




Wi-Fi Spectrum Needs Study

Final Report

Final Report to Wi-Fi Alliance, February 2017



www.QuotientAssociates.com



Contributors

*Steve Methley
William Webb*

Quotient Associates Limited

Compass House, Vision Park, Chivers Way,
Histon, Cambridge, CB24 9AD, UK
EMAIL Enquiries@QuotientAssociates.com
WEB www.QuotientAssociates.com

Version 01
Status Approved



Wi-Fi Spectrum Needs Study | Contributors

Final Report : qa1042 © Quotient Associates Ltd. 2017

Commercial in Confidence. No part of the contents of this document may be disclosed, used or reproduced in any form, or by any means, without the prior written consent of Quotient Associates Ltd.

Contents

0 Executive Summary 1

1 Traffic Predictions 3

- 1.1 External traffic predictions 3
 - 1.1.1 Cisco VNI predictions 3
 - 1.1.2 Surveyed broadband consumption 3
 - 1.1.3 Comparing sources of Internet data demand volumes 4
- 1.2 Adapting data sources for the model 5
- 1.3 Office versus consumer traffic 5
- 1.4 Internally generated traffic 5
- 1.5 Out-of-doors Traffic 6
- 1.6 Busy hour growth rates 6
- 1.7 Demand per location type 6
- 1.8 Demand scenarios 6

2 Use cases and technologies 9

- 2.1 Most important Internet devices 9
- 2.2 Device capabilities today 9
- 2.3 Projected device capabilities 10
- 2.4 Device distribution per location type 11
- 2.5 60 GHz capability 12

3 Sharing with LTE 13

4 Summary of key model inputs 14

- 4.1 Traffic demand levels 14
- 4.2 Device performance 14
- 4.3 Device distribution 14
- 4.4 Modelled environment 14

5 Model evaluation 15

- 5.1 Key Performance Metrics 15



- 5.1.1 Capacity and utilisation metrics 15
- 5.1.2 Approach to calculating utilisation 15
- 5.1.3 Reporting overall utilisation 16
- 5.1.4 Traffic and utilisation versus number of channels required 17
- 5.2 Channel selection mechanisms 17
- 5.3 Key results 17
 - 5.3.1 2020 18
 - 5.3.2 2025 20
- 5.4 Spectrum requirements summary 21
- 5.5 Power level sensitivity analysis 21
- 6 Spectrum availability and gap analysis 23**
 - 6.1 Spectrum at 5GHz 23
 - 6.1.1 Spectrum availability per region 24
 - 6.2 Gap analysis 25
- 7 Substitution potential of spectrum supply and new bands 27**
 - 7.1 Spectrum characteristics 27
- 8 Conclusions and recommendations 29**
 - 8.1 Future Wi-Fi spectrum requirements by region 29
 - 8.2 The importance of contiguous spectrum 29
- 9 Appendix A: Model parameters 31**
- 10 Appendix B: Model structure 32**
 - 10.1 Model structure (pseudo code) 32



0 EXECUTIVE SUMMARY

This report describes our predictions of the spectrum required in future, in order to satisfy the growth in the expected demand for Wi-Fi services.

Methodology

We sought evidence based traffic predictions from several perspectives. We included a Busy Hour growth scenario and an Upper Bound scenario, both with consideration of internal traffic such as in-band backhaul where appropriate. We considered location types of office, residential and mall.

In terms of use cases, we considered the shift in emphasis with respect to the most important devices used to connect to the internet. We looked at how device capabilities were expected to develop per device type over time and the distribution of device types per location type. The key performance metric we used in our airtime based modelling was Wi-Fi network utilisation. The target was 70% utilisation at the 95th percentile.

Key results

We have shown that, for the year 2025, the various regions are likely to need to find between 500 MHz and 1 GHz more spectrum than currently available to satisfy the Busy Hour scenario, which reflects the widely expected growth in traffic. If demand were to exceed the present Busy Hour predictions, our Upper Bound scenario suggests that an estimated maximum of between 1.3 and 1.8 GHz more spectrum than currently available may be needed. The Upper Bound scenario might occur due to unexpected adoption of novel applications or a further concentration of the busy hour traffic into fewer than the assumed four hours per day, for example. In other words, the Busy Hour scenario is the most likely to occur while the Upper Bound scenario is less likely, yet still plausible.

Our predictions for the new spectrum required per region are as shown in Figure 0-1.




Figure 0-1 Illustration of the spectrum shortfall per region, by year and demand level.



The amount of new spectrum required varies by geographical region, and our analysis illustrates potential effects due to spectrum which is subject to local DFS requirements¹. Our analysis assumes that new spectrum will be fully accessible by Wi-Fi.

The spectrum predictions cover the years 2020 and 2025; predicted Busy Hour and Upper Bound scenarios; two different usage levels of DFS spectrum; and three locations types consisting of office, residential and mall – of which residential was found to have the greatest spectrum requirements. Built into the predictions are technology advances, for example in terms of transmission rate, video coding improvements and device capability.

The importance of contiguous spectrum

In addition to simply needing more spectrum in total, we have shown that such spectrum needs to be assigned with sufficient contiguity such that wide channels of 160 MHz, or perhaps even wider in future, can be constructed with ease. To do otherwise would be to risk restricting the growth of Wi-Fi and the economic benefits with which it is widely associated and which was first enabled by forward-thinking spectrum regulation. Such a need for contiguity presents a significant further challenge to those with responsibility for spectrum allocation.

¹ Spectrum, access to which is determined by the constraints of Dynamic Frequency Selection procedures.



1 TRAFFIC PREDICTIONS

In this Chapter, we derive the traffic levels we will use for modelling, with reference to real-world surveys, plus predictions.

We are interested in three typical location types

- Office;
- Residential;
- Mall.

Traffic may be generated internally to the location or externally in each case. For example, traffic to and from the wider Internet is external, while traffic due, for example, to screen casting or in-band backhaul is internally generated. Both cases contribute to the loading on Wi-Fi.

1.1 External traffic predictions

External traffic relates to traffic to and from the Internet. Various surveys exist for predicted data volumes, although few, if any, go beyond five years. In order to be able to draw upon multiple sources of information, we chose to look at data from the well-known Cisco Virtual Networking Index; a regulator survey; and a commercial survey.

1.1.1 Cisco VNI predictions

The Cisco VNI forecast tool² was used, from which it was determined that North American consumer traffic at the household level would grow to 450 GB per month in 2020, with a growth factor of two over the five year period 2015-2020. Cisco do not predict beyond five years ahead. North America was chosen as representative of a developed area with significant broadband and Wi-Fi usage.

