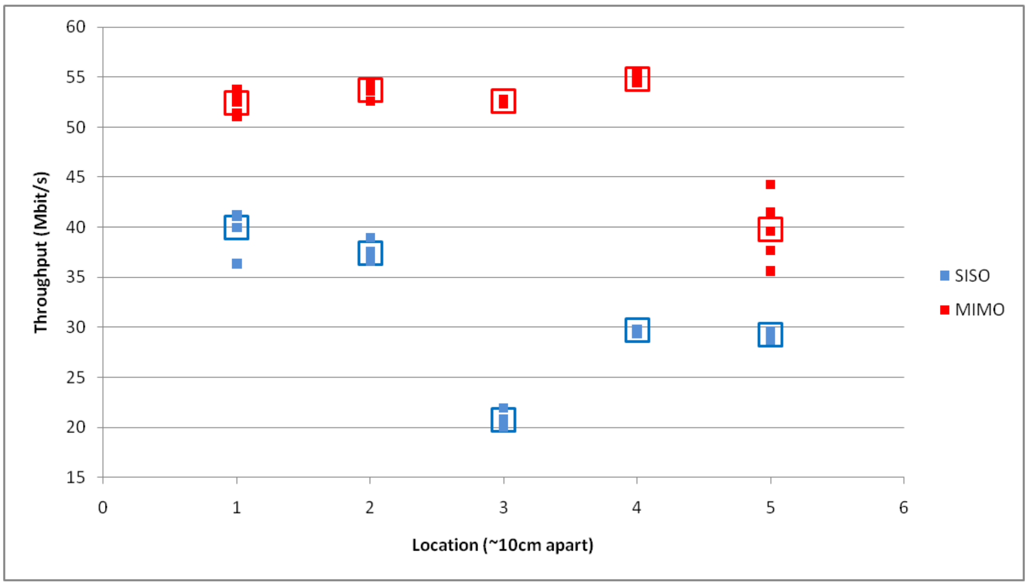

In this figure, the average throughput values, as used elsewhere in this report, are indicated by the skeleton markers. The results of the individual folder transfers are also indicated as solid markers. It can be seen that, generally, there is little variation with time. The variation with position is, however, significant and suggests that even with diversity antennas and / or MIMO operation, the laptop can be placed in a multipath null.

Similar measurements were made at the lower frequency in the 'Modern Detached' house. Here the wideband equipment was used to record throughput with the laptop at 5 cm positions along a 50 cm straight line on a desk. The propagation path was 'study' to 'bedroom 3' which, from experience, suffers from coverage problems. The measurement results for SISO and MIMO are recorded in Figures 7.16a and 7.16b respectively.

Figure 7.16a: Small scale throughput variability, SISO (MD)

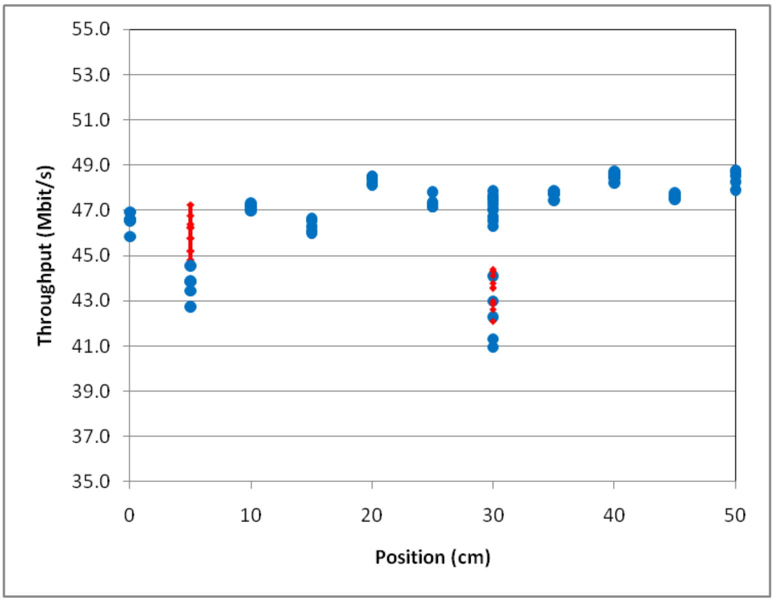

Figure 7.16b: Small scale throughput variability, MIMO (MD)

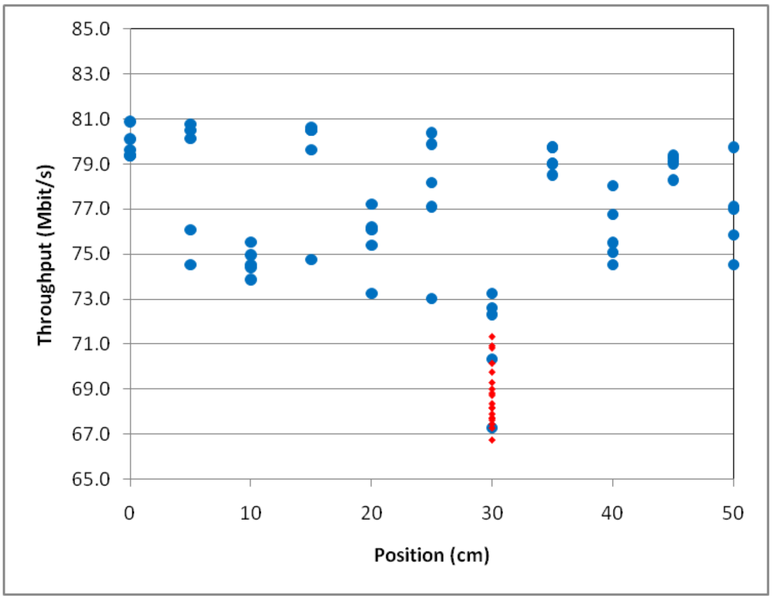

The five blue points at each position are the individual throughput values for each of the five 10 MB folder transfers. These are normally averaged to give a single result but are kept separate here to show the “random” variation that occurs between measurements separated by only a few seconds. The red points are where, at the end of the traverse, the laptop was re-positioned at a previous location (at least to within a couple of millimetres) to see whether the “nulls” were repeatable. The two graphs have been drawn with the same vertical span.

In the SISO case the spread at each position is fairly small. On the first traverse (blue points) more than one run was made at a couple of positions. At 50 cm, the spread between measurements was similar to the spread of the five points within a single measurement and all the points were not plotted. But at 30 cm there seemed to be a null, and 15 blue points have been plotted. Interestingly only 5 of these (from a single measurement) are in the null. The other 10 (collected within a few minutes of the first 5, and with no movement of the laptop) fail to show the null. However the 10 red points (collected later with the laptop repositioned at the 30 cm location) were again in the "null".

In conclusion, there is indeed a null at 30 cm although it is a bit unstable. The null at 5 cm was only seen once, but it still looks like a significant feature. Of course neither of the "nulls" is very deep in absolute throughput terms.

The MIMO result was unexpected. Not only was the variability at a given position much greater than for SISO, but there is still a null at 30 cm and if anything it is more solid than for SISO. Given that the 5 cm spacing of the measurements is less than the separation of the MIMO antennas in the laptop this is quite surprising.

7.3 Summary

7.3.1 Path-loss measurements

- Measurements are very repeatable.
- No significant difference between 'perimeter coverage' and 'whole room coverage'.
- Closing or opening doors generally has no impact on coverage, though slight additional loss (2 dB) was observed in one house at 5 GHz.
- There is almost no difference in the statistics of fading within individual rooms, either with frequency or between houses. The interdecile spread of the loss measurements is around 15 dB in all cases.

7.3.2 Throughput

- MIMO offers better throughput than SISO in almost all cases (generally by a factor between 1.6 and 1.8).
- Coverage at 5.7 GHz is always worse than at 2.4 GHz.
- Placing the hub in more open, central positions (e.g. hallway) gives better and more uniform throughput.

8 PATH-LOSS MODELLING

8.1 Introduction

The objective of the project is to compare in-home Wi-Fi coverage in different frequency bands. Coverage has been measured in terms of file-transfer speeds in the 2 GHz and 5 GHz bands, for which commercial equipment is available. To extend comparisons to lower frequencies, narrow-band measurements were made at 500 MHz and 800 MHz, as well as in the 2 GHz and 6 GHz bands. These results allow inferences to be drawn as to coverage in all four bands.

8.2 Transmission loss and propagation losses

Wi-Fi throughput is a function of transmission loss, that is, the loss between antenna terminals. Transmission loss is the attenuation experienced by radio waves between antenna terminals, including the effects of antenna gains. In the case of indoor propagation there will be multiple rays between two antennas. Each ray experiences varying propagation losses, and different antenna gains according to the take-off and arrival angles.

The narrow-band measurements show that transmission loss varies in an irregular manner with the four frequencies, particularly between 2.41 GHz and 5.76 GHz. As discussed in Annex C Section C4, this is mainly due to differences in antenna gains.

To inspect the characteristics of propagation at the four frequencies, measured transmission loss was converted to a quantity defined as 'propagation loss' in Section C4. Propagation loss, L_p, is given for our purposes by:

$$L_p = L_t + G \quad \text{dB} \qquad (8.1)$$

where G is the combined antenna gain given by the last line of Table C.4.

There are approximations in this procedure. Antenna gain G is a measured median of time-series showing wide variations, and is based on azimuthal gains only. Thus equation (8.1) ignores rays with large elevation angles at either antenna. Nevertheless, it is found that 'propagation loss' can be modelled simply and shows systematic variation with frequency, particularly if restricted to measurements on the same floor, for which there is a larger probability that the strongest ray will have small vertical angles.

8.3 Propagation loss for same-floor measurements

Figures 8.1 to 8.5 show propagation losses plotted against distance for CW measurements in the five homes. All graphs have the same scales to simplify comparison. Only same-floor measurements are shown, that is, where both transmitters and receivers are on the same floor, although always in different rooms. Where available, the data include measurements extended into an adjoining property, simulating a larger dwelling. In these graphs the symbols represent measurements, colour-coded according to the legends. The curves show free-space basic transmission loss for the same four frequencies.

Figure 8.1: Victorian terraced house, same-floor measurements

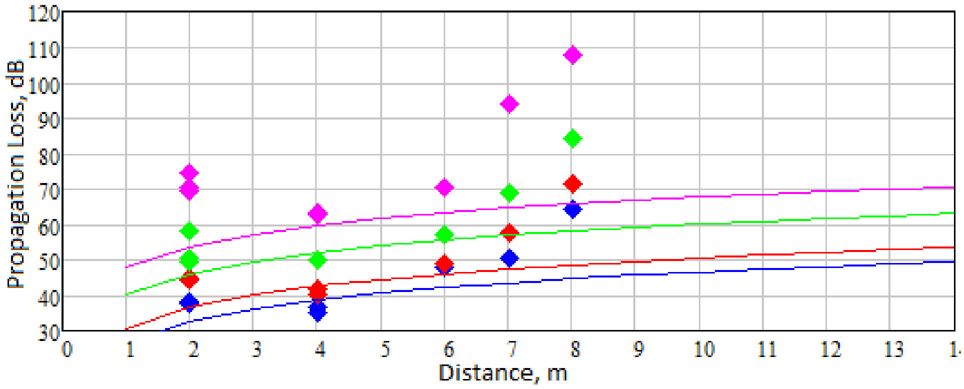

Figure 8.2: Inter-war semi-detached house, same-floor measurements

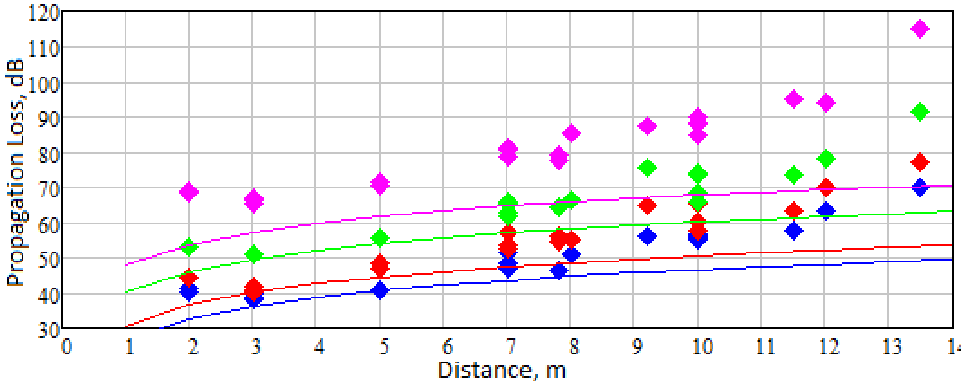

Figure 8.3: Modern bungalow

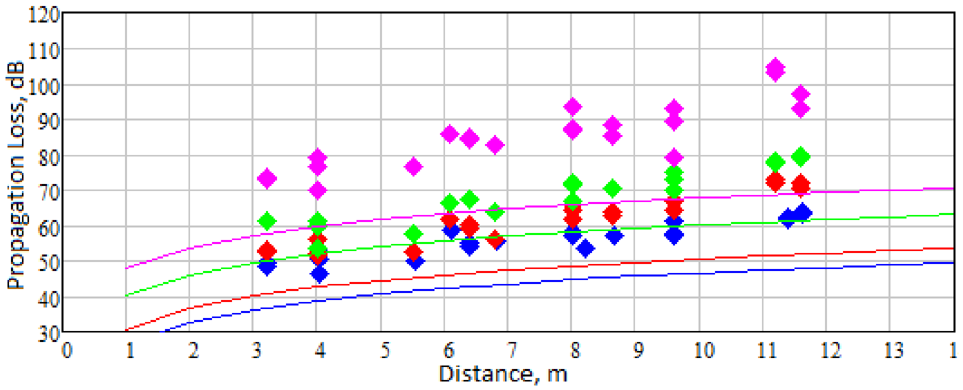

Figure 8.4: Modern detached house, same-floor measurements

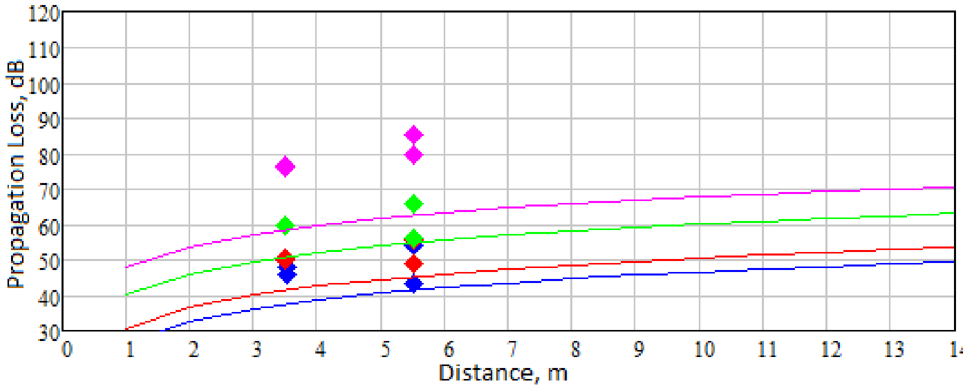

Figure 8.5: Modern terraced house, same-floor measurements

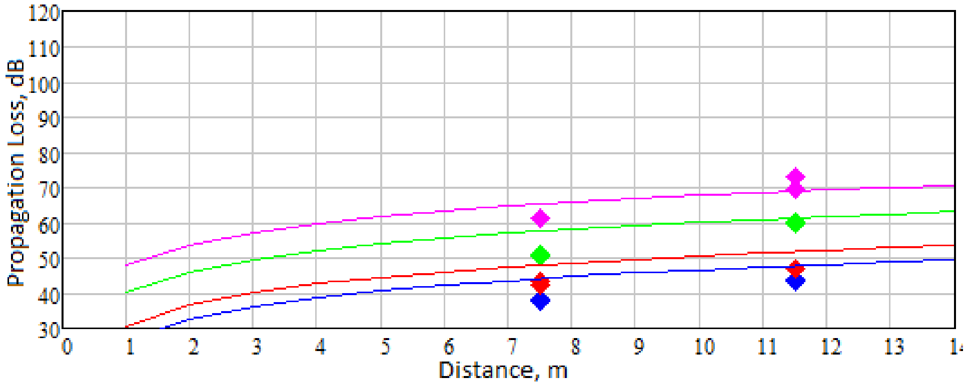

In Figures 8.1 and 8.2, and to a small extent in Figure 8.3, there is a tendency for losses relative to free-space to be at a minimum at 4 or 5 metres distance. In most cases in losses are higher in the absolute sense at 2 metres than at 4 metres.

The effect is understood as follows. At 2 metres a direct, or near-direct, ray is expected to be the strongest, even though it penetrates at least one wall, door or opening. Rays reflected from walls, floor or ceiling will travel typically twice the distance of the direct ray, for which free-space attenuation is 6 dB higher. Moreover, reflection will not be at grazing incidence, but at steeper angles where reflection coefficients are usually less. Thus most energy is transmitted via the direct ray, which needs to penetrate a wall or diffract around a door opening. Allowing for negative antenna gains for the three lower frequencies, propagation loss is expected to exceed free-space loss, as observed.

At a distance of 4 or 5 metres the ratio of indirect to direct path lengths will be smaller, incurring less relative attenuation, and some reflected paths will be closer to grazing incidence with internal surfaces and thus have higher reflection coefficients. It is now more likely that a number of rays will carry similar signal levels, with power summation reducing the net loss. Under such circumstances it is possible for the propagation loss to be less than for free space, which is observed for the lower three frequencies in Figure 8.1.

Beyond the distance of minimum loss relative to free-space, propagation losses show the effect of increasing numbers of wall penetrations, and presumably a tendency for the number of significant rays to decrease with distance. This occurs particularly strongly in the Victorian house in Figure 8.1. It is possible that this is due to the generally more solid construction, such as denser bricks. The inter-wars house is built from fairly light bricks and has mainly partition walls on the first floor, which could provide a lower-loss route for the longer paths.

Few distances are represented in Figure 8.4, but there is a tendency for losses to fall relative to free space from 3.5 to 5.5 metres.

The most conspicuous feature in Figures 8.5 is that almost all losses are less than free space, even at 11.5 metres. This is consistent with smaller penetration losses for internal walls due to lighter construction materials.

8.4 Modelling same-floor propagation losses

It was found by inspection that the propagation-loss results for which both transmitters and receivers were on the same floor could be modelled in the form:

$$L_p(d, f) = L_{bfs}(1, f) + k + (a + bf)(d - 1) \qquad \text{dB} \qquad (8.2)$$

where:

L_p is propagation loss as a function of slope distance, d (m), and frequency, f (MHz).

Distance d is measured from the receiver antenna assembly to the centre of the room being measured at the height the transmitter antennas were carried.

The terms $L_{bfs}(1, f) + k$ give the intercept at $d = 1$ m, consisting of free-space basic transmission loss plus a constant offset k (dB).

The factor $(a + bf)$ is the slope of the model in dB/m given as a linear function of frequency set by constants a and b.

Table 8.1 gives the constants when this model is fitted to all data shown in Figures 8.1 to 8.5. The values for the modern detached and modern terraced houses should be given less weight than the other three since they are based on fewer measurements and smaller ranges of distance.

Table 8.1: Model parameters for all same-floor measured results

House type	Data from Fig:	Model constants		
		k	a	b
Victorian terraced	8.1	8.3	3.00	0.00035
Inter-wars semi-detached	8.2	8.1	2.56	0.00014
Modern bungalow	8.3	16.4	1.9	0.00019
Modern detached	8.4	18.7	0.33	0.00044
Modern terraced	8.5	0.7	1.12	0.00026

Figures 8.6 to 8.10 illustrate the performance of the model for the data from Figures 8.1 to 8.5 respectively. As before the graphs have the same scales to simply comparison. The symbols represent measurements, the dashed lines are linear fits to the data for each frequency independently, and the solid lines show the model's predictions given by equation (8.2) with the parameters in Table 8.1.

