Second consultation on assessment of future mobile competition and proposals for the award of 800 MHz and 2.6 GHz spectrum and related issues

Annex 7

Consultation

Publication date:	12 January 2012
Closing Date for Responses:	22 March 2012

Contents

Annex 7

LTE technical modelling results

Introduction

A7.1 This annex presents some of the key results from our technical analysis of the downlink performance of LTE macrocell networks using paired spectrum which has been drawn upon to inform our policy analysis in Section 3 of Annex 6 of the consultation.

A7.2 The technical model we have used is an evolution of the model developed for the March 2011 consultation. For the most part the changes to the underlying model are relatively small. However, following a detailed review of responses to the March 2011 consultation and further internal analysis, a number of the parameters and assumptions have been changed. One significant change to the modelling since March 2011 is the way we approach the analysis of multi-frequency networks. For a comprehensive description of the modelling methodology and the underlying parameters and assumptions see Annex 14 and for an overview of the technical responses to the March 2011 consultation and how we have addressed these see Annex 15.

A7.3 The technical model we have used is based specifically on LTE technology and the results presented below relate to the downlink performance of LTE networks using paired spectrum. Our model only looks at the macrocell layer of LTE networks. It does not include alternative methods of dealing with dense traffic ‘hot-spots’ and/or traffic in hard to serve locations by techniques such as Wi-Fi off-load or deploying femtocells.

A7.4 The model is affected by several forms of uncertainty:

- Firstly, any model is only an approximation to reality. This is particularly true when what is being modelled is inherently non-deterministic (as per indoor coverage);
- Secondly, this model is forward looking. No one is certain what the performance of LTE will actually be;
- Thirdly, we have based our model on data which is, for the most part, in the public domain. We recognise that individual operators may have more data available, and would welcome comments from those operators on any aspect of the data, assumptions and modelling methodology we have used.

A7.5 Any attempt to derive the performance of a mobile network using a theoretical modelling approach is inevitably going to be inherently uncertain, and this is particularly the case when modelling new technologies such as LTE because the model must, by necessity, include estimates of current performance and potential improvements in performance as the technology matures for which there will be limited empirical evidence. The extent of this uncertainty is reflected in the number of comments and differing views expressed on the modelling parameters and assumptions in response to our March 2011 consultation.

A7.6 Whilst the results presented in this annex reflect the best knowledge we have as a regulator from our own research, expertise and information received from stakeholders, it is unrealistic to believe that any model can be anything more than illustrative of the real performance of actual LTE networks. We believe that the model is useful in comparing the relative variation in performance between networks operating at different frequencies. It is less useful in providing information on the absolute performance and in no sense should the results be taken as a definitive prediction of network performance.

A7.7 The results presented here do not consider questions of technology or equipment availability, e.g. how quickly it might be possible to re-farm 900 MHz and 1800 MHz spectrum. They also abstract from issues of equipment availability (including lack of support in the standards for certain combinations of bandwidth and frequency band), e.g. availability of equipment that supports a 2x15 MHz carrier in 900 MHz

Question A7.1: We would welcome comments on any aspect of the data, assumptions and modelling methodology we have used in our technical analysis, in particular our approach to serving users in a range of both easier and harder to serve locations?

Physical behaviour of networks at different frequencies

A7.8 The performance of a radio link is typically determined by its signal to interference plus noise ratio (SINR), the higher the SINR the better the performance. SINR can be used to characterise how difficult it is to serve a user at a particular location; relatively high SINRs being characteristic of easy to serve users/locations and relatively low SINRs being characteristic of difficult to serve users/locations. In a mobile network, the main source of interference in the downlink (i.e. transmissions from the base station to the user) is from cells neighbouring the serving cell which transmit on the same set of frequencies and times. For the uplink (i.e. from the user to the base station) it is from the terminals of other users which transmit on the same set of frequencies.

A7.9 Across a network there will be a wide range of users in a variety of easier and more difficult to serve locations. Difficult to serve users might typically be found deeper in-buildings, and/or further from the serving base station, and/or have an obstructed path between their location and the serving base station, etc. LTE implements an adaptive modulation and coding scheme. When SINR is high, a higher order modulation scheme is used (e.g. 64 QAM giving 6 bits per symbol) with a coding scheme that implements little or no redundancy which gives greater throughput. When SINR is low, a lower order modulation scheme is adopted (e.g. QPSK giving just 2 bits per symbol) with a coding scheme that implements a significant amount of redundancy. This allows the receiver to decode increasingly poor quality signals. However, when SINR gets very low the receiver is unable to effectively decode the wanted signal or maintain synchronisation with the network and the link is lost completely.

A7.10 As an example, as users get deeper into buildings, building penetration loss (BPL) rises. ‘On average’ this would affect the wanted signal from the serving cell just as much as the unwanted (interfering) signals from other surrounding cells[1]. So, ‘on average’, it might be expected that the signal and interference would be attenuated roughly equally; hence the signal to interference ratio (SIR) would remain constant. However, the fact that the shadowing and BPL for the wanted and unwanted signals are not completely correlated means that SINR would be likely to degrade

---

[1] The ‘on average’ here refers to the overall effect over all locations. For any specific location this will not hold true.

somewhat at greater depths indoors. In addition, the receiver noise floor[2] also needs to be accounted for. This noise floor has a fixed minimum value (due to thermal noise and the receiver's noise figure) so as the wanted signal gets weaker and weaker the SINR degrades even more quickly and performance reduces still further.

A7.11 When performance is dominated by interference (i.e. the interference power is significantly greater than the noise) a link is often referred to as 'interference limited' and conversely, when performance is dominated by noise (i.e. noise power is significantly greater than the interference) a link is often referred to as 'noise limited'. When networks are interference limited, performance differences between frequencies are minimised and when they are noise limited, performance differences between frequencies are likely to be at their greatest. This is because signals attenuate differently at high and low frequencies as they propagate from their source. Higher frequencies attenuate more rapidly with distance and as they get deeper and deeper into buildings their losses may also increase more rapidly than lower frequencies and the level of the wanted signal will approach the receiver noise more rapidly at higher frequencies than lower ones. Hence it will consequently have a lower SINR and poorer performance, i.e. it will have more locations where it is 'noise limited'. It should be stressed that there is no hard limit between 'noise limited' and' interference limited' cases, there being a smooth transition from one to the other as the network loading increases.

A7.12 A consequence of the above is that we would, in general, expect networks with relatively few sites (i.e. having large physical distances between their base stations) to have proportionately more locations that are 'noise limited' at higher frequencies than at lower ones than a network with many more base stations. Hence they would have both a poorer overall performance and a greater performance difference between high and low frequencies. Also, as users get deeper and deeper into buildings, if there is a greater building penetration loss at higher frequencies, there will be more locations inside buildings that are 'noise limited'. Hence higher frequencies will have poorer performance at greater depths within a building.

A7.13 It is perhaps worth mentioning a couple of related techniques that LTE networks can employ to improve their performance in the presence of interference. These are frequency domain packet scheduling and intelligent resource allocation and they are briefly described below:

- Frequency domain packet scheduling exploits the fact that an LTE carrier is split into multiple sub-carriers and that these sub-carriers are grouped together in frequency and split by time to form individual physical resource blocks (a 5 MHz carrier has 25 resource blocks whilst a 20 MHz carrier has 100). At any instant in time different users can be allocated a different number of physical resource blocks depending on their instantaneous demand and their signal quality. If the channel quality is significantly different for different physical resource blocks (which is typically the case for macro cellular networks with bandwidths equal or greater than about 5 MHz) then LTE can exploit this by optimally scheduling users on physical resource blocks with the best channel quality at their location. This can lead to a frequency domain packet scheduling gain which for pedestrian users can be of the order of 40% for a 10 MHz system bandwidth. To achieve this level of gain however there needs to be multiple users all demanding a service that requires a relatively small proportion of the resources available in the cell at

---

[2] The noise generated within the receiver itself above thermal noise.

any instant of time. If there are just a few users requiring a large proportion of resources then the gain is reduced.

- Intelligent resource allocation is based on adopting an intelligent approach to scheduling the physical resource blocks between neighbouring cells to avoid, as far as possible, the network scheduling the same physical resource blocks on cells that are close to each other. This reduces the amount of interference users served by one cell will see from the surrounding cells in the network. This technique is most effective when cells are relatively lightly loaded[3] which gives room for the scheduler to avoid collisions between cells but it is much less effective when cells are relatively highly loaded.

A7.14 We have simulated both of these techniques in our modelling as they have the potential to significantly impact the results. Frequency domain packet scheduling can improve average cell throughput and hence capacity. However it does not help improve coverage; this is because at the limit of coverage a user will require the vast majority of cell resources and therefore the scheduler cannot selectively schedule them on physical resource blocks with the best channel quality (it has to use them all). Intelligent resource allocation can improve both single user throughput and coverage by minimising inter-cell interference but it is most effective when the network is relatively lightly loaded becoming less and less effective as loading increases.

Summary of modelling methodology

A7.15 Underlying all the results presented in this annex are SINR distributions generated across two simulation areas. The majority of results are for the same 100 km x 100 km area we used in the March 2011 consultation (referred to as the West London area).

Figure 1: West London simulation area

---

[3] It has been suggested to us that when a cell is lightly loaded and traffic is intermittent, the network's ability to accurately estimate channel quality is likely to be reduced. This would tend to reduce the performance of the network in these circumstances to some extent – see Annex 15 for further discussion of this.

A7.16 This area was chosen as we believe it is reasonably representative of the more populous areas of the country where competition between operators will be predominantly focused.

A7.17 In addition we also report a limited set of coverage results for another area, a 100 km x 100 km area around Cambridge.

Figure 2: Cambridge simulation area

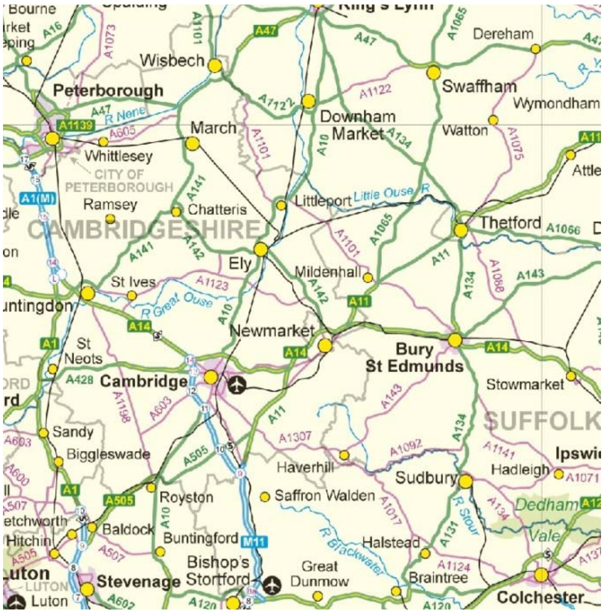

A7.18 This area was chosen to contrast the first. It is much less populous but is likely to be an area where reasonably contiguous coverage could be provided (as is the case with 2G coverage in this area today).

A7.19 In Annex 14 we give a detailed description of the model. However a high level description is as follows:

i) A synthetic base station network of a particular size (number of sites) is established covering the simulation area plus a buffer zone 20 km deep surrounding the simulation area. The base station network is constructed, as far as is possible, to have similar characteristics (in terms of site density vs. population density, antenna heights, etc.) as current or potential future mobile macrocell networks.

ii) An SINR distribution is calculated for a hypothetical test terminal (UE) positioned at the geographic location of a randomly selected sample[4] of postcode units[5] within the simulation area. SINR is calculated taking into account signals from sites within the base station network within a certain distance (20 km) of the each sample postcode unit location up to a maximum of the 20 closest sites.

iii) Using the SINR distribution generated in the step above together with an appropriate SINR to throughput mapping function, and taking into account system overheads, the average downlink single-user throughput distribution for the sample of postcode unit locations is established.

---

[4] 10,000 sample points

[5] A postcode unit is a sub-area of a postcode sector as extracted from Code-Point® data.

iv) Steps i), ii) and iii) above are repeated to establish SINR and single-user throughput distributions for a range of base station network sizes, network loadings, carrier bandwidths and building penetration depths for the frequencies under consideration (e.g. 800 MHz, 1800 MHz and 2600 MHz).

v) From the single-user throughput distribution statistics within the simulation area, the three metrics of performance are calculated:

- o *Coverage* – the proportion of domestic delivery points[6] within the simulation area to which it is technically possible to deliver a service with a particular downlink speed (if 85% of the resource blocks of the serving cell, including system overheads, were dedicated to a single customer[7]), as a function of the number of network sites and the loading on the wider network.
- o *Speed* – for a given number of sites and network loading, the downlink single-user throughput (if 85% of the resource blocks of the serving cell, including system overheads, were dedicated to a single customer) attained or exceeded by a particular proportion of domestic delivery points within the simulation area.
- o *Capacity* – for a particular wider network loading, the capacity to simultaneously serve a particular proportion of domestic delivery points within the simulation area with a given downlink speed and number of sites.

A7.20 The outdoor propagation model we are using is the Extended Hata model[8].

A7.21 Our modelling of mobile service provision within buildings is detailed in Annex 14. This adopts an approach that is consistent with our earlier work. This is because:

- The February 2009 consultation, Application of spectrum liberalisation and trading to the mobile sector[9], provided a detailed description of the issues around propagation into buildings, consideration of the literature and available evidence, and justification for the values used (see Annex 13, paragraphs A13.220 to A13.261[10]). It also included a consideration of information requested from mobile operators on their own approach to planning for indoor customers. It summarised the conclusions on appropriate modelling and parameters and consulted on that issue.

- For the March 2011 consultation, Assessment of future mobile competition and proposals for the award of 800 MHz and 2.6 GHz spectrum and related issues[11], we adopted the same basic modelling approach and parameter values consistent with the February 2009 consultation.

- For this consultation we conducted a further review of the literature on building penetration loss and relevant responses to the earlier consultations and

---

[6] Each postcode unit has associated with it a number of domestic delivery points: each delivery point will generally correspond to one residential address.

[7] 85% is considered, for the purposes of this analysis, to be a practical upper bound to loading on average

[8] ERC Report 68 and http://tractool.seamcat.org/raw-attachment/wiki/Manual/PropagationModels/ExtendedHata/Hata-and-Hata-SRD-implementation_v2.pdf

[9] http://stakeholders.ofcom.org.uk/consultations/spectrumlib/

[10] http://stakeholders.ofcom.org.uk/binaries/consultations/spectrumlib/annexes/annex13.pdf

[11] http://stakeholders.ofcom.org.uk/consultations/combined-award/

concluded that there is no strong evidence to justify switching to a different overall approach.

A7.22 The majority of our earlier analysis was focused primarily on reception at relatively hard to serve locations (proxied by locations deep indoors). In the analysis presented here we have extended this to encompass analysis of easier to serve locations (proxied by shallower depths within buildings).

A7.23 As per the February 2009 consultation the simulation models signals propagating indoors as being attenuated by two components; firstly a loss at the external wall and secondly by an increasing loss as the signal propagates further and further indoors. This is discussed in detail in Annex 14. Though presented as an actual physical distance from the external wall, we exercise caution in interpreting this literally. For instance, whilst our results for a depth of 1 metre may represent someone very close to the external wall where the major influence is the external wall loss, our results for a depth of 15 metres could be taken to represent a user physically very deep within a relatively low loss building but could also represent a user who is at a shallower physical depth but subject to greater propagation losses e.g. behind several internal walls or in a building with a very thick external wall etc. So our interpretation of the analysis is one of ability to serve a distribution of easier and harder to reach locations, rather than one of serving users at absolute depths in a building.

A7.24 Our subsequent analysis of the overall characteristics of networks with different spectrum portfolios necessarily has to make an assumption about the distribution of demand over locations which are easier and harder to serve. For the purpose of this analysis we have chosen to assume an equal number of users at each of the ‘depths’ we have modelled.

A7.25 We acknowledge this further analysis has not been subject to the same degree of consultation as our previous analysis of harder to serve locations. We therefore particularly welcome stakeholder views on our modelling and analysis of serving users in a range of locations, both easier and harder to serve.

Modelling of multi-frequency networks

A7.26 A number of the capacity results presented below are for networks with spectrum in more than one frequency band. For these cases the aim of the analysis is to understand the overall capacity of a multi-frequency (or multi-carrier) network where traffic is managed between the available carriers in such a way as to ensure that users, regardless of which actual carrier they are served by, can receive the same (or similar) quality of service (i.e. can get the same specified data-rate with the same likelihood of success, irrespective of their location).

A7.27 Our analysis is limited to modelling carriers in two different frequency bands or to two carriers in the same band. We have assumed that the traffic from harder to serve locations is preferentially served on the lower of the available frequency bands and that traffic from easier to serve locations is served on the higher of the available frequency bands. Obviously, if we model two carriers in the same band the approach is adapted accordingly but we still assume that one of the carriers is dedicated to the traffic from the harder to serve locations and the other to the traffic from the easier to serve locations.

A7.28 A detailed description of the multi-frequency analysis is contained in Annex 14. However a brief outline is provided below.

A7.29 Each multi-frequency spectrum portfolio is defined in terms of the frequency and bandwidth of the carrier in the lower frequency band and the frequency and bandwidth of the carrier in the higher frequency band. As indicated above, in some instances, we model both carriers in the same band though we still differentiate the carriers by their bandwidth.

A7.30 We analyse the performance of multi-frequency networks for a range of target coverage between 60% and 100% of locations served (in 1% steps). For each target coverage level, locations are assigned to each frequency band assuming that easier to serve locations are preferentially served by the carrier in the higher frequency band and harder to serve locations are served by the carrier in the lower frequency band up to the target coverage level.

A7.31 The proportion of locations assigned to each frequency carrier and their loadings are adjusted to ensure that:

- the target coverage can be achieved (if possible); and
- the overall demand that can be served is maximised conditional on the proportion of demand that can be served by each layer being the equal (so that the users on each layer experience the same likelihood of being able to receive service).

A7.32 If it is not possible to achieve the target coverage or to equalise the proportion of demand across the two frequency bands for a given level of coverage then the capacity is conservatively assumed to be zero.

A7.33 We are not suggesting that this is exactly how an operator would, in practice, manage traffic on their network. However it is likely that an operator will want to manage traffic, to the extent possible, in a way that maximises the overall performance of their network as a whole. The extent to which they can do this will be dependent on a number of factors, including for example whether there are a significant number of terminals on the network that are not able to work at all network frequencies. There are also practical issues in optimally managing traffic across multiple carriers (e.g. avoiding ‘ping ponging’ between layers when a user’s channel quality is close to a handover threshold, unavoidable inaccuracies in estimating channel quality, etc.) that mean that our approach is almost certainly a simplification. As a consequence of these practical issues, our results are likely to over-estimate the contribution of lower frequencies to the performance of a multi-frequency network because in a real network a proportion of traffic will inevitably be carried on the less optimal frequency layer. Despite these limitations, we believe that the results are useful in illustrating the difference in relative potential performance of different multi-frequency spectrum portfolios. We discuss these limitations further in Annex 14.

Presentation of variability in our results

A7.34 To reflect the major areas of uncertainty we have chosen to model a range of values for key parameters. To illustrate this range we have chosen to group the parameter values into two cases: those that tend, in most circumstances, to minimise the relative performance variation between frequencies (‘Min var’) and those that tend, in most circumstances, to maximise the relative performance variation (‘Max var’). The model is then run twice to produce results for these two cases. Table 1 below shows the combination of parameters we have used for the ‘Min var’ and ‘Max var’ cases.

Table 1: Min var and max var parameters

Parameter	Min var	Max var	Comment
SINR cut-off	-10 dB	-5 dB	The lowest SINR value for which a viable downlink service can be received (based on the performance of the most sensitive control channels)
BPL standard deviation	800 MHz: 4.0 dB 1800 MHz: 5.4 dB 2600 MHz: 6.0 dB	800 MHz: 8.0 dB 1800 MHz: 10.8 dB 2600 MHz: 12.0 dB	The standard deviation of the propagation loss incurred in penetrating into buildings
Median BPL	Lower absolute value and zero frequency dependency	Higher absolute value and strong frequency dependency	The median value of the propagation loss incurred in penetrating into buildings. See Annex 14 for details.
Resource allocation algorithm	Random	Intelligent	The intelligent algorithm attempts to optimise the allocation of the resource blocks on the serving cell and its neighbouring cells so as to minimise interference, while the random algorithm distributes user data randomly amongst resource blocks

A7.35 The results in this annex are presented as graphical illustrations and the ‘Min var’ and ‘Max var’ sets are typically shown as separate lines on these graphs.

A7.36 Strictly speaking, there are two different types of parameter that make up our ‘Min var’ and ‘Max var’ sets. The ranges for SINR cut-off and resource allocation algorithm represent an uncertainty in our knowledge about how an operator might manage their network. In principle, with more information we could reduce or possibly eliminate these sources of uncertainty and we would welcome any additional information from operators. However, the ranges for BPL standard deviation and median BPL represent the current uncertainty around the nature of propagation into buildings due to the myriad of different paths, locations, building types, construction materials, internal layouts and the relative importance of those locations for customers who are sensitive to differences in service quality etc. Even with perfect knowledge of every possible parameter we could never build a practical model that would eliminate these uncertainties (though, potentially, with better knowledge it could be reduced somewhat).

A7.37 In this annex the gap between the ‘Min var’ and ‘Max var’ lines is an indication of the uncertainty in our model’s prediction of performance. It illustrates the extent of knowledge we, as a regulator, can have about LTE network performance, the inherent uncertainty in propagation into buildings and the choices operators may make in relation to implementation of intelligent resource scheduling algorithms and influencing SINR cut-off by, for example, increasing power in control channels[12]. For any particular location, we believe that it is more likely than not that performance will lie somewhere between our ‘Min var’ and ‘Max var’ lines, but we are not making any specific judgement as to the likely distribution of results within this range. Operators will have better knowledge about these factors but they will likely still have a degree of uncertainty and this uncertainty may influence their network planning.

A7.38 Our results rely on mapping SINR to throughput (a measure of user data-rate in Mbps) for a 2 x 2 (transmit x receive) antenna configuration. Though we have used a particular mapping function that we believe is a reasonable representation of the likely performance of real LTE networks when they are deployed, the actual performance real LTE networks will achieve over time remains uncertain. It is possible that for a particular SINR value the data-rate that a real network could support might be significantly greater or less than the value given by the function we are using. In Annex 14 we present additional sensitivity results for an alternative mapping function that gives results that show better performance than the one used here.

A7.39 It should be noted that there are many more sources of modelling variability and uncertainty than are captured in our ‘Min var’ and Max var’ sets. In Annex 14 we have a discussion of a range of these sources. However, we believe that those captured in our ‘Min var’ and ‘Max var’ cases reflect the main sources of variability that influence the relative differences in performance between different frequencies.

Question A7.2: We would welcome any additional information, in particular from current operators, on the choice of parameters making up our ‘Min var and ‘Max var’ cases?

Performance metrics

A7.40 In this annex we mainly explore three key performance metrics:

- Coverage;
- Speed; and
- Capacity

A7.41 Coverage and capacity results are presented for two different network sizes. The first, with an equivalent national site count of 18,000, represents the upper end of

---

[12] It is likely that operators will have greater knowledge (and hence less uncertainty) about many of these factors which influence LTE network performance. However, they will still face a level of uncertainty even for the factors that they have direct influence over – for instance the efficacy of intelligent resource scheduling algorithms and their ability to reduce SINR cut-off by increasing power to certain control channels is to some extent uncertain.

the macrocell network size that existing MNOs are likely to have (particularly in the longer term). The second, with an equivalent national site count of 12,000, represents the lower end of the macrocell network size that existing MNOs are likely to have (we anticipate that all operators are likely to have access to at least this many macrocell sites within 2 to 3 years).

A7.42 A number of the results presented below have been modelled on the basis of delivering a guaranteed data-rate service to users. We report results for three nominal values: 1Mbps, 2Mbps and 5Mbps[13]. For the reasons outlined above (e.g. see A7.33), these absolute numbers should be treated with some caution. However, to set these data-rates in context 1-2Mbps is likely to be adequate for the vast majority of smartphone apps whilst 5Mbps may be needed for services such as video streaming to a tablet/laptop. In addition, though LTE can offer a guaranteed data-rate service, it is more likely that the majority of users will be served with a variable data-rate 'best-efforts' service. Networks will dynamically balance the data-rate delivered to users' dependant on a number of factors such as instantaneous demand, channel quality conditions and service mix. At busy times, most users in harder to serve locations will most likely get relatively low data-rates whilst in less busy times and when they are in easier to serve locations users may receive data-rates much closer to the maximum possible. Our choice of guaranteed data-rates is a proxy for a relatively low, medium and higher speed services and for the reasons outlined above it should not be seen as a definitive prediction that these actual speeds can be delivered simultaneously to the corresponding percentage of locations.

Coverage

A7.43 To assess depth of coverage our analysis is based on modelling 5 notional 'depths': outdoors, and 1 metre, 5 metres, 10 metres and 15 metres indoors. For the purpose of this analysis we have chosen to assume an equal number of users at each of these 'depths'. As discussed in paragraph A7.23, though modelled as an actual physical distance from the external wall, 'depth' as used in our model should be interpreted less literally. Our interpretation of the analysis is one of ability to serve a distribution of easier and harder to reach locations, rather than one of serving users at absolute depths in a building.

A7.44 Our approach to modelling coverage looks only at the ability to serve users with a guaranteed data-rate, at the extremes of coverage. This means that each cell would only be able to serve a single user and would therefore have little or no capacity in order to simultaneously serve other users. In practice coverage is likely to be a trade-off between serving the demand from users in easier to serve locations and reserving sufficient resources to serve users in the hardest to serve locations. As demand increases it becomes increasingly difficult to dedicate resources to users in the hardest to serve locations and so an operator may likely either reduce the data-rate to these users, delay their traffic choose not serve them at all.

A7.45 It is important to note that macrocells are not the only means of providing in-door coverage. The results presented here are a lower limit to the coverage that is likely to be achievable, given the ability to deploy femtocells and Wi-Fi offload.

---

[13] Note these are air interface rates and don't account for higher level application protocol overheads.

Speed

A7.46 Speed in our analysis is modelled as the single user throughput. This is the maximum speed that a single user would theoretically be able to receive if they were the only user in the serving cell demanding service at any particular instant of time. For the purposes of our analysis we assume that the user can access 85% of the resources of the serving cell (this reflects our assumption that a practical upper loading limit on any particular cell is on average approximately 85%). This metric is exactly the same as used in the March 2011 consultation. Interpretation of the speed in absolute terms is subject to the same comments made in paragraph see A7.33 over the uncertainty in our ability to predict future LTE network performance.

Capacity

A7.47 In general terms, the capacity of a network is a measure of how much offered traffic it is able to serve whilst maintaining key quality of service metrics. Such metrics might include the number of connection request failures, the number of dropped connections, the ability to maintain a minimum throughput to users, the number of lost data packets, latency, etc. Different users demanding different services from the network will need a different combination of these metrics. For instance, for a streaming video user maintaining an acceptable minimum guaranteed data-rate is important to avoid interrupts; for an online gamer latency might be the most important feature; for someone surfing the web both latency and data-rate may be key. A network will try and balance all of the competing demands of its users. Moreover, if the traffic profile of the users of one network is different from the traffic profile of another network, even if they have the same number of customers and the same network and spectrum resources they might, in practice perform very differently with one network struggling to meet demand whilst the other does not.

A7.48 As a consequence of the above, it is very difficult to derive a single capacity metric that adequately addresses all of the important network quality features that an operator is likely to feel are important. However, as many of these network quality features are likely to be independent of the frequency band they have not all been addressed in this analysis.

A7.49 The approach we have taken here is to abstract from the details and look at capacity in a more heuristic way. For the purposes of our analysis we have assumed all users are to be provided with the same service – a guaranteed data-rate service of a specified speed – and that users are uniformly distributed over all modelled locations (including over all depths at each modelled address). We then assess the relative number of such users that could simultaneously be served by the network, taking account of the resources available to the network and the resources required to serve each user with the specified service. We accept that this is an artificial scenario, but it does allow us to derive an estimate of relative capacity without having to make a lot of detailed assumptions about the specifics of the services that will be demanded by different users.

The graphs

A7.50 We use four types of graph to display the results of our technical analysis:

- Coverage as a function of number of sites;
- Coverage as a function of depth in building;

- Single-user throughput as a function of location;
- Capacity as a function of locations served.

A7.51 For all of the graphs 'Min var' results are displayed with solid lines and 'Max var' results are displayed with dashed lines.

Coverage as a function of number of sites

A7.52 These graphs plot the percentage of locations at which it is possible to receive the specified minimum data-rate on the y-axis versus number of sites on the x-axis.

A7.53 The percentage of locations displayed on the y-axis represents the percentage of locations within the simulation area where our model suggests that the specified minimum data-rate service can be received (a location being the combination of a delivery address and a depth). We assume that users are uniformly distributed over all modelled locations (including over all depths at each modelled address).

A7.54 The number of sites displayed on the x-axis represents the size of an equivalent national network (rather than the actual number of sites included in the simulation area). For instance, a network with the equivalent of 4,000 sites nationally has a total of 602 sites in our West London simulation area. Likewise, a network with the equivalent of 10,000 sites nationally has a total of 2651 sites in our West London simulation area.

Coverage as a function of depth in building

A7.55 These graphs plot the percentage of locations at which it is possible to receive the specified minimum data-rate on the y-axis (see A7.48) versus depth in building on the x-axis.

A7.56 Though these results are shown as continuous smooth lines they are derived from results obtained by modelling at 5 specific depths (outdoors, and 1 metre, 5 metres, 10 metres and 15 metres indoors).