1.1.2 Surveyed broadband consumption

For other perspectives on data volume growth, we turned to surveyed data from a regulator, Ofcom UK, plus a commercial source. Ofcom reported a data usage of 82 GB/month for UK domestic households in 2015³, growing at 40% per annum. In an alternative survey, Statista looked at worldwide fixed broadband usage per capita⁴. In 2014 this varied from 9.9 GB/month in Germany to 48.6 GB/month in Korea. The UK and USA were 22.3 and 18.5 GB/month, respectively. We note that the Ofcom data is per household, as is the Cisco VNI data, but the Statista data is per capita.

If 40% annual growth is maintained then this predicts 1.125 GB per person each busy hour in 2020, using the simplification that all traffic occurs during four busy hours per day, for example 7-11pm.

Beyond 2020 there are few predictions but many believe we will be some way along the likely "S" curve of demand growth, with demand growth slowing somewhat in the period 2020 – 2025. A slowing rate of growth of demand has been reported by mobile operators

² http://www.cisco.com/c/m/en_us/solutions/service-provider/vni-forecast-highlights.html

³ "Connected Nations Report 2015", from www.ofcom.org.uk

⁴ See <http://www.statista.com/statistics/374998/fixed-broadband-data-volume-per-capita/>



over the last few years and indeed, in some extreme cases, operators such as M1 in Singapore have seen negligible growth (only 3%) in data requirements in the last year. Equally, the arrival of a new application⁵ could result in a sudden increase in demand.

If, as expected, the growth begins to slow somewhat, then this would predict around 4.5 GB per person each busy hour in 2025.

1.1.3 Comparing sources of Internet data demand volumes




Figure 1-1 Comparison of data volumes predicted by Cisco VNI, a regulator and a commercial survey.

Figure 1-1 shows that the predictions for 2025 data volumes are broadly similar for surveyed broadband consumption from Statista / Ofcom and the Cisco VNI predictions for the USA. All figures are per user per month, having been converted where necessary. None of the data sources predict beyond 2020, so the growth factors have been extrapolated to 2025.

Cisco begins with a higher estimate of today's data volumes, but then uses a smaller growth factor, as clarified in Figure 1-2. This behaviour is the key reason that the all predictions yield similar results by 2025.




Figure 1-2 Extrapolated period-on- period growth factors for Cisco Ofcom data.

⁵ For example the app-based Pokémon Go game which became extremely popular at the time of writing.



In summary, the predicted demand from the three independent sources is reasonably aligned and differences will likely not significantly affect the amount of spectrum required in 2025 since we look for multiples of 160 MHz⁶.

1.2 Adapting data sources for the model

All three data sources have required extrapolation to 2025 by assuming the respective trends will continue and this represents a very significant uncertainty in our analysis. In order to address this most simply, we have normalised to an average traffic volume rounded up to 4.5 GB per person per busy hour for 2025 (since this is when estimates converge most closely) plus we have assumed that the Cisco VNI growth rates are appropriate.

Before making use of data volumes in the modelling, it is important to note that all the surveys quoted have used average growth rates. We discuss the phenomenon of growth rates more appropriate to the busy hours in Section 1.6. The surveys also relate to residential demand and we outline different assumed demands in office and mall in Sections 1.3 and 1.5 respectively. Finally we also discuss the effect of internal traffic due to in-band backhaul and soft AP⁷ use such as screen casting in Section 1.4. The overall effect of these discussions is summarised per location type (office, residential, mall) in Section 1.7.

Finally, in Section 1.8, we introduce our use of two overall scenarios. One is based on the traffic volume growth predicted for the Busy Hour; the other considers a scenario where growth is greater than presently expected, i.e. even greater than the predictions using busy hour trends. The function of the latter scenario is to seek to estimate a plausible Upper Bound for the amount of spectrum which Wi-Fi may require in future.

1.3 Office versus consumer traffic

The Cisco VNI suggests that, on average, a reasonable assumption is that office traffic volume is a quarter of consumer traffic volume.

1.4 Internally generated traffic

This relates to traffic which is peer to peer within the location type considered. The prime example is screen casting in the home or in-band backhaul, where the links between APs are carried within the same frequency band as the AP-to-user traffic.

Based on our extensive experience of mesh networking⁸, which is typically self-backhauled⁹, we have taken a straightforward approach of doubling the residential data requirement when self-backhauling is used. Of course not all residential locations will be self-backhauled, but we have applied a pessimistic factor of two in order to also take account of those locations which may be screen casting or otherwise using soft APs.

⁶ See Section 5.3 for why multiples of 160 MHz are desirable.

⁷ Also known as a virtual router, where software is used to turn an end station into an access point, usually for the benefit of a particular application, such as screen casting.

⁸ See, for example, "Essentials of Wireless Mesh Networking", Steve Methley, Cambridge University Press, 2009.

⁹ i.e. carried in-band.



1.5 Out-of-doors Traffic

Data usage out of doors is typically much smaller than indoors. One reason is that the user is mobile and hence less likely to be using the device intensively. Measurements are difficult to find, but a recent assessment of the use of the internet out of doors has concluded that the volume of use is one tenth of the use indoors, on average¹⁰. We note that out-of-doors does not mean in remote locations, rather it means during normal daily life, including commuting etc. This is relevant to our mall location type.

1.6 Busy hour growth rates

The Cisco VNI and other survey data all relate to average traffic levels. But it is well-known that busy hour traffic is growing much faster, for the reason that more video traffic is included in the busy hours. Data on the difference between average and busy hour growth rates has been produced by Cisco and is shown in Figure 1-3. Busy hour growth is predicted to exceed average growth by a factor of 1.5 in a five year period.




Figure 1-3 Busy hour versus average growth rates from Cisco VNI.

In our model, we take the increased growth rates of the busy hour into account, since networks must be dimensioned for busy hour rather than average volumes, and this directly affects spectrum requirements.

1.7 Demand per location type

Using the data volumes gathered from the Cisco VNI and other surveys, as described earlier, and taking into account the higher busy hour growth rate, we further double residential demand to take account of self-backhaul and screen casting as already discussed; we set office use to half the consumer value in order to accommodate the lower demand yet allow for soft APs; finally we set mall usage at one tenth of consumer demand as this is out-of-doors traffic.