Figure 8.6: Data and model, Victorian terraced house

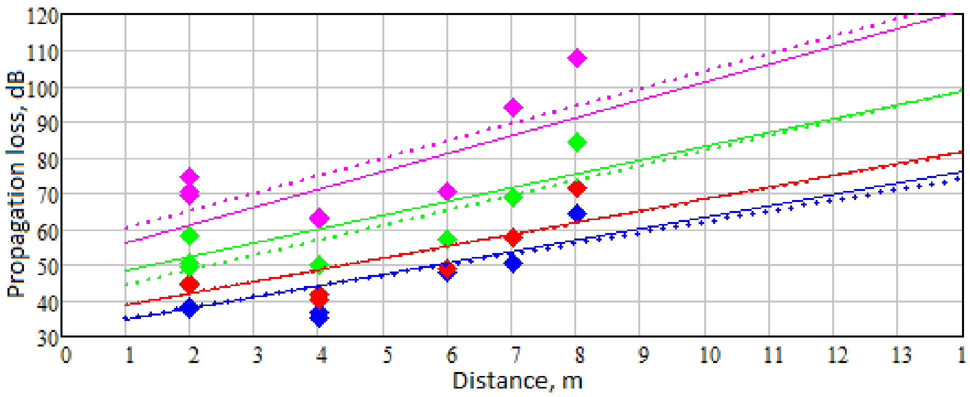

Figure 8.7: Data and model, inter-war semi-detached house

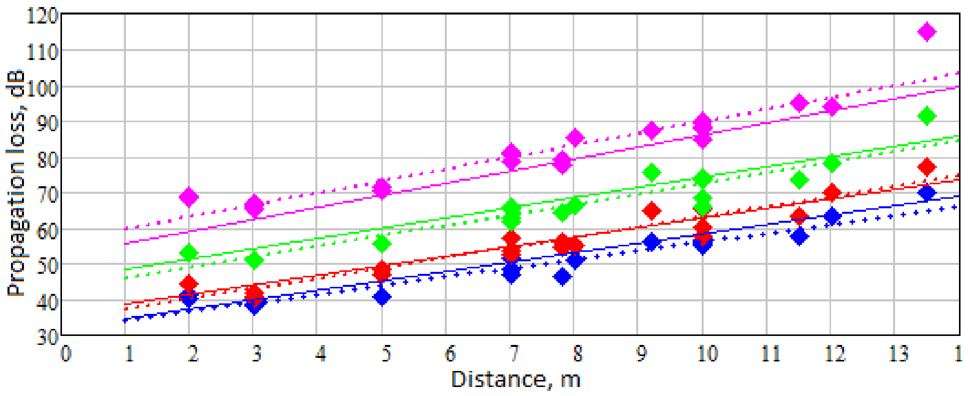

Figure 8.8: Data and model, modern bungalow

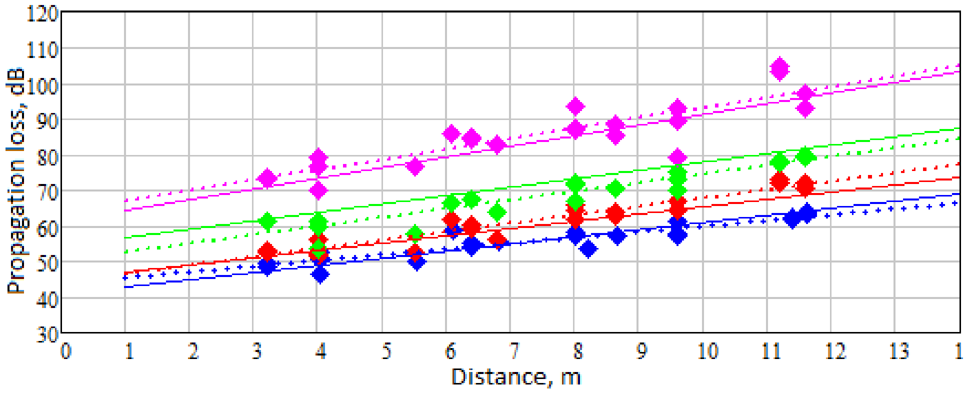

Figure 8.9: Data and model, modern detached house

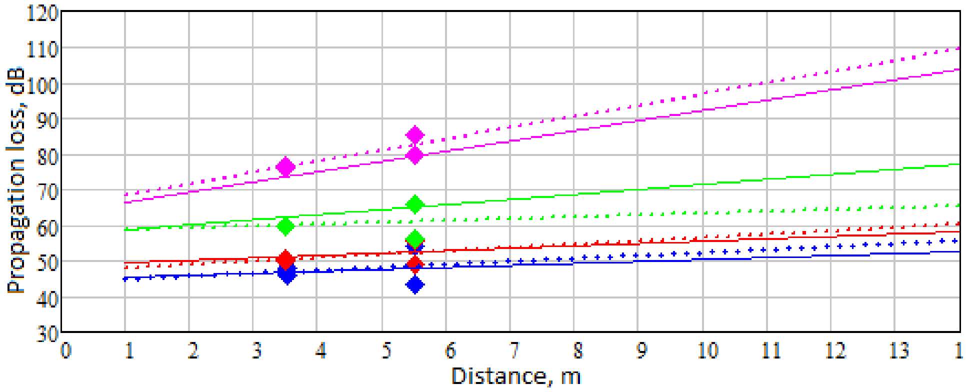

Figure 8.10: Data and model, modern terraced house

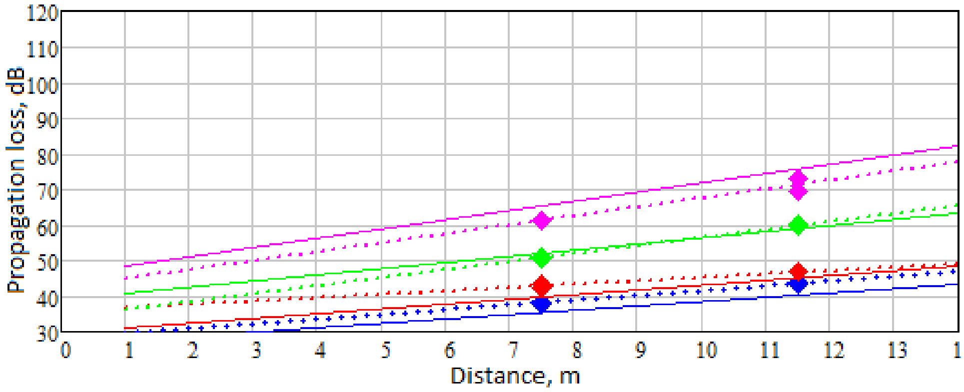

Table 8.2 gives the performance of the model in numerical terms. Although the model in each case has been fitted to the measured data shown, the fits are not perfect because an a priori decision was made that loss will vary with frequency at the 1-metres intercept proportional to 20.log(*f*).

Table 8.2: Means and standard deviations of model

Statistics of model discrepancies		Means				Standard deviations			
House type	Fig.	0.5G	0.8G	2G	5G	0.5G	0.8G	2G	5G
Victorian terraced	8.6	3.57	3.06	5.47	−1.7	6.61	6.22	6.92	10.16
Inter-wars semi-detached	8.7	1.76	0.26	2.10	−3.37	3.00	2.84	3.41	4.16
Modern bungalow	8.8	0.88	−1.31	3.64	−2.03	2.18	2.28	2.63	4.36
Modern detached	8.4	−0.67	0.33	3.24	−2.85	3.71	2.42	3.60	2.11
Modern terraced	8.5	−3.07	−2.51	0.17	4.28	0.40	0.84	1.06	1.16

The above modelling provides an indication as to the variation of propagation losses to be expected in different types of house, presented in a form which varies fairly systematically with frequency. Of particular interest is the rapid increase in loss in the Victorian property for distances beyond 7 metres, and the remarkably low losses in the modern terraced house.

9 COVERAGE

9.1 Definition of coverage

Ofcom requested that the study should provide information along the lines of "*50% of UK homes would be adequately covered at 5 GHz*", and suggested that to make such assertions would require "... *databases of known distributions of homes across the UK.*"

To achieve this we proposed to obtain coverage distributions of individual houses of various types, and to use the statistics of these house types to derive coverage figures for the UK. An obvious question is what is meant by "coverage"?

Users of domestic WLAN equipment mean at least two things by "coverage". "I can't get coverage in room X" probably indicates that a PC or laptop will not connect at all to the access point (AP). "Coverage is poor upstairs" may mean that some rooms are not connected, or that the service is intermittent or slow. So in a loose sense "coverage" means both connectivity and speed of connection.

Of course in general usage, the linkage between connection speed and coverage is subjective. Coverage for streaming of high-definition video is likely to be considered poorer than for Web browsing.

In this study we identify connection speed as throughput (given in Mbps) between an AP and a user terminal. As described in Section 6, throughput measurements have been made directly at 2.4 GHz and 5.7 GHz. At 500 MHz and 800 MHz, throughput needs to be deduced, based on the CW measurements described in Section 5.

Given a set of (measured or derived) values of throughput between one or more typical AP locations and a reasonable sampling of likely terminal locations in a house, we consider these throughput values as statistical samples. From these we derive a cumulative distribution function (CDF) which gives the probability that a given value of throughput is exceeded for that house. In this context we can interpret probability as the fraction (or percentage) of a house that is "covered" for a specified throughput requirement.

Using this definition of coverage, it is straightforward to combine the results for each of the houses measured in this study to obtain an "average" coverage distribution for the UK. The measurements for each house are simply weighted to account for the predominance of that house type in the UK housing stock (see Section 4).

A brief description of the data used is followed by a derivation of the coverage at 2.4 GHz and 5.7 GHz using this approach. The method used to obtain coverage statistics at 500 MHz and 800 MHz is then explained.

9.2 Data reduction

Sections 5 and 6 explain how the CW and throughput measurements were made and Section 7 describes the data processing.

For estimating coverage, we take the median of the CW transmission loss measurements on an AP-to-terminal path as representative of the whole room. Similarly the average throughput measured on an AP-to-terminal path is taken as representative of that room.

All the valid measurement data were used in the derivation of coverage, except for the CW measurements made between two houses. These were considered to be more relevant to interference statistics than to coverage statistics. However, the (very few) measurements between a house and a garage/shed of the same property were used, as these were considered legitimate locations that a user may require coverage.

The CW and throughput measurements come in "sets", a CW set comprising measurements at 500 MHz, 800 MHz, 2.4 GHz and 5.7 GHz on a given path, and a throughput set comprising SISO and MIMO measurements at 2.4 GHz and 5.7 GHz. Normally the measurements were made as a complete set. However it was not the case that a CW set was always made on the same path as a throughput set, and vice versa.

For some AP-room paths more than one set of CW or throughput measurements was available. Duplicates often occurred where "door open" and "door closed" measurements were made. As discussed in Section 7, door "attenuation" was always small (less than ~2 dB) and was often negative! Occasionally more than one measurement was made as a check of repeatability, and in most cases the results were very close.

Given the proposed statistical approach to coverage modelling it would be statistically unsound to include several samples from the same path without applying variable weights to the sample values. For this reason, and in order to simplify subsequent analysis, any duplicate sets were replaced by a single set whose values were the means of the values of the duplicate sets. Table 9.1 lists the datasets available, after this reduction process.

Table 9.1: Datasets used for coverage estimation

House	AP position	Number of throughput sets	Number of CW sets	Number of coincident sets
Victorian terraced	Hall	8	8	8
Victorian terraced	Study	7	7	7
Inter-war semi-detached	Bedroom 3	9	9	9
Inter-war semi-detached	Hall	2 (+ 8 partial)	12	7 (av.)
Modern bungalow	Hall	7	7	5
Modern bungalow	Reception room	10	7	4
Modern bungalow	Study	6	0	0
Modern detached	Hall	9	9	9
Modern detached	Study	8	8	8
Modern terraced	Living room	5	5	5
Large semi-detached	Hall	12	3 (partial)	1.5 (av.)
Large semi-detached	Living room	11	11 (partial)	5.5 (av.)
Basement flat	Hall	0	3	0

9.3 Coverage at 2.4 GHz and 5.7 GHz

As an example, the coverage CDF for the modern bungalow with the AP located in the hall is shown in Figure 9.1. There were 7 measurement sets in this case.

Figure 9.1: Coverage distribution, modern bungalow, AP in hall

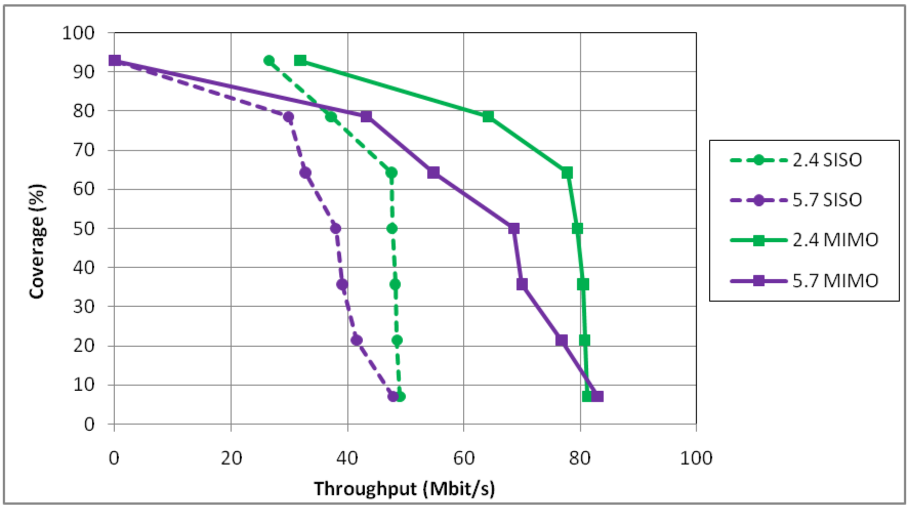

The coverage CDFs are drawn in such a way that the curves represent the percentage of rooms for which the throughput is greater than a specified throughput rate. There are relatively few points contributing to a CDF in this application, and different points will require different probability weightings when the statistics from different houses are combined (see later). The CDFs have therefore been generated directly using the data probability values rather than using an approach using equal-sized bins. To avoid bias due to the relatively large increments of probability between data points, the probability values of the plotted points are "mid-step" values.

Note that each data point has uncertainty in both its throughput value (due to measurement error) and its probability (due to limited sampling of the true distribution). The "curve" joining the data points should not therefore be taken too literally as "the" coverage CDF. The curves have not been extrapolated beyond the highest and lowest throughput values. In general the tails of the distribution can only be filled in by obtaining more samples.

The interpretation of Figure 9.1 is straightforward. For example, if your criterion for "coverage" is a requirement for more than 60 Mbps, then Figure 9.1 shows that a MIMO system will give ~60% coverage at 5.7 GHz and ~80% coverage at 2.4 GHz. A SISO system will not give this throughput at either 2.4 or 5.7 GHz. In general (though not universally) coverage at 2.4 GHz is better than at 5.7 GHz for both SISO and MIMO, reflecting the conclusions for throughput in Section 7.2.

It is important to remember that *these coverage figures are specific to the technology used in the measurements* (i.e. IEEE 802.11n in SISO or MIMO mode). More will be said below about "technology-neutral" coverage, where we extrapolate the results in frequency and remove the system limitations.

For now note one obvious feature of the system limitation. For both SISO and MIMO there is a throughput level at which the system "saturates". Figure 9.1 shows the saturation value to be ~50 Mbps for SISO and ~80 Mbps for MIMO. At 2.4 GHz, the coverage at the saturation values plummets from ~60% just below saturation to 0% above.

The 50 Mbps limit for IEEE 802.11n SISO is 77% of the theoretical physical layer limit of 65 Mbps so the measurement equipment is not significantly constraining the throughput. The observed MIMO limit of 80 Mbps is only 62% of the physical layer limit, but this presumably reflects the less than ideal MIMO channel matrix found in the sample locations. The throughput measurement system was chosen with considerable care (Section 6) and the value of 80 Mbps is likely to be better than can be obtained using average consumer-class equipment. For the purposes of this section, these saturation values will be considered to be inherent limits of the SISO and MIMO systems.

A more extreme case of saturation is shown in Figure 9.2, where coverage at 2.4 GHz SISO plummets from ~100% to 0% at 50 Mbps. The AP position in this case was placed in what was expected to be near-optimum for the house, viz. on the ceiling in a central location above the stairwell. The uniform, "saturation" throughput figures in all rooms bears out the assumption that this is indeed a good location for this AP. However, it does not mean that the rooms see the same signal levels. So if the system were to be upgraded in the future to higher bitrate technology, the coverage is unlikely to be uniform at these higher bitrates.

Figure 9.2: Coverage distribution, modern detached, AP in hall

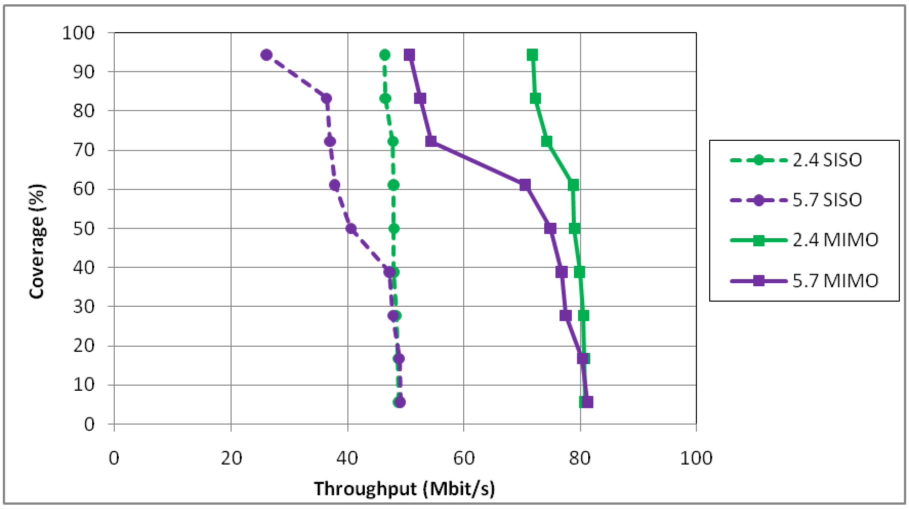

The coverage CDFs for all measured house and AP combinations are given in Annex E.1.

Obtaining representative coverage CDFs for the UK is simply a matter of including all measured data from each house and AP position in a single CDF. This assumes that the points are statistically independent, which is true. It also assumes that the measurements made during this study are representative of the UK housing stock. This is discussed in Section 4.

Of course it is not true that each data point should have an equal weighting in the statistical distribution. The various house types measured in this study represent different fractions of the UK housing stock. In addition, the different AP locations within a house may be more, or less, likely to be found in the domestic situation.

With respect to the AP locations, the measurements were generally made with the access point (or CW receiver) located near to the place where the majority of consumers would place their WLAN hub. This is rarely in an "optimum" location. It will often be next to the point-of-entry of the telecom operator's cable or ADSL connection, which of necessity will be at an extremity of the building. One exception (Figure 9.2) was where we placed the AP in a central location on the ceiling—a location that might be chosen in a commercial environment. Comparing Figures E1.8 and E1.9 for this house shows the advantage of the better AP location.

Apart from this case the AP locations are definitely *typical* rather than *optimal*. Deciding on weights to apply to the different AP locations in our dataset would be very subjective. But apart from the one case shown in Figure 9.2, the others are believed to be equally likely for a consumer to use. In the overall UK "average", the inclusion or omission of Figure 9.2 makes very little difference. In addition, a complete set of room measurements were made in most cases from each AP location. We have therefore adopted equal weighting for all AP locations within a house. So, for example, in a house with two AP locations, the weighting of each AP is 0.5.

On the other hand equal weightings cannot be applied to the different house types. Rather we have used the weightings given in Table 4.4. In the case where two houses in the same category were measured (semi-detached and large semi-detached), the houses were given equal weighting (13%). A minor complication is that throughput measurements were not made in the basement flat. At this stage we simply assume that coverage in flats will be similar to that measured in small houses, and a weighting of 28.4% is assigned to the small terraced house to represent both small terraced houses and flats. With these assumptions, a set of individual AP measurements contributes between 28.4% (semi-detached house) and 3.1% (bungalow) to the UK average statistics.

Figure 9.3 shows the UK weighted average coverage distribution.

Figure 9.3: Coverage distribution, UK weighted average

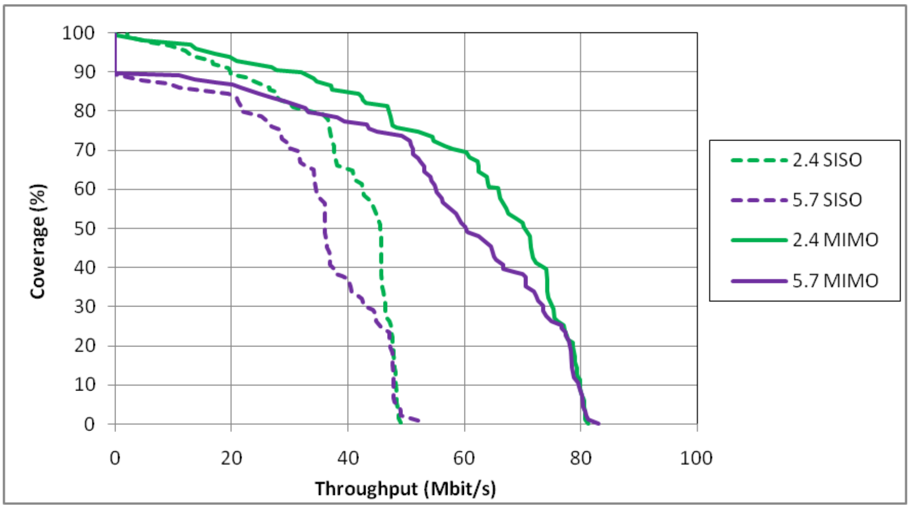

From this, inferences such as the following can be obtained:

- At 2.4 GHz, 80% of locations should get more than 30 Mbps with SISO and 45 Mbps with MIMO.
- A user requiring a throughput of more than 40 Mbps needs to use MIMO for acceptable coverage, although 20% of locations will still not be covered.