Single-user throughput as a function of location

A7.57 These graphs plot single user throughput (see A7.41) on the y-axis versus percentage of locations on the x-axis (see A7.48).

A7.58 Note that these graphs also illustrate the maximum coverage that it is possible to provide with a given network: the intercept of each line with the x-axis being the maximum proportion of locations to which it is possible to deliver a minimum level of service[14].

Capacity as a function of coverage delivered

A7.59 These graphs plot capacity on the y-axis versus percentage of locations on the x-axis (see A7.48).

---

[14] As described in Annex 14, the minimum level of service is a direct function of the SINR cut-off (i.e. the minimum SINR below which an LTE link is unable to maintain a viable communications link between base station and user terminal

A7.60 Capacity on these graphs has been displayed on a normalised scale that allows the comparison of different spectrum portfolios (and network sizes). This scale represents a ratio of the total resources available to the network to the resources needed to serve all users with the specified guaranteed data-rate service assuming that users within the simulation area are uniformly distributed over all modelled locations (including over all depths at each modelled address) (see A7.44).

A7.61 Note that these graphs also illustrate the coverage provided by a given network for the specified service speed: the intercept of each line with the x-axis being the maximum proportion of locations to which the specified service can be delivered.

Coverage

1Mbps service

A7.62 We start by considering how network coverage is predicted to vary between different networks. We first consider coverage for the provision of a relatively low-speed mobile data service – a nominal 1Mbps downlink data-rate service.

A7.63 The table below shows the data-rate achievable at the SINR cut-off values used in our ‘Min var’ and ‘Max var’ cases (which assumes 85% serving cell loading and 20% overheads).

Table 2: Data-rate achievable at the SINR cut-off

SINR cut-off	5 MHz carrier	10 MHz carrier	15 MHz carrier	20 MHz carrier
-10 dB (Min var)	0.25Mbps	0.50Mbps	0.75Mbps	1.00Mbps
-5 dB (Max var)	0.72Mbps	1.44Mbps	2.16Mbps	2.91Mbps

A7.64 Alternatively the table below shows the minimum SINR needed to achieve a notional1Mbps service.

Table 3: SINR need to achieve a 1Mbps service

SINR cut-off	5 MHz carrier	10 MHz carrier	15 MHz carrier	20 MHz carrier
-10 dB (Min var)	-3.32 dB	-6.76 dB	-8.67 dB	-10.00 dB*
-5 dB (Max var)	-3.32 dB	-5.00 dB*	-5.00 dB*	-5.00 dB*

* Coverage is limited by SINR cut-off rather than the 1Mbps data-rate

A7.65 Table 2 and Table 3 illustrate that coverage is a function of both data-rate and SINR cut-off; and in fact basing coverage on either of these is likely to be a simplification of some rather complex trade-offs as discussed in A7.39 above. Basing coverage on the ability to serve a user with a guaranteed data-rate provides a useful metric for comparison.

A7.66 We first look at how coverage varies with frequency, assuming that all networks have the same spectrum bandwidth available:

Figure 3: Coverage as a function of number of sites for a 1Mbps service, 10 MHz, 85% loading, various frequencies – West London

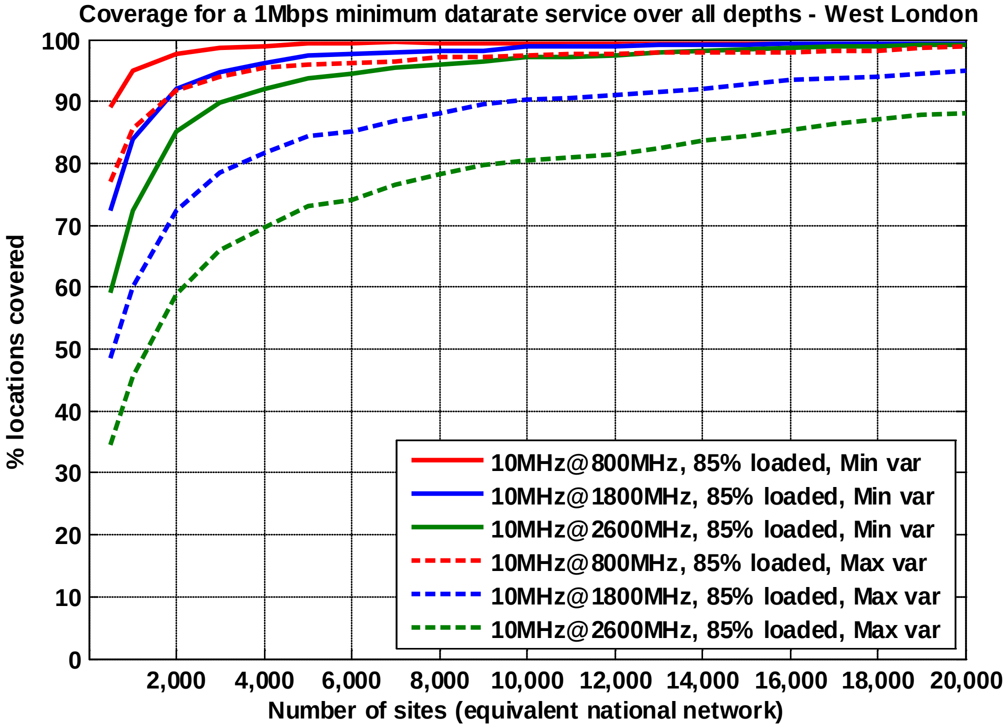

Figure 4: Coverage as a function of number of sites for a 1Mbps service, 10 MHz, 85% loading, various frequencies – Cambridge

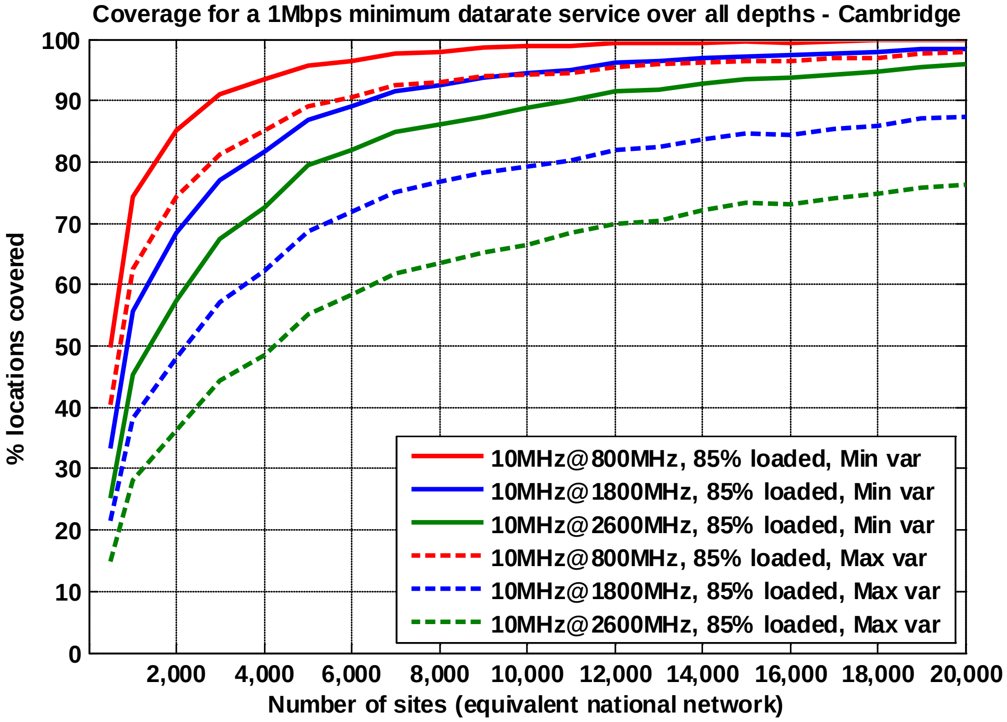

A7.67 We can see from these graphs that for networks with small numbers of sites (low thousands), coverage is predicted by our model to be better (higher proportion of locations covered) for networks using 800 MHz spectrum as compared with

1800 MHz, and with networks using 2600 MHz having poorer coverage still. The uncertainty of our model's prediction of coverage (as shown by the gap between 'Min var' and 'Max var' coverage) is less for networks using 800 MHz spectrum as compared with networks using higher frequency spectrum.

A7.68 For networks with larger numbers of sites, we see that the coverage gap is predicted to close, but the uncertainty of our model's prediction of the coverage of higher frequency networks (1800 MHz and 2600 MHz) is larger than that for lower frequency networks (800 MHz), even for quite large numbers of sites (a clear gap of around 10% of locations between 'Min var' and 'Max var' coverage remains even for 20,000 sites).

A7.69 These results are in line with expectations; results for small number of sites being more 'noise limited' and for all site counts the 'Max var' results also being more 'noise limited' especially for higher frequencies.

A7.70 It is interesting to note that for very large site numbers in the West London simulation case, the 'Min var' results converge and approach almost 100% coverage. This is because, for very large site numbers the distance between sites gets quite small and therefore outdoor propagation losses are also small; and in the 'Min var' case, though there is a difference in building penetration loss standard deviation, median building penetration losses are identical at all frequencies (see Table 1). Consequently, SINR remains relatively high at all frequencies leading to generally good coverage and minimal variation between frequencies.

A7.71 Looking in more detail at how coverage varies between frequencies in our West London simulation area (Figure 5), we see that, for networks with the equivalent of 12,000 sites nationally, there is little difference in the predicted extent of coverage outdoors across the three frequency bands. We also see that the extent of coverage is predicted by our model to be less for users inside buildings, for all frequencies, with the degradation increasing the deeper into the building the user is (e.g. in harder to serve locations)[15]. In the 'Max var' case, the predicted extent of this degradation is materially greater for 2600 MHz than for 800 MHz, with 1800 MHz in between. The degradation is also predicted by our model to be materially greater for the hardest to serve locations inside buildings (e.g. our 10m and 15m results) than it is for easier to serve locations inside buildings (e.g. our 1m and 5m results).

A7.72 Again, this is in line with our expectations, the greater building penetration losses for higher frequencies in the 'Max var' case leading to lower SINR and more 'noise limited' locations the deeper into buildings you get.

---

[15] The noticeable discontinuity in the curve at 1m in Figure 5 and subsequent figures of coverage as a function of 'depth' in-building (especially in the 'Max var') cases is due to the step change in propagation losses going from outdoors to indoors.

Figure 5: Variation of coverage with 'depth' in building for a 1Mbps service, 10 MHz, 85% loading, 12,000 sites, various frequencies – West London

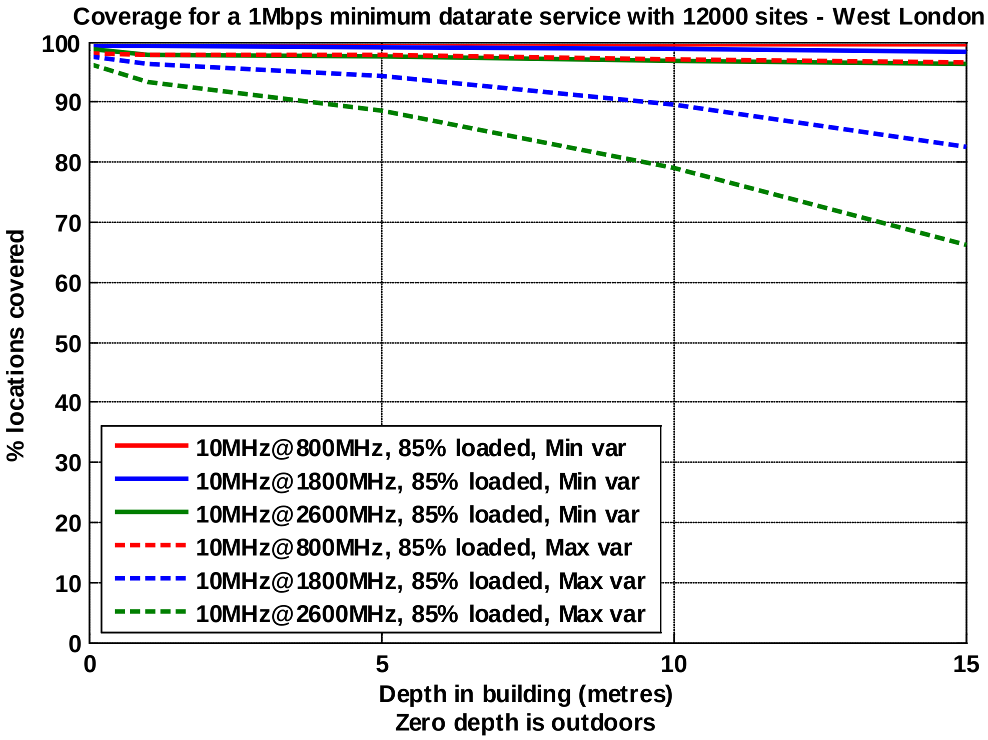

A7.73 In our Cambridge simulation area (Figure 6), we see that, for networks with the equivalent of 12,000 sites nationally, there is a small difference in the predicted extent of coverage outdoors across the three frequency bands and that for users inside buildings, for all frequencies, the degradation in coverage increases faster than for the West London simulation area.

A7.74 This is as expected as there are relatively fewer base stations in the Cambridge simulation area; hence they are geographically more widely separated than in the West London simulation area. The greater distances between base stations means that signals are generally weaker and hence performance is poorer for all frequencies but especially so for higher frequencies.

Figure 6: Variation of coverage with 'depth' in building for a 1Mbps service, 10 MHz, 85% loading, 12,000 sites, various frequencies - Cambridge

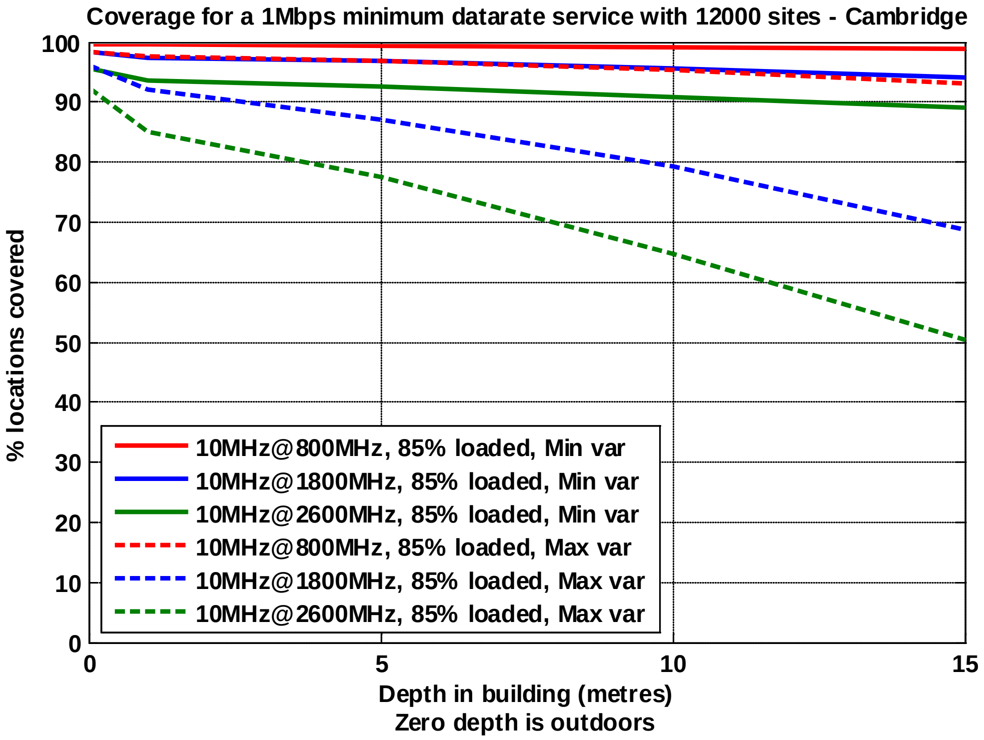

A7.75 It should be stressed that the above results are not a prediction of the nationwide coverage achievable with 12,000 sites. Rather they indicate the coverage achievable with a network with the equivalent of 12,000 sites nationally in the two particular simulation areas.

A7.76 Looking at the same result for a larger number of sites (Figure 7 and Figure 8) we see that the difference in coverage predicted by our model between outdoor and the hardest to serve indoor locations (e.g. our 15m results) has reduced for both the West London and Cambridge simulation areas but for our 'Max var' case there remains a noticeable degradation in the 1800 MHz and 2600 MHz results for those locations deepest inside buildings (e.g. our 10m and 15m results[16]). And that for the Cambridge simulation area there is a noticeable degradation for our 'Max var' case even at the easier to serve locations in buildings (e.g. our 1m results) for the higher frequencies.

---

[16] Though actual physical depths of 10m to 15m may be relatively rare particularly in residential areas, as discussed in A7.23 above, depth as used in our model should be interpreted less literally

Figure 7: Variation of coverage with 'depth' in building for a 1Mbps service, 10 MHz, 85% loading, 18,000 sites, and various frequencies – West London

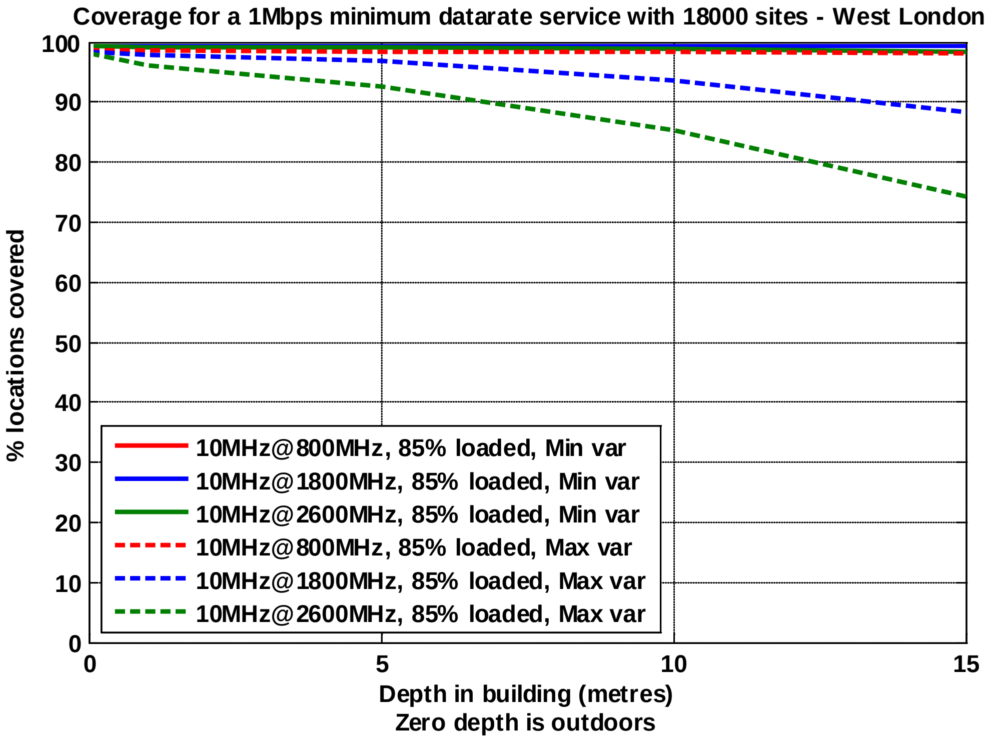

Figure 8: Variation of coverage with 'depth' in building for a 1Mbps service, 10 MHz, 85% loading, 18,000 sites, and various frequencies - Cambridge

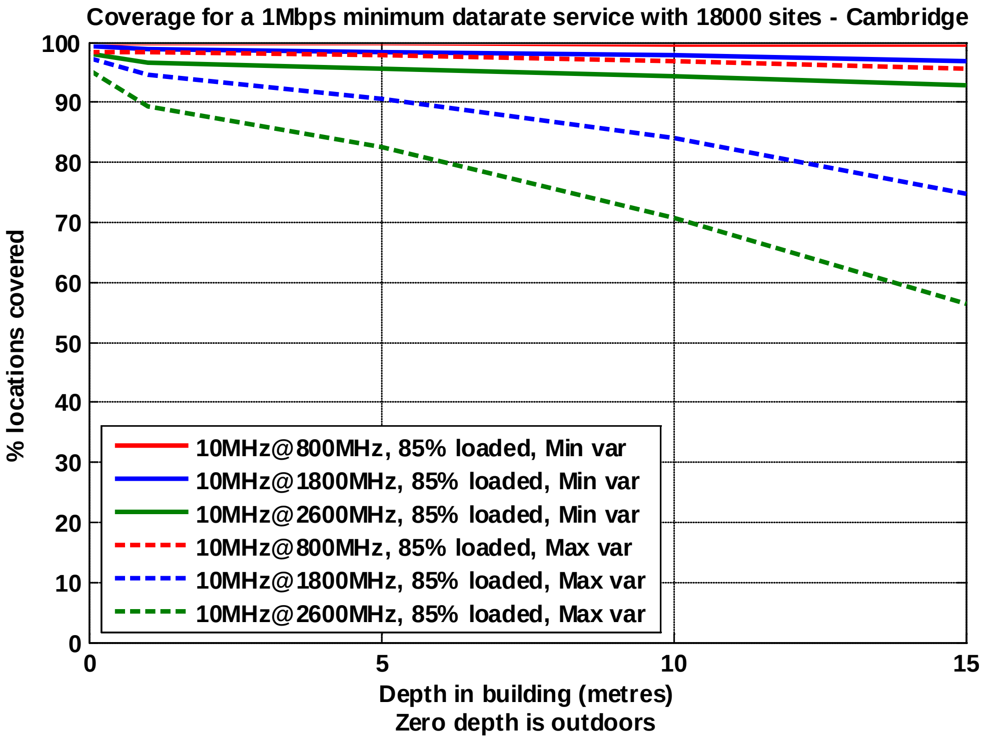

A7.77 The above results are all for equal quantities of spectrum. We next look at how coverage varies with the amount of spectrum used (Figure 9 to Figure 14).

Figure 9: Coverage as a function of number of sites for a 1Mbps service, 800 MHz, 85% loading, various bandwidths – West London

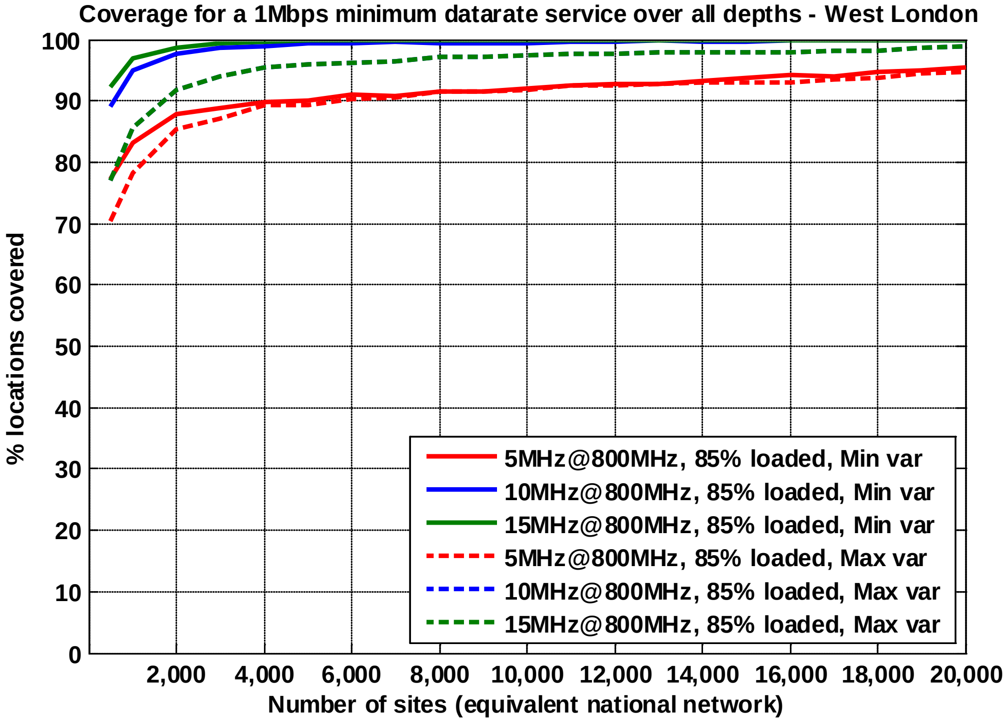

Figure 10: Coverage as a function of number of sites for a 1Mbps service, 800 MHz, 85% loading, various bandwidths – Cambridge

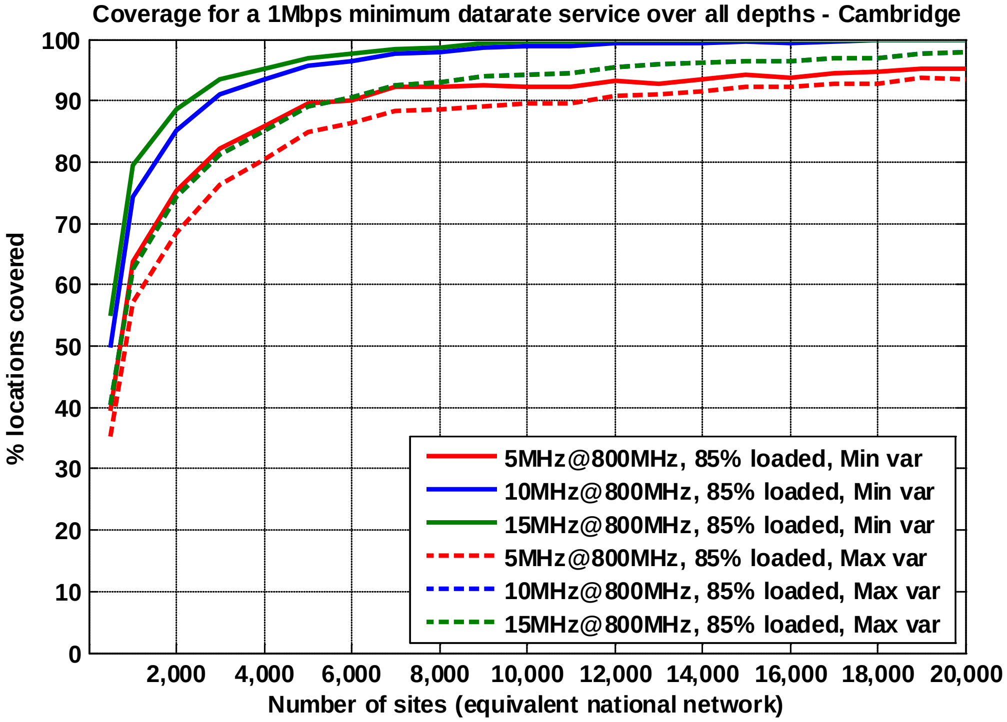

Figure 11: Coverage as a function of number of sites for a 1Mbps service, 1800 MHz, 85% loading, various bandwidths – West London

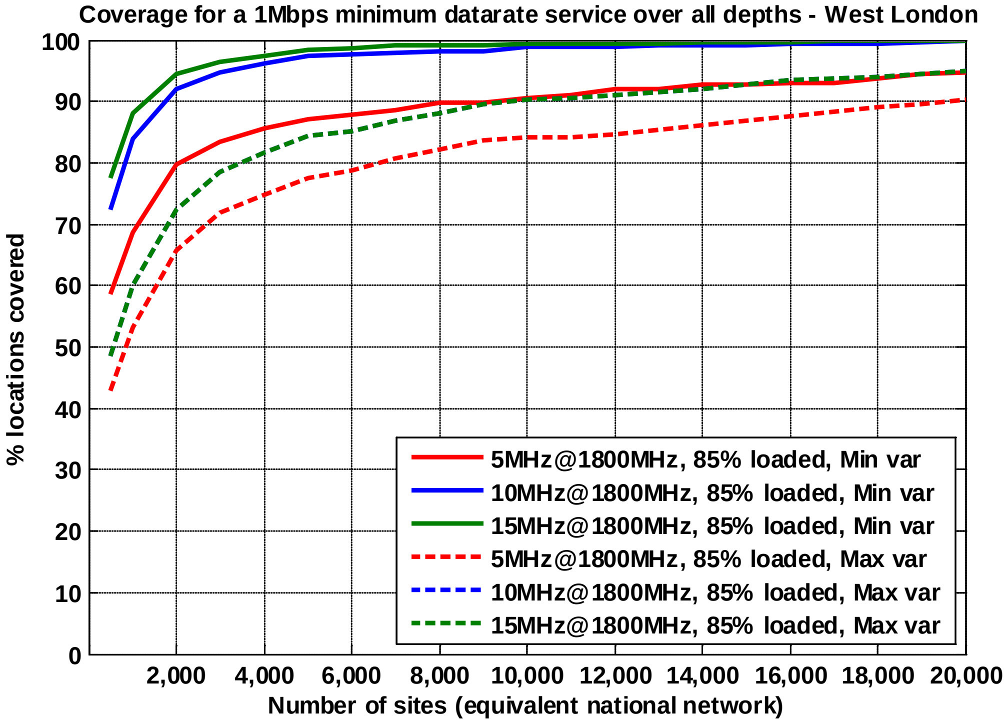

Figure 12: Coverage as a function of number of sites for a 1Mbps service, 1800 MHz, 85% loading, various bandwidths – Cambridge