1.8 Demand scenarios

We employ two scenarios, a Busy Hour Scenario based on the predictions described earlier, plus we add a scenario with higher demand. The latter demand level is greater than the predicted busy hour and was chosen to be equal to 2x the average traffic level in 2020 and

¹⁰ "Out-of-home use of the internet", Broadband Stakeholder group, September 2014.



4x the average traffic level in 2025¹¹. The higher scenario is intended to reflect what could happen if usage increased or was concentrated into fewer busy hours per day, for example. As it is intended to reflect an upper bound to spectrum requirements, we refer to it as the Upper Bound scenario¹². Having some notion of an upper bound is likely to be helpful to those in charge of the allocation of spectrum resources.

We show the relationship of the Busy Hour scenario and the Upper Bound scenario to predicted average data volumes in Table 1-1.

<i>Scenario</i>	<i>2020</i>	<i>2025</i>
Busy Hour	150%	225%
Upper Bound	200%	400%

Table 1-1 The two scenarios used in the study, with their relationship to predicted average traffic volumes, for the relevant year.

In selecting an upper bound we have chosen a data volume which is less than twice the predicted busy hour volume, so it seems far from being out of the question that such volumes might occur. On the other hand it is significantly higher than any prediction we have seen reported elsewhere. Alternative estimates of an upper bound are possible, but we feel the one we have selected is reasonable.

It is worth noting that possible reasons for higher than expected data volumes include not only the obvious case of increased data demand by users, but also a case where data demand is flat on average yet becomes more concentrated into fewer busy hours per day. For example the busy hour period could decrease from the present 7-11pm to, say, 8-10pm. Networks would then have to be dimensioned to deal with peak traffic volumes which are higher by a factor of 2. Moreover, concentration effects could happen at the same time as an overall increased demand, leading to compound growth behaviour.

The data volumes used in the modelling are as follows. The detailed traffic volumes levels used for the Busy Hour scenario are shown in Table 1-2.

<i>Busy Hour Scenario</i>	<i>2020</i>	<i>2025</i>
Office demand	1688	5063
Residential demand	6750	20250
Mall demand	338	1013

Table 1-2 Busy Hour demand parameters (Mbytes per person during busy hour)

The detailed traffic volumes levels used for the Upper Bound scenario are shown in Table 1-3.

¹¹ For the purposes of comparison, this is also equivalent to a demand 30% higher than the busy hour prediction in 2020 and 78% higher than the busy hour prediction in 2025.

¹² Of course, for the avoidance of doubt, this is simply our estimate of where an upper bound might lie. In other words, it does not represent a true upper bound in the mathematical sense.



<i>Upper Bound Scenario</i>	<i>2020</i>	<i>2025</i>
Office demand	2250	9000
Residential demand	9000	36000
Mall demand	450	1800

Table 1-3 Upper Bound demand parameters (Mbytes per person during busy hour)



2 USE CASES AND TECHNOLOGIES

In the preceding Chapter we looked at the data volume demands expected in the future. In this Chapter we investigate the capabilities of the devices used. Again we are interested in performance expected in the future.

2.1 Most important Internet devices

The mix of devices which people use is not static, as shown in Figure 2-1. There is clearly an increase in the use of smartphones and tablets and a decrease in the use of desktops and laptops. For example, desktop use has halved over the three years of the survey and smartphone use has more than doubled. A key conclusion for our study is that smartphones have overtaken laptops as the most important way to connect to the Internet. This is true worldwide, and is even more pronounced in Asia, for example.




Figure 2-1 Changes in the type of device used to connect to the internet¹³.

These trends are important since each type of device possesses a different level of Wi-Fi capability, which may be dictated by device size, power supply and price. We illustrate this next.

2.2 Device capabilities today

Makers of Wi-Fi access points have an interest in knowing what capabilities the clients are likely to have. Results from such a study have been published¹⁴, which we reproduce in Figure 2-2.

¹³ Ofcom Communication Market Report, 2015, available from www.ofcom.org.uk

¹⁴ "Competitive Test Report Executive Summary - Spring 2016", available from www.mojonetworks.com



Client Make	Model	802.11 Version	Streams	Chipset Vendor	Chipset	Driver Version	OS Version
Google	Nexus 5x	11ac Wave 2	2x2	Qualcomm	QCA6174	-	Android 6.0.1
Acer	Aspire F15	11ac Wave 2	1x1	Qualcomm	QCA9377	12.0.0.203	Windows 10
OnePlus	OnePlus Two	11ac Wave 2	1x1	Qualcomm	QCA9377	-	Android 6.0.1
Apple	Macbook Pro	11ac Wave 1	3x3	Broadcom	BCM4360	7.21.94.136.1a1	MacOS 10.11.2
Dell	Inspiron 7000	11ac Wave 1	2x2	Intel	7265	18.40.0.9	Windows 10
HP	Elitebook	11ac Wave 1	2x2	Intel	7265	18.33.0.1	Windows 7
Apple	Macbook Air	11ac Wave 1	2x2	Broadcom	BCM4360	7.21.94.136.1a1	MacOS 10.11.2
Apple	iMac	11ac Wave 1	3x3	Broadcom	BCM4360	7.21.94.136.1a1	MacOS 10.11.3
Apple	iPad Pro	11ac Wave 1	2x2	Broadcom	-	-	iOS 9.2.1
Samsung	Tab S2	11ac Wave 1	2x2	-	-	-	Android 6.0.1
HP	ProBook	11n	2x2	Broadcom	BCM943228Z	6.30.223.255	Windows 10
Apple	iPad Mini	11n	1x1	Broadcom	-	-	iOS 9.3.1

Figure 2-2 Client device capabilities today¹⁴.

Reading the data in Figure 2-2, it can be seen that low end phones and small tablets are currently 1x1 MIMO, whereas higher end phones and full size tablets are 2x2. Many laptops are 2x2, with one high end example having 3x3 MIMO today.

2.3 Projected device capabilities

We can project the device capabilities of Figure 2-2 into the future by incorporating the expected development of Wi-Fi over time.

Table 2-1 shows rates (Mb/s) versus bandwidth for different versions of 801.11 from legacy to future examples, against bandwidth used and number of spatial streams used.