Some general conclusions can be made:

- On average, MIMO is always better than SISO.
- On average 2.4 GHz is better than 5.7 GHz
- The saturation effect is less of a limitation in the average coverage than for the examples shown in Figure 9.1 and Figure 9.2. It affects SISO more than MIMO.

There are also some caveats:

- "On average" in the above is important. Looking through the individual coverage distributions in Annex E.1, cases will be found where 5.7 GHz is better than 2.4 GHz and even where SISO is better than MIMO.
- The results reflect the particular hardware used for the measurements (Section 6). Current consumer-grade equipment may well have lower performance and hence coverage. And of course a consumer's hardware and software applications may reduce data rates below the throughput figures. In one test case, the throughput measured using a consumer hub and PC collocated with the test kit gave file transfer rates that were only ~35% of the throughput measurement.

It is important to remember that the coverage distributions are *percentage of locations covered*, not *percentage of houses that have 100% coverage*. The individual house coverage distributions in Annex E.1 can be used to say something about how good coverage might be within a single house. So for example, in most of the house types, there was at least one room where the 5.7 GHz kit failed to connect, giving strictly 0 Mbps throughput. Even at 2.4 GHz there were often rooms at which throughput was less than 20 Mbps.

The houses where throughput was high in all rooms were, not unexpectedly, the smaller houses (modern terraced house, Figure E1.10) and larger houses where the AP was centrally located (modern detached house with AP in hall, Figure E1.8).

Note however that the number of houses measured does not allow accurate statistical figures to be given for an individual house type. This would need more measurements from a number of houses of that type.

9.4 Predicted coverage at 500 MHz, 800 MHz, 2.4 GHz and 5.7 GHz

9.4.1 Relationship between transmission loss and throughput

We want to plot coverage distributions at 500 MHz and 800 MHz, and this requires throughput figures for these frequencies. We cannot do this directly as no measurements were made of throughput (indeed no commercial equipment using IEEE 802.11n is available) at UHF. But we do have CW measurements at all four frequencies.

In order to make progress some relationship between transmission loss and throughput needs to be established between the measurements taken at 2.4 GHz and 5.7 GHz, and this used in some way to extrapolate the results to the UHF frequencies.

Initial results were quite encouraging. Figure 9.4 and Figure 9.5 show throughput plotted against transmission loss for the Victorian terraced house, for SISO and MIMO respectively. The throughput and transmission loss show a remarkable degree of correlation for both datasets. The SISO and MIMO trends are different, but importantly the 2.4 GHz and 5.7 GHz results appear to follow the same trend lines.

Figure 9.4: Throughput scatterplot, Victorian terraced house, SISO

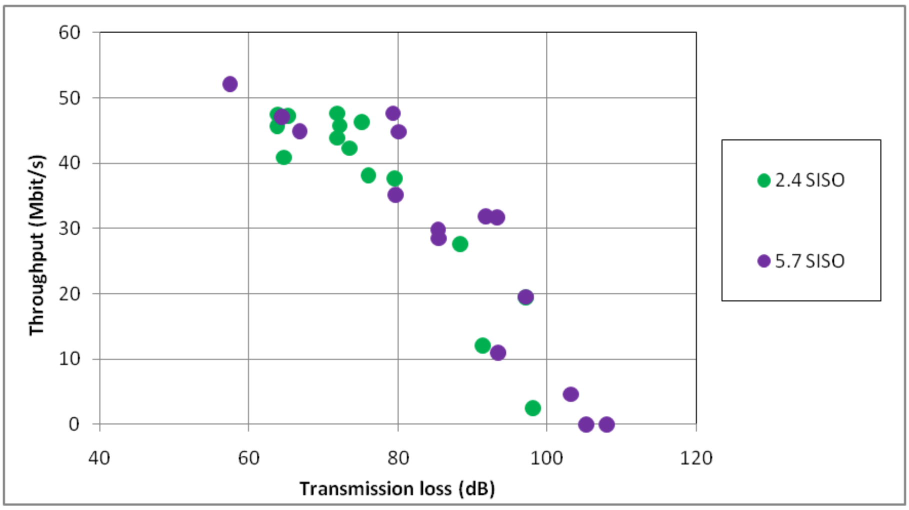

Figure 9.5: Throughput scatterplot, Victorian terraced house, MIMO

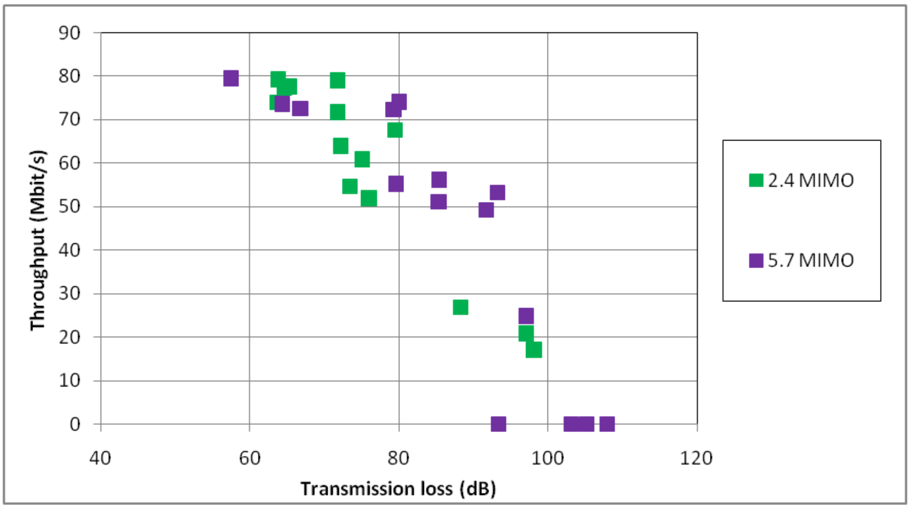

There is some scatter, of course, but this is not unexpected given the different equipment and data processing methods used for the two measurements. In particular, the transmission loss measurement represents the room median value, while the throughput measurement is a point value.

These graphs suggest that we are sampling the performance curve (throughput versus carrier-to-noise ratio) of the IEEE 802.11n equipment. At a given frequency, the points would be expected to lie on a line if the CW and throughput measurements had been made simultaneously using the same equipment. The different equipment and data processing methods used for the two measurements would be expected to introduce the observed scatter.

For transmission loss to adequately represent the same performance curve at both frequencies two things must hold:

1. The multipath environment must be the same at 2.4 GHz and 5.7 GHz.
2. The antennas used for the CW and throughput measurements must have the same antenna patterns.

The scope of this study did not allow for channel sounding, so no definite conclusions could be made on point 1. However the lack of frequency-dependence in the shape of the individual room transmission loss CDFs suggests there is no major difference in the multipath environment at 2.4 GHz and 5.7 GHz (see Section 7.1).

The assumption of point 2 is highly questionable, given the discussion in Section 8 and Annex C.3. We consider this further below.

In principle we could relate our measurements to the theoretical performance curve of the IEEE 802.11n system. Published results are hard to interpret. The protocol is adaptive and switches modulation scheme depending on the signal strength. To the accuracy of the scatter within the data, the trends of Figure 9.4 and Figure 9.5

appear to follow a straight line. We will therefore define an “effective C/N”, *C* (dB), in terms of transmission loss, *L* (dB), as:

$$C = K - L \tag{9.1}$$

The constant *K* may be different for different frequencies – we propose to check this using the data. The value of *K* is related theoretically to the calibration constants used to derive transmission loss from received signal strength, and to the noise floor in the receiver bandwidth. In fact the absolute value of *K* is not important for subsequent analysis, only its relative value for different frequencies.

Throughput is plotted against effective C/N in Figure 9.6 and Figure 9.7 for all locations at which both CW and throughput measurements were made. (The data from the large semi-detached and basement flat were not included in this analysis as they were not available at the time the analysis was done and in any case they had a very limited number of coincident CW-throughput datasets—see Table 9.1.)

Figure 9.6: Effective C/N, all data, SISO

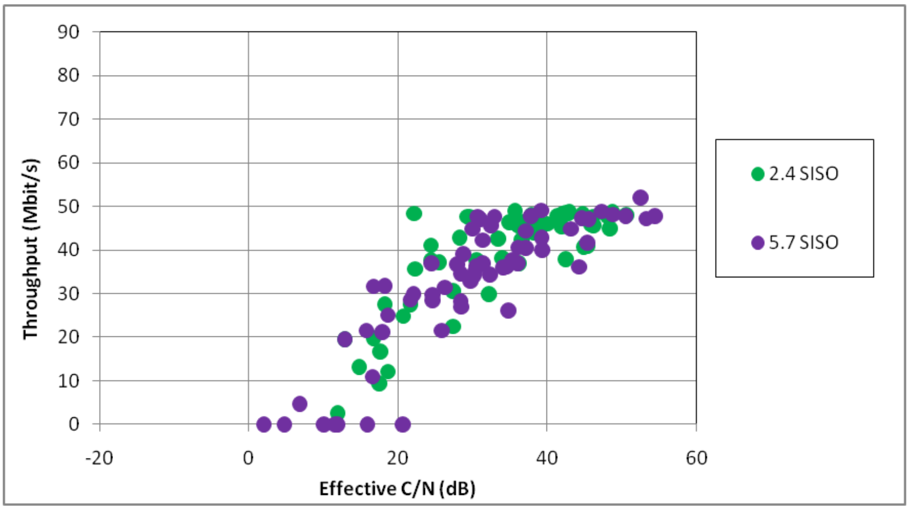

Figure 9.7: Effective C/N, all data, MIMO

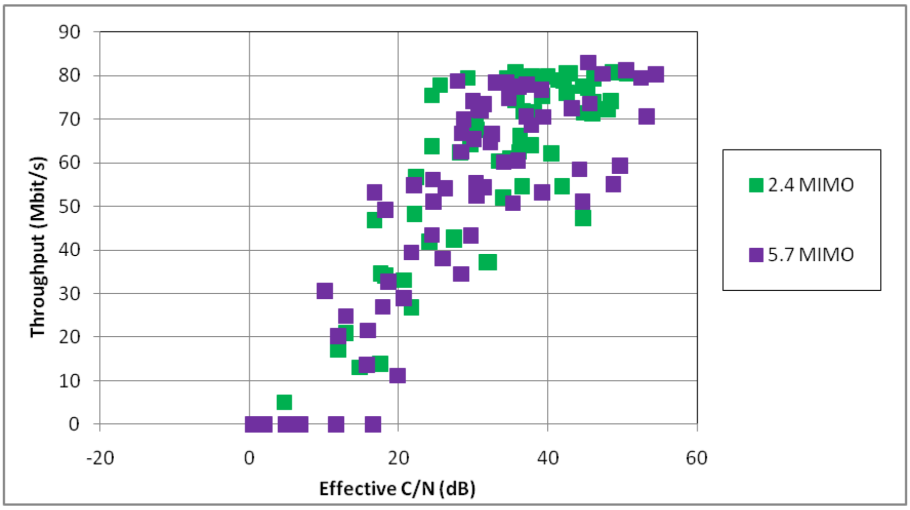

Again there is a good degree of correlation in both the SISO and MIMO data although the scatter is greater than for the single house. In these graphs, the constant *K* was set to 110.0, making the throughput drop to zero when the effective C/N is between 0 and 20 dB, which is typical.

Close inspection of the graphs suggests that a simple linear fit to all the data may not be the best "model" for the performance curve. One issue, most clearly visible in the SISO graph of Figure 9.6, is the saturation effect referred to in the previous section. Above ~40 dB C/N the performance levels off at 50 Mbps (SISO). For MIMO (Figure 9.7) the effect is similar—a levelling off at 80 Mbps above 40 dB C/N.

A second issue, although less serious, is evidence of a threshold effect at low C/N values. There appears to be a "step" at ~10 dB C/N, below which the throughput drops suddenly to zero. A threshold effect is to be expected theoretically—a minimum value of C/N being required for association to take place.

In order to get the best estimate of the "slope" portion of the performance curve, we therefore fitted the data between the threshold and saturation regions. C/N values of 10 dB and 40 dB were used to define the slope region. The results are shown in Figure 9.8 and Figure 9.9.

Figure 9.8: Effective C/N, slope region, all data, SISO

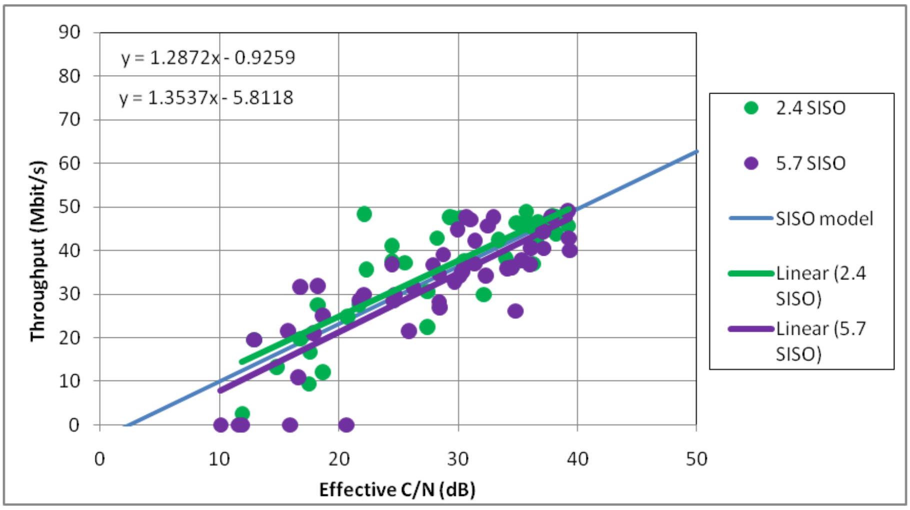

Figure 9.9: Effective C/N, slope region, all data, MIMO

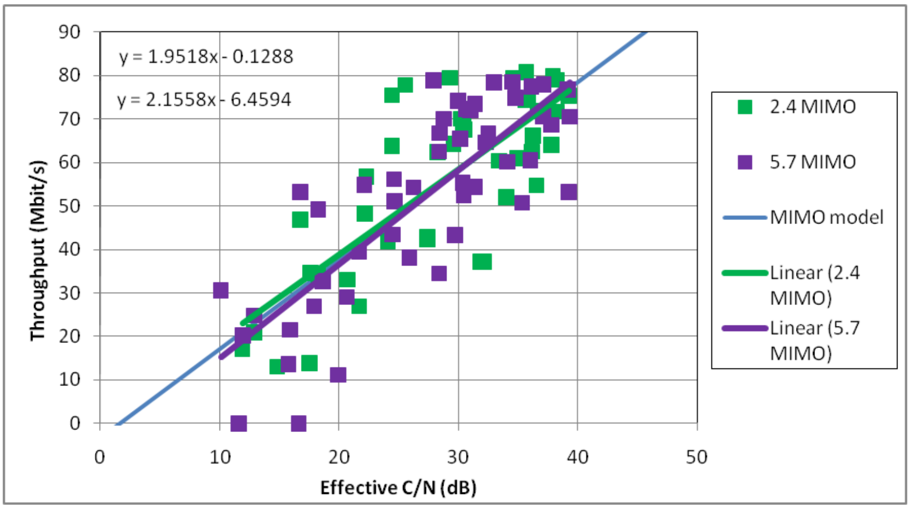

In these figures the (short) green and purple lines show the best least squares linear fit to the 2.4 GHz and 5.7 GHz data, respectively. The equations of the best fit lines are shown at the top of the figures. It is remarkable that, for both SISO and MIMO, the best fits for the two frequencies agree very well.

It is possible to get even better agreement between the fits by small frequency-dependent adjustments to the constant K in the definition of effective C/N (Equation 9.1). This has not been done as the changes required are well within measurement errors, and it seemed more prudent to assume true frequency-independence, rather than introducing a weak dependence that had little statistical significance.

The independence from frequency is a non-trivial result. As stated above, the antennas used for the CW and throughput measurements were different, and differences in the antenna patterns might be expected to introduce frequency biases in the results. In Section 8 and Annex C.3 it was shown that the individual antennas did deviate significantly from isotropic. However the results derived here suggest either that the biases in the CW and throughput measurement antennas are very similar, or that any differences cancel out when aggregated over the ensemble of houses.

So a model for the slope of the performance curve was obtained by simply taking a line mid-way between the 2.4 GHz and 5.7 GHz best fits. These are shown as the (long) blue lines in Figure 9.8 and Figure 9.9, defined as:

$$T = \alpha C - 3.3 \tag{9.2}$$

where

T is throughput (Mbps)

C is effective C/N (dB), given in terms of transmission loss, *L*, by:

$$C = 110 - L \tag{9.3}$$

The slope, α, is different for SISO and MIMO:

$$\alpha = \begin{matrix} 1.32 & \text{Mbit/s/dB} & \text{for SISO} \\ 2.05 & \text{Mbit/s/dB} & \text{for MIMO} \end{matrix} \tag{9.4}$$

Note that Equations 9.2 and 9.3 are independent of frequency and that the constant offset –3.3 in Equation 9.2 is the same for SISO and MIMO. The fact that such a simple formula can be used to predict throughput based on transmission loss supports the idea that the trend line represents the system performance curve.

The implications are very significant for the task in hand. If Equation 9.2 does indeed represent the performance curve of the SISO and MIMO systems, then *we can apply it at any frequency*, and so obtain throughput/coverage results at 500 MHz and 800 MHz.

On the other hand, even although Equation 9.2 is independent of frequency, *it does not follow that the throughput and coverage is the same at all frequencies*. The measured transmission losses are different at each frequency, and this will give different coverage.

9.4.2 Predicted coverage

We can now use Equation 9.2 to make predictions of coverage at the four frequencies. If we can ask what coverage would be available using current IEEE 802.11n equipment, Equation 9.2 needs to be modified slightly to add the threshold and saturation limits into Equation 9.2:

$$\begin{array}{ll} \text{If C < 10} & T = 0 \\ \text{Otherwise} & T = \min(\alpha C - 3.3, S) \end{array} \tag{9.5}$$

where S = 50 for SISO and 80 for MIMO.

Applying Equations 9.5 and 9.3 to the measured transmission losses at all four frequencies, the predicted average coverage for the UK is shown in Figure 9.10 and Figure 9.11 for SISO and MIMO respectively. Here we include all the houses measured in Table 9.1 with the weightings given in Table 4.4: only CW measurements are required to apply the model.

Figure 9.10: Predicted coverage distribution, UK weighted average, SISO

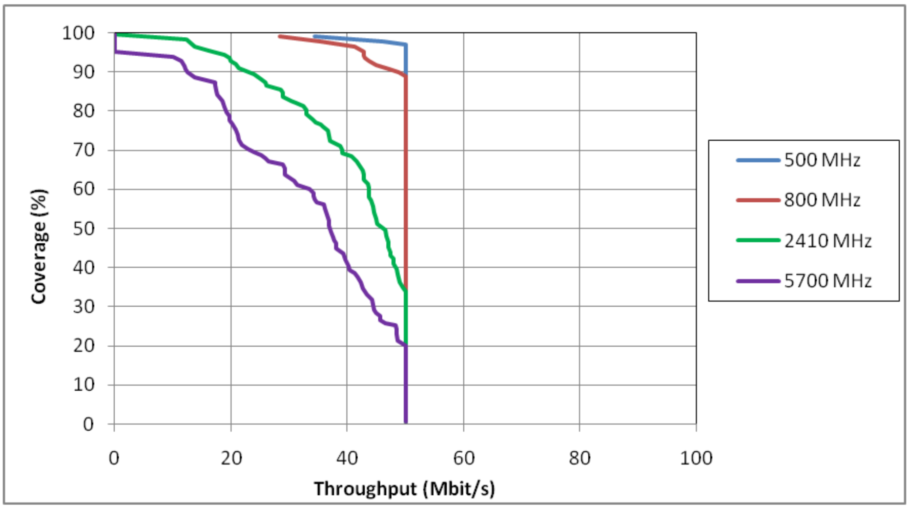

Figure 9.11: Predicted coverage distribution, UK weighted average, MIMO

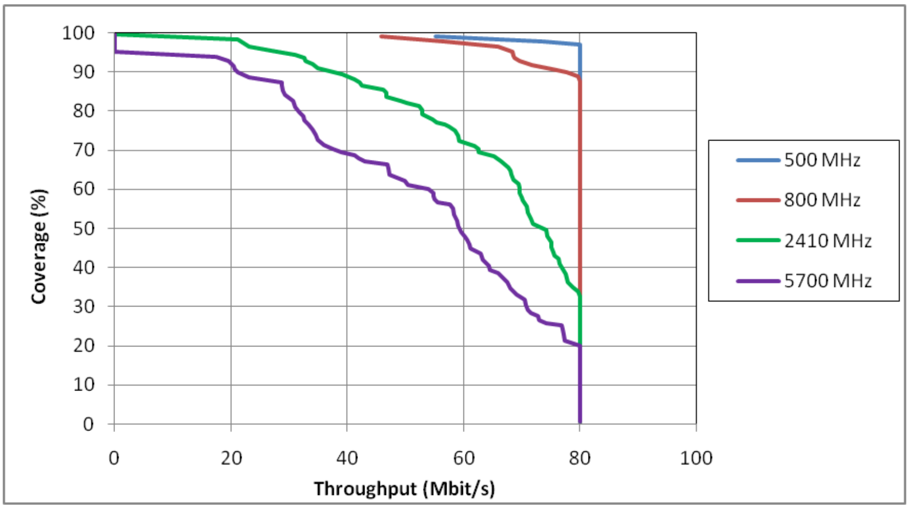

These figures look rather strange, but only because the predicted coverage at 500 MHz and 800 MHz is so high that throughput reaches the saturation level for ~90% of locations. The coverage at UHF is essentially binary: a throughput requirement less than the threshold value give coverage almost everywhere, while naturally there is no coverage at all for a throughput requirement greater than the saturation level.