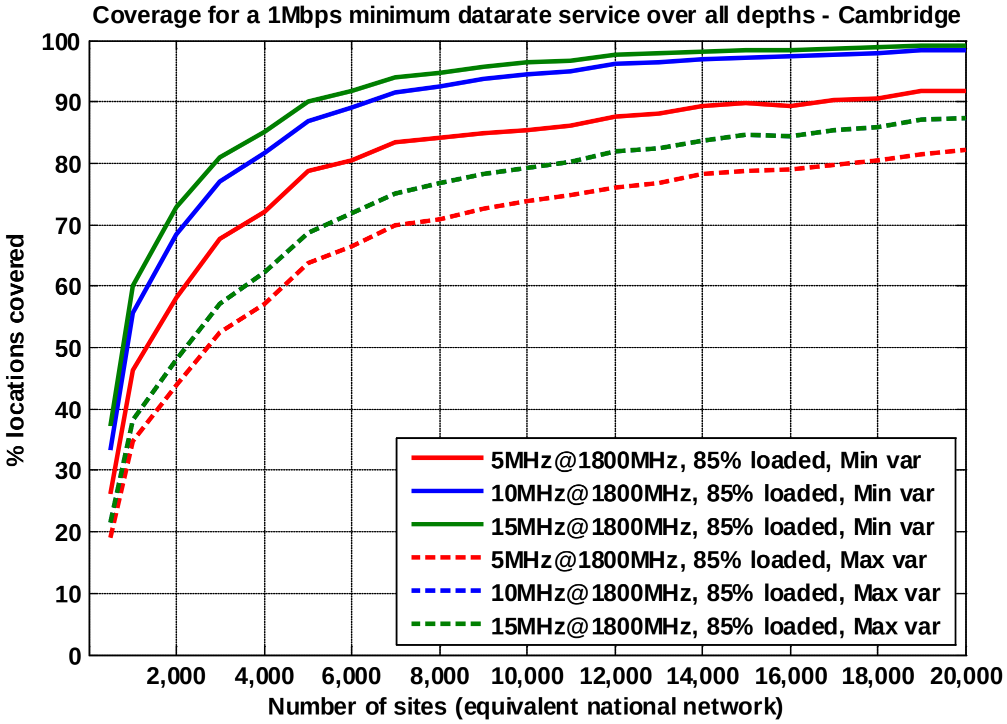

Figure 13: Coverage as a function of number of sites for a 1Mbps service, 2600 MHz, 85% loading, various bandwidths - West London

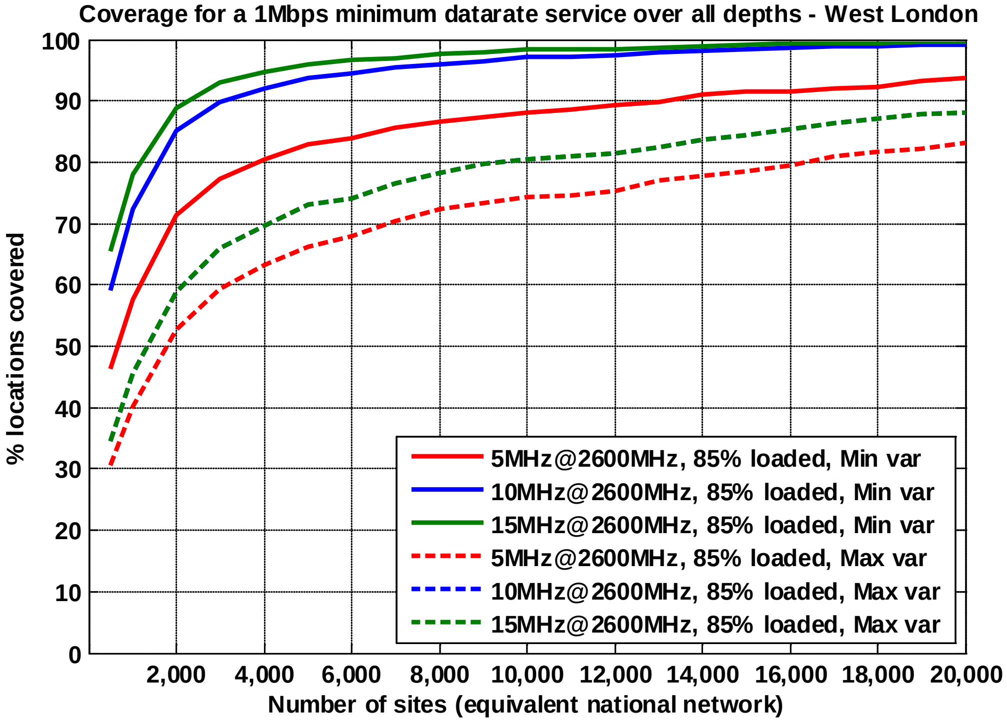

Figure 14: Coverage as a function of number of sites for a 1Mbps service, 2600 MHz, 85% loading, various bandwidths - Cambridge

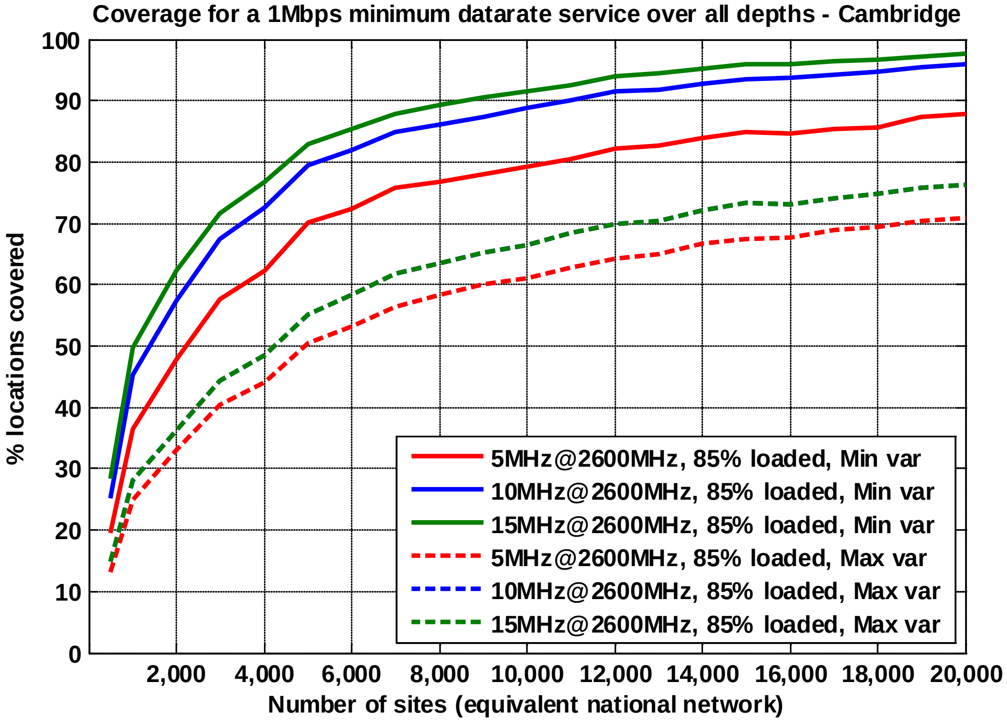

A7.78 We see from these results that, for a notional 1Mbps downlink data-rate service, for all frequencies, coverage is predicted by our model to be noticeably worse with only 5 MHz of spectrum, and the difference between frequencies greater, as compared with 10 MHz Coverage is not however predicted to be materially better with 15MHz as compared with 10 MHz of spectrum, at least in this case where we only require a notional 1Mbps service to be provided[17].

A7.79 To reinforce our conclusion about coverage with only 5 MHz being noticeably worse than 10 MHz, the difference in coverage at all frequencies is clear when comparing Figure 15 and Figure 16 below with Figure 3 and Figure 4 above.

Figure 15: Coverage as a function of number of sites for a 1Mbps service, 5 MHz, 85% loading, various frequencies – West London

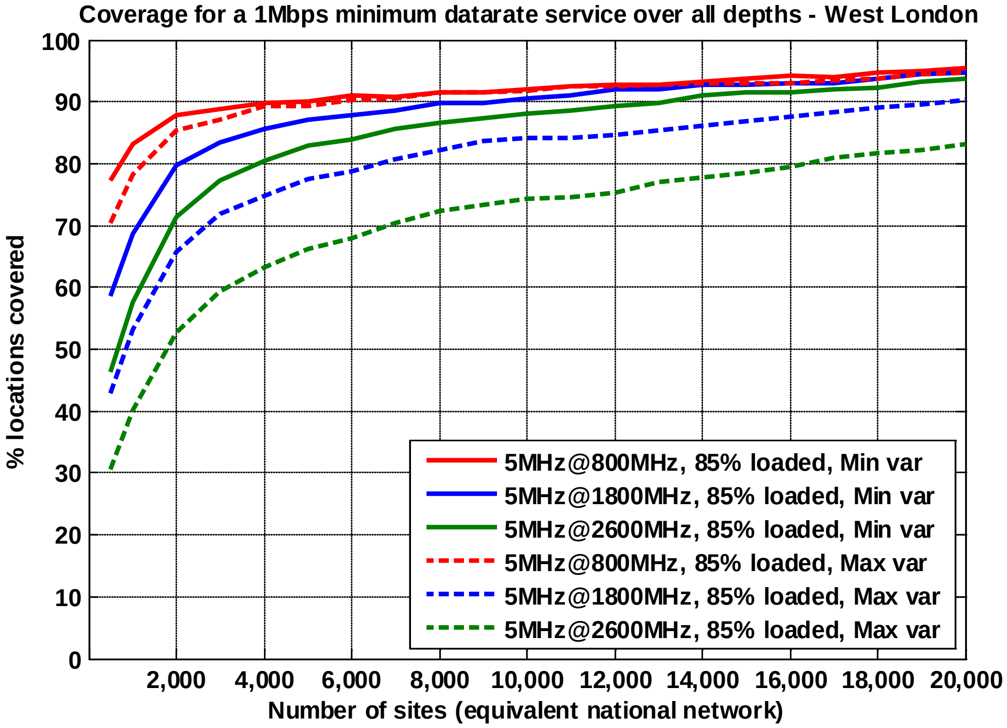

---

[17] The ‘Max var’ coverages for 10 MHz and 15 MHz of spectrum are actually predicted by our model to be identical which is why only the 15 MHz line can be seen on the graphs – the 10 MHz line is hidden underneath. This is because coverage is actually limited by the SINR cut-off rather than the 1Mbps data-rate for carrier bandwidths 10 MHz and greater – see Table 3 above.

Figure 16: Coverage as a function of number of sites for a 1Mbps service, 5 MHz, 85% loading, various frequencies – Cambridge

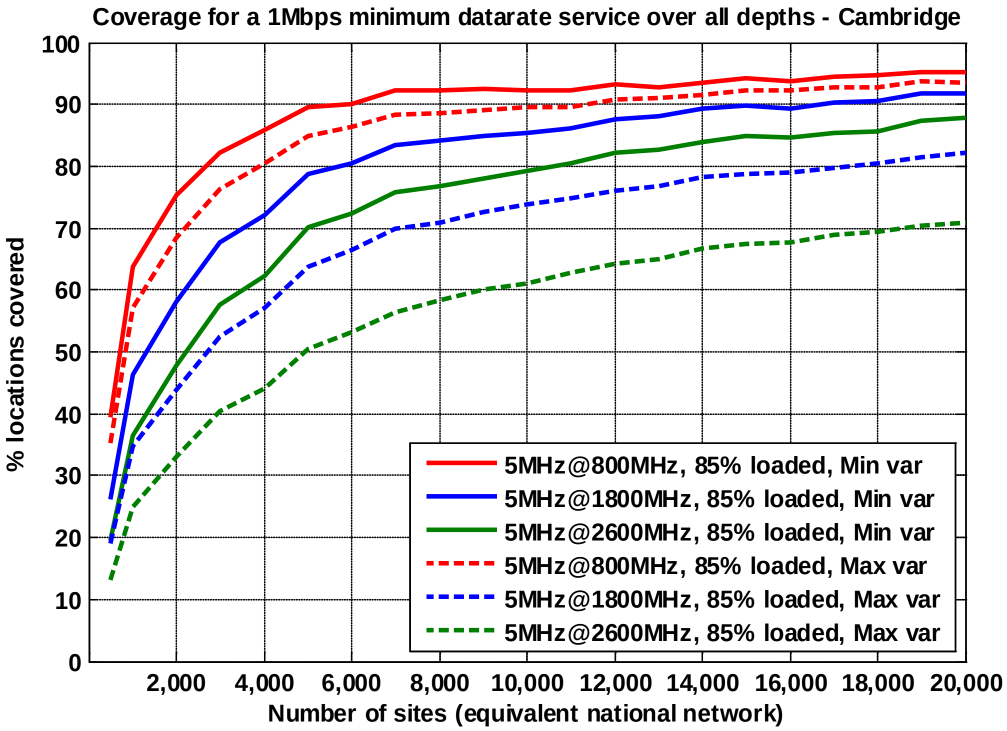

A7.80 All of the above results are for fully loaded networks (85% loading). As illustrated by the next graphs (Figure 12 and Figure 13), a network may be able to provide coverage to more locations if it is less loaded, although this will also reduce the capacity of the network.

Figure 17: Coverage as a function of number of sites for a 1Mbps service, 800 MHz, various loadings, 5 MHz – West London

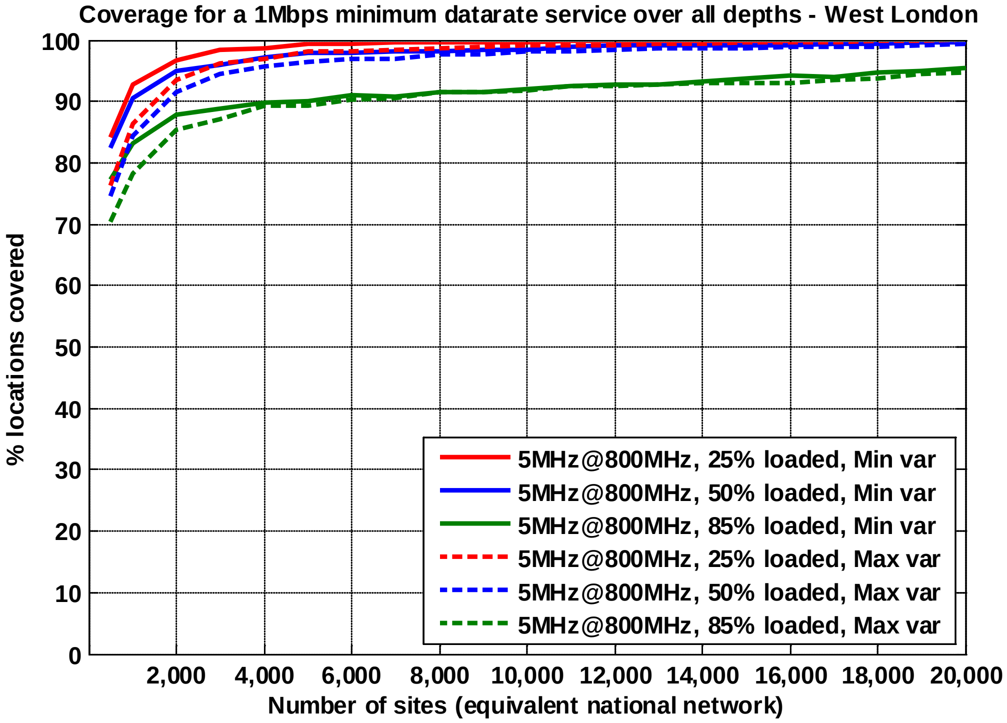

Figure 18: Coverage as a function of number of sites for a 1Mbps service, 800 MHz, various loadings, 5 MHz – Cambridge

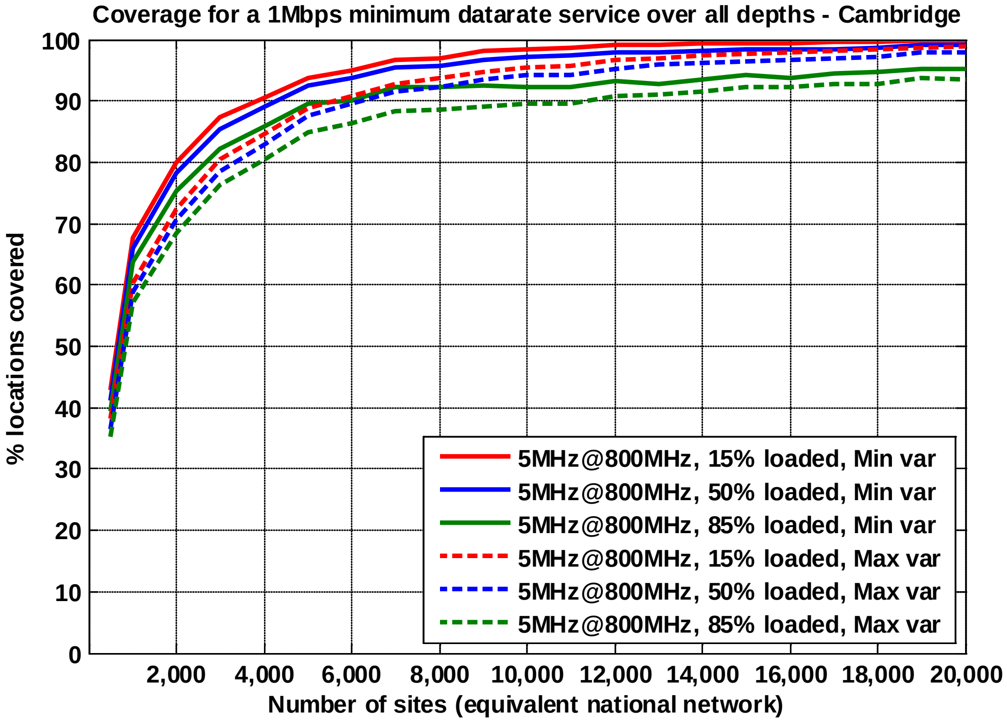

A7.81 It turns out, however, that the coverage benefits of lighter loading are predicted to be fairly limited for the simulation areas considered when the service of interest is only around 1Mbps[18]. In this case the benefits of lighter loading are only really material when the bandwidth available is 5 MHz; if the bandwidth available is 10 MHz or more then the coverage benefits are almost negligible (e.g. comparing Figure 19 with Figure 3 and Figure 20 with Figure 4). We will see later, however, that lighter loading can have a more material impact on the ability to deliver higher speed services.

Figure 19: Coverage as a function of number of sites for a 1Mbps service, 10 MHz, 25% loading, various frequencies – West London

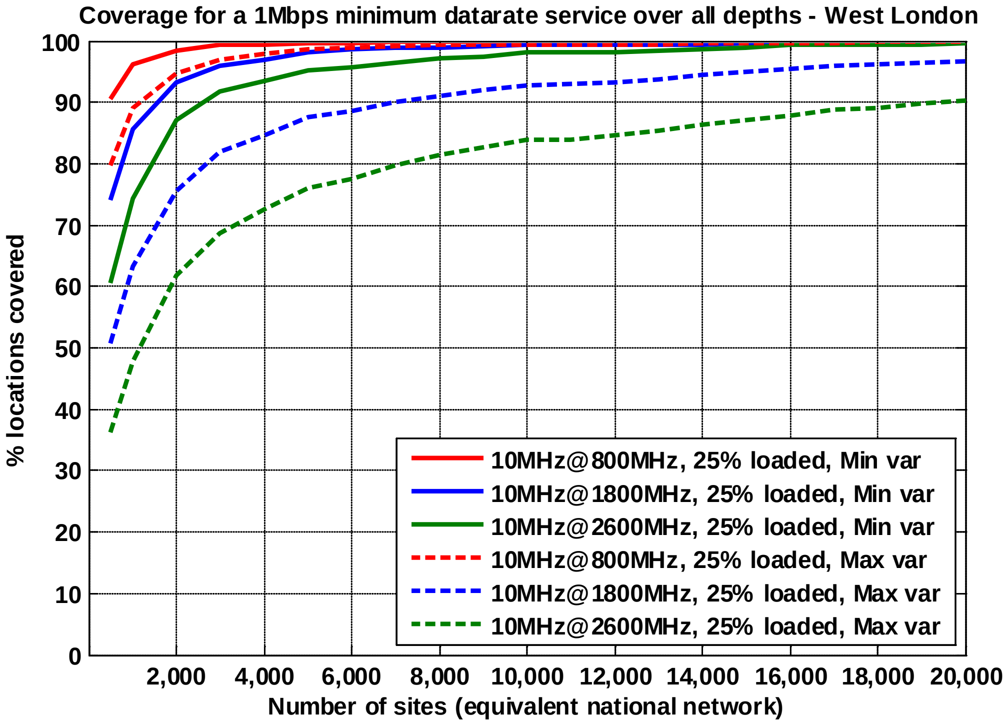

---

[18] It should, however, be remembered (as discussed in paragraph A7.37) that a guaranteed data-rate of 1Mbps in our modelling is a proxy for a relatively low speed service and should not be seen as a definitive prediction.

Figure 20: Coverage as a function of number of sites for a 1Mbps service, 10 MHz, 25% loading, various frequencies – Cambridge

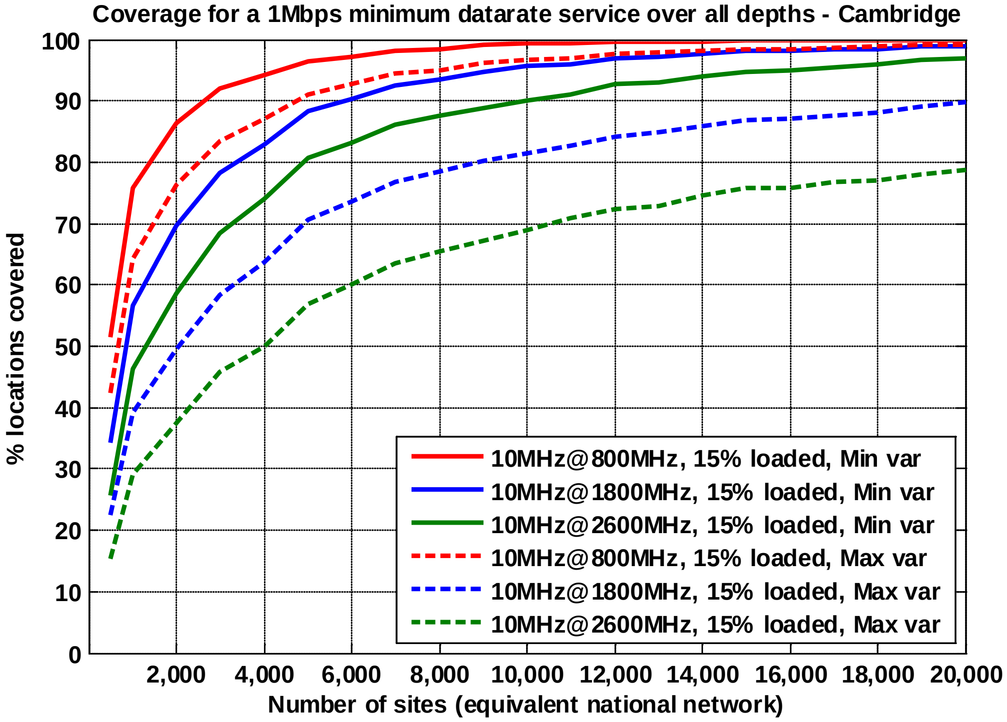

A7.82 In summary therefore, it would appear that a network using 800 MHz spectrum may have two advantages as compared with networks without 800 MHz spectrum:

- The ability to deliver good coverage using a smaller number of sites initially (albeit such a network will have correspondingly less capacity and consequently may need to be expanded to meet growth in demand, and ultimately may need to deploy a similar if not larger number of sites – see the discussion of capacity below);
- Potentially better coverage even in the longer term, particularly in harder to serve locations such as deep inside buildings, in particular if reality turns out to be more like our 'Max var' rather than our 'Min var' case.

A7.83 However, equally importantly, the gap in coverage between networks with and without 800 MHz spectrum, in particular between networks with 800 MHz spectrum and networks with 1800 MHz spectrum, does not appear to be that large for networks with large numbers of sites – for example 18,000 sites. Figure 7 for example indicates that for our 'Max var' case users of such a network with 10 MHz of 1800 MHz spectrum are predicted by our model to be able to receive a notional 1Mbps service at more than 95% of locations outdoors and at easier to serve locations inside buildings (e.g. our results up to 5m), and more than 85% of hardest to serve locations in buildings (e.g. our 15m results). For the 'Min var' case coverage for a notional 1Mbps service is predicted by our model to be greater than 95% of locations for both the easiest and hardest to serve locations inside buildings (e.g. up to 15m).

A7.84 Finally, the coverage of a fully loaded network (85% loaded) with only 5 MHz of spectrum is predicted by our model to be less than the same network with 10 MHz of spectrum, irrespective of frequency or number of sites. This loss in coverage can be mitigated to some extent through operation of the network at lighter loading, but at the cost of reduced capacity (and the capacity of 5 MHz of spectrum is already less than half that of a network with 10 MHz).

A7.85 These results are important because, as noted above, they provide an indication of the likely limit of good quality mobile data coverage by networks. In subsequent graphs we will see how coverage declines as the target speed of service increases. Nevertheless, when interpreting these later results, it is important to remember that, just because a network is unable to provide a higher-speed mobile data service to a particular location, doesn't mean that that network cannot provide a mobile data service at all; the network may well be able still to provide a good quality mobile data service to that location just at a lower speed until the SINR cut-off is reached.

Higher speed services

A7.86 We now turn our attention to the ability of networks with different spectrum portfolios to deliver higher speed services to various locations.

A7.87 The coverage that a network can provide at different speeds can be assessed by looking at the speed that the network can deliver to a single user at various locations, assuming that there are no other users in the same cell[19]demanding service at the same time (single-user throughput as a function of location).

Figure 21: Single-user throughput as a function of location for various quantities of 800 MHz spectrum – 18,000 sites

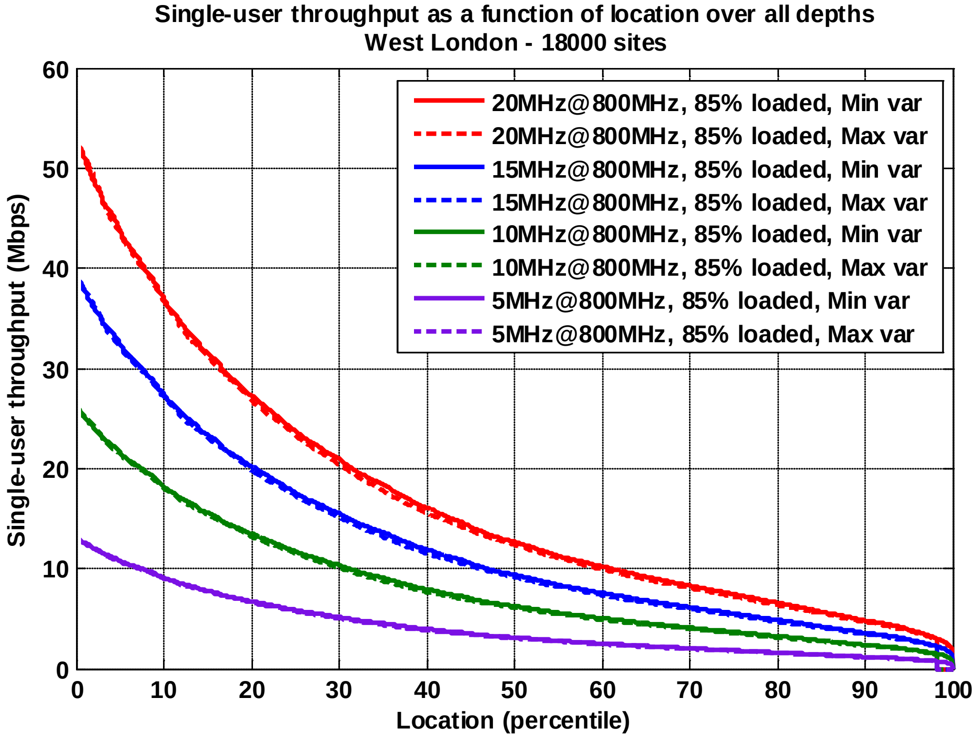

---

[19] The same sector.