<i>Spatial streams</i>	<i>Variant</i>	<i>20</i>	<i>40</i>	<i>80</i>	<i>160</i>
1	11b	11			
1	11g	54			
1x1	11n	72	150		
	11ac	87	200	433	867
	11ax	143	287	600	1.2 Gb/s
2x2	11n	144	300		
	11ac	173	400	867	1.7 Gb/s
	11ax	300	600	1.2 Gb/s	2.4 Gb/s
3x3	11n	216	450		
	11ac	289	600	1.3 Gb/s	2.3 Gb/s
	11ax	450	900	1.8 Gb/s	3.6 Gb/s
4x4	11n	288	600		
	11ac	346	800	1.7 Gb/s	3.5 Gb/s
	11ax	600	1200	2.4 Gb/s	4.8 Gb/s

Table 2-1 802.11 rates (Mb/s) versus bandwidth (MHz) and spatial streams.



Table 2-2 constructs a timeline of capabilities for each device type based on our understanding of when the different variants of 802.11 will be in the majority in the marketplace. For example 11n is here today, 11ac will come to dominate in the next few years and we have assumed that 11ax will become mainstream by 2025.

In terms of devices, smartphones are typically single antenna devices in the short and medium terms, apart from the larger phones, where two antennas become realistic. Laptops and tablets form a natural second group, mostly based on size where 2x2 will be most common until 2020. Premium devices are 3x3 today and are likely to become 4x4 in the next few years. We do not anticipate that client devices will have the maximum 8 antennas allowed in 11ax, even if access points do.

As far as channel width is concerned, 11n is limited to 40 MHz but we expect later 802.11 versions will concentrate on the mandated 80 MHz and the optional 160 MHz.

<i>Device / Year</i>	<i>2015</i>	<i>2020</i>	<i>2025</i>
Smartphone	11n 1x1 40 MHz (150 Mb/s)	11ac 1x1 80 MHz (433 Mb/s)	11ax 2x2 80 MHz (1.2 Gb/s)
Tablet/Laptop and high end phone	11n 2x2 40 MHz (300 Mb/s)	11ac 2x2 80 MHz (867 Mb/s)	11ax 3x3 80 MHz (1.8 Gb/s)
High end laptop	11ac 3x3 80 MHz (1.3 Gb/s)	11ac 4x4 80 MHz (1.7 Gb/s)	11ax 4x4 80 MHz (2.4 Gb/s)

Table 2-2 Projected 802.11 device capabilities over time.

We have made several simplifications, which include that 11ax is similar to 11ac, but with the addition of 1024 QAM (MCS 10, 11). We take into account that all rates are maxima; many are possible only at very short distances/low interference situations. The model will allow channel widths up to 160 MHz to be used where propagation conditions allow.

To complete the picture, we need to understand the distribution of device types over location types.

2.4 Device distribution per location type

We do not have precise data on which to base Table 2-3. Instead we have a number of points such as the residential data of Figure 2-1, from which we extrapolate to the office and mall location types. We expect more high-end laptops to be used in office locations and more smartphones to be used in the mall location.

<i>Device</i>	<i>Office %</i>	<i>Residential %</i>	<i>Mall %</i>
Smartphone	50	70	90
Tablet/Laptop and high end phone	25	30	10
High end laptop	25	0	0

Table 2-3 Projected distributions of device types as percentage per location type.



2.5 60 GHz capability

Over both scenarios in all years, we assume that the distribution of 60GHz capable users is 10%, 20% and zero in office, residential and mall locations respectively. Although all APs are assumed to be 60 GHz capable, user devices will connect at 60 GHz only if propagation conditions are suitable, leading to a lesser proportion of 60 GHz devices actually connected in office and residential location types.

The detail of this Chapter forms some of our model inputs, as we summarise in Chapter 4.



3 SHARING WITH LTE

For purposes of this study we have assumed that LTE in 5 GHz will not bring a significant amount of new traffic, but will rather take some traffic share away from Wi-Fi. Another way to say this is that we have no evidence that users will generate any more traffic because they could choose LTE over Wi-Fi in 5 GHz. Given that the proportion of traffic presently carried over LTE is dwarfed by Wi-Fi, which carries 80% and is increasing, this appears a reasonable assumption.

In terms of modelling, we do not specifically model LTE. If required it could be dealt with as a potential reduction in the number of channels available to Wi-Fi when evaluating the spectrum required.



4 SUMMARY OF KEY MODEL INPUTS

Our model requires inputs of traffic demand; device capabilities over time; distributions of device type over location types; and the dimensions of the modelled environment. This Chapter provides an easy reference to all the necessary model parameters.

We believe that where we have needed to make assumptions, they have been reasonable for the purposes of this spectrum estimation study, yet we are aware that other assumptions are possible. With this in mind our approach has been to make it clear where we have made assumptions and what these assumptions comprise.

4.1 Traffic demand levels

We have defined two scenarios of traffic demand, which are firstly a projected Busy Hour scenario and secondly an Upper Bound scenario. The latter is intended to reflect what could happen if usage increased or concentrated into fewer busy hours per day, for example. We illustrated these scenarios in Table 1-1 on page 7.

We gave modelled data volumes in Table 1-2 and Table 1-3 on page 8.

4.2 Device performance

We gave our assumed timeline of device capabilities in Table 2-2 on page 11.

4.3 Device distribution

We gave the distribution of device types in Table 2-3 on page 11.

4.4 Modelled environment

We have taken the dimensions of buildings from the IEEE 802.11 'TGax simulation scenarios'¹⁵ for office and residential locations, plus we constructed a mall with dimensions 30m x 300m x 2 floors. Further details, including AP densities, are given in Appendix A.

We assume the office and mall location types are managed and hence use a centralised channel selection procedure, whereas this is not appropriate for residential, see Chapter 5.

¹⁵ 11-14-0980-16-00ax-simulation-scenarios, downloaded from http://www.ieee802.org/11/Reports/tgax_update.htm



5 MODEL EVALUATION

Note that we provide a summary of the model results in terms of a table of total predicted spectrum demand in Section 5.4. A walkthrough of model operation is provided in Appendix B.

5.1 Key Performance Metrics

In evaluating the performance of a Wi-Fi system, some thought is needed as to the best metrics to use. For some users the maximum speed will be important, with too low a speed preventing some applications such as video streaming. For others, latency will be critical, causing delays in applications such as web browsing or voice. In some cases, congestion levels might impact on applications such as video conferencing by causing erratic delays.

Equally, in some cases high speeds cannot be achieved because the user is too far from the access point (AP) and so does not have a sufficiently high SNR to use higher-order modulation. This is not an issue of capacity or adequate spectrum but one of AP layout, with a denser grid of APs providing higher signal levels. The actual density used will vary from deployment to deployment and indeed more APs might be installed if sufficient users experience problems (subject to AP-AP interference considerations).