Comparing the predicted coverage distributions at 2.4 GHz and 5.7 GHz with those based directly on the measurements in Figure 9.3 give confidence in the modelling. The curves agree very well, the differences being mainly at the two ends of the curves where the threshold and saturation effects dominate. The simplistic treatment of these effects could perhaps be improved.

Of course applying the threshold and saturation limits for throughput means that we are building in the system-dependent limitations of the IEEE 802.11 protocols. Since saturation is clearly limiting the system performance at UHF it is worth asking what sort of performance could be achieved if the limitation were removed.

It is not entirely consistent to simply remove the threshold and saturation effects from the observed SISO and MIMO performance curves. To achieve throughputs higher than the saturation levels would require higher modulation schemes and/or higher bandwidths and so would no longer be IEEE 802.11n. The performance curves may therefore be different. But extrapolating our modelled performance curves beyond the threshold will at least give some idea of the “propagation limited” throughput achievable.

The extrapolation of the performance curve requires a minor change to Equation 9.2 to prevent the throughput being less than zero for small or negative values of C/N:

$$T = \max(\alpha C - 3.3, 0) \tag{9.6}$$

The extrapolated coverage for the UK average is shown in Figure 9.12 and Figure 9.13 for SISO and MIMO respectively.

Figure 9.12: Predicted coverage, UK weighted average, SISO extrapolated

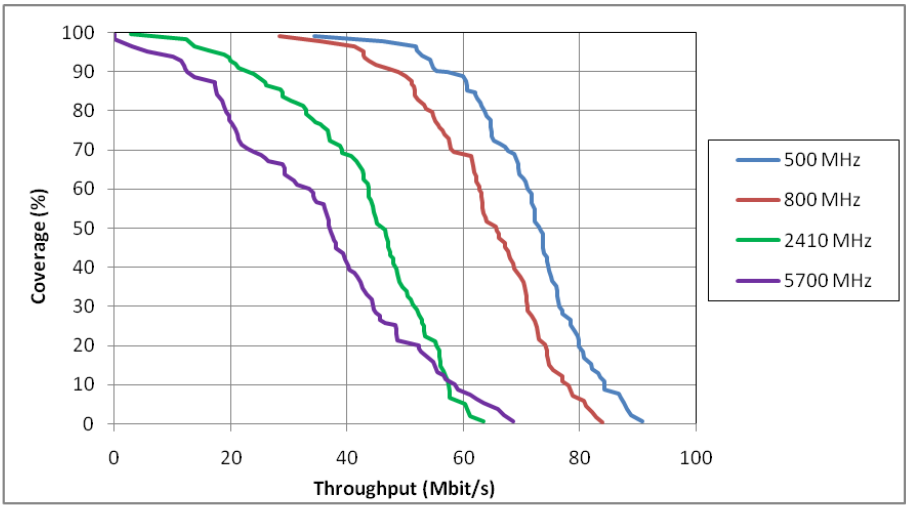

Figure 9.13: Predicted coverage, UK weighted average, MIMO extrapolated

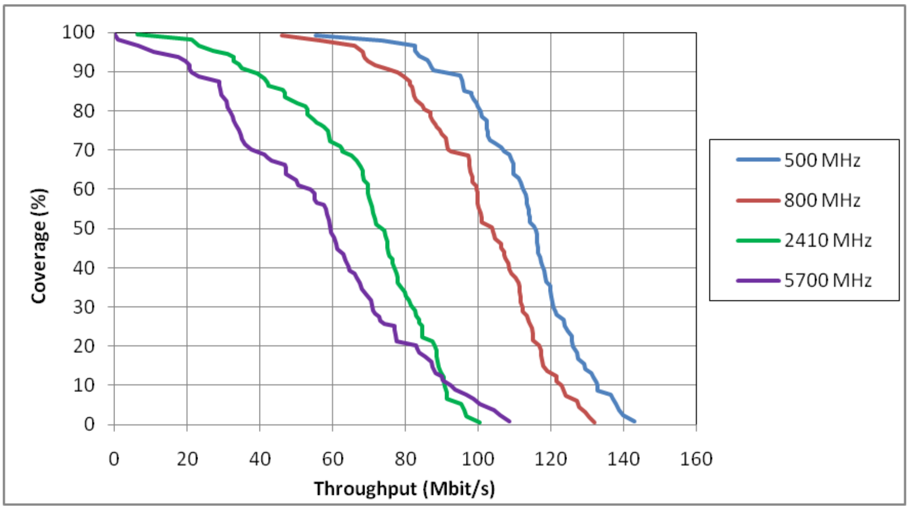

Figure 9.13 shows that throughput at UHF can theoretically exceed 120 Mbps for 30% of locations (note the extended throughput scale in this figure). Of course achieving this in practice requires adequate spectrum to be available.

Above saturation, the 2.4 GHz and 5.7 GHz coverage distributions are predicted to cross. The significance of this is unknown, but may be a receiver saturation effect in the CW measurements for the points giving the lowest transmission losses.

As expected MIMO coverage is better than SISO coverage. For each of MIMO and SISO, general conclusions are:

- Coverage decreases with frequency.
- The shape of the coverage distributions is broadly similar at all four frequencies.
- The distributions are offset along the throughput axis by an amount that reflects differences in the transmission losses at the four frequencies. The main influence is the 20 log(frequency) dependence of the basic free space transmission loss, although there are some propagation effects from the building structure (see Section 8.4).

Even although we have removed the saturation and threshold constraints of the SISO and MIMO systems from the above coverage distributions, they are still system-specific in one important aspect. Since they were derived using the performance curves based on an effective C/N, if the carrier level at the receiver is increased, the system would provide greater throughput and coverage. There is one easy way to increase the carrier level and that is to increase the transmitted power.

Predicted coverage distributions therefore depend on the transmitted power used by the system whose coverage is being considered. *The coverage distributions shown here have assumed the same ERP at all four frequencies*. This is not unreasonable. The transmitter power used for the throughput measurements at 2.4 GHz and 5.7 GHz was 20 dBm, which is common in consumer equipment of this type. This power level would not be unreasonable at UHF, and so *the power level assumed at all four frequencies was 20 dBm*.

The coverage distributions can be drawn for different transmitted powers by simply translating the distribution along the throughput axis by an amount corresponding to the slope parameter α. Specifically the throughput increases by 13.2 Mbps (SISO) or 20.5 Mbps (MIMO) for each 10 dB increase in power, independent of frequency, subject of course to any saturation implied by the protocol used.

So although Figure 9.12 and Figure 9.13 suggest that coverage at UHF is “better” than at 2.4/5.7 GHz, another way of looking at this is to say that the coverage at UHF is similar to that obtained at 2.4/5.7 GHz when a significantly lower power level (in this case ~ 20 dB lower, that is ~0 dBm) is used at UHF.

To summarise, there is one major caveat in interpreting the coverage distributions:

- The distributions assume the same ERP of 20 dBm at each frequency. The coverage at any frequency can be enhanced or reduced by increasing or decreasing the transmitted power. The distributions are shifted along the throughput axis by 13.2 Mbps (SISO) or 20.5 Mbps (MIMO) for each 10 dB increase in power, independent of frequency, subject of course to any saturation implied by the protocol used.

For completeness the predicted, unconstrained coverage distributions for each house (averaged over all AP positions) are given in Section E.2. Coverage distributions that include the saturation and threshold effects can be deduced from these by applying the constraints of Equation 9.5. This essentially means limiting the throughput to 50 Mbps for SISO and 80 Mbps for MIMO—the threshold effect at low throughput levels is minimal.

10 CONCLUSIONS

Measurements have been carried out in a selection of houses broadly representative of the UK housing stock.

CW measurements of path loss were made simultaneously at four frequencies, with continuous logging of signals received at a 'hub location' from a transmitter carried around each room of interest.

Measurements of throughput were made using standard 802.11n equipment. Significant problems were encountered in configuring 802.11n equipment to run at anything close to the rates expected; these problems were largely due to device driver issues, and were only resolved when an integrated solution in the form of a high-end laptop was adopted.

In the course of the measurements, the following observations were made:

- Coverage, defined either in terms of 'points where connection was possible' or 'throughput speed attained' was invariably better at 2.4 GHz than at 5.7 GHz.
- The statistics of CW signal variation within rooms were similar at all four frequencies.
- No statistical difference was observed between the coverage to points around the room perimeter and the coverage within the entire room.
- In general, there was no significant difference in coverage with interior doors open or closed, although one house did exhibit a slight excess loss at 5 GHz with doors closed.

The throughput measurements alone, when weighted according to the UK distribution of the house types in which they were made, allow the coverage of homes at 2.4 GHz and 5 GHz to be estimated. The results of the throughput and path-loss measurements showed a clear correlation, broadly similar at both 802.11n frequencies. This allowed the results for the two higher frequencies to be extrapolated to 500 and 800 MHz.

Cumulative distribution functions have been derived for which predict the area coverage that can be expected in the typical UK house at the four frequencies, at bitrates corresponding to the use of 2x2 MIMO or SISO technology in 20 MHz channels.

For the example of MIMO coverage, 80% of a typical house would receive a throughput of 30 Mbps at 5 GHz, 55 Mbps at 2.4 GHz, 85 Mbps at 800 MHz and 100 Mbps at 500 MHz. These coverage figures assume that the same system configuration (including transmitted power) is used at all frequencies, and that adequate spectrum is available. Furthermore, it should be noted that these figures take no account of possible interference from unrelated, neighbouring systems.

A REFERENCES

[Behzad, 2008] Behzad, A., “*Wireless LAN radios*”, Wiley-Interscience, 2008.

[Cheung, 2002] Cheung, D., Prettie, C., “*A Path Loss Comparison Between the 5 GHz UNII Band (802.11a) and the 2.4 GHz ISM Band (802.11b)”, Intel Labs, January 2002.*

[EHS, 2008] “English Housing Survey—Housing stock report 2008”, Department for Communities and Local Government, October 2010 (available at: http://www.communities.gov.uk/publications/corporate/statistics/ehs200809headlinereport).

[Erceg, 1999] Erceg et al., "*An Empirically Based Path Loss Model for Wireless Channels in Suburban Environments*" *IEEE J. on Selected Areas in Communication*, vol. 17, pp. 1205-1211, July 1999.

[Mikas, 2008] “Measurement and prediction of signal propagation for WLAN systems”, Czech Technical University in Prague, 2008 (available at http://www.urel.feec.vutbr.cz/ra2008/archive/ra2003/papers/205.pdf).

[Novarum, 2009] “*Enterprise 802.11n Wireless LAN Access Point Benchmark*”, Novarum Inc., 2009 (available at: http://www.novarum.com/publications.php).

[Perahia, 2008] Perahia, E. and Stacey, R., “*Next Generation Wireless LANs*”, Cambridge University Press, 2008.

Saleh, 1987] Saleh, A.A.M. and Valenzuela, R. A., *"A Statistical Model for Indoor Multipath Propagation*", IEEE J. on Selected Areas in Communications, vol. SAC-5, pp. 128-137 (February 1987).

[Suzuki, 1977] Suzuki, H., “*A statistical model for urban radio propagation”,* IEEE Transactions on Communications, 25(7), 673-680, 1977.

[Turin, 1972] Turin et al., "*A statistical model of urban multipath propagation*", IEEE Trans. on Vehicular Technology, vol. VT-21, pp. 1-9, February, 1972.

B IEEE 802.11 STANDARDS

This annex provides an overview of the aspects of the 802.11 standards that are relevant to the present project.

B.1 Overview & historical development

The IEEE 802.11 working group created the first Wireless LAN standard in 1997. In its first incarnation, the physical layer operated in the 2.4 GHz band and supported 1 Mbps and 2 Mbps data rates using either direct sequence (DSSS) or frequency hopping spread spectrum (FHSS) techniques.

In 1999 the 802.11b extension was published, extending data rates to 5 Mbps or 11 Mbps through the substitution of Complementary Code Keying (CCK) for the original Barker spreading codes. This extension also operated in the 2.4 GHz band.

In the same year, the 802.11a amendment was also ratified, which introduced an entirely new modulation type (OFDM in 20 MHz channels) and a new operating frequency (5 GHz). In 2003, the same approach was adopted for use at 2.4 GHz, being standardised as 802.11g. The delay in this standardisation was largely due to the need to ensure compatibility with legacy 802.11 and 802.11b systems operating in the same band.

The latest amendment to the standard is 802.11n, which embraces both the 802.11a and 802.11g standards, but also allows the use of techniques such as 'channel bonding' (operation in 40 MHz channels) and MIMO.

B.2 802.11 a & g

These standards essentially differ only in their frequency bands of operation (although 802.11g also includes features to allow interworking with legacy equipment).

The physical layer uses OFDM modulation, with 52 subcarriers (four of which are pilots) in a nominal 20 MHz channel (carrier spacing of 312.5 kHz and guard interval of 0.8 μs). The modulation and code rate can be changed dynamically to make optimum use of the radio channel. The options available are shown in Table B.1 below.

Table B.1: 802.11a,g modulation and coding

Data rate (Mbps)	Modulation	Code rate	Required sensitivity
6	BPSK	1/2	−82 dBm
9	BPSK	3/4	−81 dBm
12	QPSK	1/2	−79 dBm
18	QPSK	3/4	−77 dBm
24	16-QAM	1/2	−74 dBm
36	16-QAM	3/4	−70 dBm
48	64-QAM	2/3	−66 dBm
54	64-QAM	3/4	−65 dBm

The sensitivities achieved by modern devices typically exceed the requirement by some 10 dB.

B.3 802.11n

The new 802.11n standard essentially takes the existing a and g standards and adds spatial multiplexing (MIMO), the ability to use 40 MHz channels (channel bonding) and other options to improve throughput.

Spatial multiplexing, with support for a minimum of two streams, is mandatory except for handheld terminals, but channel bonding is optional (and often not permitted by the firmware at 2.4 GHz so as to avoid interference).

The spatial multiplexing allows up to four parallel streams, and this, in combination with the use of 40 MHz channels, would allow a maximum possible throughput approaching 600 Mbps.

Table B.2: 802.11n modulation and coding (800 ns GI, 20 MHz channel)

MCS	Data rate (Mbps)	Spatial streams	Modulation	Code rate	Required sensitivity
0	6.5	1	BPSK	1/2	*−82 dBm*
1	13	1	QPSK	1/2	*−81 dBm*
2	119.5	1	QPSK	3/4	*−79 dBm*
3	26	1	16-QAM	1/2	*−77 dBm*
4	39	1	16-QAM	3/4	*−74 dBm*
5	52	1	64-QAM	2/3	*−70 dBm*
6	58.5	1	64-QAM	3/4	*−66 dBm*
7	65	1	64-QAM	5/6	*−65 dBm*
8	13	2	BPSK	1/2	*−82 dBm*
9	26	2	BPSK	3/4	*−81 dBm*
10	39	2	QPSK	1/2	*−79 dBm*
11	52	2	QPSK	3/4	*−77 dBm*
12	78	2	16-QAM	1/2	*−74 dBm*
13	104	2	16-QAM	3/4	*−70 dBm*
14	117	2	64-QAM	2/3	*−66 dBm*
15	130	2	64-QAM	3/4	*−65 dBm*
-	-	-	-	-	-
31	260	4	64-QAM	5/6	

The shaded schemes are those for which the same modulation type and code rate combinations are available in 802.11a or 802.11g type equipment.

Note that the bit rates achieved are somewhat different, as 802.11n has 52 data subcarriers compared to the 48 used in 802.11a/g.

B.4 Throughput

The bit rates given in Tables B.1 and B.2 relate to the rates available at the input/output of the PHY layer (i.e. they include the impact of guard intervals and convolutional coding, but no other overheads).

The other layers of a real-world communications protocol will add significant additional overheads, associated with, for example, flow control and security. A common rule of thumb is that the useful bitrate achieved from an 802.11g system will not exceed around 22 Mbps.

C CALIBRATION OF CW EQUIPMENT

This annex describes the calibration procedure used with the CW measuring equipment. The equipment itself, the measurement method, and the results, are given in Annex D.

C.1 Arrangement of CW equipment

Figure C.1 is a schematic diagram of the measurement arrangement. The four receiving antennas are in the same square layout as the transmitting antennas. They are mounted with vertical polarisation on an aluminium base-plate which can be oscillated through 180° around their vertical axis, each cycle taking of the order of 10 seconds. The overall assembly is mounted on a stable tripod.

Figure C.1: Measuring equipment schematic

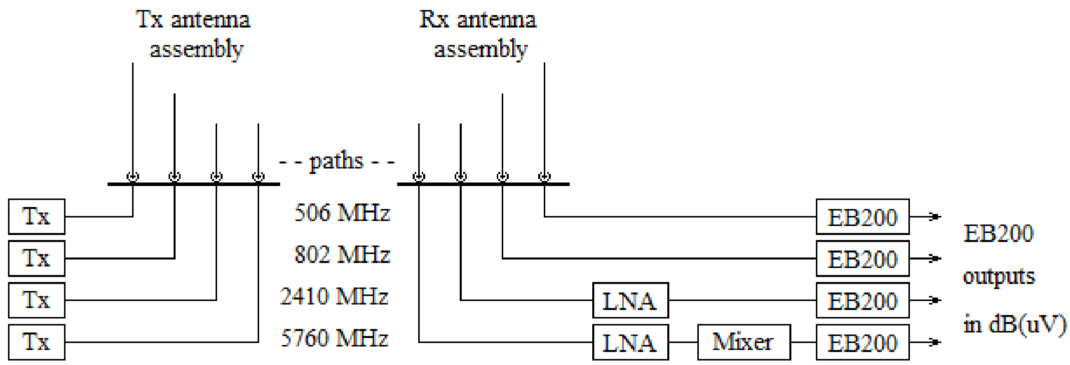

Low-noise amplifiers precede the EB200 measuring receivers for the upper two frequencies. A down-converting mixed is also needed at the highest frequency.

C.2 Calibration parameters

Figure C.2 shows the parameters considered in the calibration process in relation to one representative transmit-receive chain.

Figure C.2: Calibration parameters

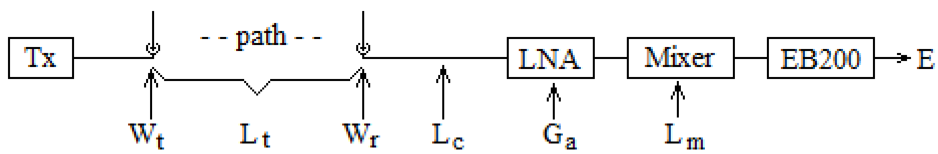

Table C.1 provides further details on these parameters.

Table C.1: Calibration Parameters

Symbol	Units	Comments
W_t	dB(mW)	Transmitter power at the transmitting antenna terminals.
W_r	dB(mW)	Received power at the receiving antenna terminals.
L_t	dB	Transmission loss, equal to basic transmission loss minus the combined antenna gains.
L_c	dB	Cable losses between the receiving antenna and EB200 input. Separate values are needed for two lengths of cables used.
G_a	dB	Low noise amplifier gain. Zero in 506 and 802 MHz channels.
L_m	dB	Mixer insertion loss. Zero in 506, 802 and 2410 MHz channels.
E	dB(µV)	EB200 indication.

The parameters in Table C.1 are related by the following expressions:

$$L_t = W_t - W_r \quad \text{dB} \qquad \text{(C1)}$$

where

$$W_r = E - 107 + L_m - G_a + L_c \quad \text{dB(mW)} \qquad \text{(C2)}$$

The constant 107 converts from dB(µV) to dB(mW) on a 50-Ohm transmission line.

Thus

$$L_t = W_t - E + 107 - L_m + G_a - L_c \quad \text{dB} \qquad \text{(C3)}$$

It is convenient to rearrange Equation C3 in the form

$$L_t = C - E \quad \text{dB} \qquad \text{(C4)}$$

where

$$C = W_t + 107 - L_m + G_a - L_c \quad \text{dB} \qquad \text{(C5)}$$

The parameter *C* thus incorporates all calibration information required to convert an EB200 indication *E* in dB(µV) to transmission loss L_t in dB.

Two sets of values for the final calibration value C applicable to the use of 'short' and 'long' cables between the receiver antennas and the rest of the receiving equipment. In fact 'long' cables consisted of the 'short' cables plus a plugged-in extension, primarily to permit oscillation of the receiving antennas. In practice, after some initial tests without rotation, long cables were normally used for all CW measurements.