Figure 22: Single-user throughput as a function of location for various quantities of 800 MHz spectrum – 12,000 sites

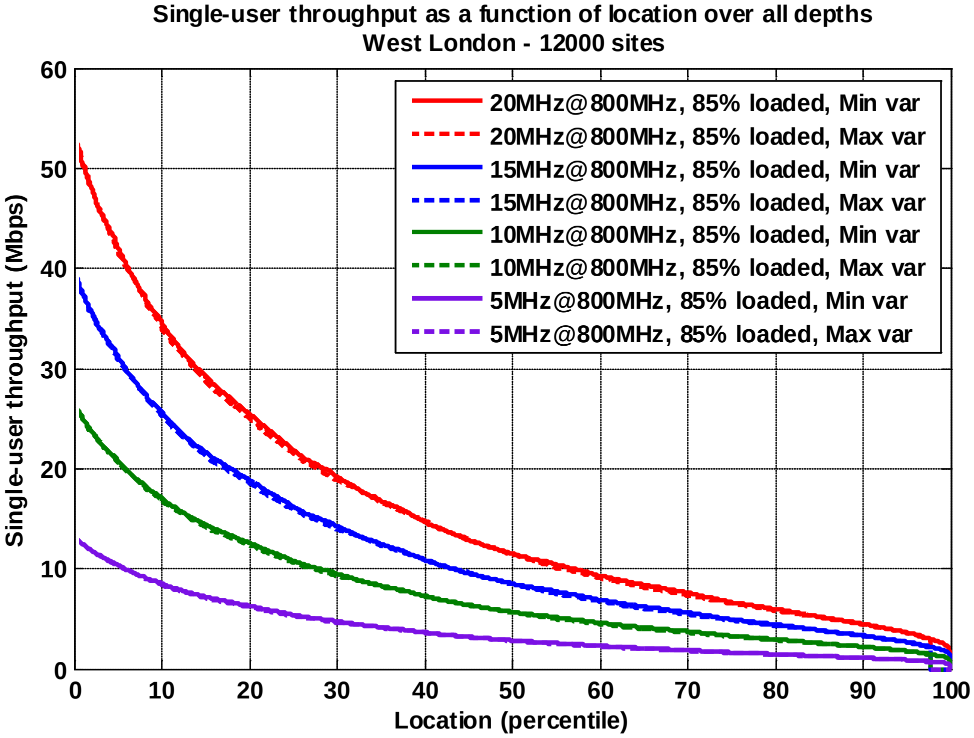

Figure 23: Single-user throughput as a function of location for various quantities of 800 MHz spectrum – focus on speeds up to 10Mbps – 18,000 sites

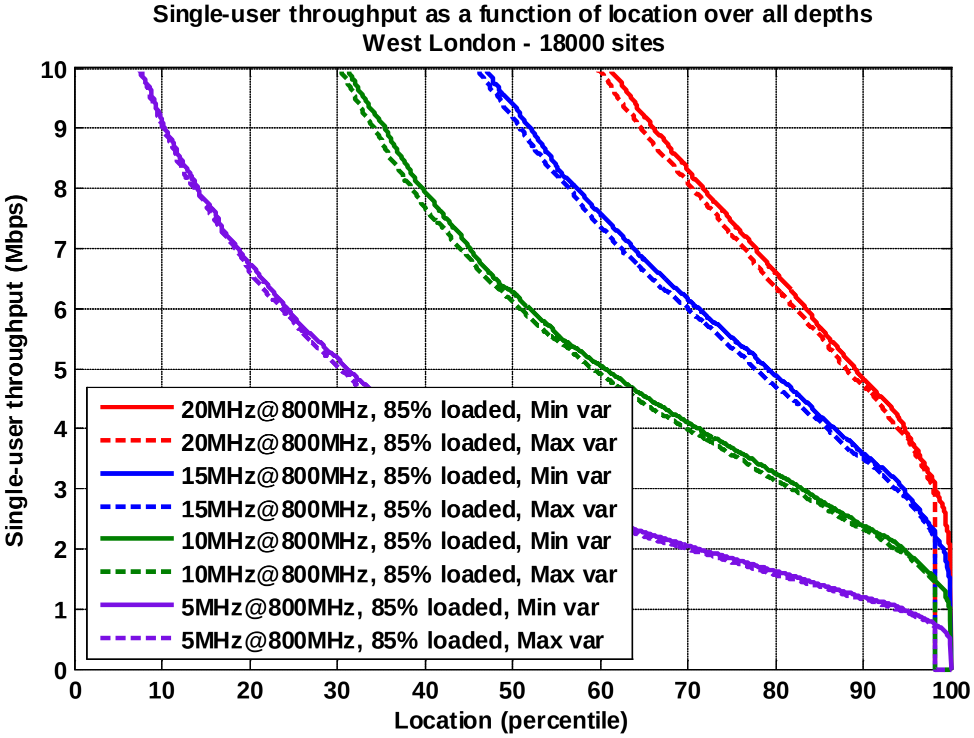

Figure 24: Single-user throughput as a function of location for various quantities of 800 MHz spectrum – focus on speeds up to 10Mbps – 12,000 sites

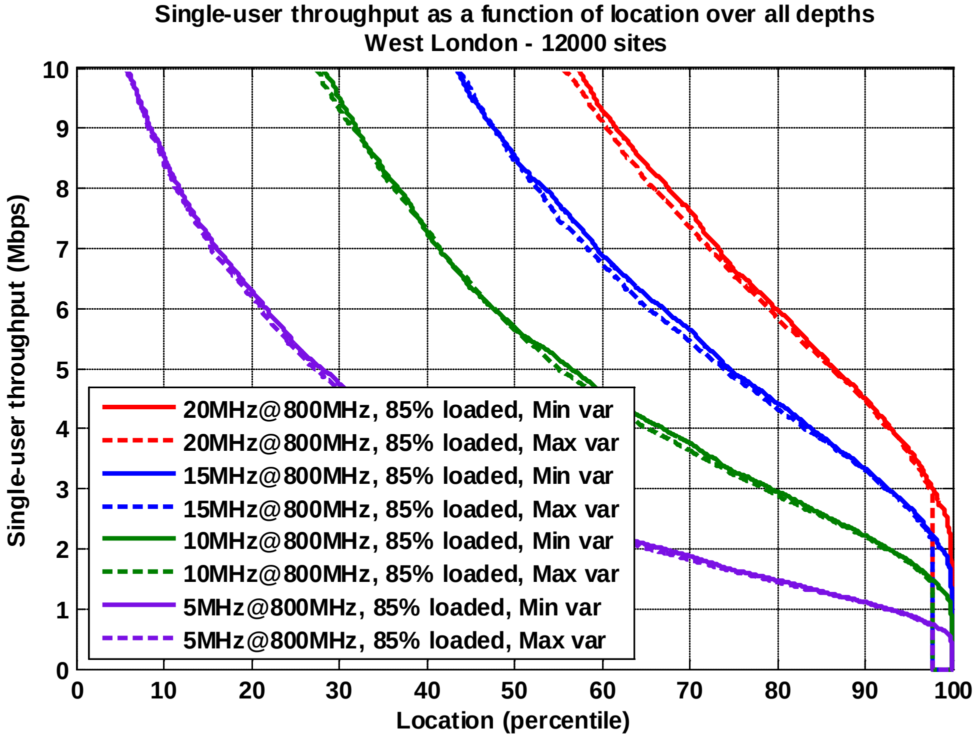

A7.88 From this we can see immediately that the proportion of locations to which a network with only 5 MHz of 800 MHz spectrum is predicted by our model to be able to deliver a particular speed of service declines rapidly as the required speed increases: from approximately 95% of locations for a notional speed of 1Mbps, to only 70% of locations for a notional speed of 2Mbps, and to approximately 30% of locations for a notional speed of 5Mbps. These results are relatively insensitive to site count with the network with the equivalent of 12,000 sites nationally performing only very slightly worse than the network with the equivalent of 18,000 sites nationally.

A7.89 By contrast, a network with 10 MHz of 800 MHz spectrum is predicted by our model to be able to deliver a notional speed of 2Mbps to approximately 95% of locations, and to approximately 60% of locations for a notional speed of 5Mbps.

A7.90 Please bear in mind that the specific data-rates reported here (e.g. 1Mbps, 2Mbps, 5Mbps) should not be seen as a definitive prediction but an indication of relative performance.

A7.91 So far as variation with frequency is concerned, we can see from Figure 25 and Figure 26 that there is greater uncertainly of our model's prediction of coverage for any given speed with higher frequencies (greater difference between 'Min var' and 'Max' var' coverage) but that the variation is not particularly large.

Figure 25: Single-user throughput as a function of location for various frequencies, 85% loading, 10 MHz bandwidth – 18,000 sites

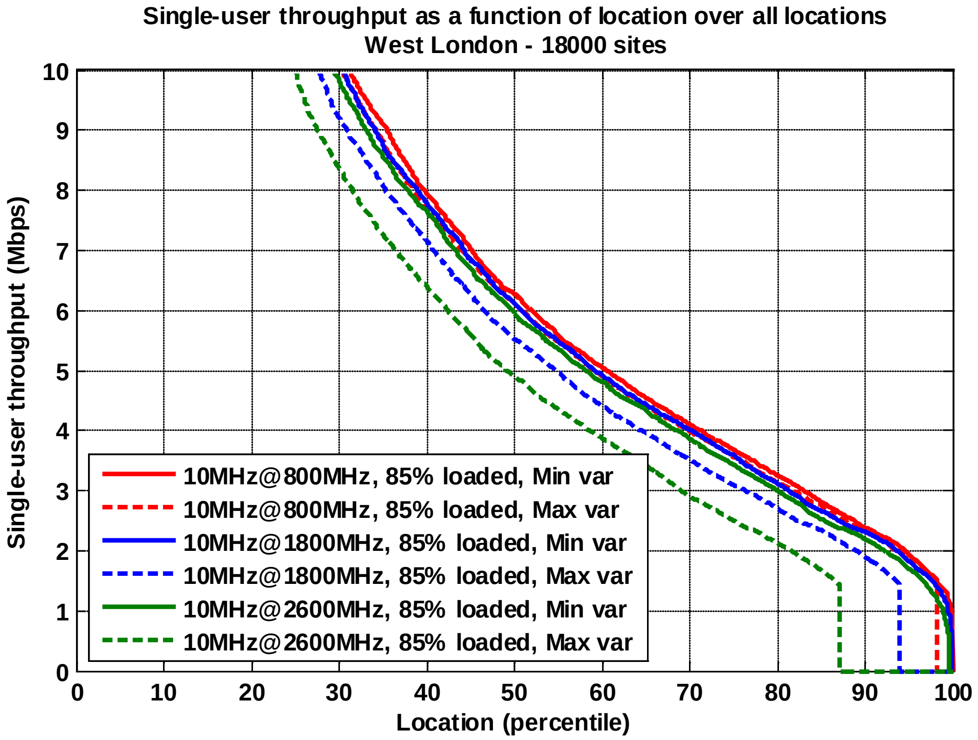

Figure 26: Single-user throughput as a function of location for various frequencies, 85% loading, 10 MHz bandwidth – 12,000 sites

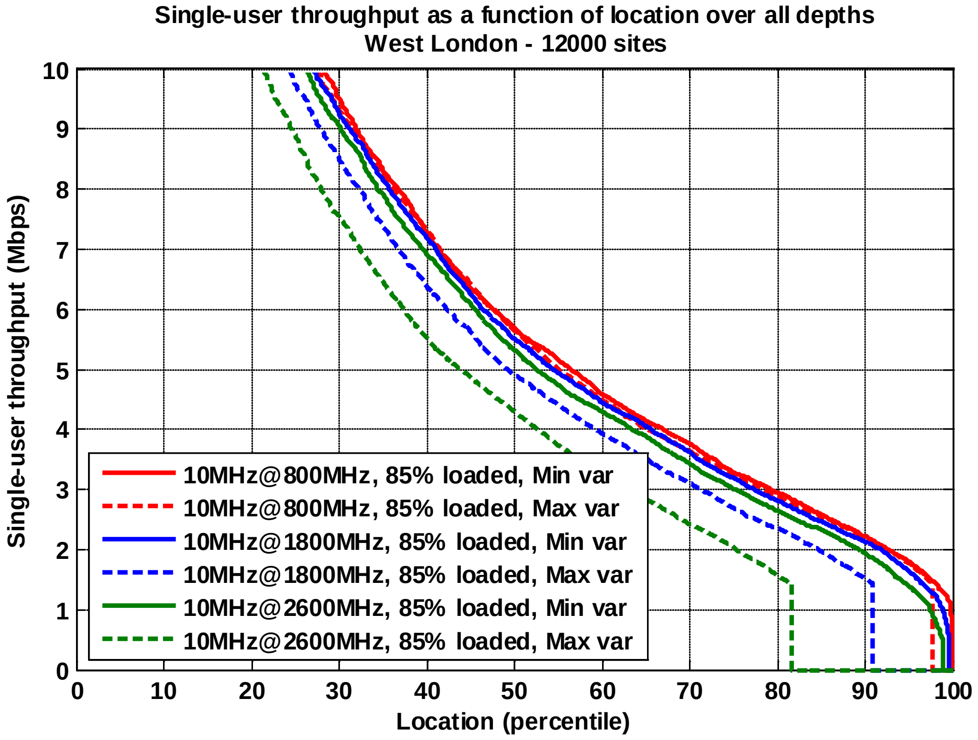

A7.92 If we look at the proportion of locations at each depth to which our model predicts a notional 5Mbps service can be delivered using 10 MHz of spectrum at 800 MHz, 1800 MHz and 2600 MHz, using 18,000 sites (Figure 27 and Figure 28), we see a level of coverage below around 60% as predicted by our model for all three frequencies at any depth. This level of coverage is very unlikely to be seen as acceptable to a mobile operator. However, in reality it is more likely that service would be offered at lower data-rates rather than abruptly cutting-off as implied by the graph.

A7.93 As far as the uncertainly of our model's prediction of coverage (the difference between 'Min var' and 'Max var' coverage) is concerned, Figure 27 and Figure 28 follow the same pattern as Figure 7 and Figure 8 but with a lower level of overall coverage. This is because greater data-rates require greater SINR values hence can only be achieved over a smaller proportion of locations – outdoor propagation losses dominate (i.e. the network is 'interference limited') and hence overall there is a lower level of coverage than can be achieved than for a notional 1Mbps service; however, building penetration losses affect the results in almost the same way for a notional 5Mbps or a notional 1Mbps services and it is these losses that dictate the relative differences between the frequencies in the 'Min var' and 'Max var' cases.

Figure 27: Variation of coverage with depth in building for a 5Mbps service, 10 MHz, 85% loading, 18,000 sites, various frequencies

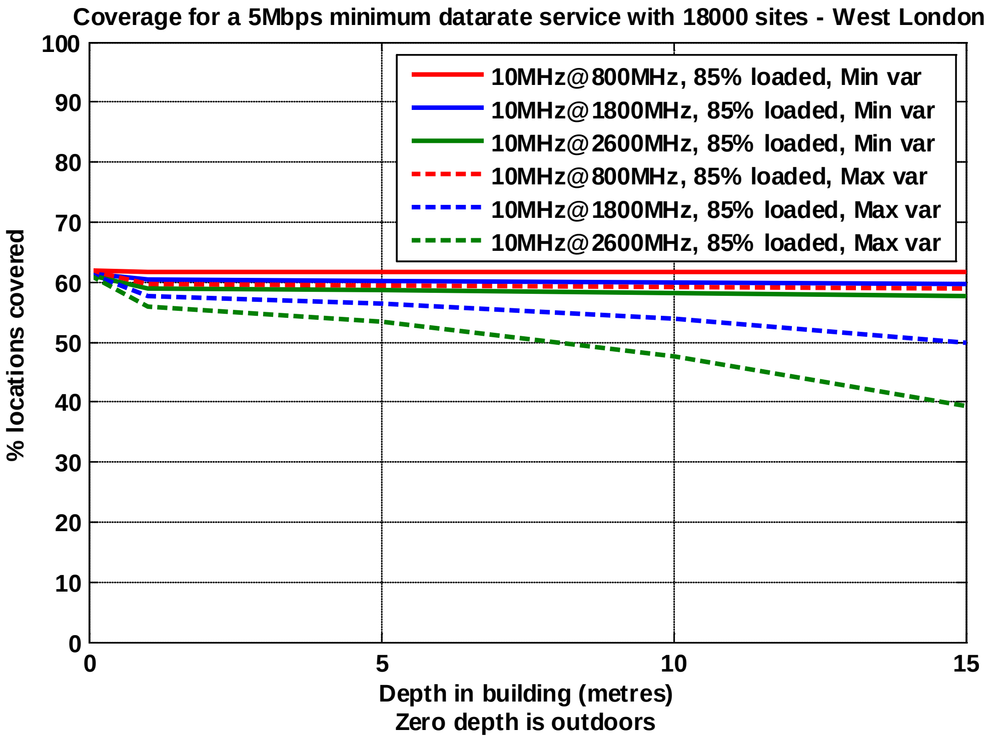

Figure 28: Variation of coverage with depth in building for a 5Mbps service, 10 MHz, 85% loading, 12,000 sites, various frequencies

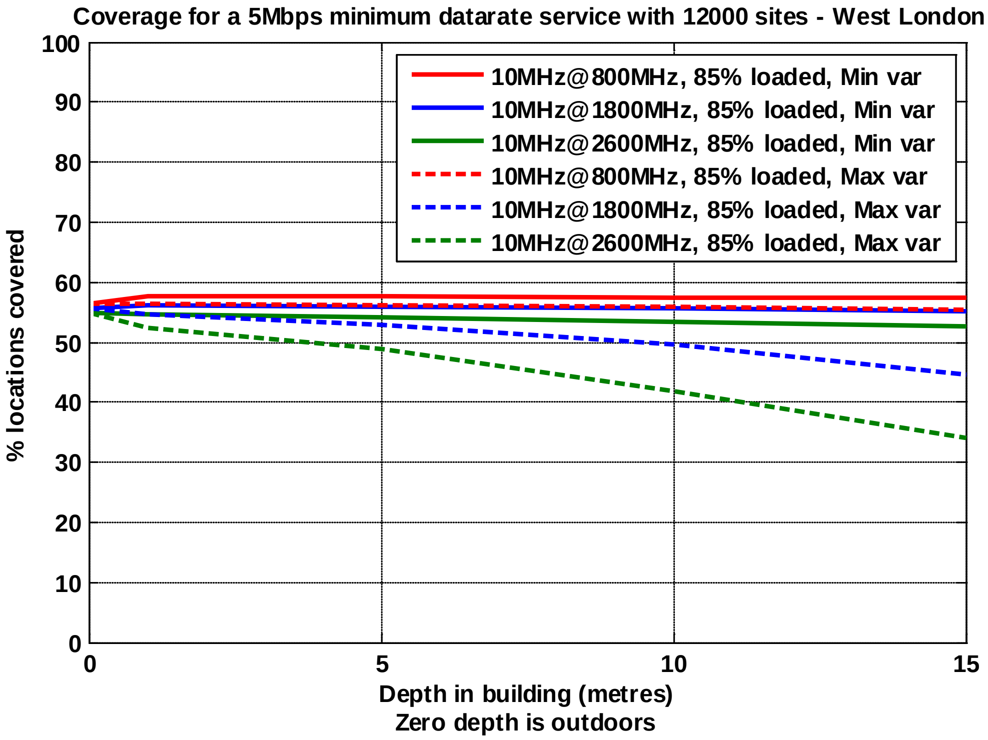

A7.94 All of the above results are, however, are dependent on the assumptions we make as to how heavily loaded each network is (what proportion of the available capacity is used, and hence how much interference is generated by neighbouring base stations). All of the results above assume that the network is close to fully loaded (85% of resources used). If we look at results for less loaded networks then we see an increase in the speed that can be delivered at each location, and consequently an increase in the proportion of locations to which any given speed can be delivered (albeit at the cost of reduced capacity).

Figure 29: Variation of single user throughput with throughput in building, 5 MHz, 85% loading, 18,000 sites, various loading levels

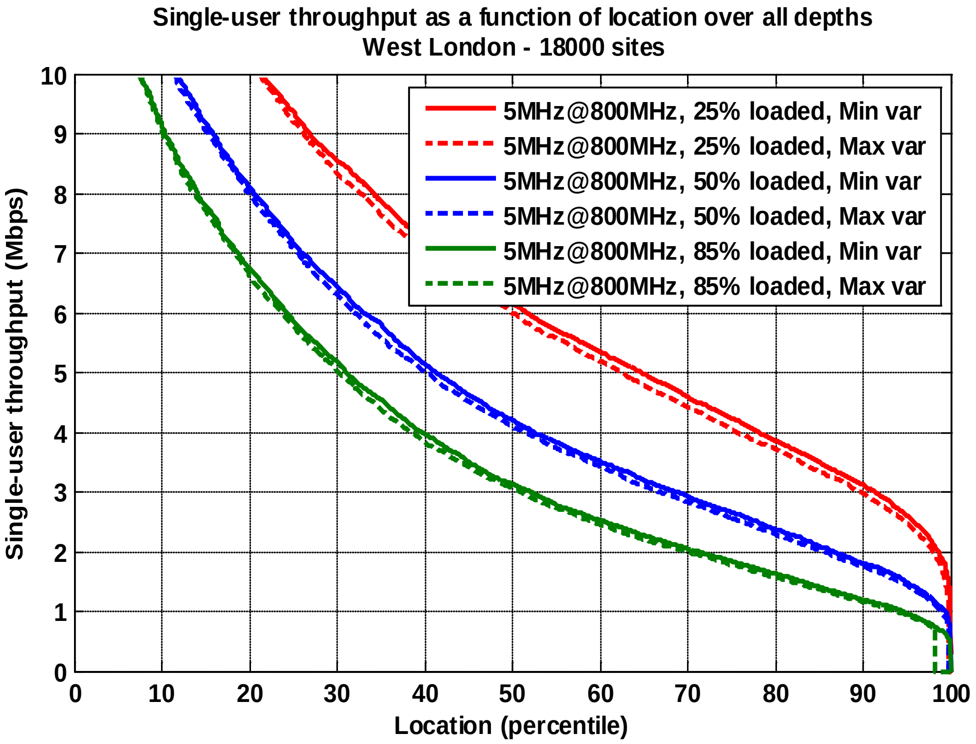

Figure 30: Variation of single user throughput with throughput in building, 5 MHz, 85% loading, 12,000 sites, various loading levels

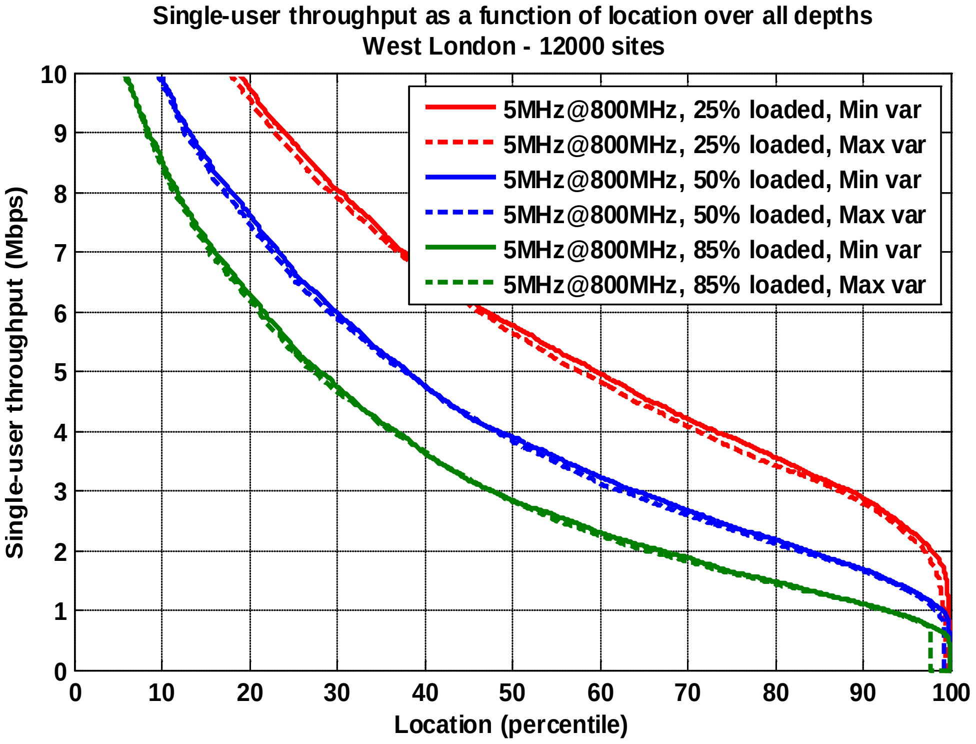

A7.95 Any proper consideration of the capability of networks to provide higher speed services to particular locations therefore needs to have regard to the overall capacity of those networks (relative to the comparator networks of interest) and hence the (relative) loading under which those networks will likely have to operate. This we have done in our multi-frequency analysis, the results of which are reported below.

A7.96 Even so, it can be seen from Figure 29 and Figure 30 that, even when lightly loaded (25% loading[20]), a network with only 5 MHz of 800 MHz spectrum (or in fact any frequency) is predicted by our model to be likely to struggle to provide a notional 5Mbps service to more than 60% of locations, albeit it may be able to provide a notional 2Mbps service to more than 95% of locations at the same depth[21]. By contrast (looking at Figure 31 and Figure 32) networks with 10 MHz of spectrum (again of any frequency) are predicted by our model to be able to provide a notional 5Mbps service to between 65% and 80% of locations when moderately loaded (50% loading)[22].

Figure 31: Single-user throughput as a function of location for various frequencies, 50% loading, 10 MHz bandwidth – 18,000 sites

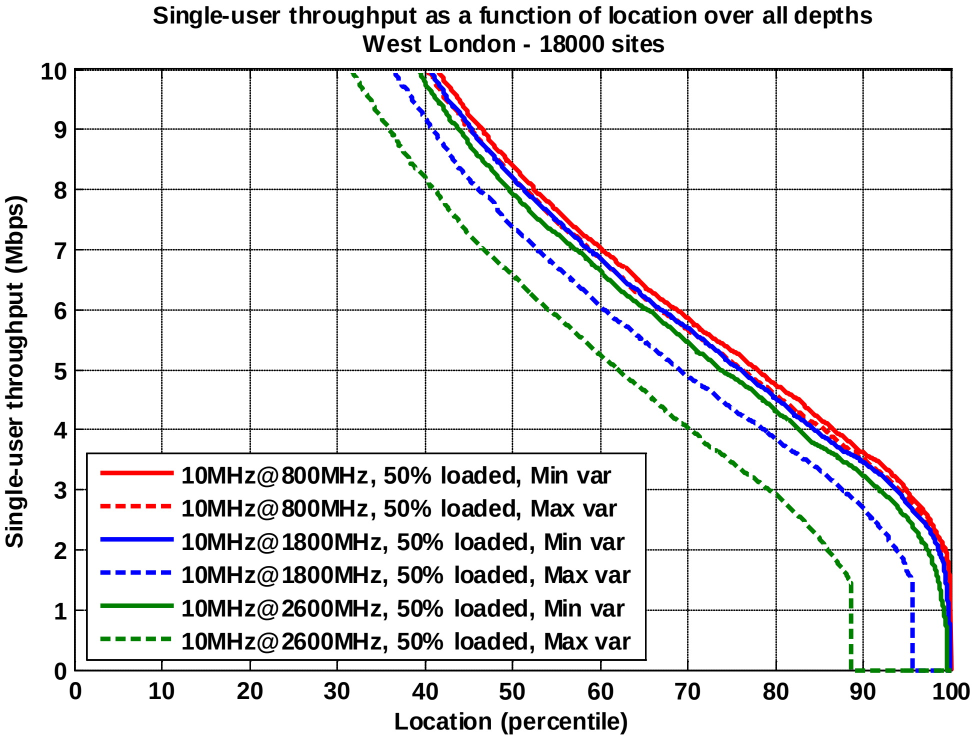

---

[20] Maintaining a loading as low as 25% is likely to be seen by most operators as undesirable.

[21] Coverage will be better at shallower depths, but worse at deeper depths.

[22] It is perhaps worth reiterating (as indicated in paragraph A7.37 above) that the actual data-rates predicted by our modelling should not be seen as a definitive predictions but rather as indications of likely relative performance.

Figure 32: Single-user throughput as a function of location for various frequencies, 50% loading, 10 MHz bandwidth – 12,000 sites

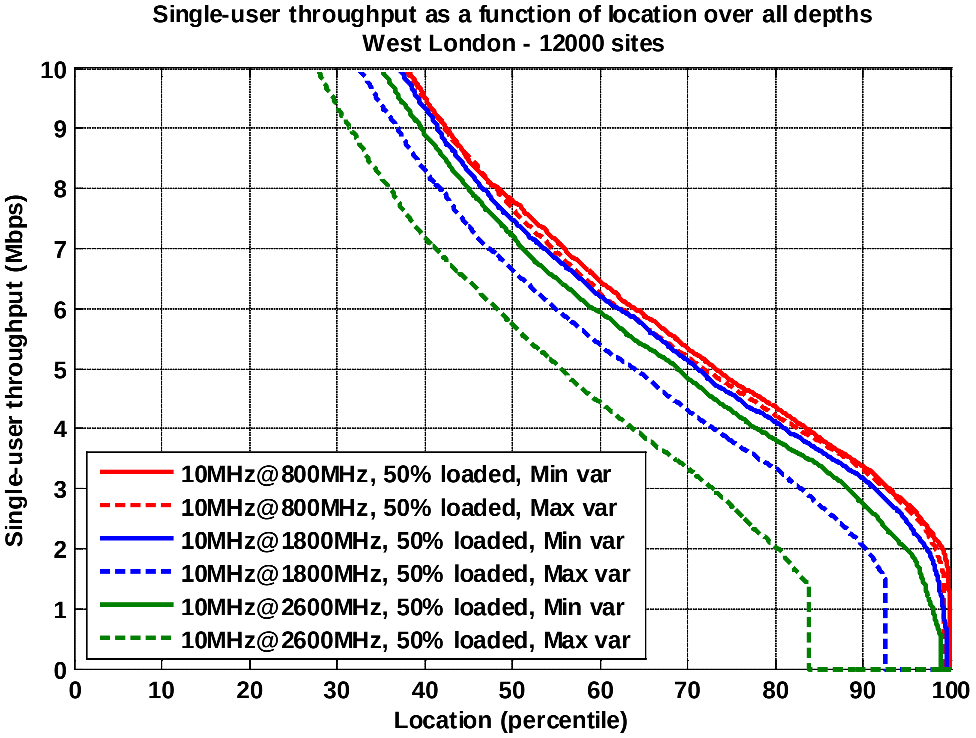

Capacity

A7.97 As we have just noted, our modelling indicates that the set of locations to which a particular speed of service can be delivered will depend to a material extent upon the level to which the network is loaded, at least as regards speeds in excess of about 1Mbps. Conversely, this implies that the maximum loading at which a network can operate, and hence its overall capacity to deliver services to users, will depend to a material extent on the coverage that the network is aiming to provide. Hence, if we are to compare the capacity of different networks consistently, we need to ensure that those networks are operating at a loading which still leaves them capable of delivering a specified level of coverage for some specified speed of service. Our refined capacity analysis does this. However, it is worth noting that macro cellular LTE may not be a cost effective alternative to small cell solutions such as Wi-Fi hot spots or femtocells. It will be an important option where Wi-Fi is not available (such as indoor public spaces), and in areas where Wi-Fi is available but contended.