5.1.1 Capacity and utilisation metrics

For these reasons we have chosen not to focus on high speed but instead on metrics related to capacity. Specifically, we look at two key metrics; the percentage of offered traffic that is carried plus a measure of AP utilisation, based on airtime usage¹⁶. Both relate to data rates and latency. A network that is able to carry all the traffic offered and where the AP utilisation is within normal bounds will be able to deliver the maximum data rates that a device can access according to its signal level.

Even with these metrics there is no hard failure point at which performance changes from acceptable to unacceptable. As we have all experienced, the effect of congestion is a soft failure: Congested networks may be adequate if a little frustrating, then as loading builds the performance may degrade to become very frustrating for users and then finally degrading to the point where using the network becomes impractical.

Taking this into account, we suggest that networks should be

1. Able to carry at least 95% of the offered traffic;
2. Have an AP loading ideally below 70%.

APs with a utilisation around 70% tend to be suitable for data traffic. A lower utilisation may sometimes be suitable for more sensitive specific traffic such as voice. On the other hand, newer forms of Wi-Fi are expected to achieve lower latency at higher utilisations. Our target is 70%, which hence represents a modestly conservative approach.

5.1.2 Approach to calculating utilisation

We note that our results are specific to the location types we have simulated, which are based on IEEE 802.11 simulation scenarios. In general, adding more APs can typically

¹⁶ We define utilisation in our airtime based model as that percentage of airtime that an AP observes as being utilised, both by itself and other neighbouring co-channel networks.



improve the situation because (1) devices are closer to an AP and so can use more efficient modulation and coding schemes and (2) more APs generally add more capacity. However, adding more APs also increases interference and so incremental improvements steadily decrease. More pragmatically, the number of APs may be limited by factors such as room size and access to backhaul. We have opted for numbers of APs that we consider practical but relatively dense. This results in a lower bound for spectrum requirements. The model is described in the Appendices.

The reporting of utilisation is complicated by our consideration of multiple frequency bands. An AP may work across the 2.4GHz, 5GHz and 60GHz bands, while devices may work across some or all of these bands. A situation could be envisioned, for example, where an AP had 5GHz and 60GHz capability but the devices within its coverage operated only on 5GHz. The AP might then appear to be 50% utilised, yet from the device viewpoint it might appear completely congested.

We have opted for an approach where we have weighted the utilisation in each band by the percentage of devices using that band. So if the 5GHz band was 80% utilised and the 60GHz band 10% utilised and if 50% of devices access each band, we would measure utilisation as $(80\% * 50\%) + (10\% * 50\%) = 45\%$. On the one hand, in this situation, the overall AP performance might be judged acceptable but for the 50% of devices in the 5GHz band they might find usage unacceptable. On the other hand, a more complex set of multiple metrics could aim to address this, but a likely problem is that too complex a set of metrics may quickly make the results difficult to interpret.

In summary we use the approaches of

- moving terminals towards their less-preferred band when their preferred band is congested; and
- using summary measures of utilisation composed of multiple underlying measures.

5.1.3 Reporting overall utilisation

Because each deployment includes multiple APs, and because the loading can vary across them¹⁷, then some measure of overall AP utilisation is needed. The simplest approach would be to take an average, but this might mask a situation where some APs are working well and others completely congested. To avoid this issue, we favour a cumulative distribution approach where we take the 95th percentile point. If, for example, the 95th percentile is 40% this means that 95% of all APs have a loading of 40% or less. We believe this strikes a good balance between requiring every AP to be lightly loaded and ensuring that the vast majority of users have a good experience.

¹⁷ For example APs on the edge of a building can be subject to more external interference.



5.1.4 Traffic and utilisation versus number of channels required

In summary, we seek to identify outcomes where both

- at least 95% of the offered traffic is carried, and
- the 95th percentile of AP utilisation is below 70%.

We vary the spectrum available at 5GHz to try to satisfy these targets.

We choose 5GHz rather than other bands because

- there is limited space, if any, to expand at 2.4GHz; and
- many devices either cannot support 60GHz or will not be within range of a 60GHz AP.

Finally we note that although our choice of variable is the spectrum available at 5GHz, we could equally well have chosen bands from around 2 GHz to 10 GHz, since propagation will be similar¹⁸. Hence, we are not restricting the applicability of our results solely to 5GHz.

5.2 Channel selection mechanisms

We assume the office and mall location types are managed and hence use a centralised channel selection procedure¹⁹. We assume the residential situation is unmanaged and hence we use a random channel assignment procedure. The random assignment generally leads to a less efficient selection and an associated increase in spectrum required.

5.3 Key results

We present results relating to two time horizons and three location types

- 2020 and 2025, reflecting our expectation of two demand scenarios at each of these times as described in Chapter 1;
- office, residential and mall locations.

Offices are assumed to be multi-storey and open plan. Residential is assumed to consist of multi-story apartments which are relatively small with two occupants. The mall is assumed to be two-storey with a large central atrium. The model is described in the Appendix.

For each of the three years selected we model each environment. For each outcome we examine 95th percentile AP utilisation against the amount of 5 GHz spectrum required. The point at which our target metrics are met then determines the spectrum requirements. We assume 60 GHz penetration in devices of 10% in office, 20% in home and 0% in mall in all cases. Additionally we adopt the engineering assumption that in the future it will be more desirable to use wider channels for both speed and lower battery use in mobile devices²⁰.

All results presented here carried at least 95% of presented traffic. We do not model situations with fewer than 4 channels in use as this leads to excessive interference. Hence, in the following graphs, the 80 MHz channels results begin at a minimum of 320 MHz; and 160 MHz results begin at a minimum of 640 MHz of total spectrum.

¹⁸ See Section 7.1.

¹⁹ For a discussion of management and co-ordination possibilities see our earlier report "Study on the use of Wi-Fi for Metropolitan Area applications", available from www.ofcom.org.uk

²⁰ We note however that 160 MHz channels are not presently allowed in China.



5.3.1 2020

Busy Hour scenario 2020

Our busy hour scenario uses predictions of traffic demand, device capabilities over time, and distributions of device type over location for the year 2020, as summarised in Chapter 4.