Table C.2 shows values obtained from manufacturer's data sheets and measurements conducted by the project to obtain the final values of parameter C for short and long receiver cables.

Table C.2: Calibration results

Description	Symbol					Units
Frequency		506	802	2410	5760	MHz
RG58A specific cable loss		0.4	0.5	1.1	2.1	dB/m
Cable length (short)		1.0	1.0	1.0	1.0	m
Cable length (long)		2.5	2.5	2.5	2.5	m
Cable loss (short)	L_c	0.4	0.5	1.1	2.1	dB
Cable loss (long)	L_c	1.0	1.3	2.7	5.1	dB
TX power (after 30-min warm-up)	W_t	18.0	16.0	20.4	15.9	dB(mW)
LNA gain	G_a	0.0	0.0	28.2	30.9	dB(mW)
Mixer loss	L_m	0.0	0.0	0.0	10.8	dB
Calibration (short cables)	C	124.7	122.5	154.5	141.0	dB
Calibration (long cables)	C	124.1	121.7	152.9	137.9	dB

As a check on the above calibration results, EB200 indications were noted for two levels of received power injected directly into the receiving antenna sockets. This test was conducted with short receiver cables. Table C.3 shows the results.

Table C.3: Calibration tests with injected received power

Tables values are signal levels in dB(mW) at receiver inputs	Injection level dB(mW)	Frequency, MHz			
		506	802	2410	5760
Measured levels using calibration data from Table C.2	−20.0	−20.0	−18.5	−21.03	−22.5
	−80.0	−80.1	−79.5	−80.8	−80.0
Discrepancy (measured minus injected levels)	−20.0	−0.0	1.51	−1.0	−2.5
	−80.0	−0.1	0.51	−0.8	0.0

The level of agreement shown in Table C.3 is considered adequate for the purposes of the required measurements. The errors at −20 dBm for the higher frequencies are due to LNA gain.

C.3 Antenna characteristics

The 2 GHz and 6 GHz antennas used for the CW measurements were each a dual-band unit manufactured by Antenna Factor. Figures C.3 and C.4 give the manufacturer's information for vertically-polarised azimuthal gain, and Figures C.5 and C6 give one cut through the elevation pattern[9]. In each case there are separate figures for 2.45 GHz and 5.30 GHz.

Figure C.3: Azimuthal gain at 2.45 GHz

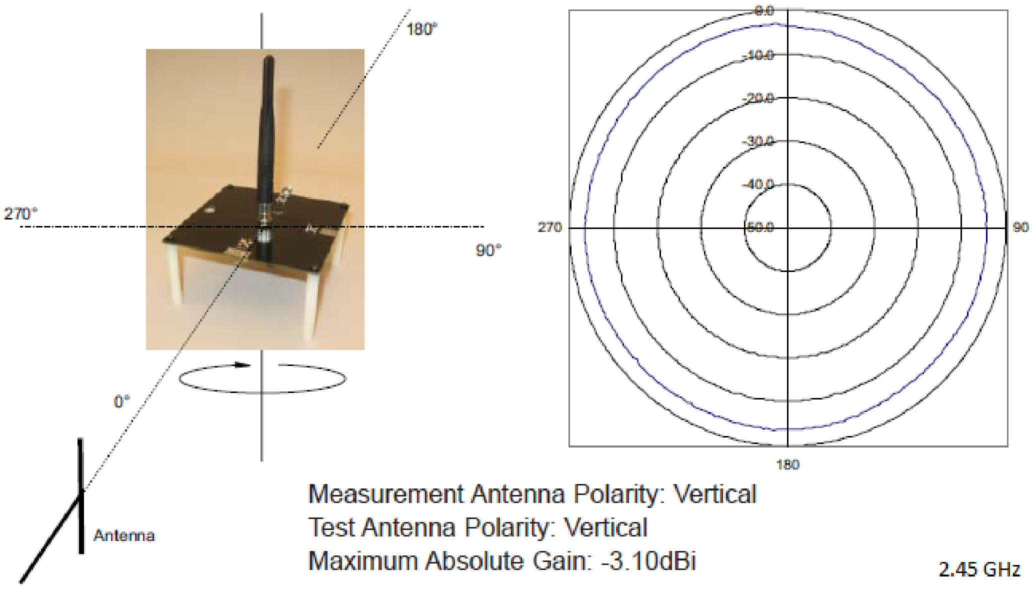

Figure C.4: Azimuthal gain at 5.30 GHz

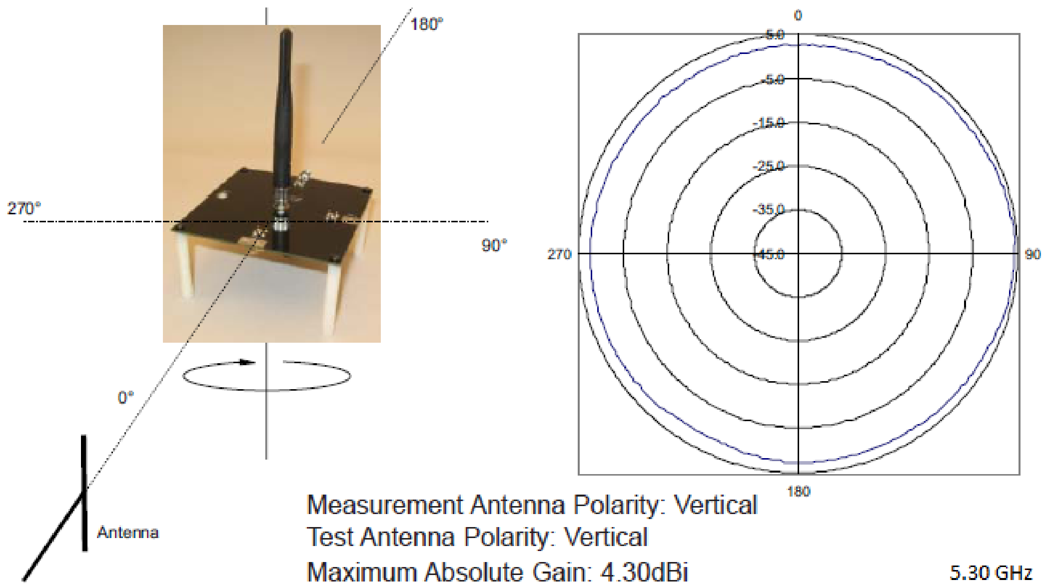

[9] There is an apparent contradiction in the horizon gain values shown in figures C.3 (~−3 dBi at all azimuths) and C.5 (<−10 dBi at the two horizon points). The manufacturer was contacted, but no explanation was forthcoming.

Figure C.5: Elevation pattern at 2.45 GHz

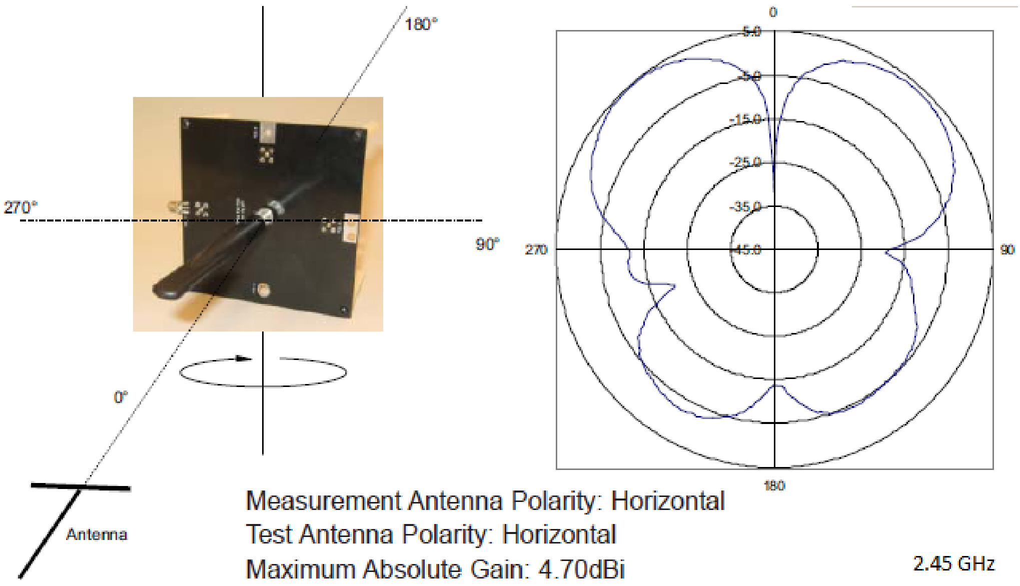

Figure C.6: Elevation pattern at 5.30 GHz

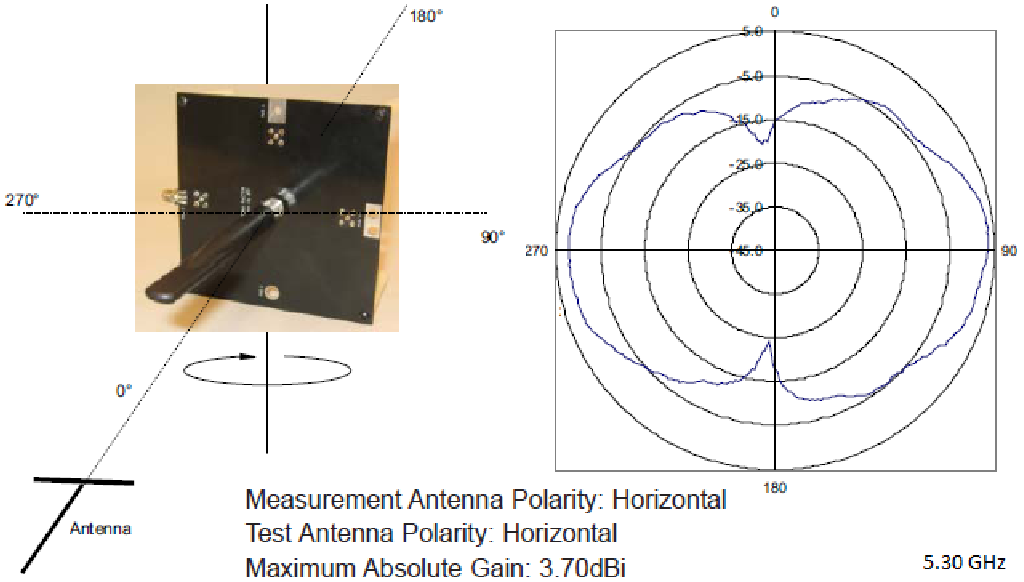

The variations in azimuthal gain are estimated from Figures C.3 and C.4 as follows:

2.45 GHz	−4.5 to −3.0	dBi
5.30 GHz	+2.0 to +4.0	dBi

The variation with elevation angle both above and below the mounting base-plate are estimated from Figures C.5 and C.6 as follows:

2.45 GHz	−33 to +5	dBi
5.30 GHz	−24 to +3	dBi

The antennas used for 500 MHz and 800 MHz were sleeve dipoles constructed within the project, and detailed characterisation of each individual antenna in isolation was not practicable within the resources available.

The manufacture's information on the UHF antennas is useful, although not given at the actual frequencies used in the project. A further point of difference is that the CW measurements were made with 4-antenna assemblies for both the transmitters and receivers. Some interaction between the antennas on each common baseplate can be assumed.

To provide more information on the actual antenna gains during the CW measurements the equipment was used in an anechoic chamber.

Most of the anechoic-chamber results were between the two groups of four vertical antennas separated horizontally by 6.4 m. The receiving antenna group was oscillated through 180° as normal during in home-measurements. Measurements were run for eight rotational positions of the transmitting antenna assembly around a vertical axis at 45° intervals.

The long receiver cables were in use. Transmission losses were calculated using the calibration corrections used for other narrow-band measurements as given in Table C.2.

Figure C.7 shows the concatenated time-series for transmission losses measured at the four frequencies. The receiving antenna assembly is oscillating through 180°, and there are 200 points at each of 8 orientations of the transmitting antenna assembly. These are referred to as the azimuthal antenna measurements.

Transmission losses at all frequencies show considerable variation with antenna-assembly orientation. Some of these variations, particularly at 2.4 GHz, are characteristic of multi-path propagation. In an anechoic chamber this is attributed to passive re-radiation from different antennas in each assembly.

Figure C.7: Anechoic chamber transmission loss time series

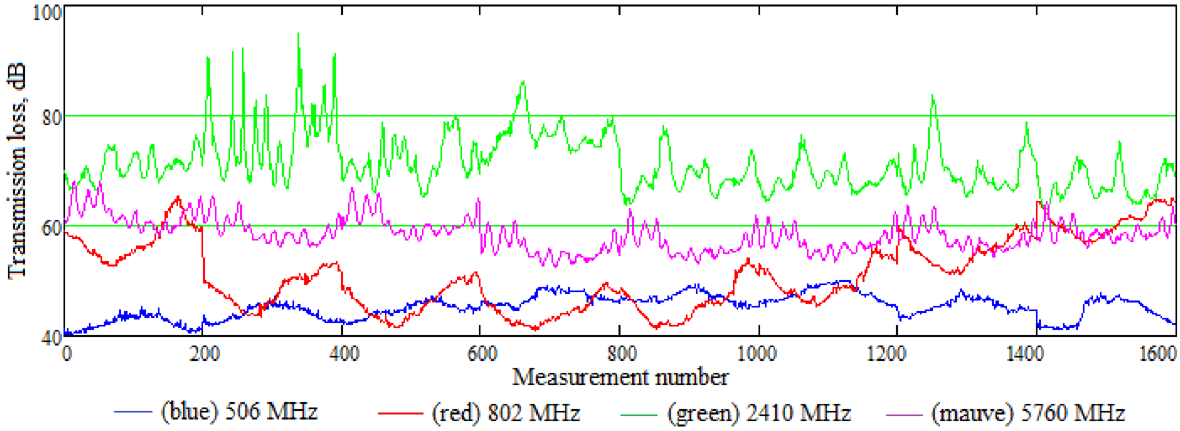

To assess the net effect of antenna gains, the following analysis process was followed. At each frequency:

a) all results from the eight series were aggregated and the median values of transmission loss, L_t in dB, was extracted

b) the free-space basic transmission loss, L_{bfs} in dB, was calculated for the 6.4 m range

c) combined antenna gains in dBi were calculated given by $L_{bfs} – L_t$.

Table C.4 shows the results for these horizontal antenna-gain measurements. They indicate clearly that the net effect of antenna gain varies irregularly with frequency. In the last row the 'net' antenna gains are inferred from the measurements calculated for each frequency as the calculated free-space basic transmission loss over a distance of 6.4 m minus the median measured transmission loss. This row shows the 'combined' gains on this basis, that is, the sum of the two inferred gains.

Table C.4: Antenna-gain results

Frequency	MHz	506	802	2410	5760
Median measured L_t	dB	45.5	50.3	70.2	58.3
Calculated L_{bfs}	dB	42.6	46.6	56.2	63.8
Combined antenna gains	dBi	−2.9	−3.7	−14.0	+5.5

Table C.5 compares the manufacturer's azimuthal gains with half of the combined gains from the Table C.4 for the upper two frequency bands.

Table C.5: Comparison of manufacturer's and measured antenna gains

Band	Antenna gain, dBi		
	Manufacturer's		Measured
	Minimum	Maximum	
2 GHz	−4.5	−3.0	−7.0
6 GHz	+2.0	+4.0	+2.7

Both manufacturer's information and the anechoic measurements show that the dual-band antennas have significantly higher gain in the higher band, by about 6.7 dB and 9.7 dB respectively.

The discrepancies in the above table are considered acceptable when it is taken into account that:

i) the anechoic chamber measurements were not at the same frequencies as the manufacturer's results

ii) both transmitting and receiving antennas were not isolated, but in four-antenna assemblies for which some degree of interaction is inevitable.

No similar gain comparisons are available for the two lower-frequency antennas.

C.4 Transmission and propagation losses

Transmission loss is defined as the loss between transmitter and antenna terminals. It thus includes antenna gains and propagation losses.

The CW measurements were used to derive transmission losses at all four frequencies simultaneously. In general transmission losses were seen to increase with frequency, except for the 2.41 GHz and 5.76 GHz results, where either could show the larger transmission losses for different cases. It is assumed from the foregoing antenna information that his is primarily due to the lower antenna gains at 2.41 GHz.

To obtain information on how propagation losses vary with frequency, it is desirable to separate the effects of antenna gains.

For a single radio path, or at least where any multipath propagation falls within narrow ranges of take-off and arrival angles at the antennas, the antenna gains can be quantified. Under these conditions, transmission loss L_t and the antenna gains G_t and G_r, are related according to:

$$L_t = L_{bt} - G_t - G_r \quad \text{dB} \qquad \text{(C6)}$$

where L_{bt} is defined as basic transmission loss.

This approach is not applicable for indoor propagation, where propagation is usually multipath with wide ranges of take-off and arrival angles. The project did not undertake channel sounding, and thus no identification of individual rays is possible.

To remove the effects of antenna gain as far as practicable a quantity referred to as "propagation loss" is calculated as the median measured transmission loss for each room and receiver position plus the net combined antenna gain as inferred above. "Propagation loss" in this sense would be equivalent to basic transmission loss for single-ray propagation. "Propagation loss" is used in Section 8.1 of this report.

D DETAILS OF MEASUREMENT LOCATIONS

This annex presents details of the homes in which measurements have been made. The descriptions of both throughput and CW measurements refer to the following information.

Table D.1: Summary building data

Description	Abbreviation	Largest horizontal extent*	Floor area*
Victorian terraced	VT	12.5 m	91 m^2
Inter-war semi-detached	SD	15.0 m	107 m^2
Modern bungalow	MB	18.0 m	165 m^2
Modern detached	MD	11.5 m	83 m^2
Modern terraced	MT	11.2 m	66 m^2
Large Semi Detached	LSD	15.5 m	180 m^2
Basement Flat	BF	8 m	14 m^2

* excluding garages, bathrooms, hallways

D.1 Victorian terraced house

Figure D.1 shows floor plans for this Victorian terraced house, with brick external walls and most internal walls of lath-and-plaster. The second floor is a study which has been built into the roof space.

Figure D.1: Victorian terraced house

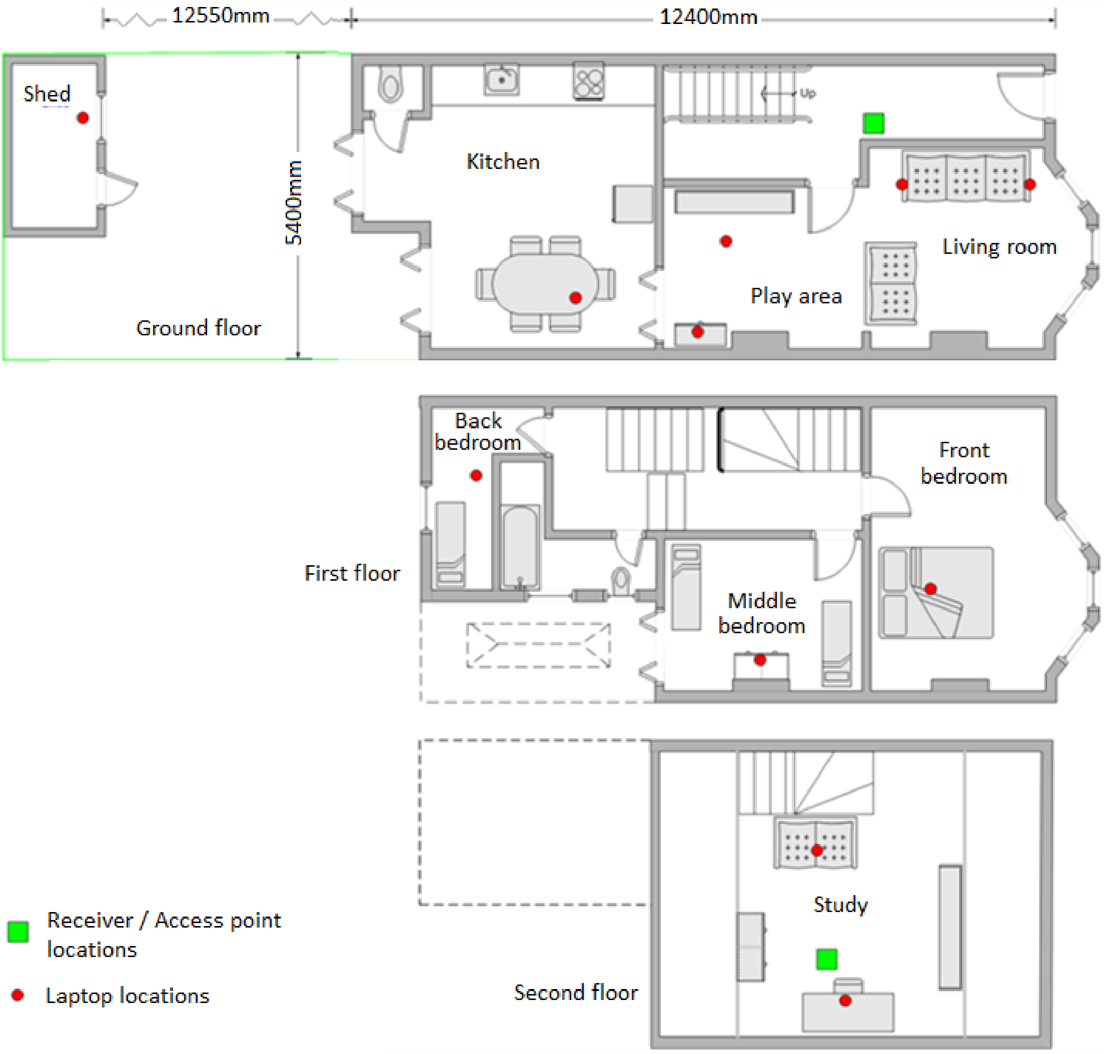

CW measurements were also made in neighbouring houses. These are identified in the results as 'No.22' and 'No.39'. The relative building locations are shown in Figure D.2.