A7.98 For example, Figure 33 and Figure 34 show that a network with 10 MHz of 800 MHz spectrum can operate at 85% loading and provide a notional 5Mbps service to 60% of locations, but that if the network operator wanted to provide a notional 5Mbps service to 85% of locations the loading on the network would have to be 40% or less[23]. By contrast, a network with 15 MHz of 1800 MHz spectrum would be able to provide a notional 5Mbps service to 85% of locations whilst operating at somewhere between 50% and 70% loading with the equivalent of 18,000 sites nationally, and between 40% and 60% loading with the equivalent of 12,000 sites nationally. A network with 20 MHz of 2600 MHz spectrum would be able to provide a 5Mbps service to 85% of locations whilst operating at somewhere between 45% and 85%

---

[23] Maintaining a loading as low as this is likely to be seen by most operators as undesirable.

loading with the equivalent of 18,000 sites nationally, but could not reach 85% at any load in the 'Max var' case with the equivalent of 12,000 sites nationally. A network with only 5 MHz of 800 MHz spectrum has to maintain a loading at 15% to provide a notional 5Mbps service to between about 70% and 80% of locations.

A7.99 Looking at the variation between 'Min var' and 'Max var' results in another way, at a constant loading of 40%, a network with 10 MHz of 800 MHz spectrum is predicted by our model to be able to deliver a notional 5Mbps service to almost 85% of locations with the equivalent of 18,000 sites nationally, and about 80% with the equivalent of 12,000 sites nationally. Whereas a network with 15 MHz of 1800 MHz spectrum is predicted by our model to be able to provide a notional 5Mbps service to somewhere between 84% and 95% of locations, and a network with 20 MHz of 2600 MHz spectrum somewhere between 85% and 96% of locations with the equivalent of 18,000 sites nationally, and between 78% and 94% with the equivalent of 12,000 sites nationally. This again exemplifies the greater uncertainty of our model's prediction of coverage delivered by networks using higher frequencies. But in this case it shows how networks with larger quantities of higher frequencies may still be able to operate at higher loading, and hence have higher capacity, for relatively high coverage levels.

Figure 33: Maximum loading as a function of coverage delivered: 5Mbps, 18,000 sites

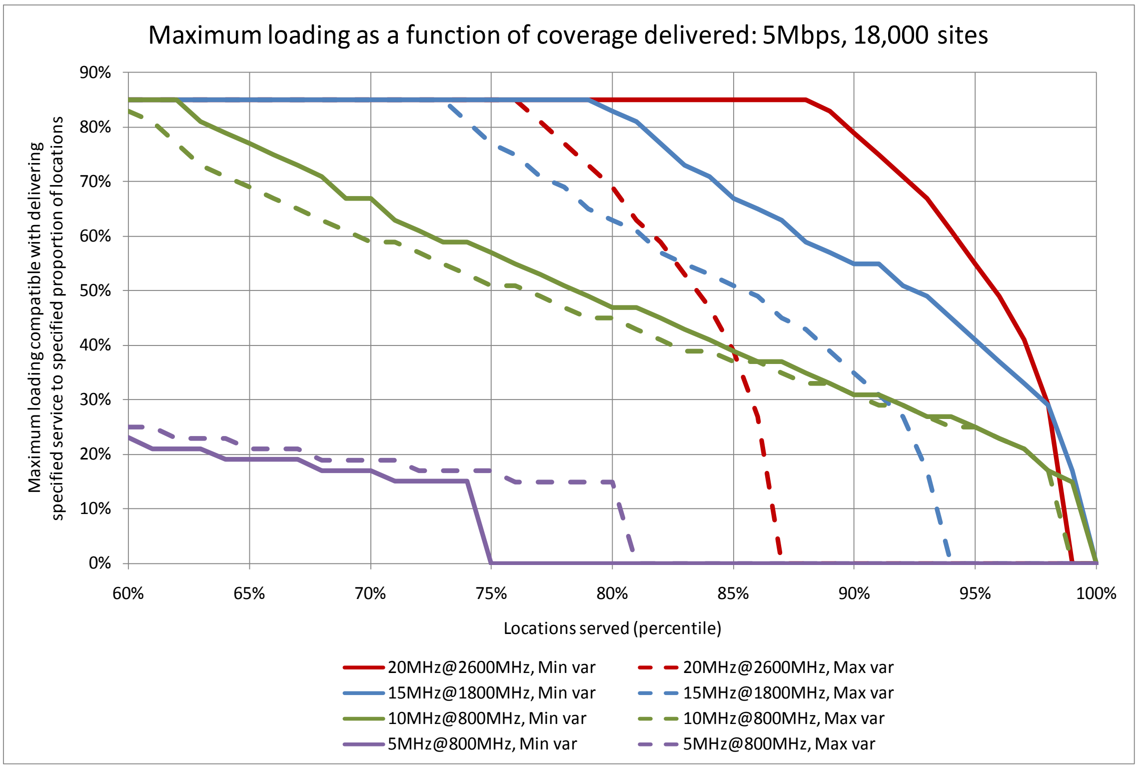

Figure 34: Maximum loading as a function of coverage delivered: 5Mbps, 12,000 sites

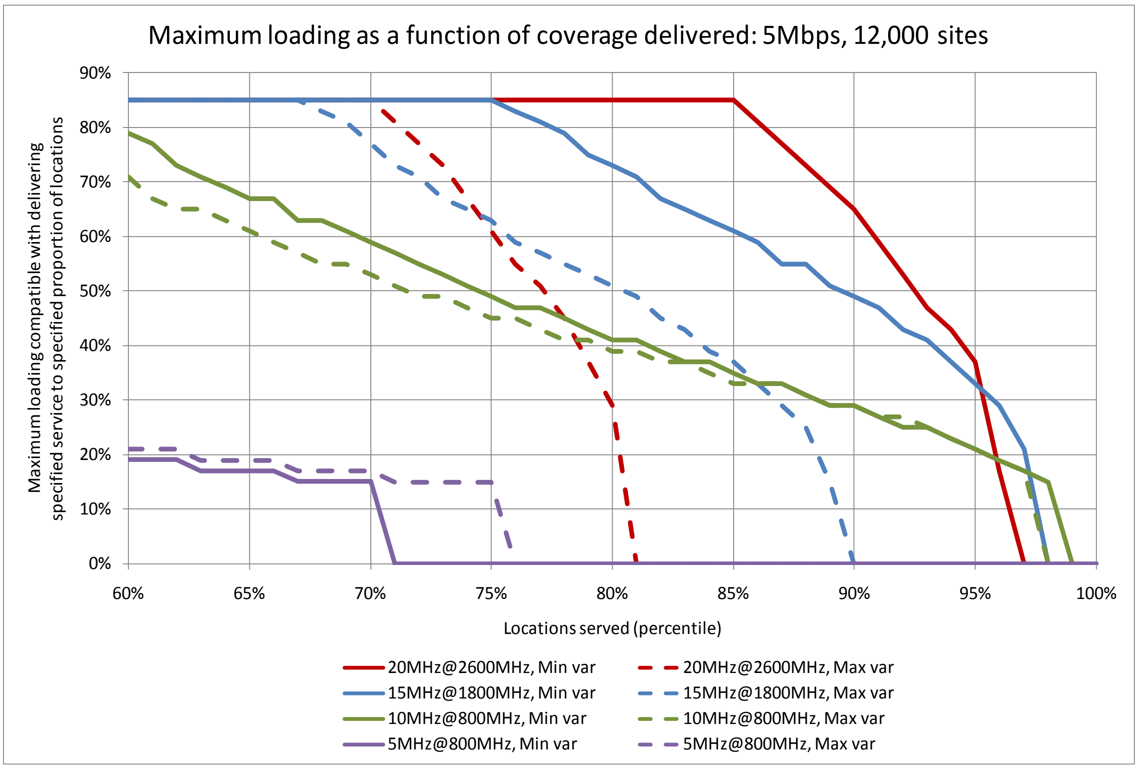

A7.100 Turning now to Figure 35 and Figure 36, we see the impact of the different loadings and other factors (in particular the fact that to deliver the same service to users in harder to serve locations requires more network resources) on the capacity of the networks to serve users (the relative number of simultaneous users of the specified service – in this case a notional 5Mbps service – that each network can support).

A7.101 Here we can see how the capacity of a network with only 10 MHz of 800 MHz spectrum is predicted by our model to be materially less than that of networks with 15 MHz of 1800 MHz spectrum or 20 MHz of 2600 MHz spectrum for coverage up to about 85% of locations. But we also see how the network with 800 MHz spectrum is predicted by our model to be able to provide a notional 5Mbps service for all coverages up to almost 100% of locations (with little difference between ‘Min var’ and ‘Max var’ results). For the networks with higher frequencies the uncertainty of our model’s prediction of results at all coverages is greater, and also there appears a sudden drop-off in the ’Max var’ capacity results (the ‘knee’) if the network tries to provide coverage beyond a certain critical point – e.g. beyond about 92% for the network with 15 MHz of 1800 MHz spectrum.

A7.102 The point at which the capacity of a network drops to zero (the intercept on the x-axis), and the variability of this point between the ‘Min var’ and ‘Max var’ results, also provides us with an indication of the extent and uncertainty of the prediction of the ultimate limit of coverage for the specified service, not forgetting however that the network may still be able to provide a lower speed of service beyond this point.

Figure 35: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

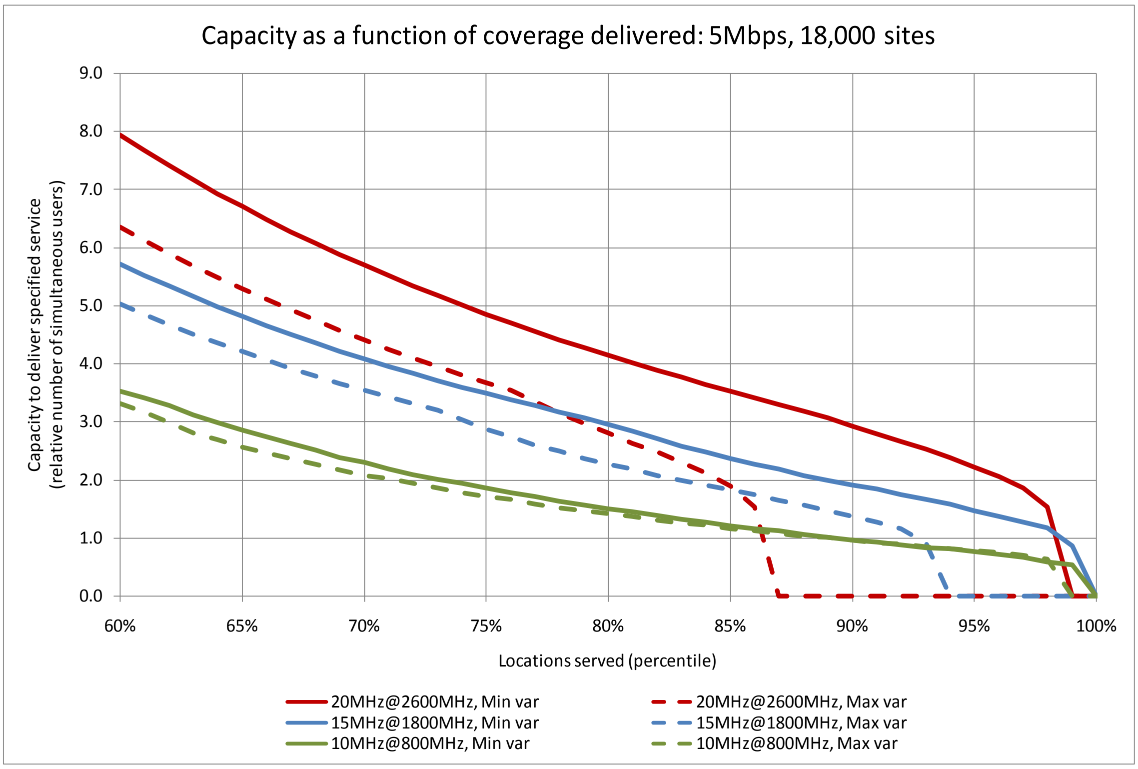

Figure 36: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

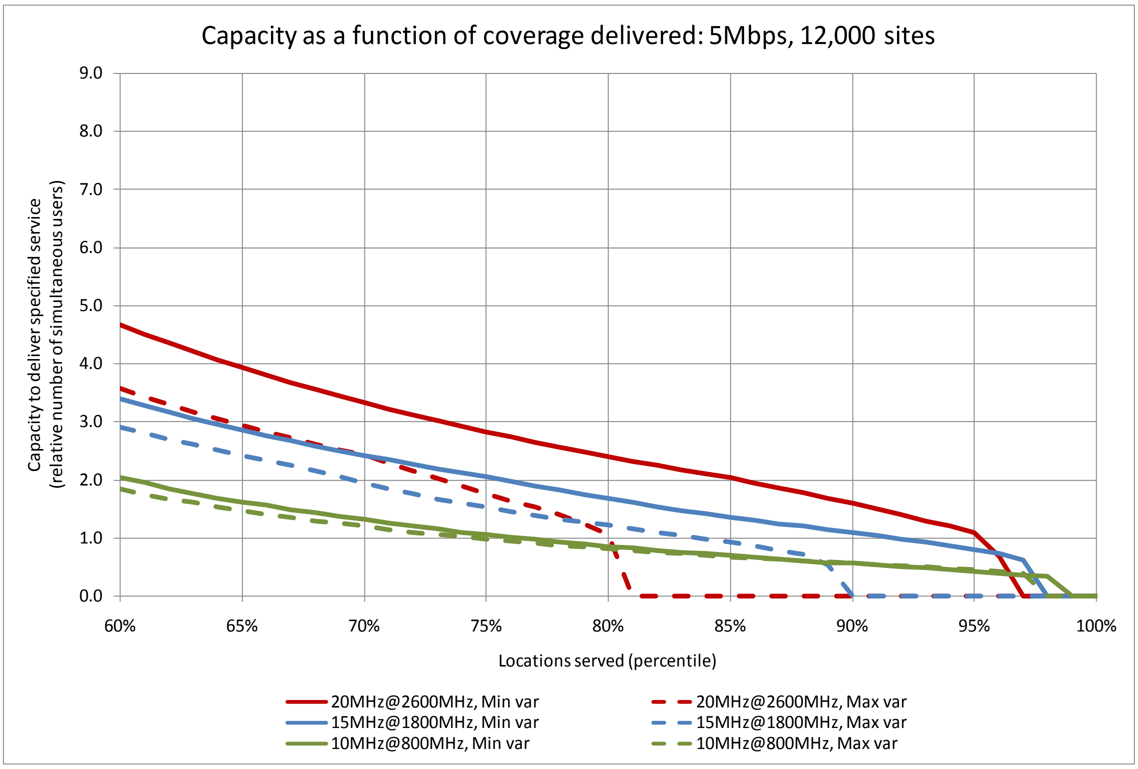

A7.103 (Note: in this graph and similar graphs that follow that show capacity as a function of coverage, the y axis uses a normalised scale indicating the relative number of simultaneous users the networks could support using the relevant bandwidth and frequency. It is not a prediction of the actual number of simultaneous users that the networks could support.)

A7.104 Indeed, Figure 37 shows how the maximum speed of service that the three networks (with the equivalent of 18,000 sites nationally) can provide to customers varies with location, for the case where each network is just able to provide a notional 5Mbps service to users at approximately the 83% percentile of locations in the 'Max var' case. This clearly shows how the networks with larger amounts of higher frequency spectrum are able to offer higher speed services to users at the majority of locations, albeit with greater uncertainty over the potential speed our model predicts to be available (difference between 'Min var' and 'Max var' results). But also how those networks may not be able to offer the same speed of service, or any service at all, to users in the hardest to serve locations – beyond the 83% of locations to which they can deliver a notional 5Mbps service – as compared with the speed of service that the network with only 10 MHz of 800 MHz spectrum can deliver.

Figure 37: Single-user throughput as a function of location for various frequencies, loadings and bandwidths – 18,000 sites

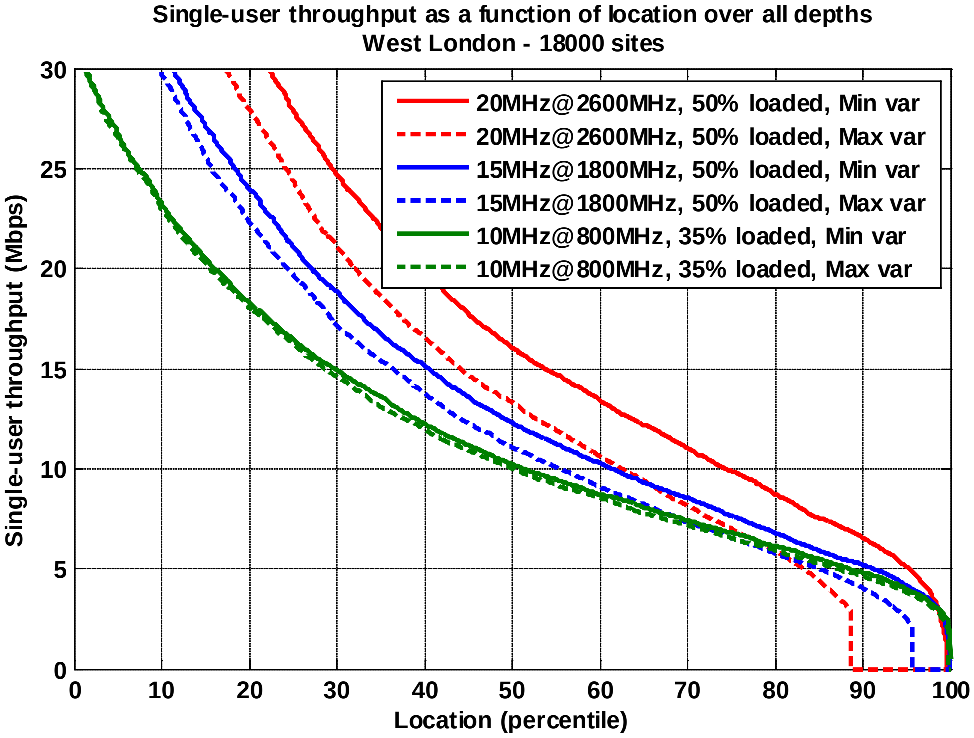

A7.105 However, Figure 38 shows that for networks with the equivalent of 12,000 sites nationally, even with larger amounts of higher frequency spectrum cannot achieve coverage beyond about the 80% percentile of locations. But they are still able to offer higher speed services to users at the majority of locations.

Figure 38: Single-user throughput as a function of location for various frequencies, loadings and bandwidths – 12,000 sites

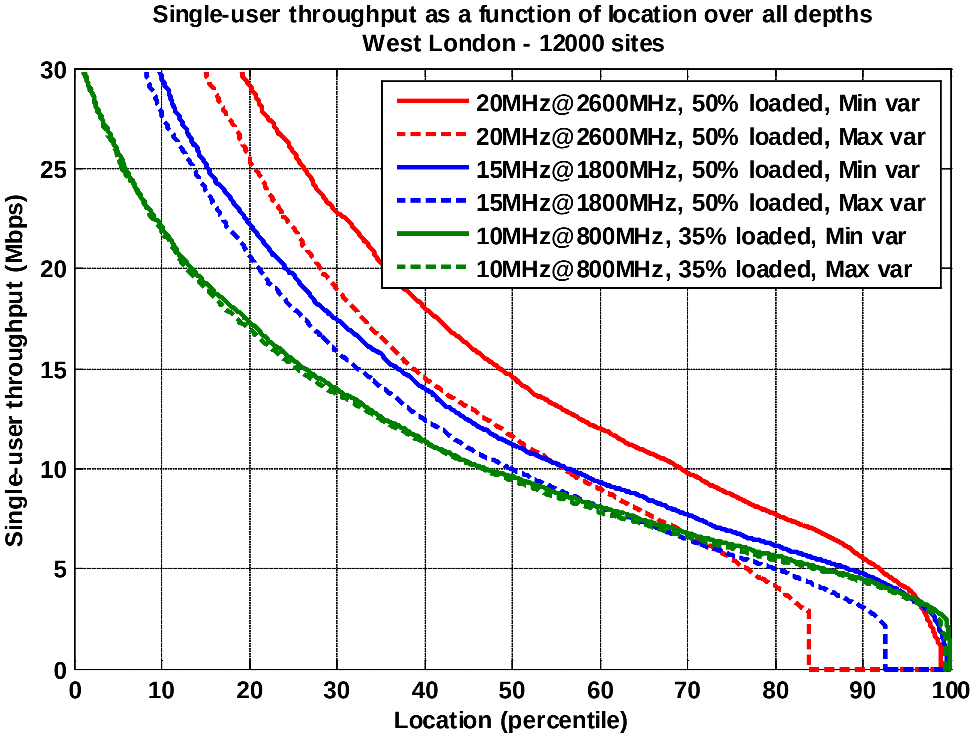

A7.106 In Figure 39 and Figure 40 we look at the impact that additional spectrum at 2600 MHz is predicted by our model to have on the relative capacity of each of the three single-frequency networks we have just been looking at (using our multi-frequency analysis). Comparing Figure 39 and Figure 40 with Figure 35 and Figure 36 we see a clear increase in overall capacity, and also a closing of the gap in relative capacity between the network with 10 MHz of 800 MHz spectrum and those with larger amounts of higher frequencies. The ability of the higher frequency networks to maintain capacity up to the point where they become coverage limited is also improved as a result of the additional 2600 MHz spectrum (the ‘knees’ in the curves are sharper), but the absolute limit of coverage for each of the networks remains essentially unchanged.

Figure 39: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

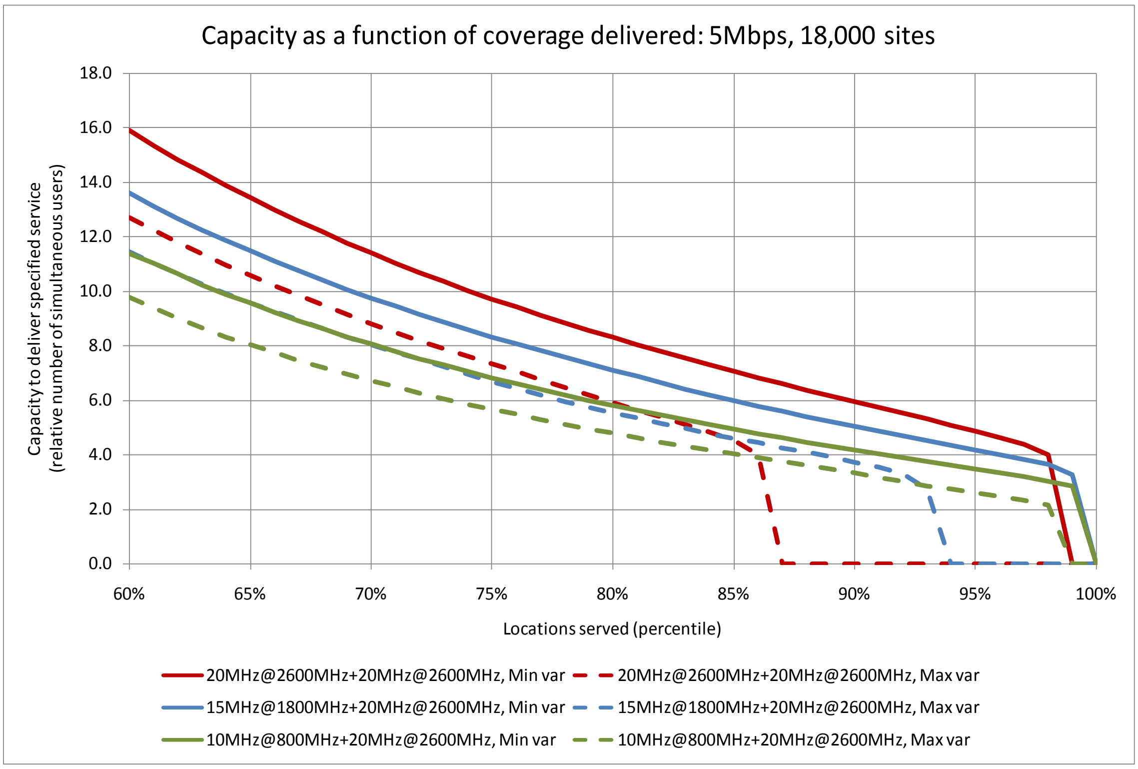

Figure 40: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

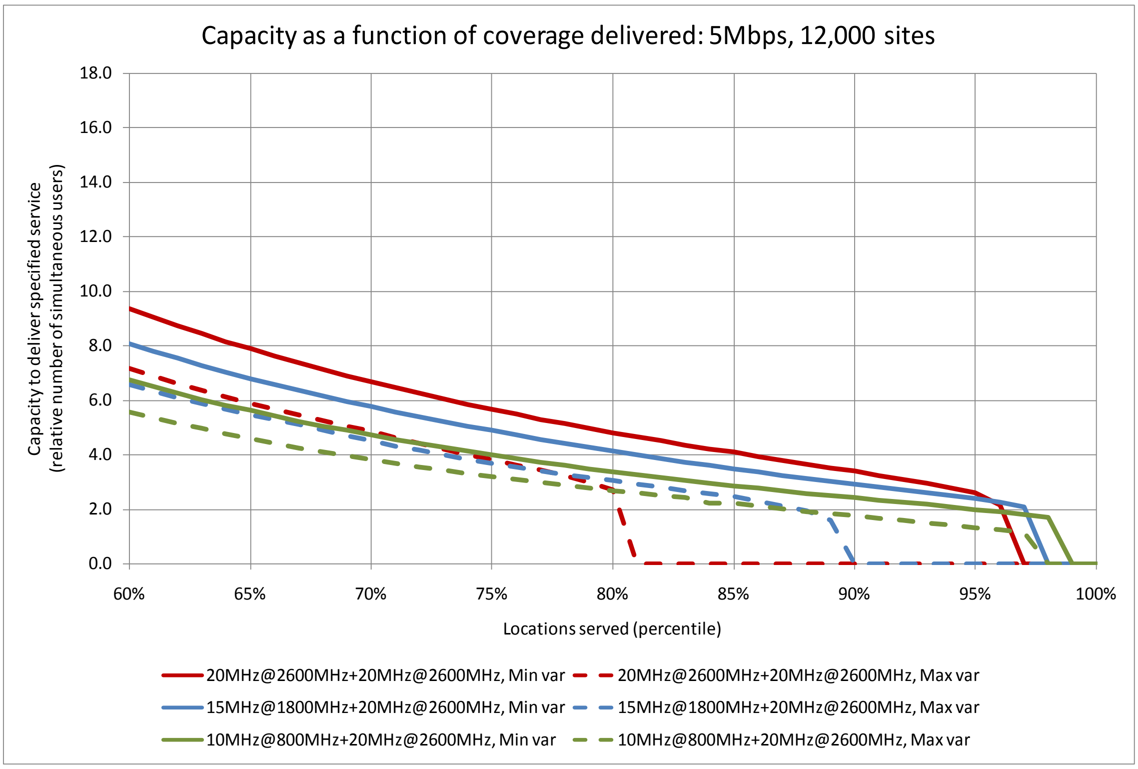

A7.107 So what combinations of spectrum are predicted by our model to have similar capacity? In Figure 41 and Figure 42 we look at the same three frequencies as before, but now with the same bandwidth available in each case (20 MHz). Whilst the 'Min var' capacities for these three spectrum portfolios are similar, the higher frequency networks have consistently less capacity in the 'Max var' case, albeit they still have more capacity than an 800 MHz network with 15 MHz of spectrum (as shown for comparison on the same graph).