Figure 5-1 Busy Hour scenario, year 2020, using 80 and 160 MHz channels (target utilisation 70%)

From Figure 5-1 we can see that as the amount of spectrum is increased on the x-axis, then the utilisation of the AP reduces on the y-axis. In other words, greater spectrum provision leads to a less heavily loaded Wi-Fi network. Our chosen utilisation point is 70%, which was chosen as this level will facilitate normal data communication²¹. However, we note that some networks may need to run at lower utilisation in order to ensure latency requirements. In this sense our study provides a lower bound of spectrum requirements and is one reason why we also have a scenario with higher demand.

In the busy hour scenario for 2020, 800 MHz spectrum is required²². This is set by the residential location type which has the highest spectrum requirement.

²¹ A utilisation of 70% at the 95th percentile means that 95% of APs are operating at or below 70% utilisation.

²² All results are rounded to the next higher multiple of 160 MHz.



Upper Bound Scenario 2020




Figure 5-2 Higher Demand scenario, year 2020, using 80 and 160 MHz channels (target utilisation 70%)

Our upper bound scenario increases demand level to twice that predicted for average traffic²³ growth in 2020 and four times in 2025 (30% and 78% more traffic than the busy hour prediction), but is otherwise the same, i.e. device performance improvements and distribution over time are the same. As this higher demand level is greater than the predicted busy hour levels, it provides the potential to account for situations where more spectrum is needed because

- future demand simply outstrips the industry-based estimates we have used;
- future demand concentrates traffic further into fewer busy hours per day;
- future demand includes applications where lower network utilisation is necessary, e.g. to reduce latency.

Any or all of these situations could occur in the future.

From Figure 5-2 we can see that 960 MHz spectrum is required whether 80 or 160 MHz channels are used. This could be 12 channels of 80 MHz or 6 channels of 160 MHz. Wider channels have a potential for higher data rates. An implication is that applications which need to burst data faster will benefit from wider channels as will battery operated devices, whose lifetime will be extended by the shorter on-times facilitated by faster, wider channels.

It follows that that spectrum should be made available contiguously, specifically to suit wider channels; in this case offering spectrum in multiples of 160 MHz should be considered best practice as this offers the greatest speed potential for a user's device. For this reason, we have given spectrum demand figures which have been rounded to the next higher multiple of 160 MHz, in all cases.

Some users may have higher or lower speeds, depending on the signal quality which will depend on distance from the AP, obstructions and interference. The maximum speeds are set by the technology predicted in 2020 and summarised in Chapter 2.

²³ Surveys and forecasts generally report average traffic levels, rather than busy hour.



5.3.2 2025

Busy Hour scenario 2025

We repeat the graphs of the previous section, but this time for predicted demand, device performance and location distribution appropriate to 2025.




Figure 5-3 Busy Hour scenario, year 2025, using 80 and 160 MHz channels (target utilisation 70%)

In the busy hour scenario for 2025, 1120 MHz of spectrum is required to achieve an AP utilisation of 70%. Once again 160 MHz channels give the greatest opportunity for bursty high speed data usage, yet do not demand any more total spectrum than 80 MHz channels in this case.

Upper Bound scenario 2025




Figure 5-4 Higher Demand scenario, year 2025, using 80 and 160 MHz channels (target utilisation 70%)

In the upper bound scenario for 2025, 1920 MHz of spectrum is required to achieve an AP utilisation of 70%.



We have given all spectrum estimates as multiples of 160 MHz, in order to maximise the highest potential transmission speed²⁴. We note that, although not in the 802.11 standards, in principle even wider channels may be considered in the future. This would increase the need for contiguous spectrum - although not necessarily total spectrum - and the benefit would be a higher maximum transmission speed potential.

5.4 Spectrum requirements summary

While spectrum requirements may appear modest based on busy hour growth predictions, Table 5-1 shows that the upper bound scenario calls for substantially more spectrum, especially in 2025. To avoid doubt the values in Table 5-1 represent total spectrum demand, in other words including spectrum already allocated to Wi-Fi at 5GHz²⁵.

Scenario	2020	2025
Busy Hour	800 MHz	1120 MHz
Upper Bound	960 MHz	1920 MHz

Table 5-1 Total spectrum required by year, for base and higher demand scenarios.

The difference between the two demand scenarios consists of relatively modest uplifts of 30% in 2020 and 78% in 2025, which we feel is far from being beyond the realms of possibility. The residential location is the most demanding in each case and hence this sets the spectrum requirements. Looking ahead to 2025 it would seem reasonable to consider investigating channels which are wider than the current maximum of 160 MHz. This would call for a greater degree of contiguous spectrum, although not necessarily more total spectrum.

5.5 Power level sensitivity analysis

Considering a general case, Figure 5-5 shows the effect of a 10dB increase or decrease in power level. This is relevant since different power restrictions exist in different regions.




Figure 5-5 Effect of increasing or decreasing transmit power level.

²⁴ We note that 160 MHz channels are not presently allowed in China.

²⁵ See Chapter 6 for our analysis of the amount of new spectrum required, per region.



The result shows that back-off behaviour will affect throughput. In this case, reducing power will cause back-off to be triggered less, but the danger is that signal quality and hence throughput will also decrease²⁶. Reducing power is thus not a panacea.

At the time of writing, the 802.11ax Task Group is discussing various methods to avoid the back-off issue, such as BSS colouring. This reduces time spent in back-off; although it does still suffer a degraded signal quality relative to a clear channel.

²⁶ Since noise will eventually become dominant in the signal to interference plus noise (SINR) metric.



6 SPECTRUM AVAILABILITY AND GAP ANALYSIS

6.1 Spectrum at 5GHz

Spectrum at 5GHz is of great interest. This is because the current generation of Wi-Fi, 11ac, targets 5 GHz exclusively. In turn this is driven by the exhaustion of 2.4 GHz. 11ac (and 11ax to be standardised by 2019²⁷) offer a combination of range and rate that is unmatched by 2.4 GHz or 60 GHz.

There has been some movement in both Europe and the USA to extend the bandwidth available at 5GHz. However, extending the spectrum is not without its issues, especially with respect to incumbent usage in the 535-5470 MHz gap, notably EESS²⁸ as we evaluated in our earlier work which was submitted to ITU JTG 4-5-6-7²⁹. At the top end of the band the challenge comes from co-existence with radars and potentially lower power limits.

Furthermore, some 5 GHz channels are subject to DFS (Dynamic Frequency Selection), as a way to limit interference to other band users, specifically the radio location service. This has led to an observation that channels subject to DFS may be less used than other channels. By way of an informal example, we noted an effect in Wi-Fi channel usage during simple walking surveys in central London and San Francisco³⁰, see Figure 6-1 and Figure 6-2.