Figure D.2: Victorian terraced house—neighbouring measurement locations

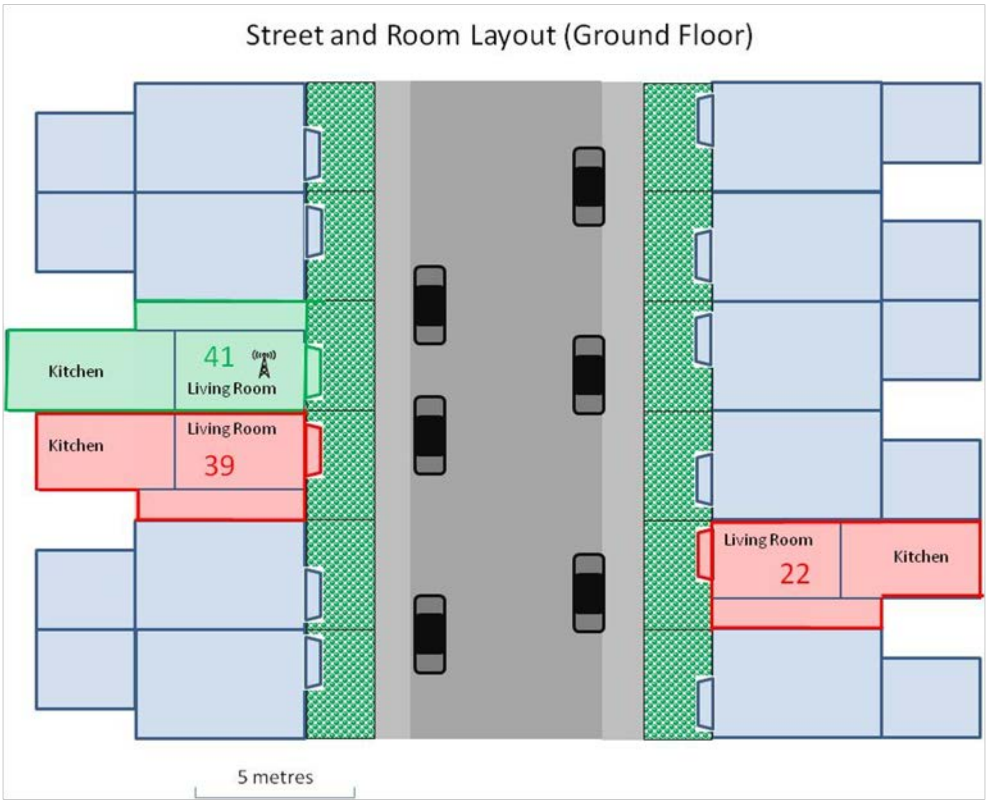

D.2 Inter-war semi-detached house

Figure D.3 shows floor plans for this 1930s chalet-style semi-detached house. The roof is dual-pitch with the eves at ground floor level. On the south side the first-floor walls are either rafters-and-tiles, or tile-hung dormers.

Figure D.3: Inter-war semi-detached house

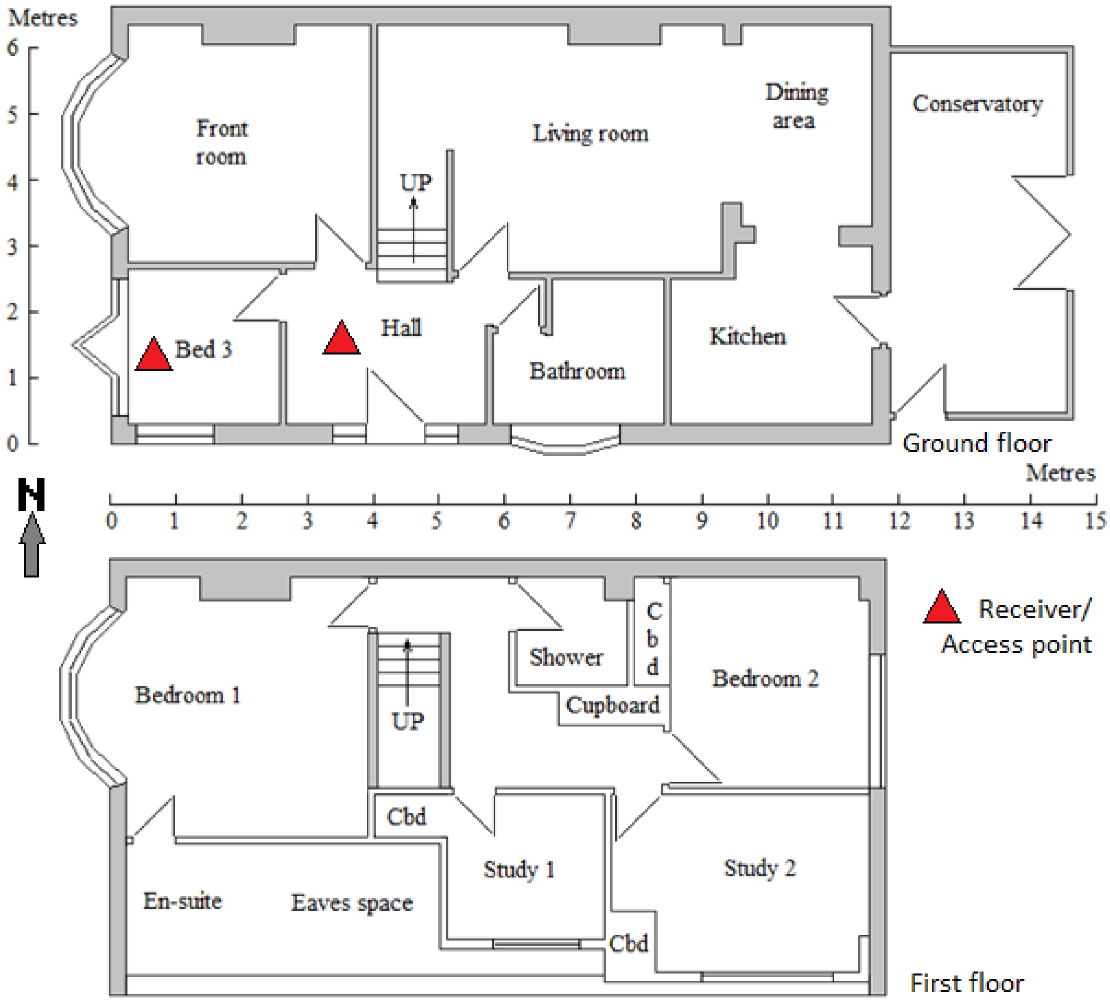

D.3 Modern bungalow

Figure D.4 shows the floor plans of this 1950s bungalow. Although this is not exactly a recent build, the construction is modern. External walls are cavity with external brick and internal block construction. Internal walls are of stud and plaster-board. The first-floor room is in the roof, aligned with the right-hand end of the floor plan.

Figure D.4: Modern bungalow

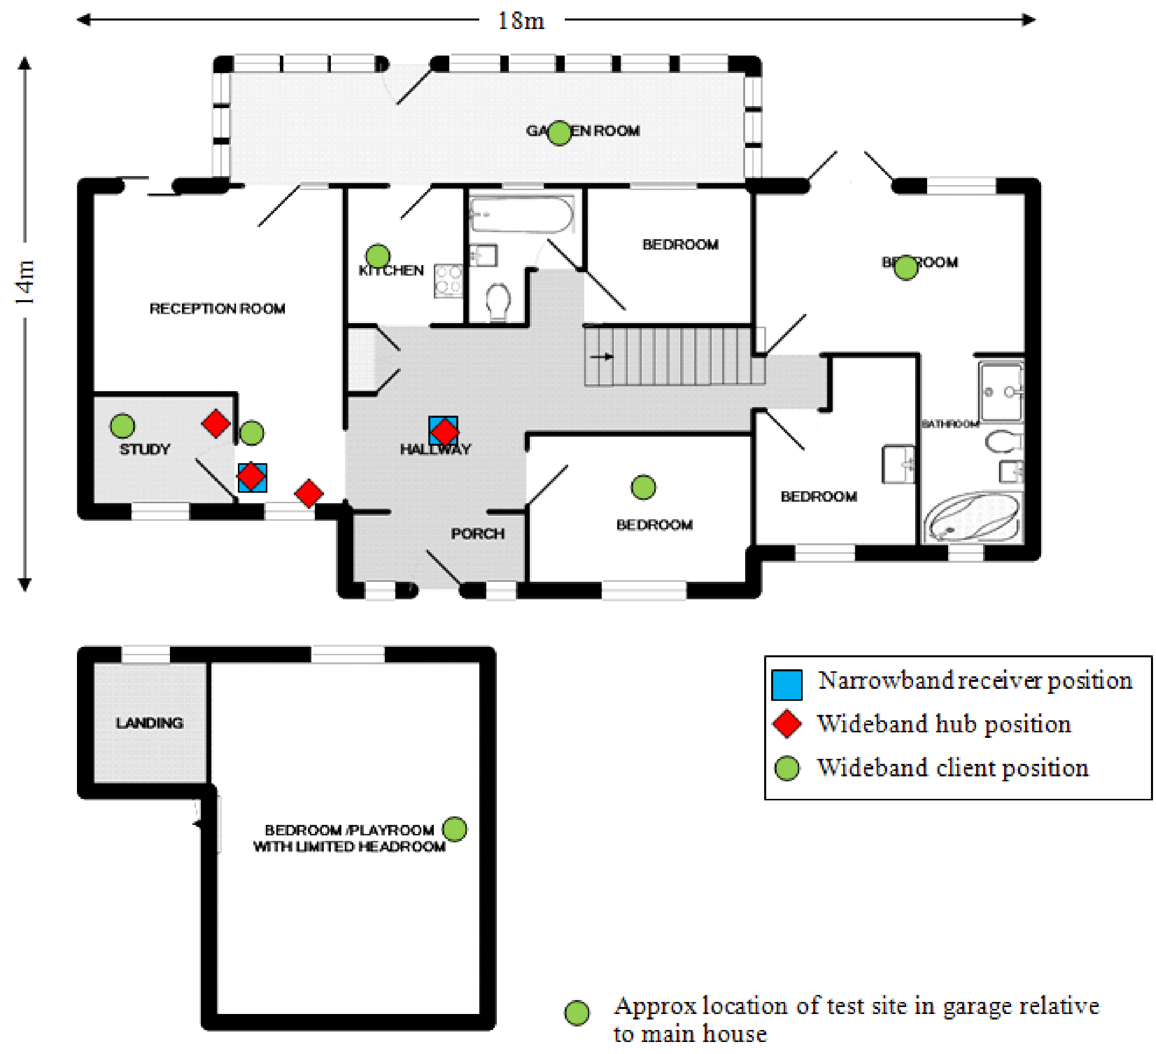

D.4 Modern detached house

Figure D.5 shows floor plans for this modern detached house.

Figure D.5: Modern detached house

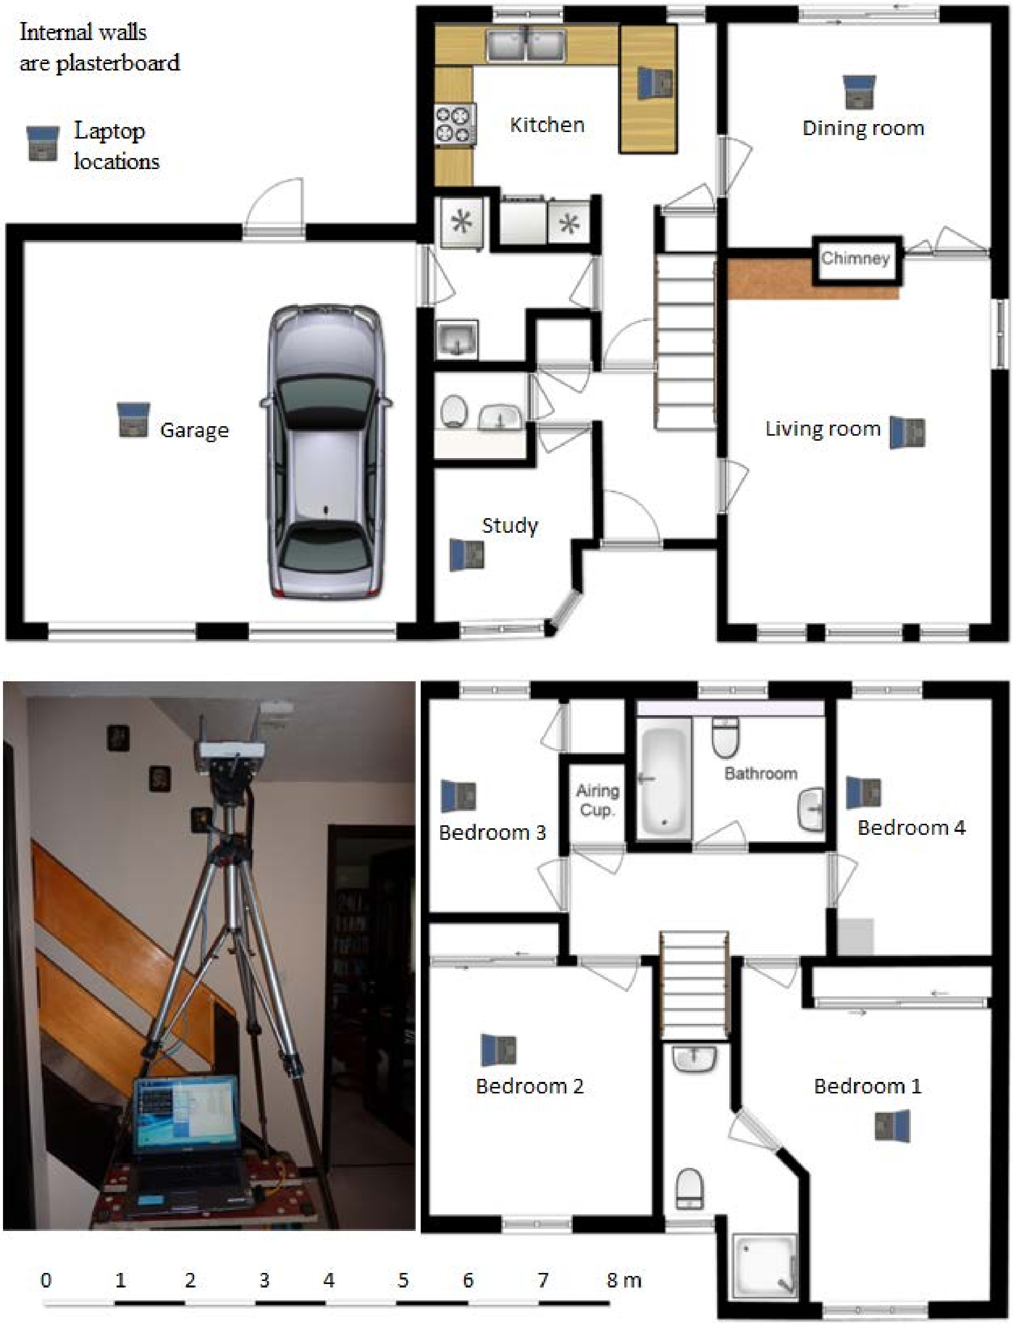

Two hub locations were used:

1. In the study, near the point of entry of the cable broadband connection, at shelf height (1.8 m). This is shown approximately by the laptop symbol in the study.

2. Downstairs hallway at ceiling level close to the stairs newel post, as shown in the insert photograph. This simulates a ceiling-mounted hub in a location that intuitively might be expected to give better coverage than the study.

D.5 Modern terraced house

Figure D.6 shows floor plans for this modern terraced house.

Figure D.6: Modern terraced house

In this house only one hub location was used. This was in the living room, located 0.85 m from the E wall and 0.65 m from the window in the S wall. This was near the main ADSL connection point. The antenna height was 1.2 m.

D.6 Large semi-detached house

Figure D.7 shows the floor plans (not including the cellars) of this spacious Victorian semi-detached house. The floor-plan is 14.5 m x 10 m.

Figure D.7: Large semi-detached house

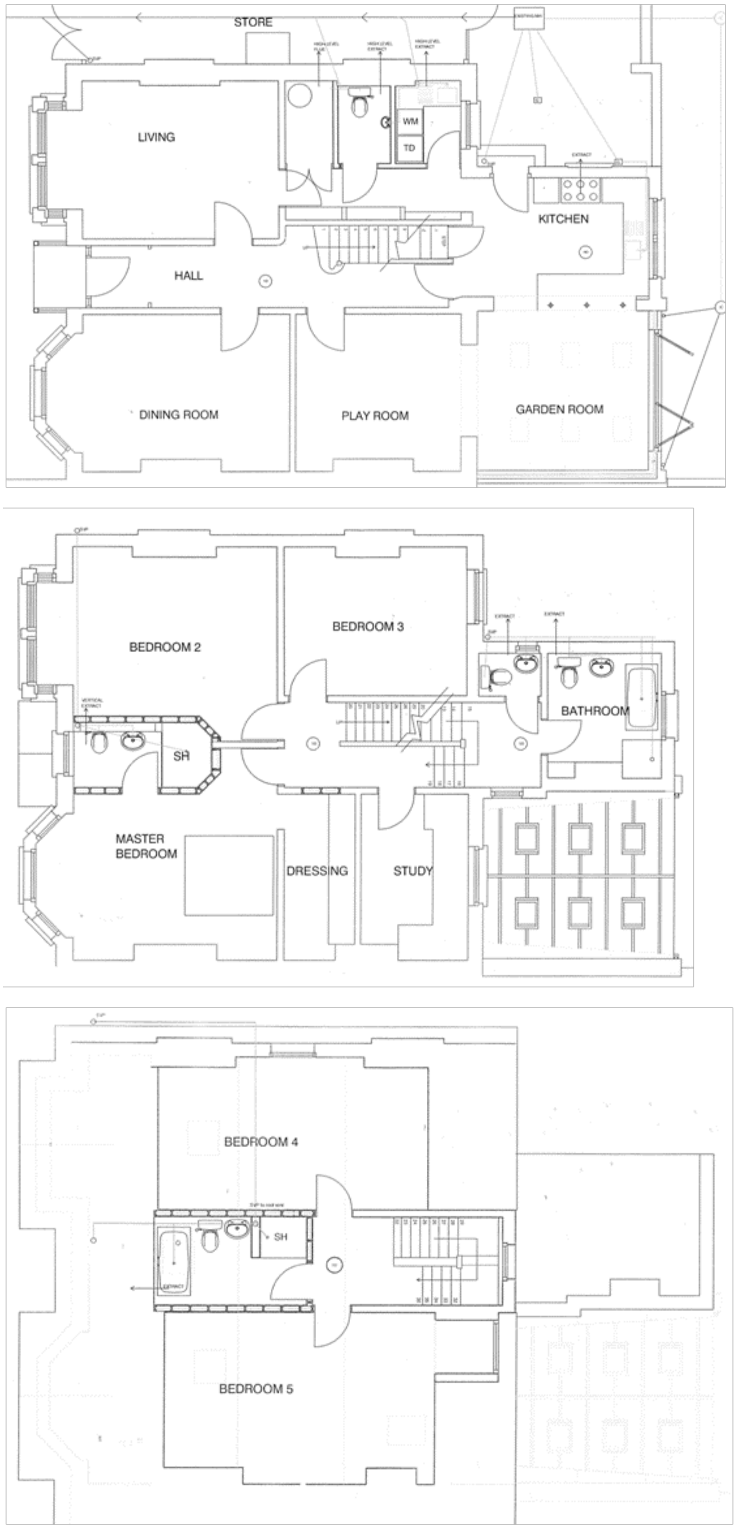

D.7 Basement flat

The layout of this small flat is sketched in Figure D.8, below.