Figure 41: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

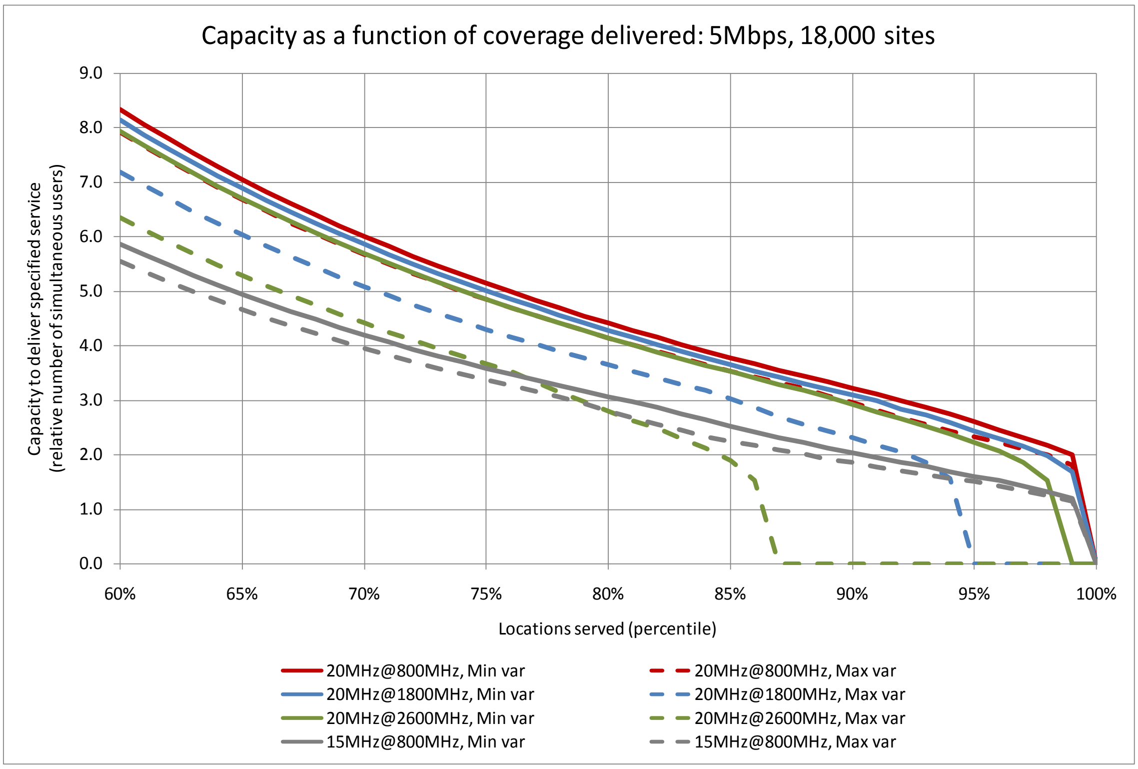

Figure 42: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

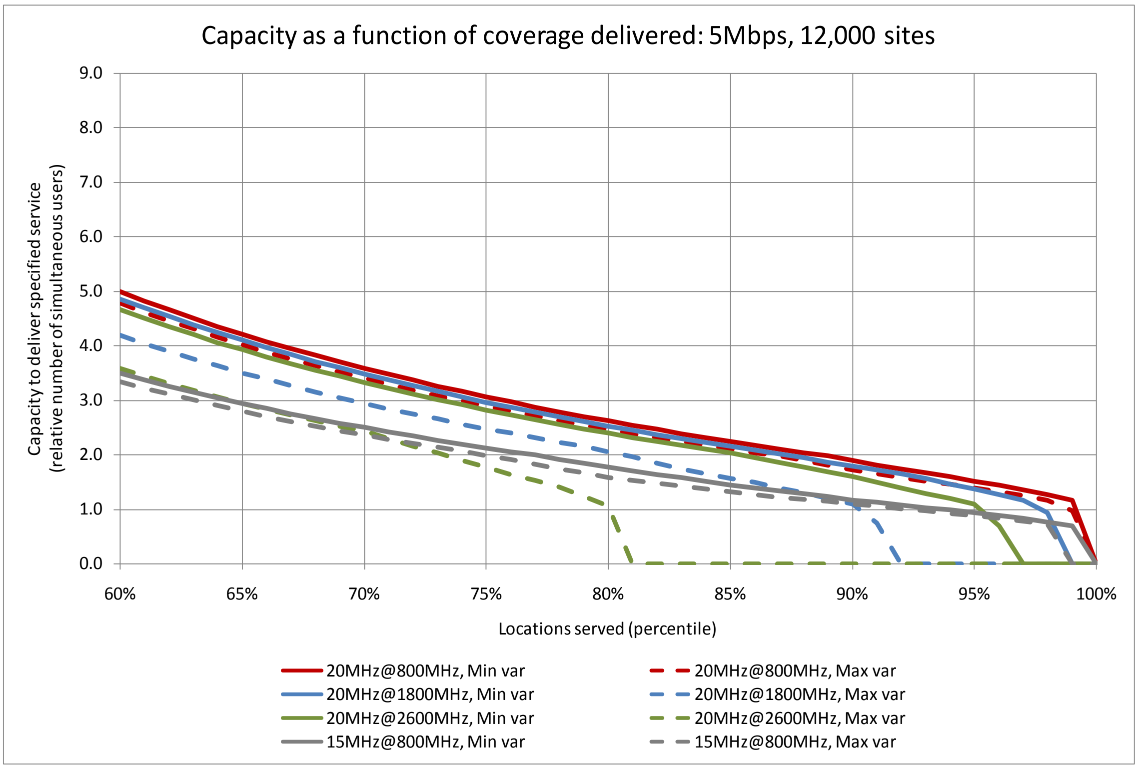

A7.108 Considering multi-frequency spectrum portfolios of a similar size (Figure 43 and Figure 44), we find that a range of portfolios including a mix of 800 MHz, 1800 MHz and 2600 MHz spectrum, with a total bandwidth of 25 MHz, are all predicted by our model to provide a very similar capacity to a single-frequency network with 20 MHz of 800 MHz spectrum.

Figure 43: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

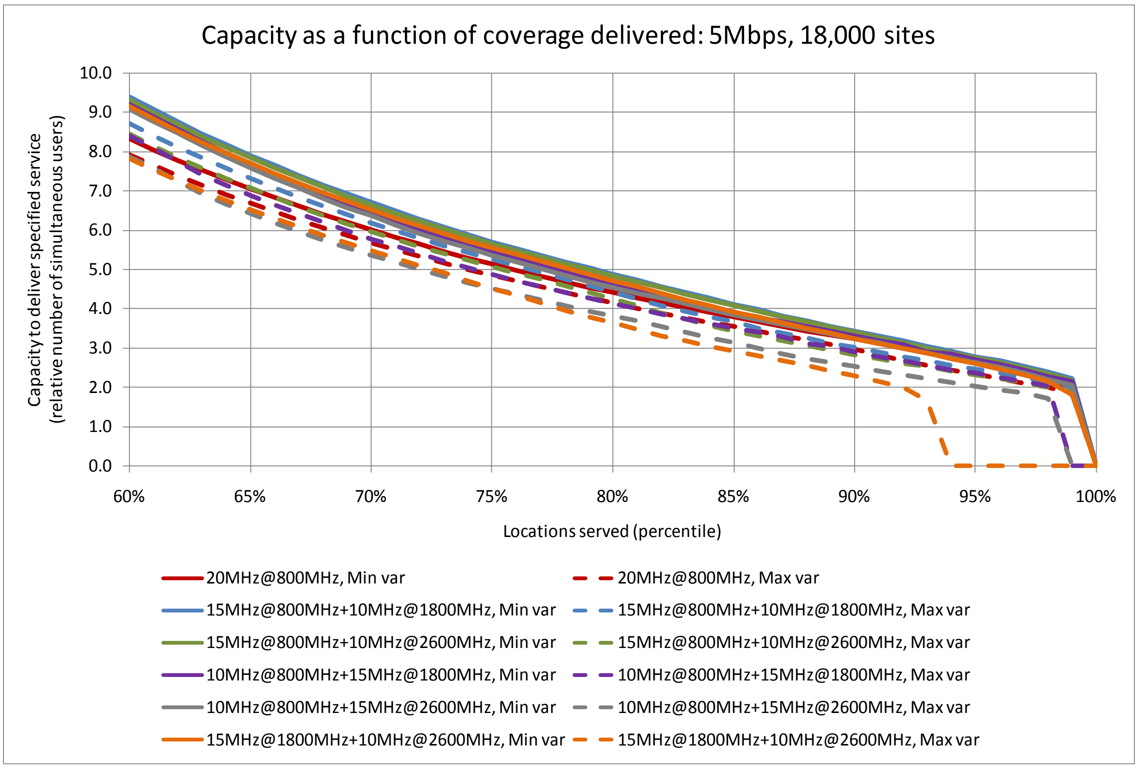

Figure 44: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

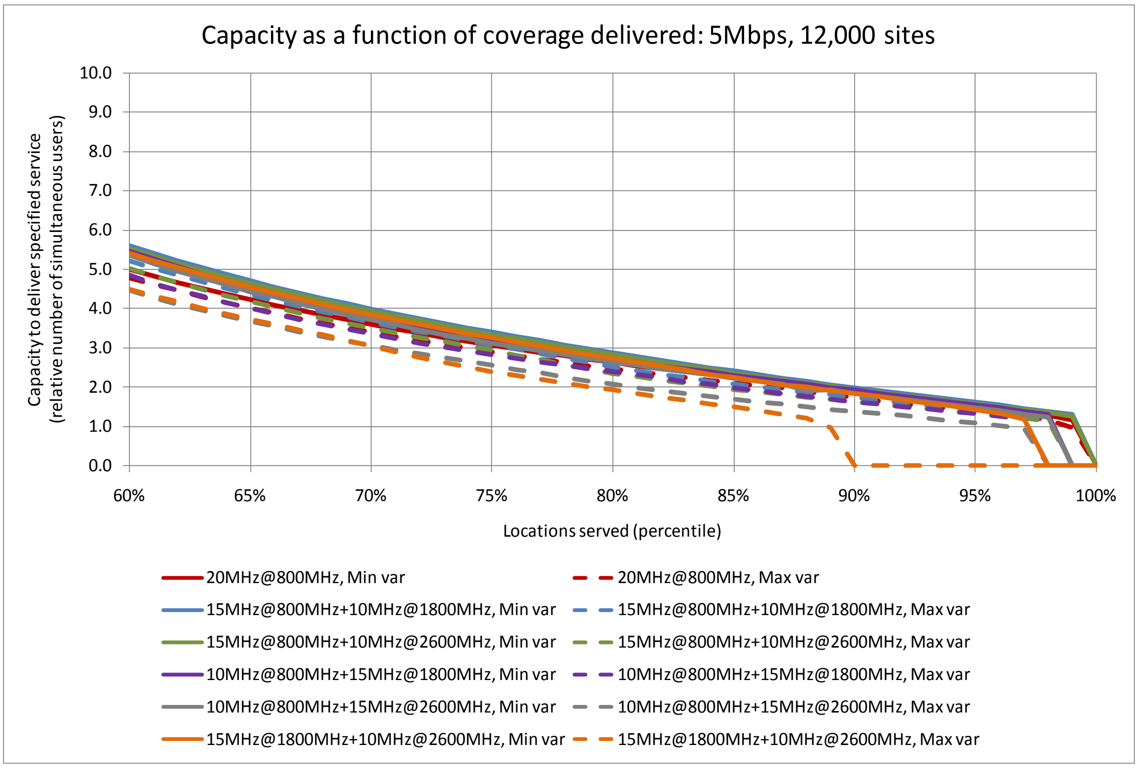

A7.109 However for larger spectrum portfolios (Figure 45 to Figure 50) we find that multi-frequency portfolios are not predicted by our model to need additional spectrum bandwidth overall in order to be able to provide similar capacity to a single-frequency 800 MHz network. But the uncertainty our model's prediction of capacity provided across locations (range of results between 'Min var' and 'Max var') is larger for networks that rely on large quantities of 2600 MHz spectrum, or do not have any 800 MHz spectrum, as compared with those that don't.

Figure 45: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

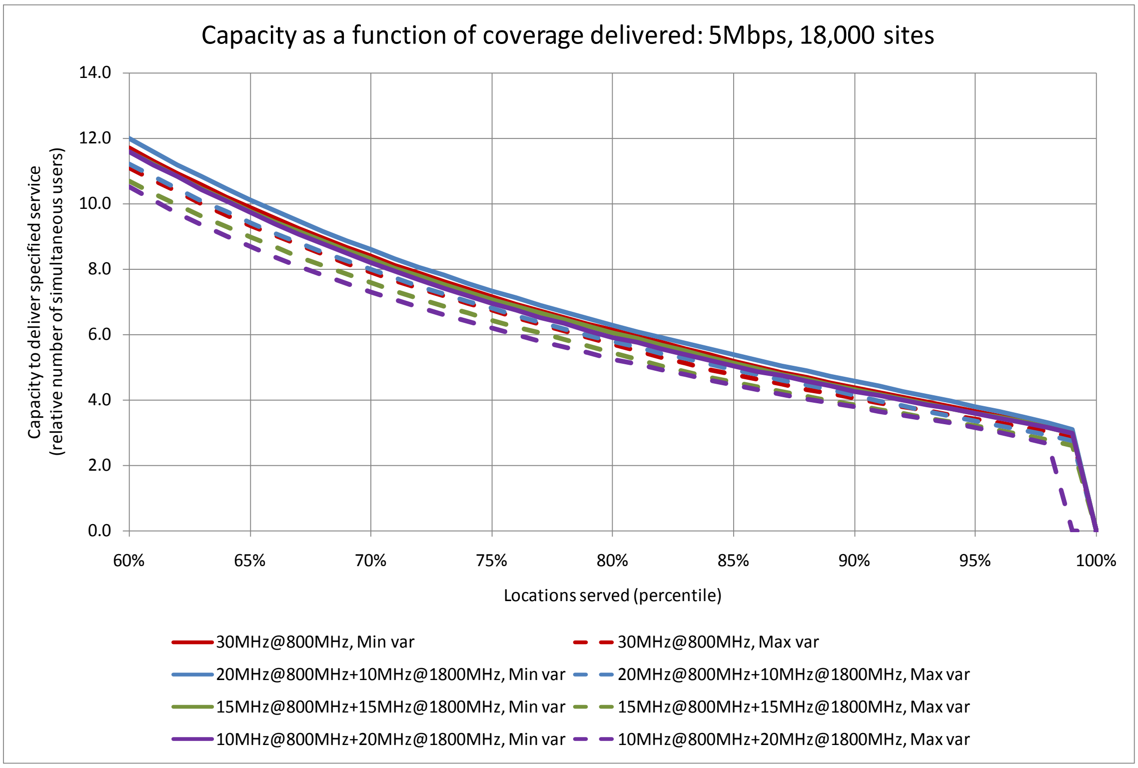

Figure 46: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

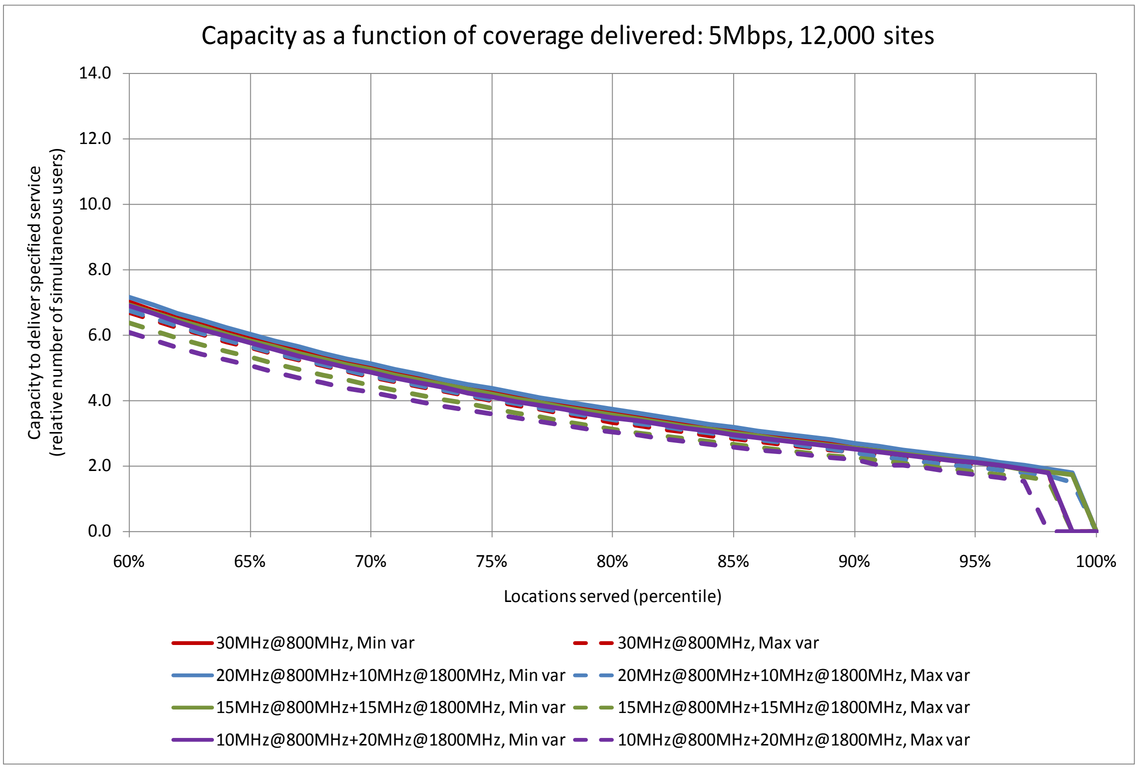

Figure 47: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

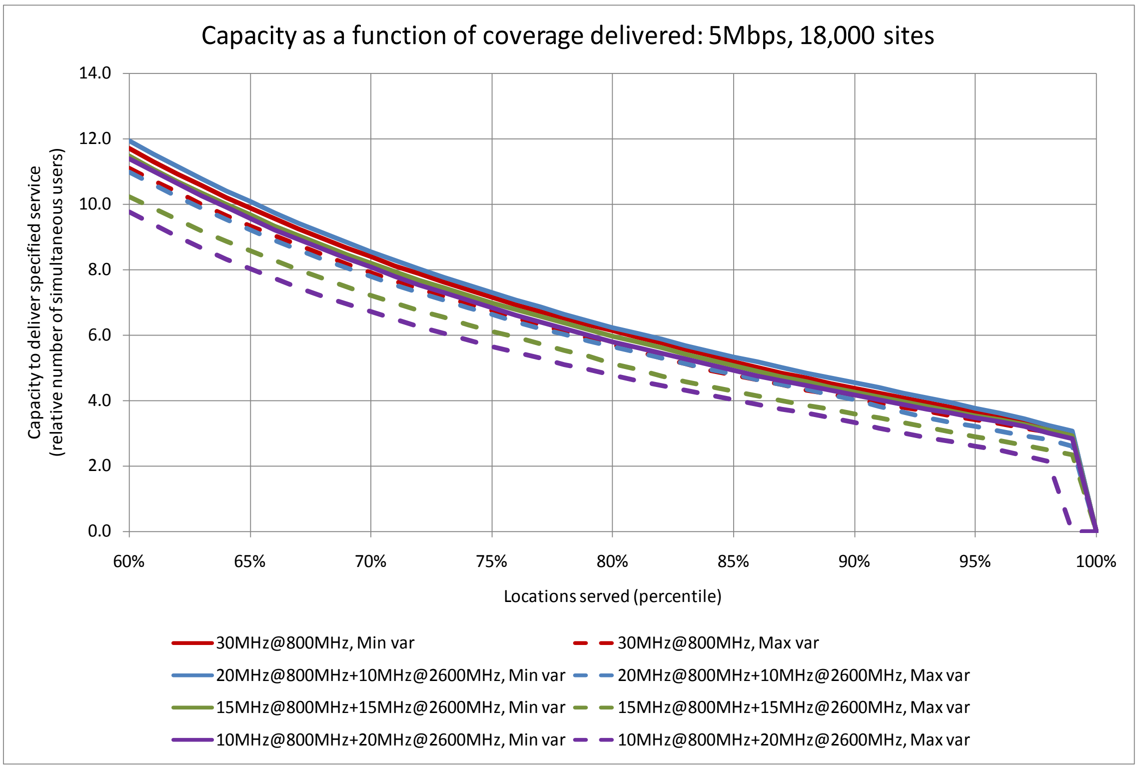

Figure 48: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

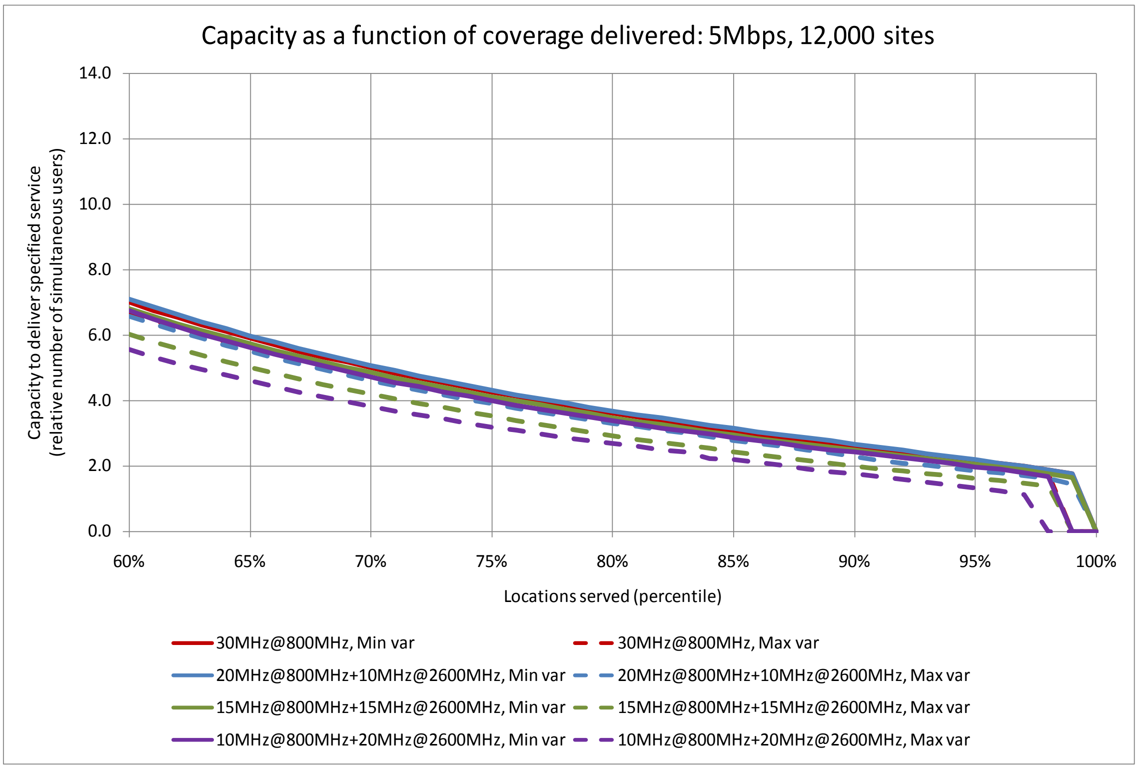

Figure 49: Relative capacity as a function of coverage delivered: 5Mbps, 18,000 sites

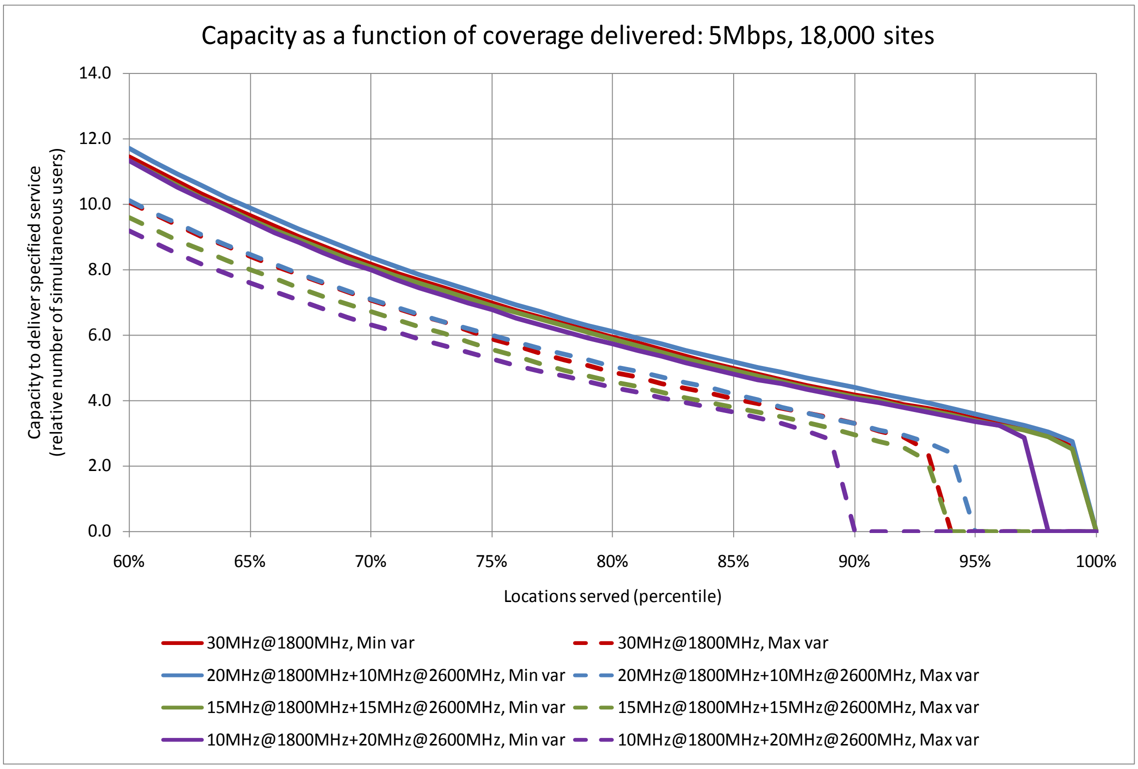

Figure 50: Relative capacity as a function of coverage delivered: 5Mbps, 12,000 sites

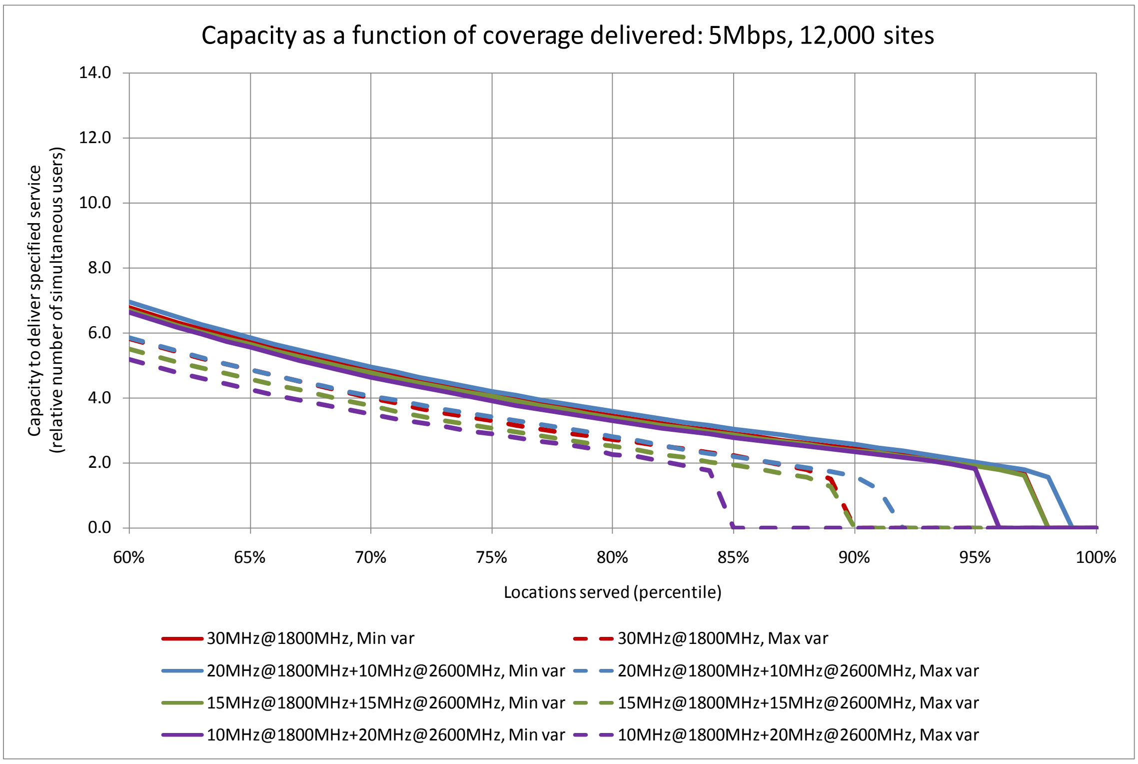

A7.110 We therefore conclude that, at least for networks with more than 20 MHz of spectrum bandwidth in total, capacity is largely a function of the total bandwidth of spectrum available to a network, rather than the specific mix of frequencies. The uncertainty of our models prediction of capacity over locations, however, can vary materially according to the mix of bandwidths between frequencies.

Comparison of portfolios

A7.111 In this section we look at the predicted performance of the networks that the existing MNOs could theoretically build in both the nearer and the longer term, using their existing spectrum holdings, and how the performance of networks using a range of other spectrum portfolios might compare.

A7.112 In all cases we look at the capacity to deliver a notional 5Mbps services (since we consider this to be a reasonable basis for the comparison of network capacities in this context) for networks with:

- 12,000 sites: since we consider this to be a reasonable estimate of the total number of macrocell sites that an operator is likely to have access to in the next 2 to 3 years; and

- 18,000 sites: since we consider this to be a reasonable estimate of the total number of macrocell sites that an operator might be able to use in the longer term).

A7.113 For the purposes of this analysis we have assumed that 900 MHz performs identically to 800 MHz and that 2100 MHz performs identically to 1800 MHz This is obviously an approximation. However, in Annex 14 we have looked at the differences in performance our model predicts for LTE networks operating at 900 MHz as opposed to 800 MHz and 2100 MHz as opposed to 1800 MHz and have concluded the differences are relatively minor and therefore using 800 MHz and 1800 MHz results as a proxy for 900 MHz and 2100 MHz respectively is unlikely to have a material impact on the results.

A7.114 Looking first at the capacity and coverage that networks using the MNOs' existing spectrum portfolios might be able to deliver in the longer term (once the majority of existing spectrum has been re-farmed for LTE or LTE-Advanced), Figure 51 and Figure 52, we see that:

- Everything Everywhere (EE) is predicted by our model to have significantly more capacity than any of the other existing MNOs, but with uncertainty over their coverage affecting the hardest to serve 5% of locations;
- Vodafone and Telefónica are predicted by our model to have about half the capacity of EE, but with coverage to almost 100% of locations;
- Hutchison 3G (H3G) is predicted by our model to have the least capacity, less than half that of Vodafone and Telefónica, and less than a quarter that of Everything Everywhere; the uncertainty of their coverage is also larger even than Everything Everywhere's.