Figure 6-1 Walk survey results from London, number of APs detected per channel




Figure 6-2 Walk survey results from San Francisco, number of APs detected per channel

²⁷ http://www.ieee802.org/11/Reports/802.11_Timelines.htm

²⁸ Earth Exploration Satellite Service.

²⁹ "5 GHz Co-existence Investigations: Final Report", Quotient Associates for Ofcom UK, February 2014.

³⁰ "Study on the use of Wi-Fi for Metropolitan Area applications", Aegis and Quotient for Ofcom UK.



Figure 6-1 from the London walk survey shows that there are more channels used in the low band of 5 GHz (channel 36 to 64) than in the high band (channel 100 and above). The effect is more dramatic in the San Francisco walk survey, Figure 6-2, where no high band usage of 5 GHz was seen at all.

We note that DFS requirements are not uniform; rather they vary by geographical region in response to the different radars used. Where a Wi-Fi AP is used globally, it must comply with each and every local DFS requirement. This adds complexity to the design.

6.1.1 Spectrum availability per region

The amount of spectrum presently available to Wi-Fi at 5GHz varies by region, see Table 6-1.

<i>Band (MHz)</i>	<i>Bandwidth (MHz)</i>	<i>Europe</i>	<i>USA</i>	<i>Japan</i>	<i>China</i>
5150-5250	100	Yes	Yes	Yes	Yes
5250-5350	100	Yes	Yes	Yes	Yes
5470-5725	255	Yes	Yes	Yes	No
5725-5850	125	No	Yes	No	Yes
Total bandwidth (MHz)	580	455	580	455	325

Table 6-1 Wi-Fi spectrum availability per sub-band and region.

Within these overall spectrum ranges, regions also differ on the requirement for DFS, see Table 6-2. In summary, while Europe and Japan are similar, other regions have not only different amounts of spectrum allocated at 5GHz, but also different proportions which are subject to DFS. We note that 5725-5850 is not currently used for Wi-Fi in Europe due to a 25mW power restriction, although FWA is allowed. This band may be opened to Wi-Fi in Europe in future.

<i>Band (MHz)</i>	<i>Europe</i>	<i>USA</i>	<i>Japan</i>	<i>China</i>
5150-5250	No	No	No	No
5250-5350	Yes	Yes	Yes	Yes
5470-5725	Yes	Yes	Yes	N/A
5725-5850	N/A	No	N/A	No
Total - No DFS (MHz)	100	225	100	225
Total - DFS (MHz)	355	355	355	100

Table 6-2 Wi-Fi DFS requirement per sub-band and region.

We are interested in both the total spectrum available and that proportion constrained by DFS. We may then compare this available spectrum with the total spectrum requirements we predicted for Wi-Fi in the years 2020 and 2025, in Table 5-1 on page 21.



6.2 Gap analysis

By comparison with the total spectrum requirement predictions of Table 5-1 on page 21 and the available spectrum with and without DFS constraints in Table 6-2, we can predict the spectrum shortfall per region. To be clear, this spectrum shortfall is the amount of new spectrum which will need to be found and made accessible for Wi-Fi use, if Wi-Fi is to meet the demand we predicted in the years 2020 and 2025.

We illustrate this graphically in Figure 6-3 and more precisely in tabular form in Table 6-3.




Figure 6-3 Illustration of the spectrum shortfall per region, by year and demand level.

In summary, for the year 2025, the various regions are likely to need to find between 500 MHz and 1 GHz more spectrum to satisfy the Busy Hour scenario, which reflects the widely expected growth in traffic. If demand exceeds the present Busy Hour predictions, the Upper Bound scenario suggests that an estimated maximum of between 1.3 and 1.8 GHz more spectrum may be needed. Such additional spectrum might be required due to unexpected adoption of novel applications or a further concentration of the busy hour traffic into fewer than the assumed four hours per day, for example.

While the predicted total amount of spectrum required is the same across regions (Table 5-1), the shortfall per region (Figure 6-3) depends on available spectrum. It also depends on the proportion of that spectrum which is assumed to be usable due to the effects of DFS constraints per region, as follows.

To account for the apparent under-use of DFS spectrum (see Figure 6-1 and Figure 6-2), we have assumed two cases. Firstly where all DFS spectrum is actually used and secondly where only 30% of the spectrum constrained by DFS is used. As expected, where less DFS spectrum is usable, more new spectrum is required.



It is interesting to note that, as the amount of the spectrum shortfall increases with year and demand level, the relative influence of DFS spectrum reduces. This occurs because we have chosen to assume that new spectrum will not be constrained by DFS. Clearly if any new spectrum were identified that also required DFS, then the influence of DFS constraints would no longer diminish as we have shown. In other words our assumption is that new spectrum made available for Wi-Fi will be readily accessible by Wi-Fi.

This is a very important point and we note that DFS conditions, some of which were set more than ten years ago are presently being re-examined in international fora. Given the apparent difficulties of using DFS spectrum as suggested by the walk tests, we note that removing or reducing existing DFS constraints might reduce the need to make more spectrum available for use by Wi-Fi at 5 GHz³¹.

Table 6-3 contains the source data which was used to construct the graphical illustration of Figure 6-3.

<i>Spectrum Shortfall (MHz)</i>	<i>2020 Busy Hour</i>	<i>2020 Upper Bound</i>	<i>2025 Busy Hour</i>	<i>2025 Upper Bound</i>
Europe & Japan - all DFS spectrum in use	345	505	665	1465
Europe & Japan - 30% DFS spectrum in use	593	753	913	1713
USA - all DFS spectrum in use	220	380	540	1340
USA - 30% DFS spectrum in use	468	628	788	1588
China - all DFS spectrum in use	475	635	795	1595
China - 30% DFS spectrum in use	545	705	865	1665

Table 6-3 Spectrum shortfall per region, by year and demand level.

³¹ An analysis of DFS conditions and their practical effects is beyond the scope of the present study.



7 SUBSTITUTION POTENTIAL OF SPECTRUM SUPPLY AND NEW BANDS

7.1 Spectrum characteristics

For alternative spectrum to offer a substitute for 5 GHz spectrum from the point of view of an application, it would need to offer similar propagation and data carrying ability.

For the example of 2.4 GHz, atmospheric loss is broadly similar to 5GHz; see Figure 7-1, although building penetration loss may be much higher in some materials, such as red brick for example. However, the bandwidth available is too low in practice, totalling only 83 MHz.