Figure D.8: Basement flat

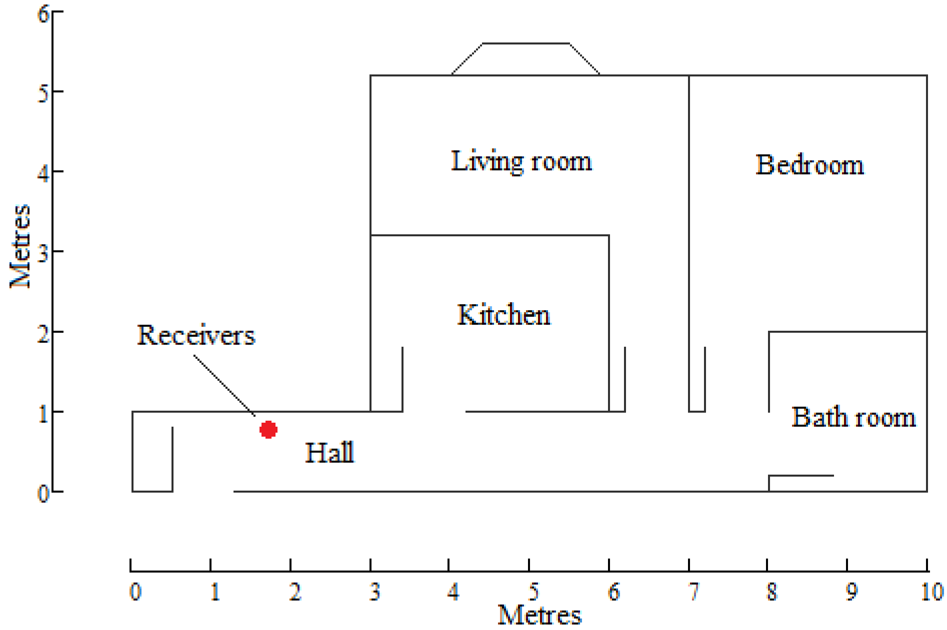

E COVERAGE RESULTS

This Annex is a catalogue of coverage distributions for the individual houses.

Section E.1 gives coverage distributions based directly on measured throughput at 2.4/5.7 GHz, for both SISO and MIMO. Each house/access point measurement set is given separately.

Section E.2 gives the predicted coverage distributions for the four frequencies, 500 MHz, 800 MHz, 2.4 GHz and 5.7 GHz. The SISO and MIMO results are given separately. Note the different axis scale for the MIMO results. These distributions assume an ERP of 20 dBm at each frequency, and ignore the system-dependent threshold and saturation effects of IEEE 802.11n. See Section 9.4 for details.

E.1 Coverage based directly on measured throughput at 2.4/5.7 GHz

E.1.1 Victorian terraced

Figure E1.1: AP in hall

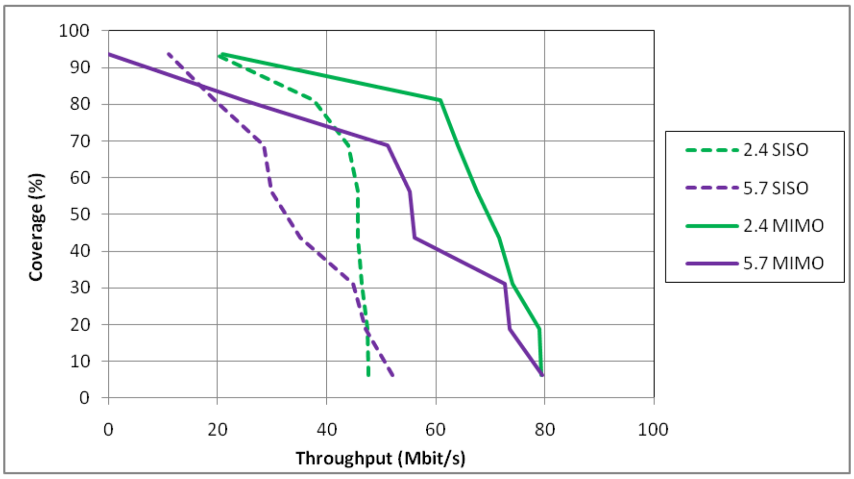

Figure E1.2: AP in study

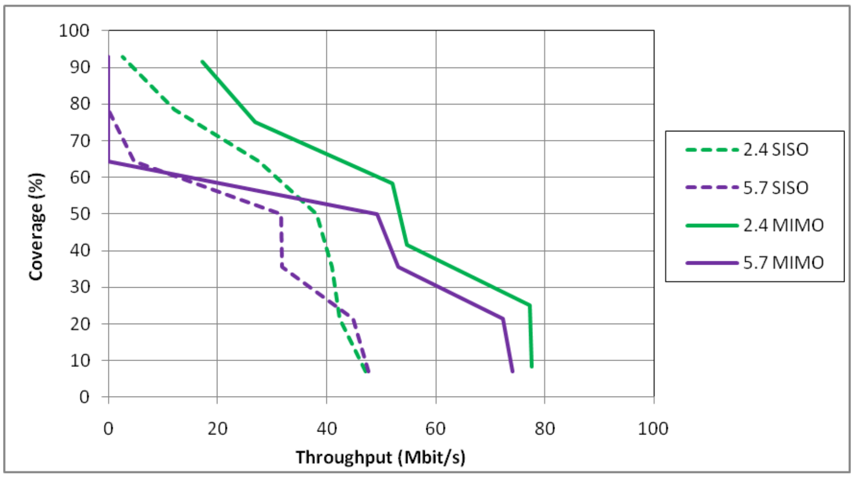

E.1.2 Inter-war semi-detached

Figure E1.3: AP in bedroom 3

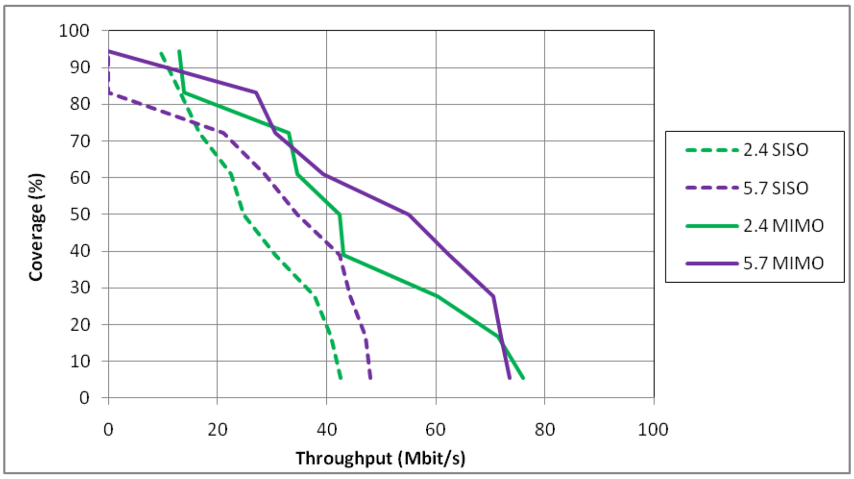

Figure E1.4: AP in hall

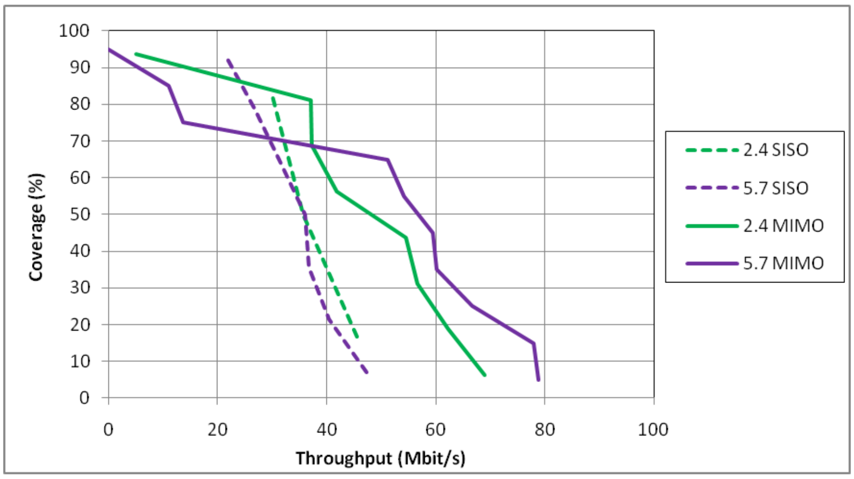

E.1.3 Modern bungalow

Figure E1.5: AP in hall

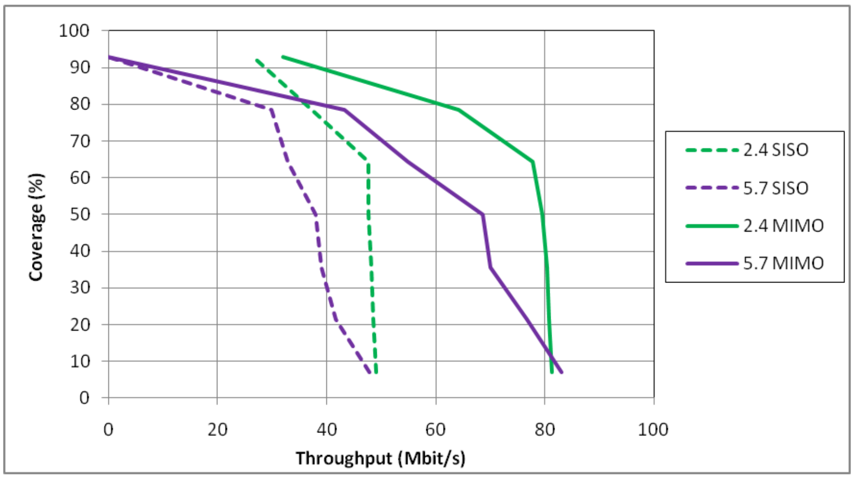

Figure E1.6: AP in reception room

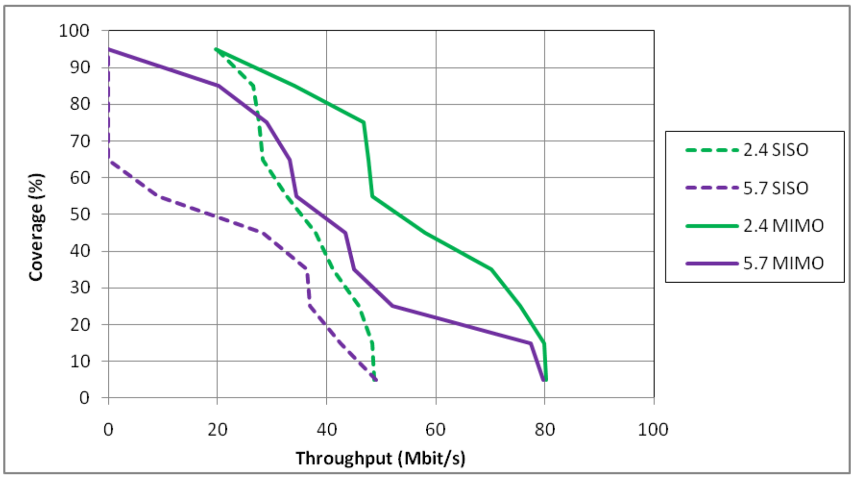

Figure E1.7: AP in study

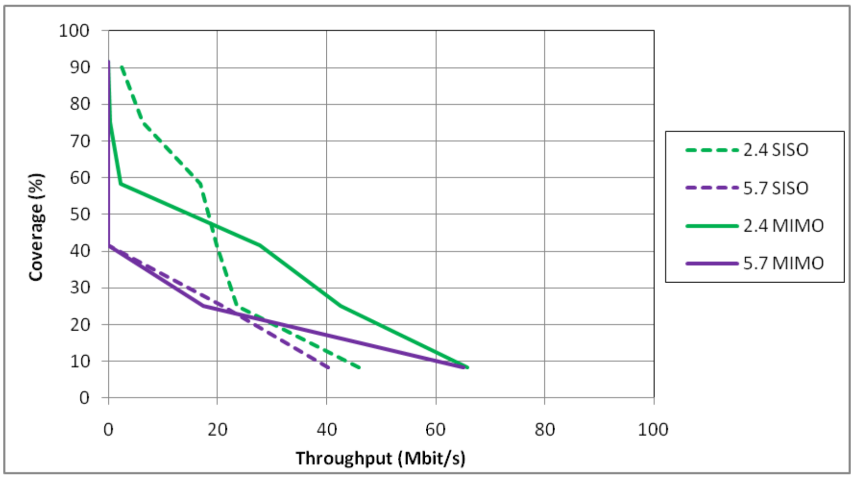

E.1.4 Modern detached

Figure E1.8: AP in hall

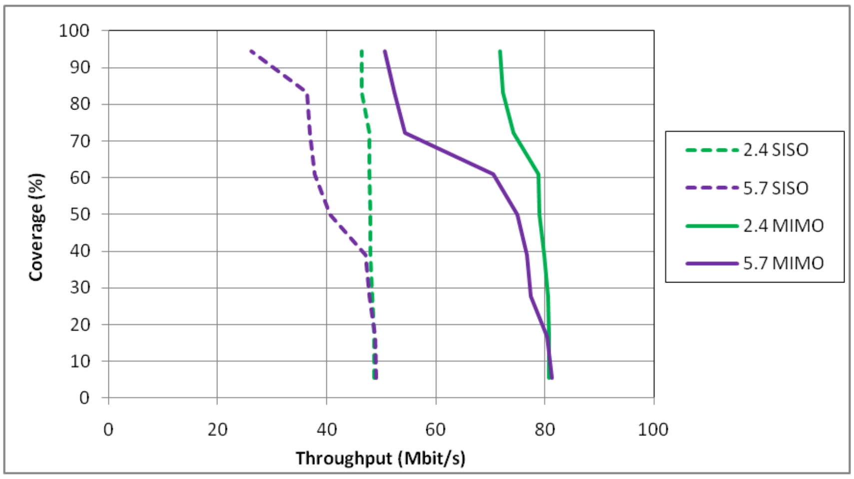

Figure E1.9: AP in study

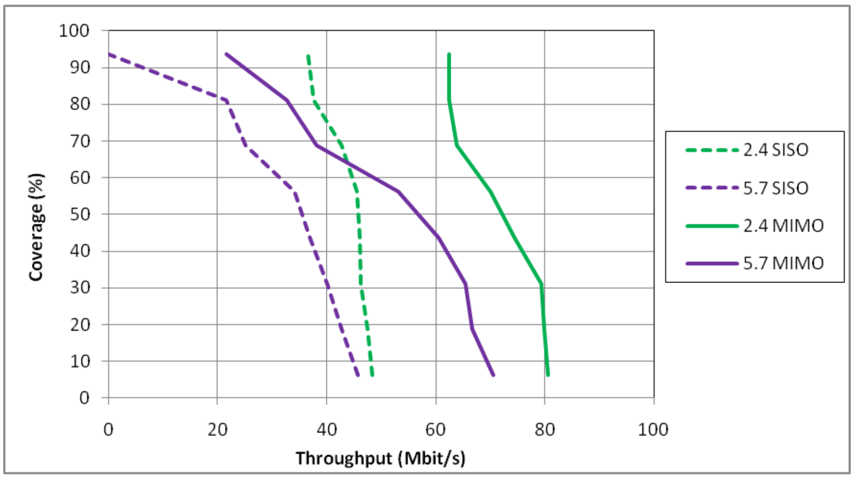

E.1.5 Modern terraced

Figure E1.10: AP in living room

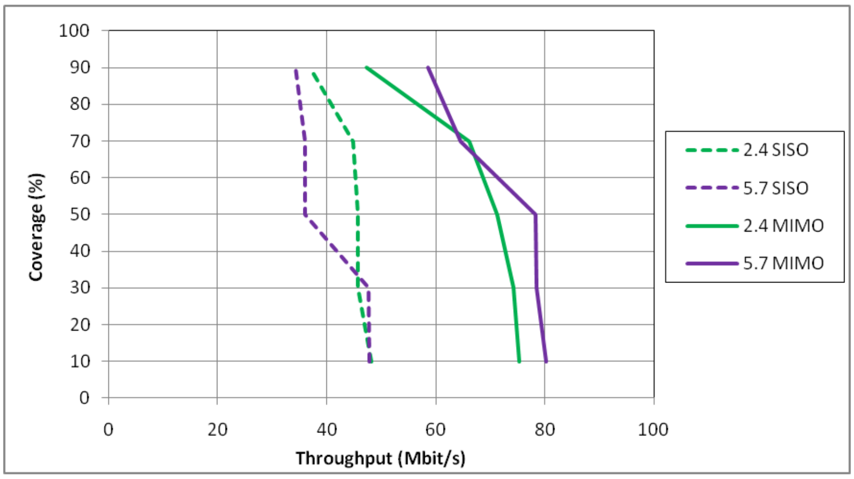

E.1.6 Large semi-detached

Figure E1.11: AP in hall

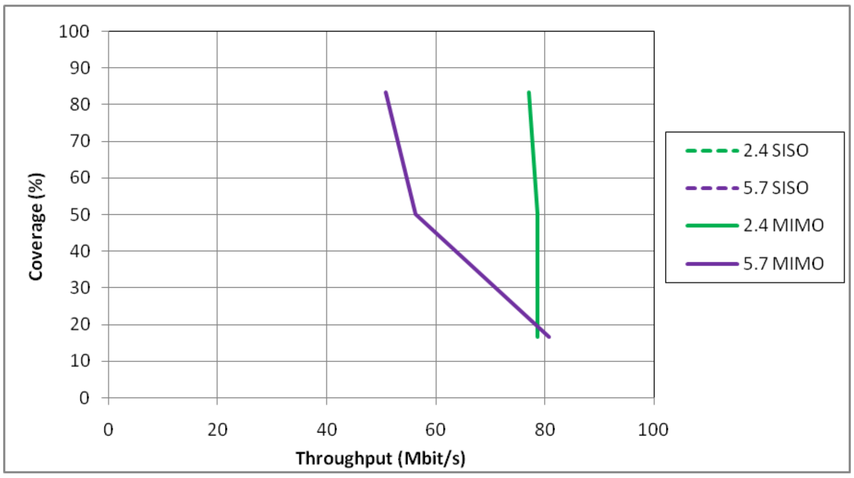

Figure E1.12: AP in living room

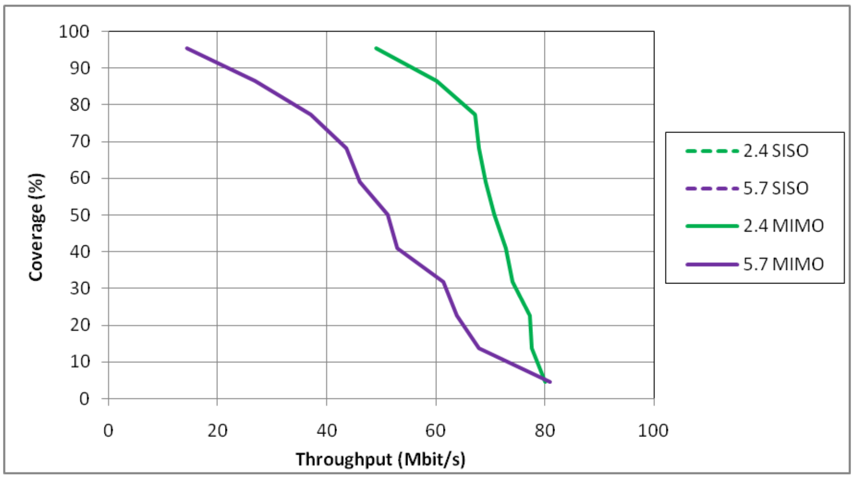

E.1.7 Basement flat

No throughput data taken.