Figure 51: Longer term relative capacity and coverage of MNOs' existing spectrum portfolios: 5Mbps, 18,000 sites

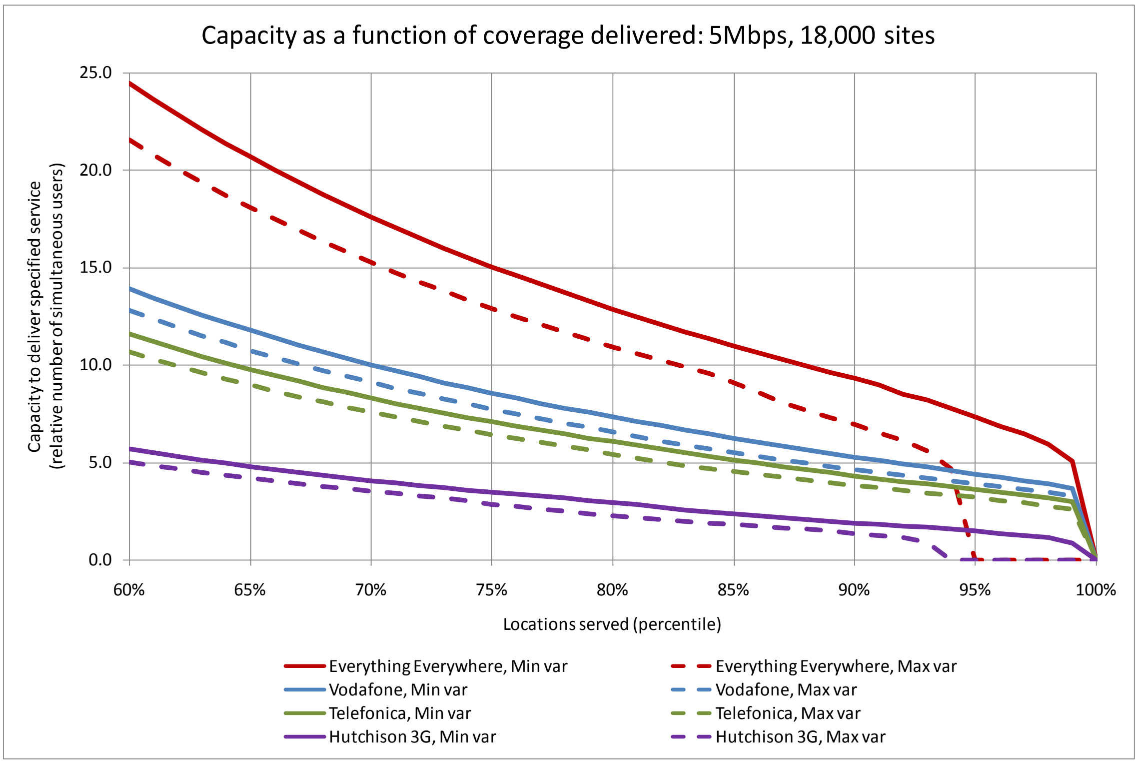

Figure 52: Longer term relative capacity and coverage of MNOs' existing spectrum portfolios: 5Mbps, 12,000 sites

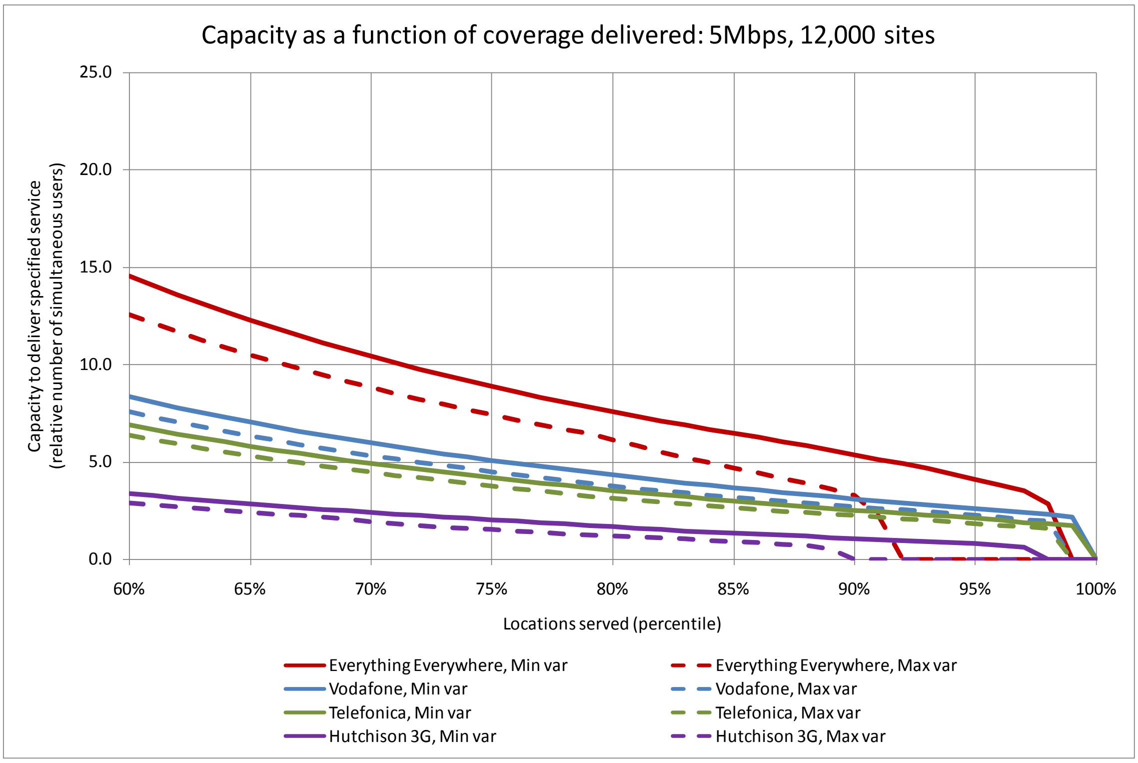

A7.115 Note that for the purpose of these charts we have used 800 MHz as a proxy for 900 MHz, and 1800 MHz as a proxy for 2100 MHz We have also used the following quantities of spectrum to represent the existing spectrum holdings of the various existing MNOs:

- For Everything Everywhere: 40 MHz@1800 MHz + 20 MHz@2100 MHz.
- For Vodafone: 15 MHz@900 MHz + 5 MHz@1800 MHz + 15 MHz@2100 MHz.
- For Telefónica: 15 MHz@900 MHz + 5 MHz@1800 MHz + 10 MHz@2100 MHz.
- For H3G: 15 MHz@2100 MHz.

A comparison of the capacity and coverage of various spectrum portfolios relative to the predicted longer term capacity of Vodafone and Telefónica's existing spectrum holdings (excluding 2100 MHz spectrum)

A7.116 We now consider how a variety of spectrum portfolios, which might for example be acquired by H3G or a new entrant through the auction, compare with those of the other existing MNO's.

A7.117 For simplicity and due to the uncertainty around LTE equipment availability and re-farming from 3G at 2100 MHz, we focus on capacity derived from spectrum other than the existing 2100 MHz spectrum. By way of comparison, Figure 53 and Figure 54 present the same results as Figure 51 and Figure 52 but excluding 2100 MHz spectrum.

Figure 53: Longer term relative capacity of existing spectrum holdings, excluding 2100 MHz: 5Mbps, 18,000 sites

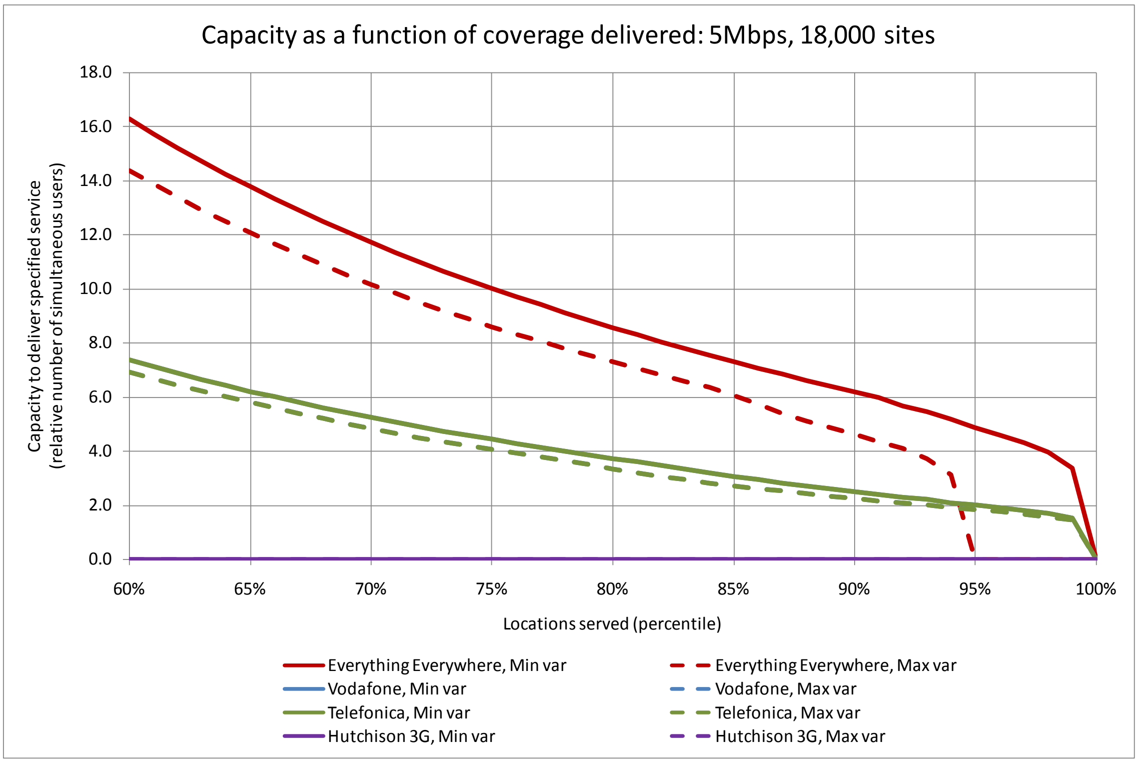

Figure 54: Longer term relative capacity of existing spectrum holdings, excluding 2100 MHz: 5Mbps, 12,000 sites

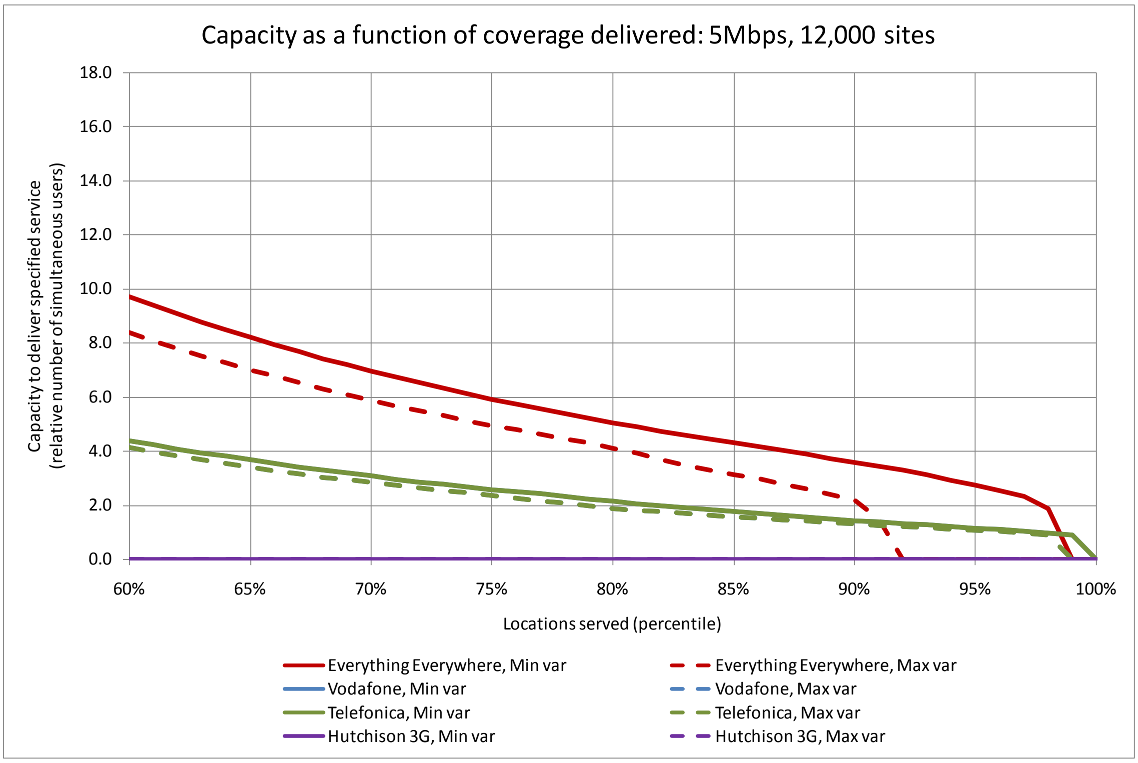

A7.118 (Note that Vodafone and Telefónica hold the same amount of non-2100 MHz spectrum – which is why there appears to be only one line on the graph for these two operators – and H3G holds none – which is why their line is at zero.)

A7.119 We start by comparing the predicted capacity and coverage of some relatively small spectrum portfolios with those of Vodafone and Telefónica, specifically:

- 10 MHz@800 MHz;
- 15 MHz@1800 MHz; and
- 20 MHz@2600 MHz.

Figure 55: Relative capacity and coverage of some small spectrum portfolios compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 18,000 sites

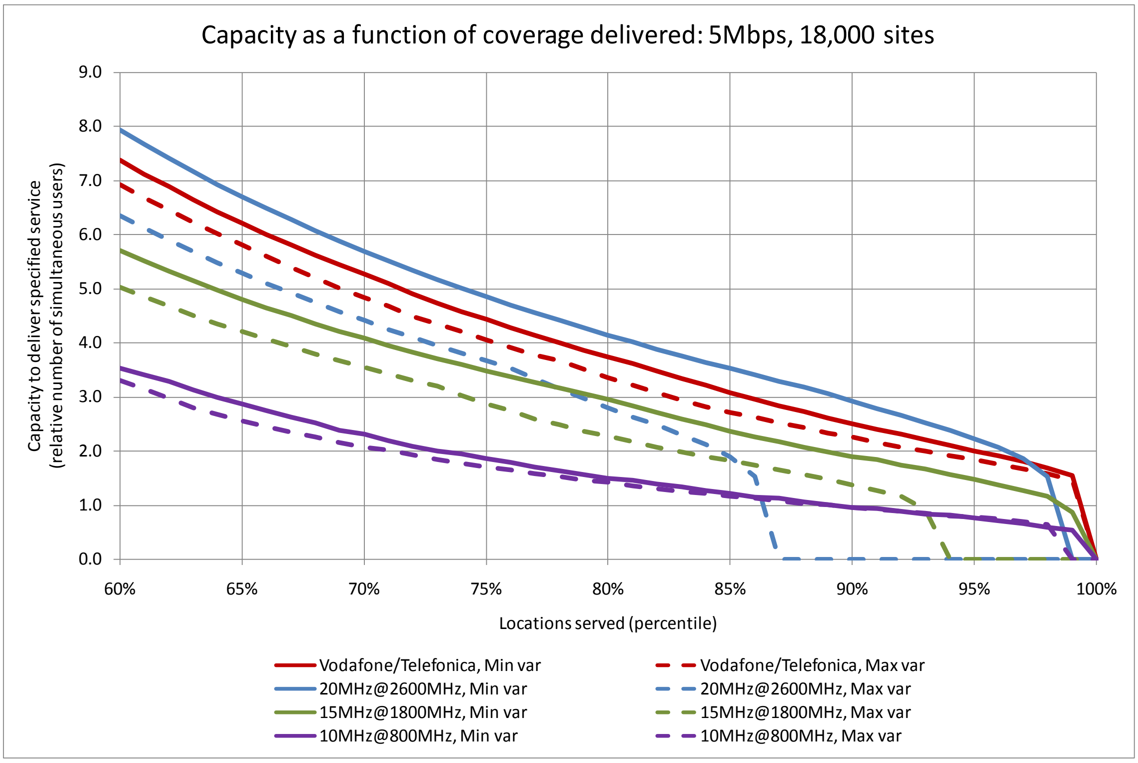

Figure 56: Relative capacity and coverage of some small spectrum portfolios compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 12,000 sites

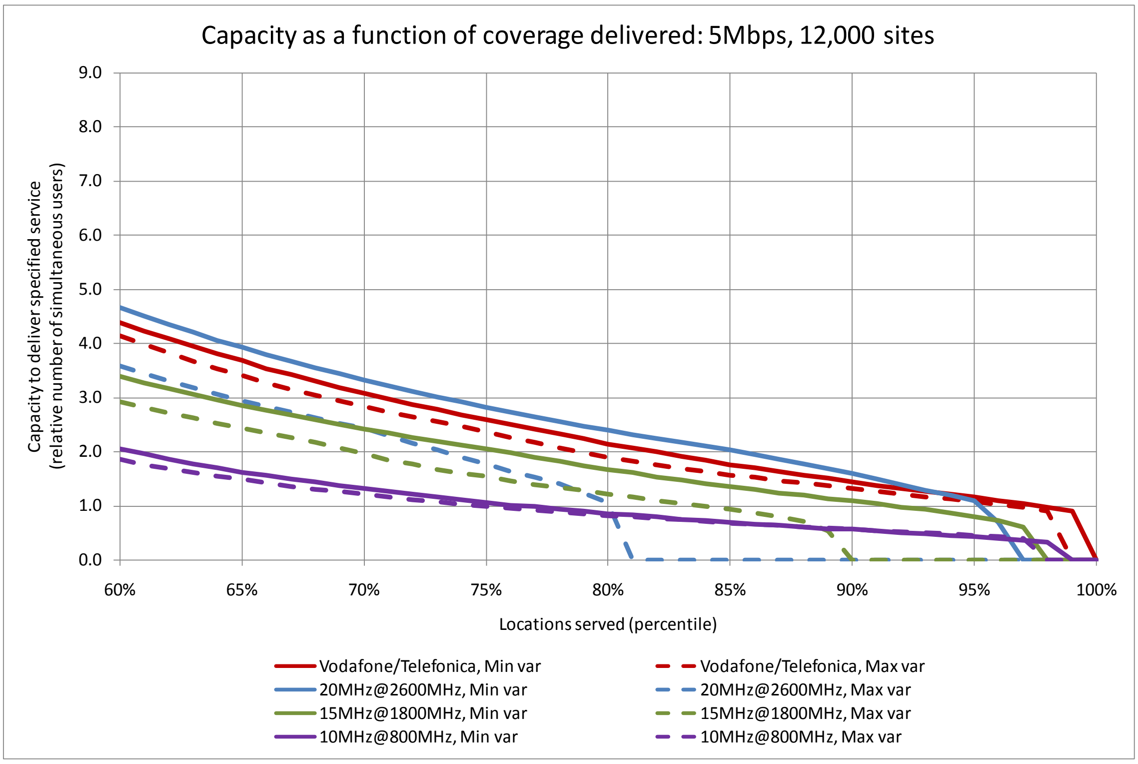

A7.120 We see, unsurprisingly, that none of these portfolios is predicted by our model to match both the capacity and coverage predicted for Vodafone and Telefónica in the longer term:

- the portfolio with 10 MHz of 800 MHz spectrum is predicted by our model to have comparable coverage but at most half the capacity of Vodafone and Telefónica;

- the portfolio with 15 MHz of 1800 MHz spectrum is predicted by our model to come closer to matching the capacity, albeit still being some way short of doing so, but coverage of the hardest to serve 5% of locations is uncertain; and

- the portfolio with 20 MHz of 2.6 GHz spectrum is predicted by our model to have comparable capacity to that of Vodafone and Telefónica in the longer term, but is predicted not to provide coverage at all to some hardest to serve locations, and coverage is uncertain for more than 10% of locations.

A7.121 We next consider a group of somewhat larger spectrum portfolios:

- 15 MHz@800 MHz;

- 10 MHz@800 MHz + 10 MHz@2600 MHz;

- 15 MHz@1800 MHz + 10 MHz@2600 MHz; and

- 30 MHz@2600 MHz.

A7.122 In Figure 57 and Figure 58 we see that both the capacity and coverage of the first two of these portfolios are not dissimilar to that which might be achieved by

Vodafone and Telefónica in the longer term using their existing spectrum portfolios, although capacity is slightly less.

Figure 57: Relative capacity and coverage of some medium spectrum portfolios compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 18,000 sites

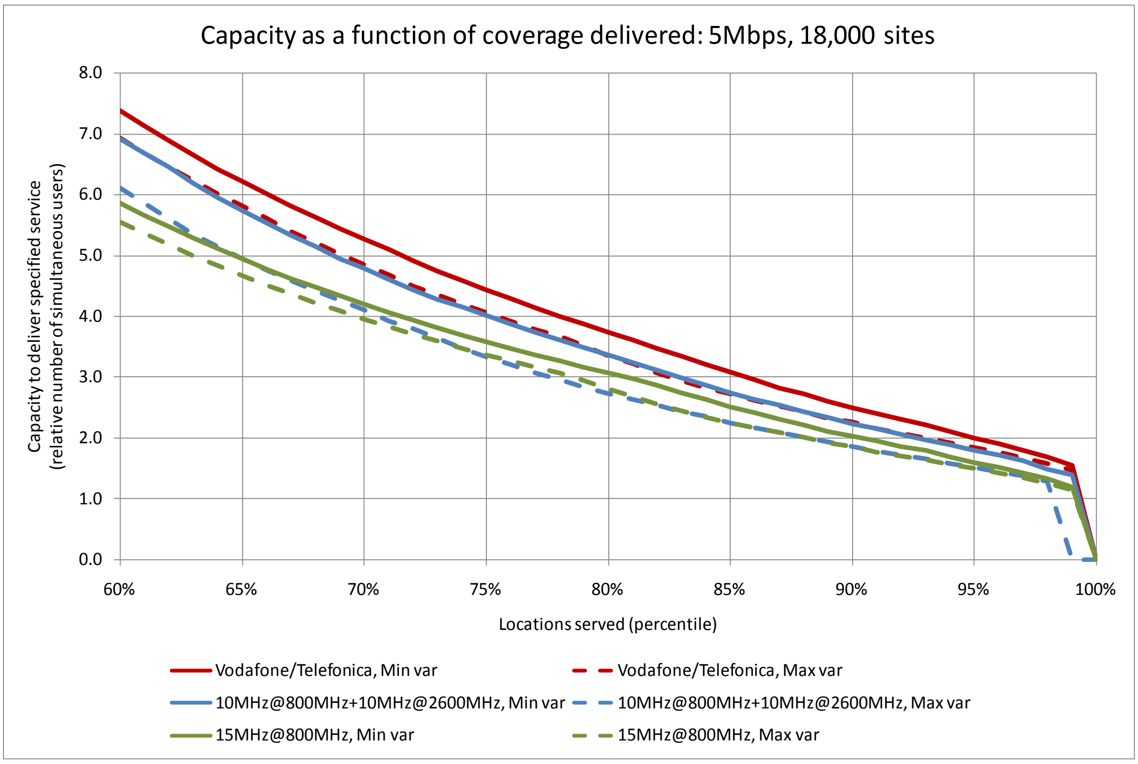

Figure 58: Relative capacity and coverage of some medium spectrum portfolios compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 12,000 sites

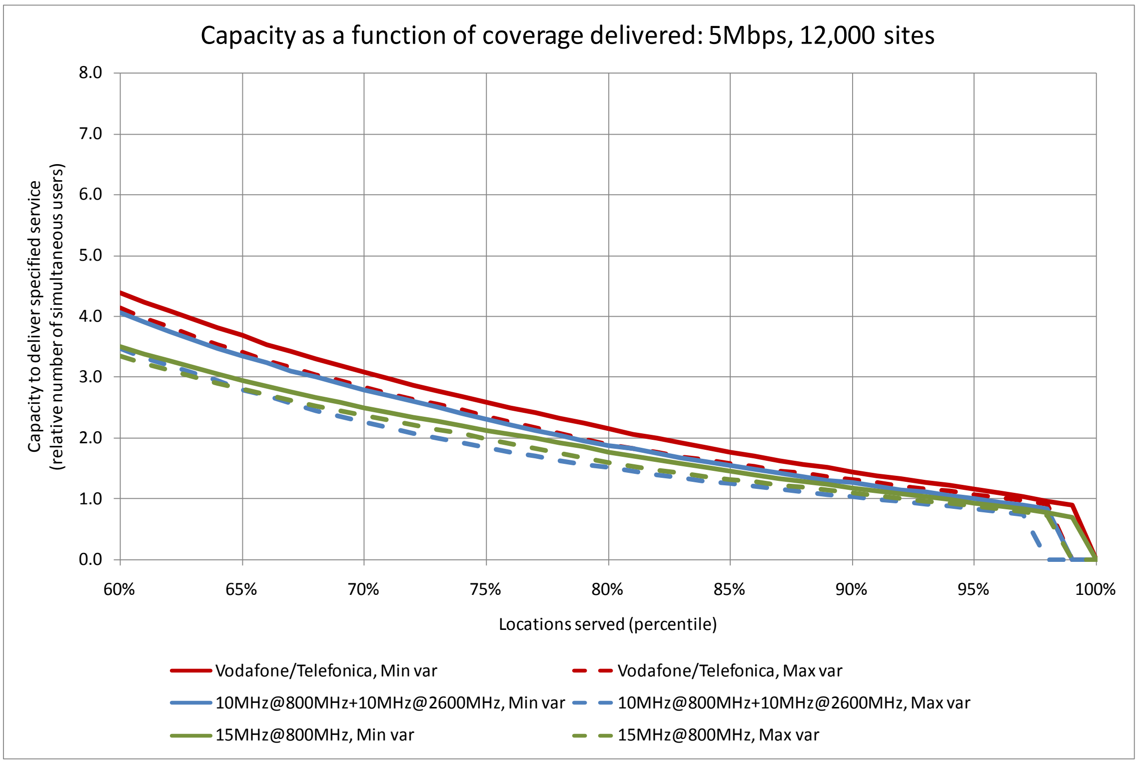

A7.123 In Figure 59 and Figure 60 we see that the latter two portfolios are predicted by our model to have somewhat more capacity than the existing spectrum portfolios of Vodafone and Telefónica, a little more in the case of the 15 MHz@1800 MHz + 10 MHz@2600 MHz portfolio, and quite a bit more in the case of the 30 MHz@2600 MHz portfolio. However in both cases there is greater uncertainty over the coverage that can be provided, significantly more so in the case of the latter portfolio (coverage to the last 14% of locations being uncertain).