Figure 7-1 Loss due atmospheric absorption and rain fade versus frequency (FCC OET).

For the case of UHF spectrum (e.g. 700 MHz) the propagation loss is significantly smaller and self-interference would be more likely, unless powers were kept low and/or cell sizes increased. However the main problem at UHF is the limited bandwidth available compared to 5GHz. This is more a fundamental problem with UHF since it is simply not realistic to expect to use 1 GHz bandwidth at only 700 MHz.

It is an often used engineering rule of thumb that when the bandwidth required exceeds 10% of the given carrier frequency, then the design is classed as wideband and becomes significantly more challenging. We can see from the following list that more bandwidth is more easily achieved at higher carrier frequencies.

- 10% of 700 MHz = 70 MHz
- 10% of 2.4GHz = 240 MHz
- 10% of 5.5GHz = 550 MHz
- 10% of 60 GHz = 6GHz



We note that 60 GHz offers a great deal of bandwidth due to its higher carrier frequency. However the propagation and penetration with respect to 5GHz is much poorer (Figure 7-1). This makes 60 GHz suitable for small areas or short distances, plus it will not usefully penetrate walls. It is also subject to significant rain attenuation if used outdoors. 60 GHz thus cannot normally be considered a good substitute for 5 GHz. Nonetheless, there may well be a case for expanding the use of 60 GHz in order to satisfy those new and as yet uncharacterised demands such as Virtual Reality and Augmented Reality, especially if these future applications prove to be ones which tend to be used within the confines of a room (as might VR gaming, for example).

In summary, spectrum in the range 2-10 GHz may be expected to offer a reasonable degree of substitutability for 5 GHz in terms of propagation and bandwidth. This is subject of course to consideration of other users in this spectrum range, which are many and varied. The need for contiguous spectrum in order to support wide Wi-Fi channels provides a further constraint.



8 CONCLUSIONS AND RECOMMENDATIONS

8.1 Future Wi-Fi spectrum requirements by region

We have shown that between 500 MHz and 1 GHz of new spectrum will be needed in 2025 to satisfy the anticipated busy hour, with between 1.3 and 1.7 GHz required if demand exceeds the busy hour prediction by a relatively modest 78%, for example due to novel and as yet un-anticipated applications, or the further concentration of traffic into fewer busy hours than the present four hours per day. The amount of new spectrum required varies by geographical region, and our analysis illustrates potential effects due to spectrum which is subject to local DFS requirements³². Our analysis assumes that new spectrum will be fully accessible by Wi-Fi.

Our predictions for the new spectrum required per region are as shown in Figure 8-1³³.




Figure 8-1 Illustration of the spectrum shortfall per region, by year and demand level.

The spectrum predictions cover the years 2020 and 2025; predicted Busy Hour and Upper Bound demand levels; two different usage levels of DFS spectrum; and three locations types consisting of office, residential and mall – of which residential was found to have the greatest spectrum requirements. Built into the predictions are technology advances, for example in terms of transmission rate, video coding improvements and device capability.

8.2 The importance of contiguous spectrum

In addition to simply needing more spectrum in total, we have shown that such spectrum needs to be assigned with sufficient contiguity such that wide channels of 160 MHz, or

³² Spectrum, access to which is determined by the constraints of Dynamic Frequency Selection procedures.

³³ See also Table 6-3 on page 27.



perhaps even wider in future, can be constructed with ease³⁴. To do otherwise would be to restrict the growth of Wi-Fi and the economic benefits with which it is widely associated and which was first enabled by forward-thinking spectrum regulation. Such a need for contiguity presents a significant further challenge to those with responsibility for spectrum allocation.

³⁴ Where regulation allows.



Wi-Fi Spectrum Needs Study | Conclusions and recommendations

Final Report : qa1042 © Quotient Associates Ltd. 2017

Commercial in Confidence. No part of the contents of this document may be disclosed, used or reproduced in any form, or by any means, without the prior written consent of Quotient Associates Ltd.

9 APPENDIX A: MODEL PARAMETERS

Number of channels: 2.4GHz – 3; 5GHz – variable; 60GHz – 3.

Transmit power: 20dBm.

Path loss coefficients: see Appendix B.

Channel bandwidths: 2.4GHz – 20MHz, 5GHz – variable, 60GHz – 2160MHz.

Percentage use of 60GHz: Office – 10%, Residential – 20%, Mall – 0%.

Machines per person = 3.

Noise rise above thermal: 2.4GHz – 15dB, 5GHz – 10dB, 60GHz – 10dB.

Dimensions of buildings: IEEE 802.11 office and residential location types³⁵, plus the mall with dimensions 30m x 300m x 2 floors.

Access points per floor: AP density is set in m^2 and the user density also in m^2 .

1 AP per $100m^2$ in office and residential and 1 per $150 m^2$ in mall.

1 person per $10 m^2$ in office, 1 per $25 m^2$ in residential and 1 per $7 m^2$ in mall. So e.g. 4 people per AP in residential

Data volumes: variable.

³⁵ 11-14-0980-16-00ax-simulation-scenarios, downloaded from http://www.ieee802.org/11/Reports/tqax_update.htm



10 APPENDIX B: MODEL STRUCTURE

10.1 Model structure (pseudo code)

Define environment

- Set x and y dimensions of floor;
- Set number of floors;
- Set average distance between interior walls;
- Set average wall loss and average floor loss;
- Set access points per floor;
- Set users per floor;
- Set data rate requirements per user in MB over busy hour.

Place access points

- Place APs on even rectangular grid across floor;
- For each band in use;
 - For office and mall assign frequencies across access points based on available number of channels using algorithm where successive APs look for lowest interference channel and grab it.
 - For residential assign frequencies randomly.

Place users

- Place users on even rectangular grid across floor.

Place machines

- Place machines on even rectangular grid across floor.

Calculate max data rates based on SINR per user

- For each band in use
 - Find closest access point on same floor;
 - Calculate path loss to serving access point and hence SNR;
 - Calculate path loss to every other access point using same frequency and add interference levels;
 - For those APs where the interference level is above the threshold then assume that terminals will back-off.
 - Add in assumed levels of non-Wi-Fi interference;



- Use look-up tables to convert to data rates.

Calculate transmission time per user as % of total time assuming perfect balancing of load across band³⁶.

Add total transmission times per access point.

Build histogram of congestion.

³⁶ We define utilisation in our airtime based model as that percentage of airtime that an AP observes as being utilised, both by itself and other neighbouring co-channel networks.