E.1.8 UK average

Figure E1.13: All APs

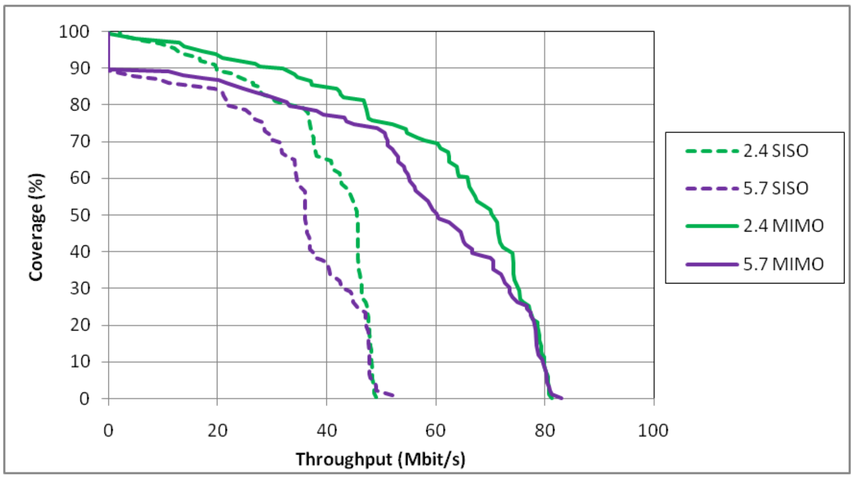

E.2 Predicted coverage at four frequencies

E.2.1 SISO

Figure E2.1: Victorian terraced, SISO

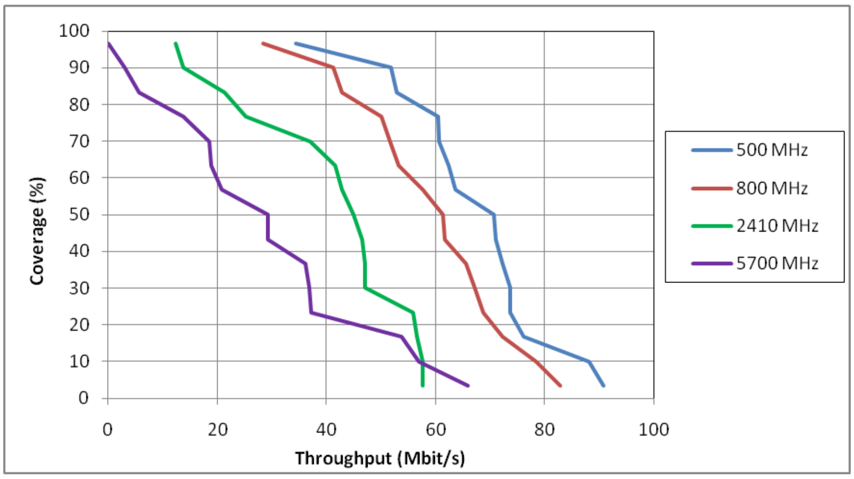

Figure E2.2: Inter-war semi-detached, SISO

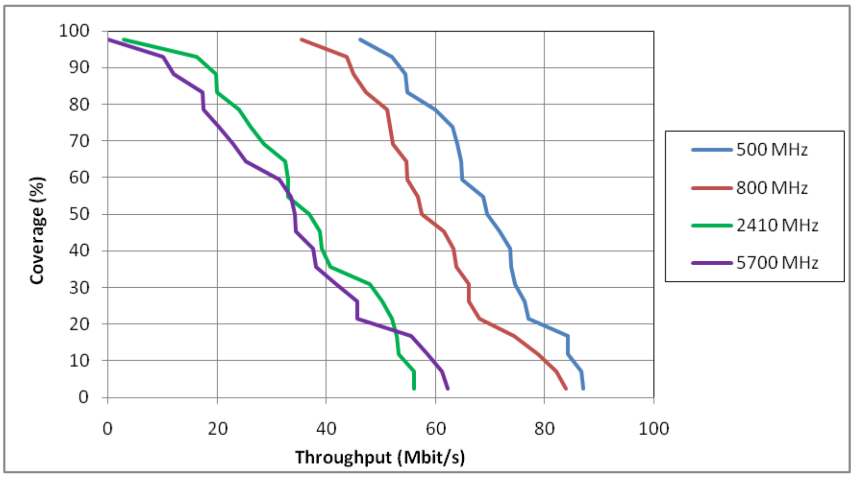

Figure E2.3: Modern bungalow, SISO

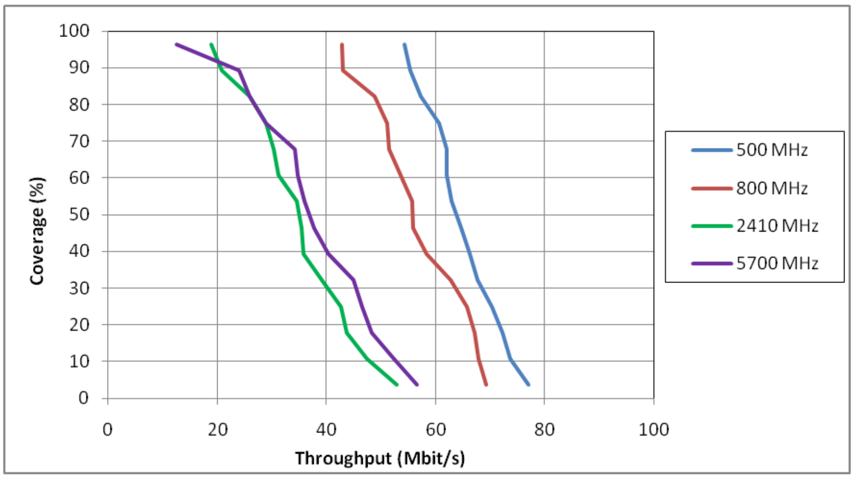

Figure E2.4: Modern detached, SISO

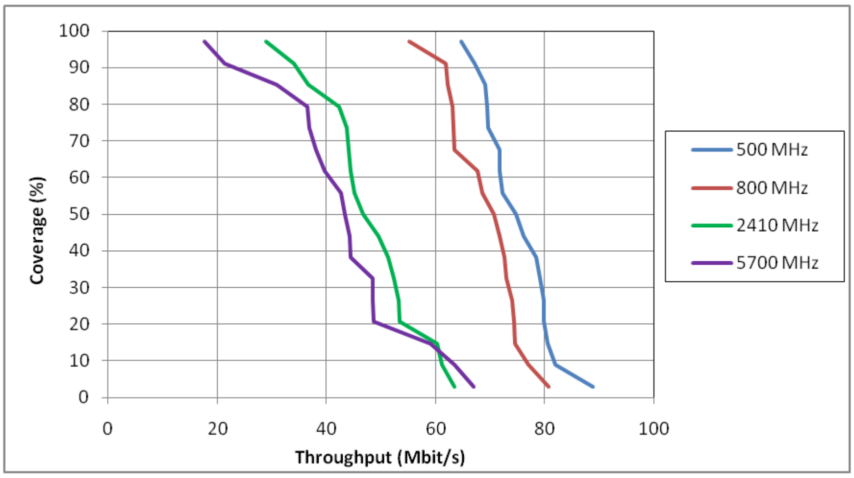

Figure E2.5: Modern terraced, SISO

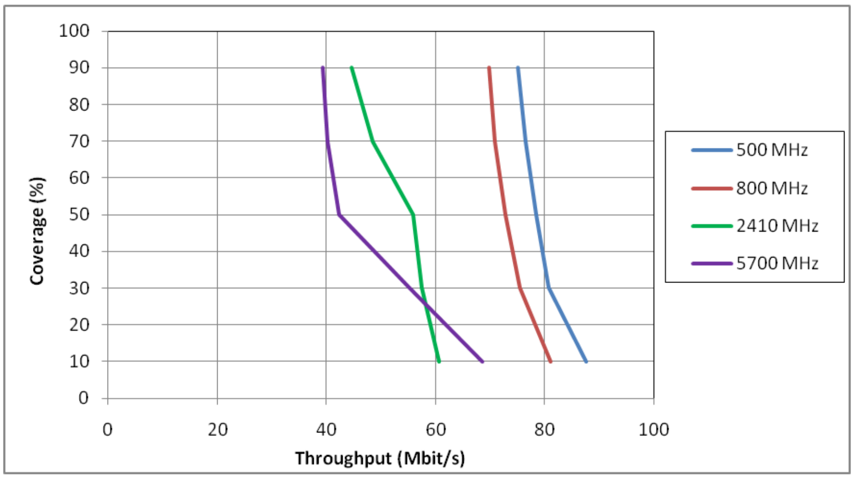

Figure E2.6: Large semi-detached, SISO

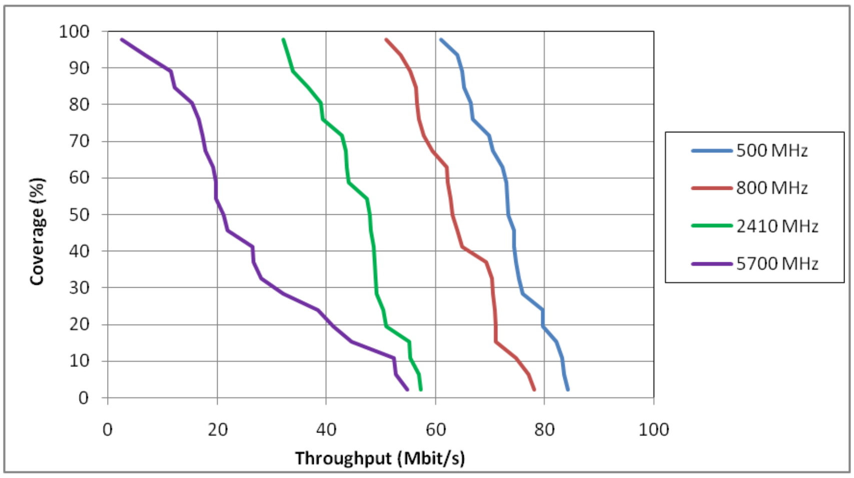

Figure E2.7: Basement flat, SISO

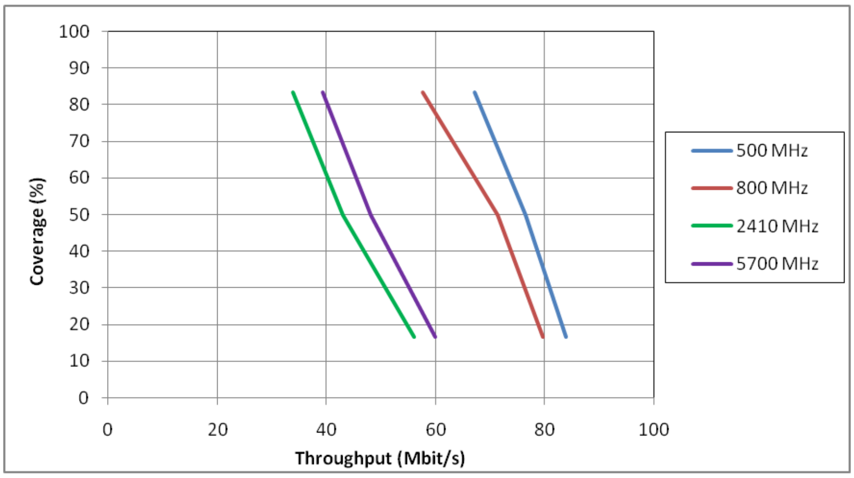

Figure E2.8: UK average, SISO

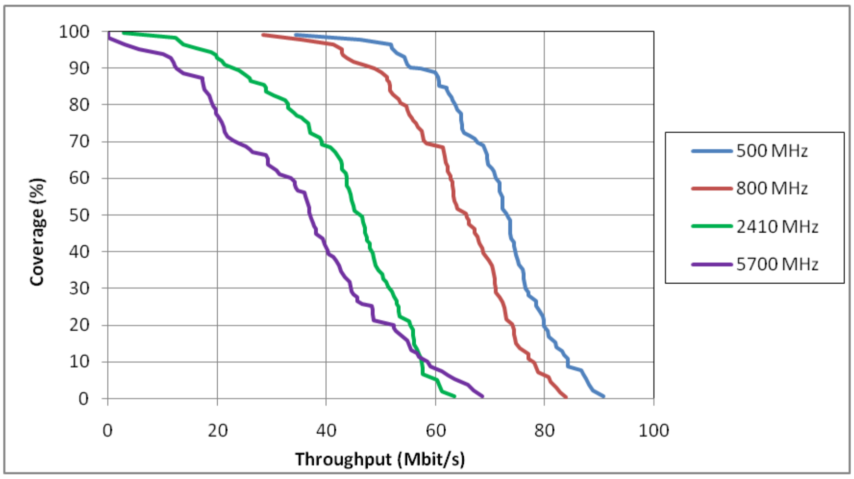

E.2.2 MIMO

Figure E2.9: Victorian terraced, MIMO

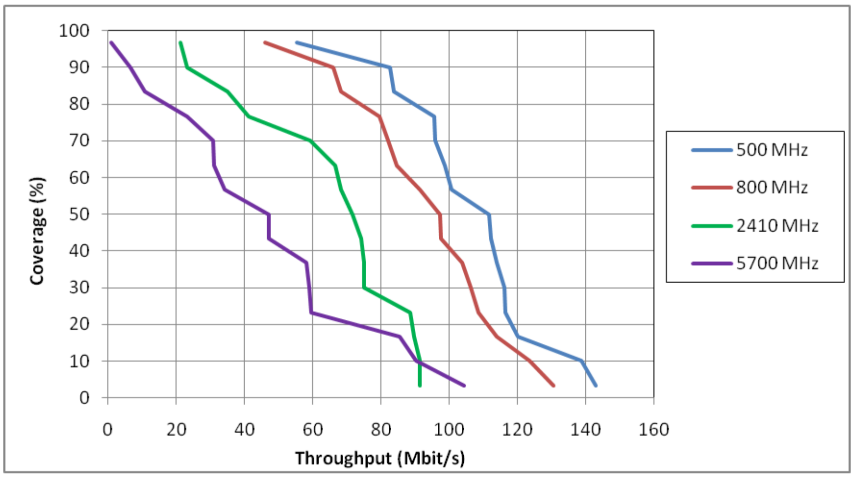

Figure E2.10: Inter-war semi-detached, MIMO

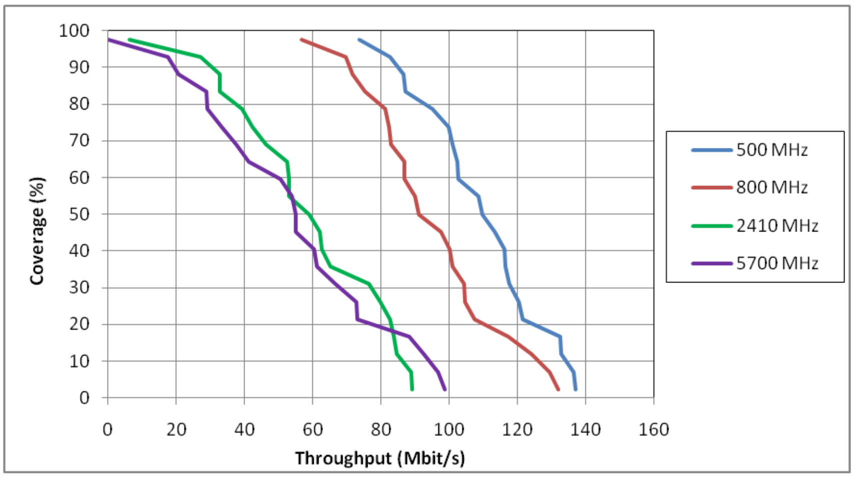

Figure E2.11: Modern bungalow, MIMO

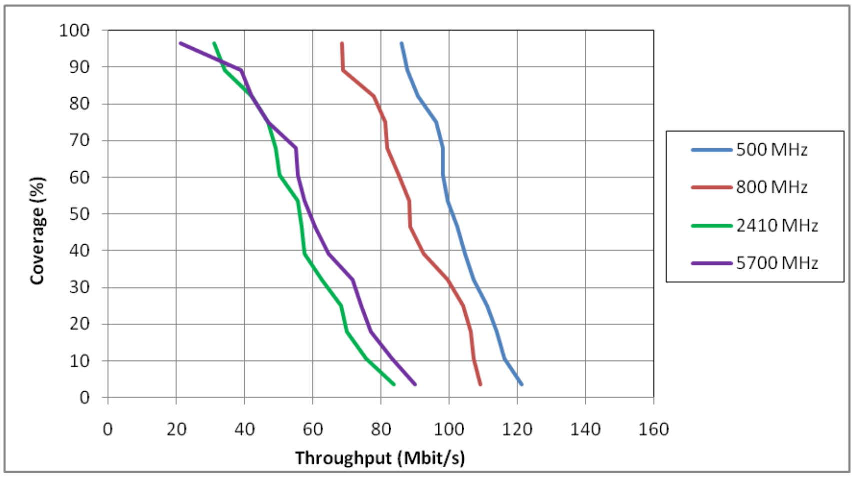

Figure E2.12: Modern detached, MIMO

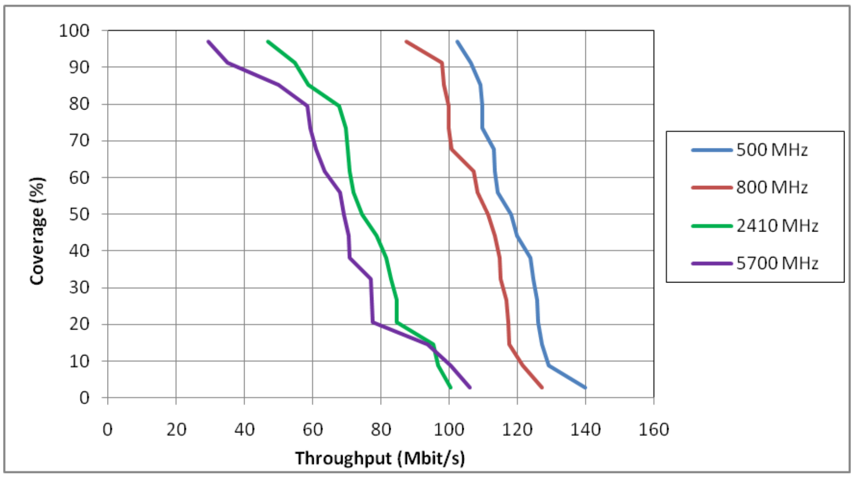

Figure E2.13: Modern terraced, MIMO

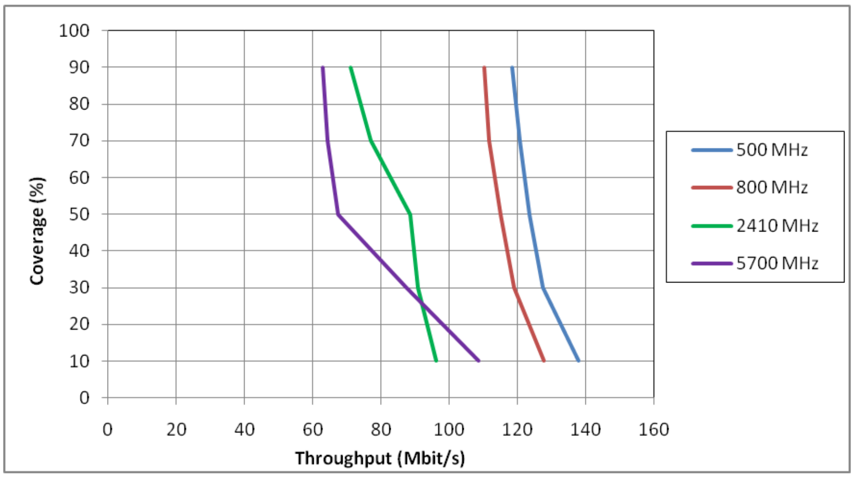

Figure E2.14: Large semi-detached, MIMO

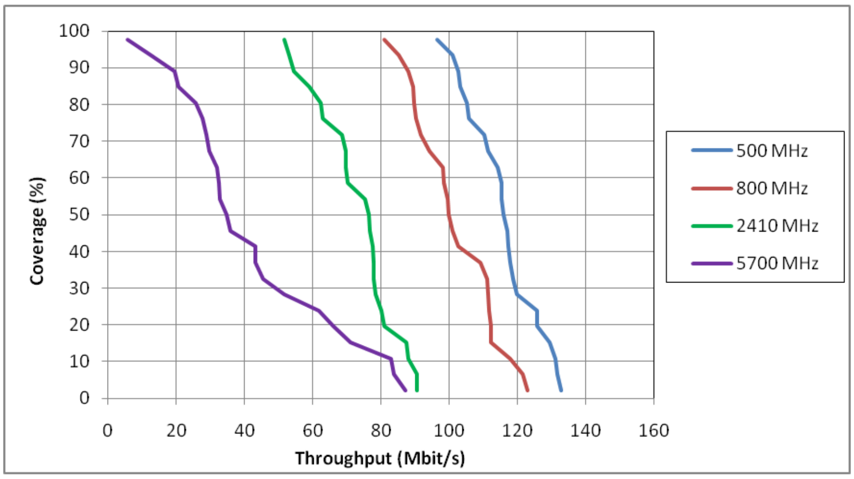

Figure E2.15: Basement flat, MIMO

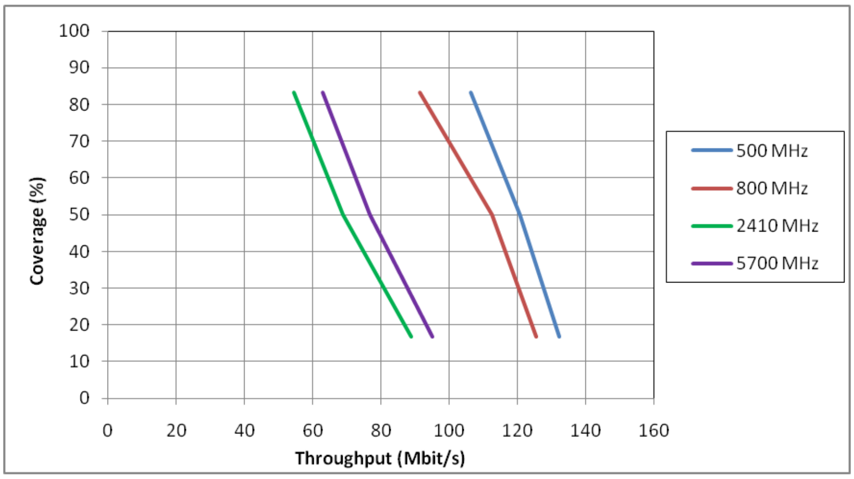

Figure E2.16: UK average, MIMO