Figure 59: Relative capacity and coverage of some medium spectrum portfolios compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 18,000 sites

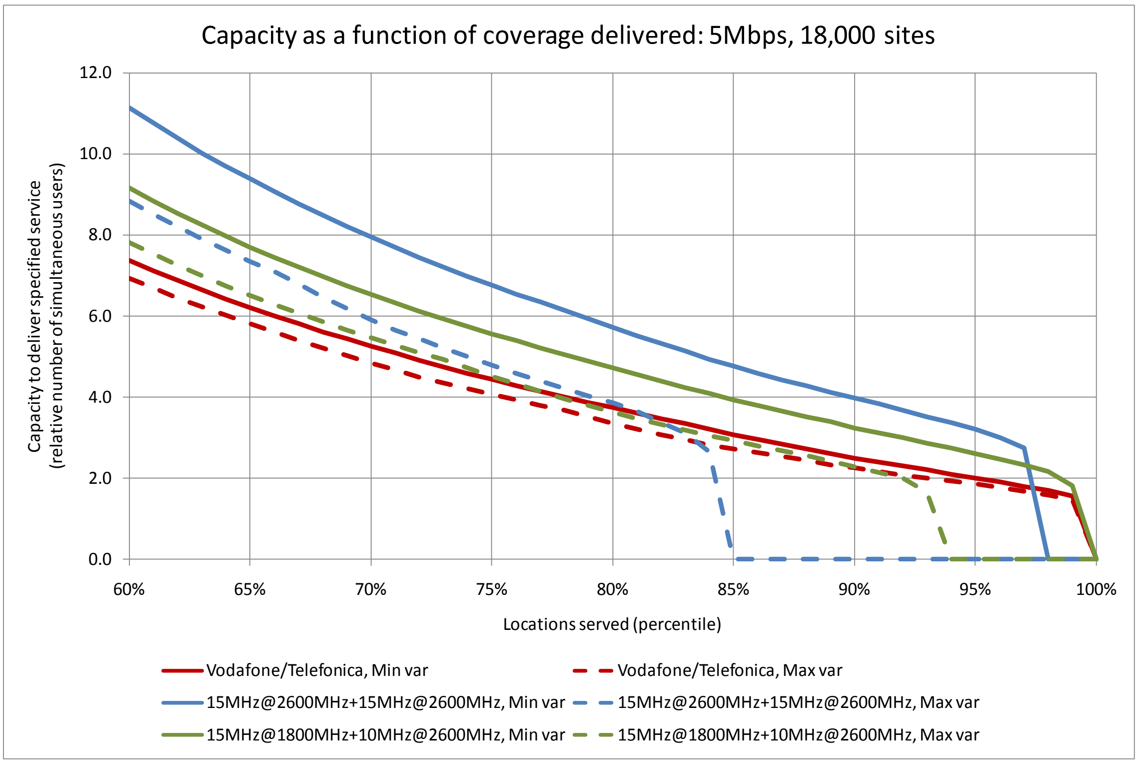

Figure 60: Relative capacity and coverage of some medium spectrum portfolios compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 12,000 sites

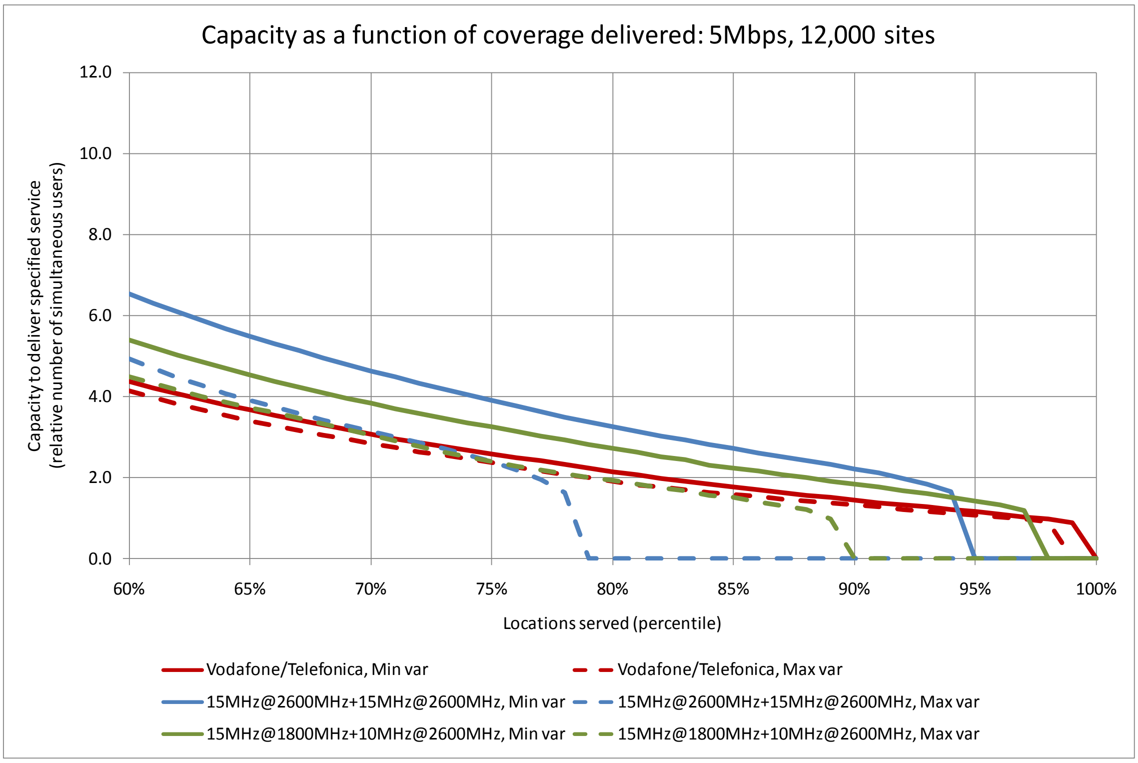

A7.124 Our final set of portfolios are larger still:

- 20 MHz@800 MHz;
- 15 MHz@800 MHz + 10 MHz@2600 MHz;
- 10 MHz@800 MHz + 15 MHz@1800 MHz;
- 10 MHz@800 MHz + 15 MHz@2600 MHz;
- 15 MHz@1800 MHz + 15 MHz@2600 MHz; and
- 40 MHz@2600 MHz.

A7.125 We see from Figure 61 to Figure 64 that all of these portfolios are predicted by our model to have greater capacity than that predicted for Vodafone and Telefónica in the longer term, using their existing spectrum portfolios; moderately so in the case of the first four portfolios, more so in the case of the last two, with the last portfolio have a predicted capacity approaching that of Everything Everywhere's existing spectrum holding.

A7.126 In the case of the first four portfolios, coverage is predicted by our model to be comparable to, if not better than, that predicted to be achievable in the longer term by Vodafone and Telefónica using their existing spectrum holdings. However, in the case of the last two portfolios, there is material uncertainty over the quality of coverage achievable. In the case of the 15 MHz@1800 MHz + 15 MHz@2600 MHz portfolio uncertainty over the quality of coverage is comparable to, but slightly worse than that predicted for Everything Everywhere. In the case of the last portfolio, of 40 MHz@2600 MHz, the uncertainty is materially greater, with coverage of the last 10% of locations being uncertain (for a notional 5Mbps service).

Figure 61: Relative capacity of some larger spectrum portfolios as compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 18,000 sites

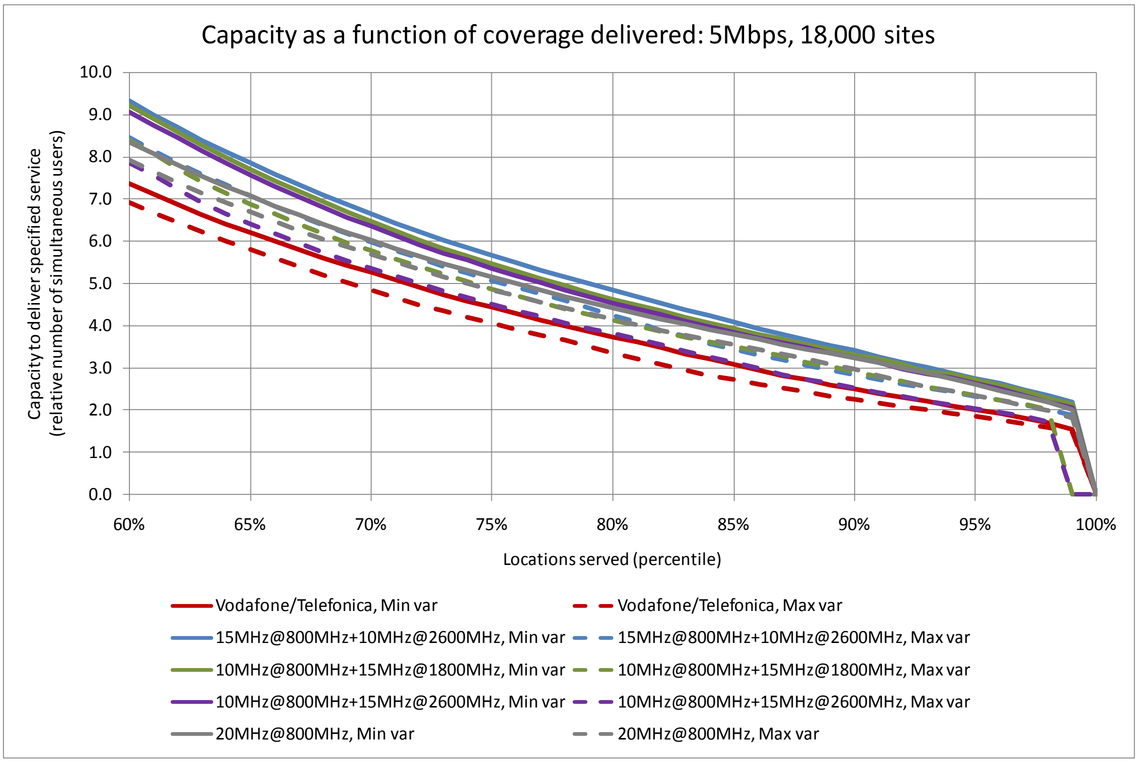

Figure 62: Relative capacity of some larger spectrum portfolios as compared with Vodafone/ Telefónica existing spectrum holdings (excl 2100 MHz): 5Mbps, 12,000 sites

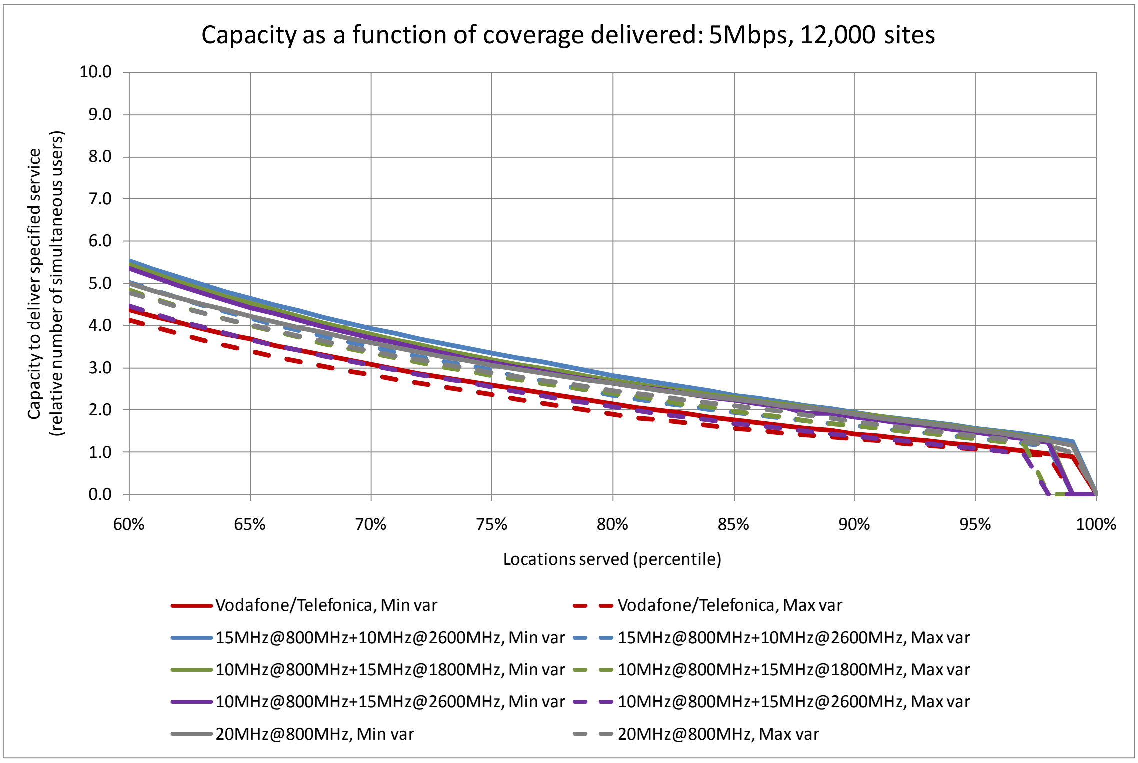

Figure 63: And some portfolios with even more bandwidth, but lower coverage quality: 5Mbps, 18,000 sites

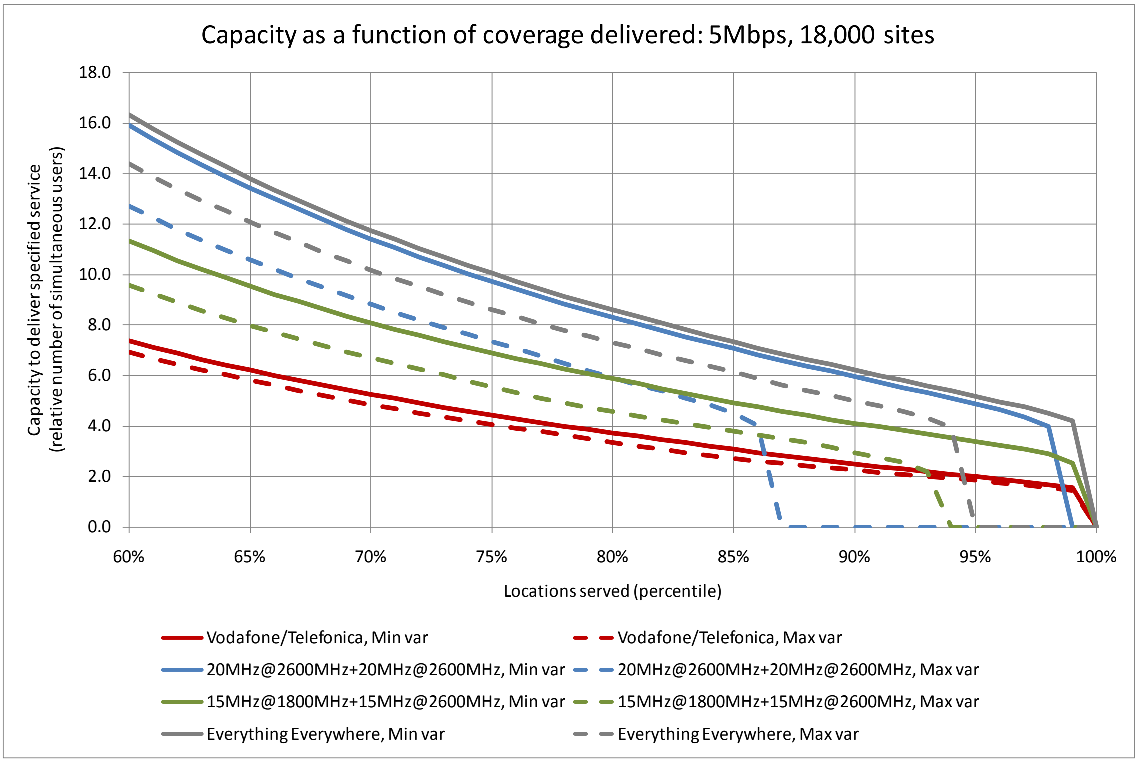

Figure 64: And some portfolios with even more bandwidth, but lower coverage quality: 5Mbps, 12,000 sites

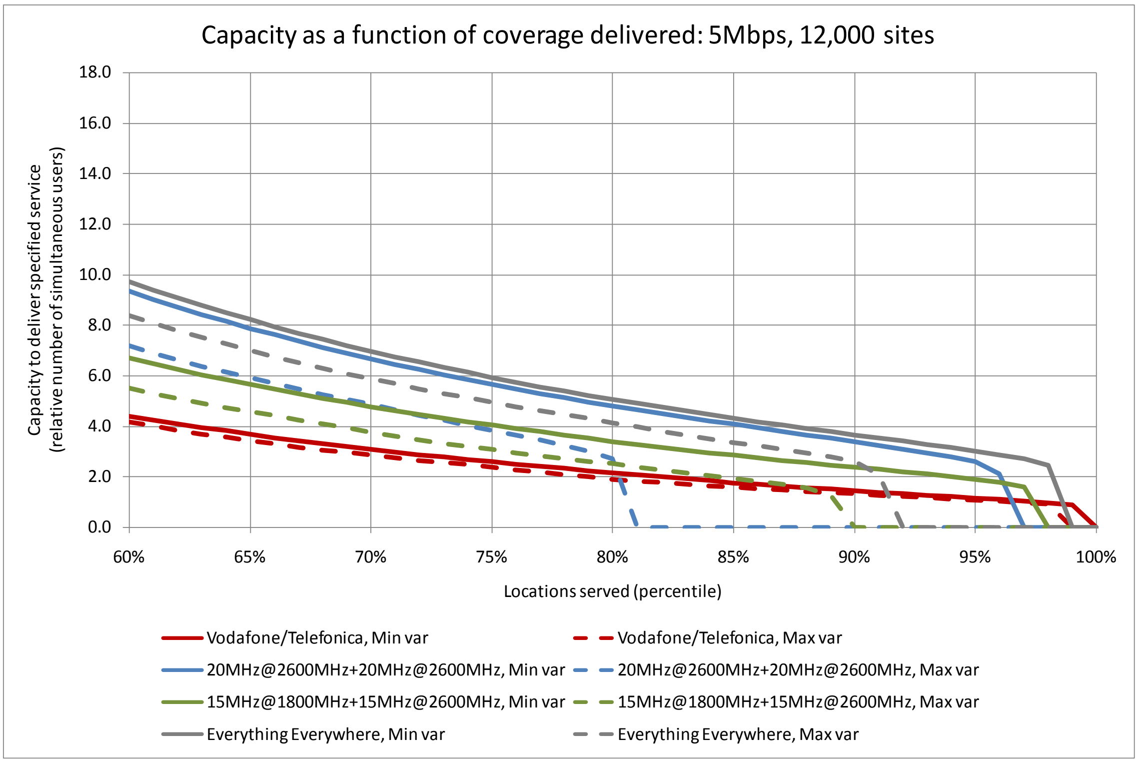

Potential additions to Everything Everywhere’s existing spectrum portfolio to improve quality of coverage

A7.127 Next we consider what additions to Everything Everywhere’s existing spectrum portfolio might help them to achieve a higher quality of coverage than their existing spectrum portfolio alone.

A7.128 If the service of interest was to be just a notional 2Mbps service then, as shown in Figure 65 and Figure 66, the addition of just 5 MHz of 800 MHz spectrum to Everything Everywhere’s existing spectrum portfolio would reduce if not entirely eliminate the uncertainty of Everything Everywhere’s coverage, ensuring that its coverage was comparable to that achievable by Vodafone and Telefónica.

A7.129 However, if the service of interest was to be a notional 5Mbps service, our model shows no improvement in the uncertainty of Everything Everywhere’s coverage as a result of the addition of just 5 MHz of 800 MHz spectrum (Figure 67 and Figure 68). However an additional 10 MHz of 800 MHz spectrum would once again reduce the uncertainty that Everything Everywhere’s coverage would match that achievable by Vodafone and Telefónica. In this case a further increase in the quantity of 800 MHz spectrum would not materially reduce the uncertainty of Everything Everywhere’s coverage further for delivery of a notional 5Mbps service, although it would increase the overall capacity.

Figure 65: Impact on Everything Everywhere's network coverage and relative capacity of 5 MHz of 800 MHz spectrum: 2Mbps, 18,000 sites

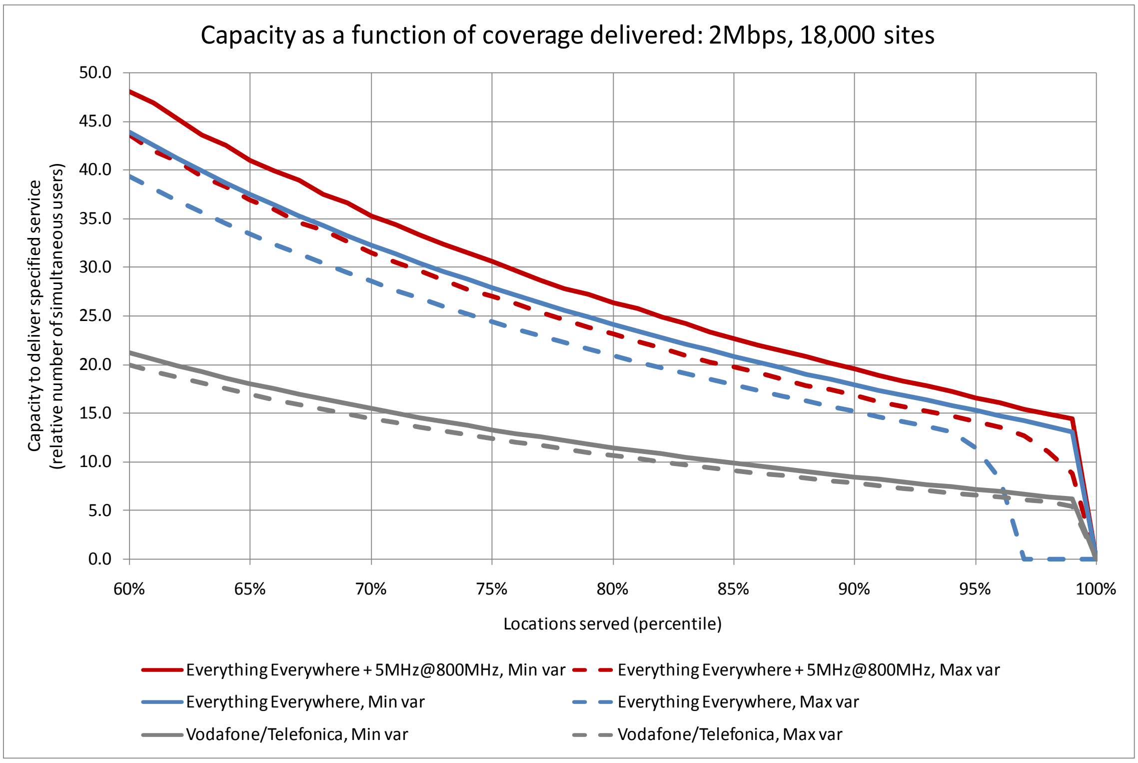

Figure 66: Impact on Everything Everywhere's network coverage and relative capacity of 5 MHz of 800 MHz spectrum: 2Mbps, 12,000 sites

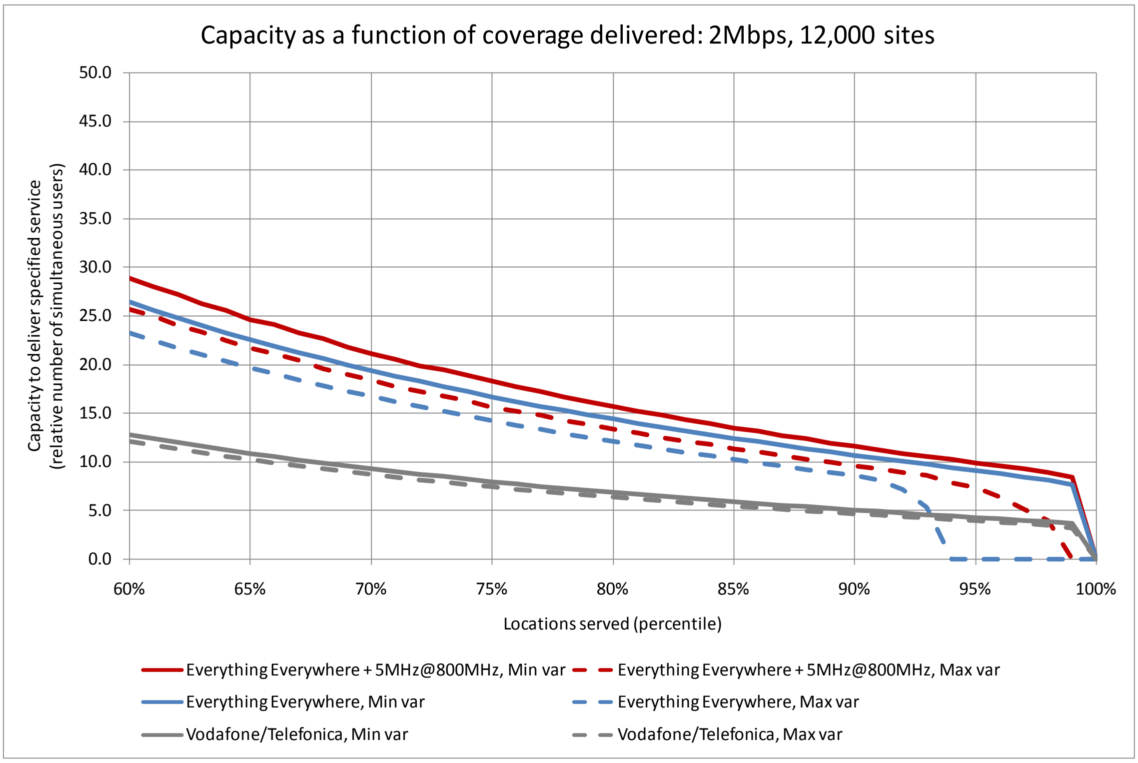

Figure 67: Impact on Everything Everywhere's network coverage and relative capacity of various amounts of 800 MHz spectrum: 5Mbps, 18,000 sites

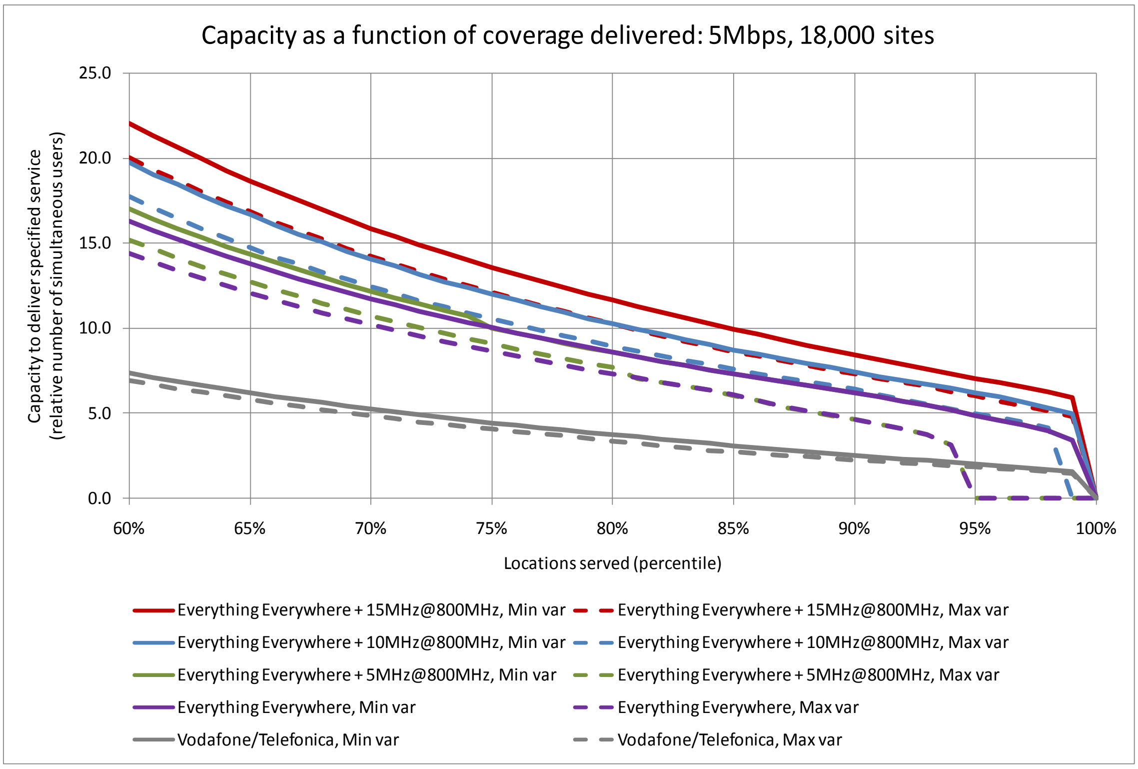

Figure 68: Impact on Everything Everywhere's network coverage and relative capacity of various amounts of 800 MHz spectrum: 5Mbps, 12,000 sites

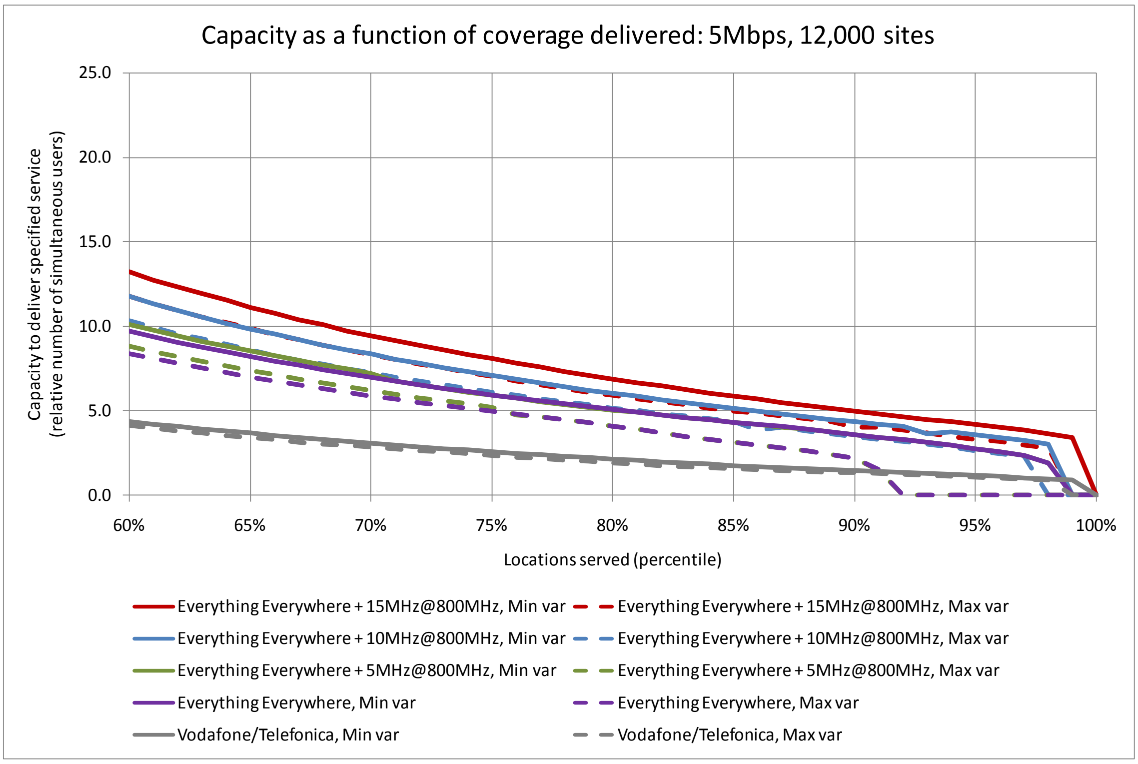

Potential additions to Vodafone and Telefónica’s existing spectrum portfolios to increase capacity

A7.130 Vodafone and Telefónica’s existing spectrum portfolios are predicted to provide materially less capacity than Everything Everywhere’s existing spectrum portfolio in the long term on a like-for-like basis (same service and same number of sites), for target coverages up to about 95% of locations. In this section we consider what additional spectrum would be required to bring the predicted capacity for each of Vodafone and Telefónica up to the same level as Everything Everywhere.

Figure 69: Impact on Vodafone/ Telefónica’s network relative capacity of 30 MHz of 2600 MHz spectrum: 5Mbps, 18,000 sites

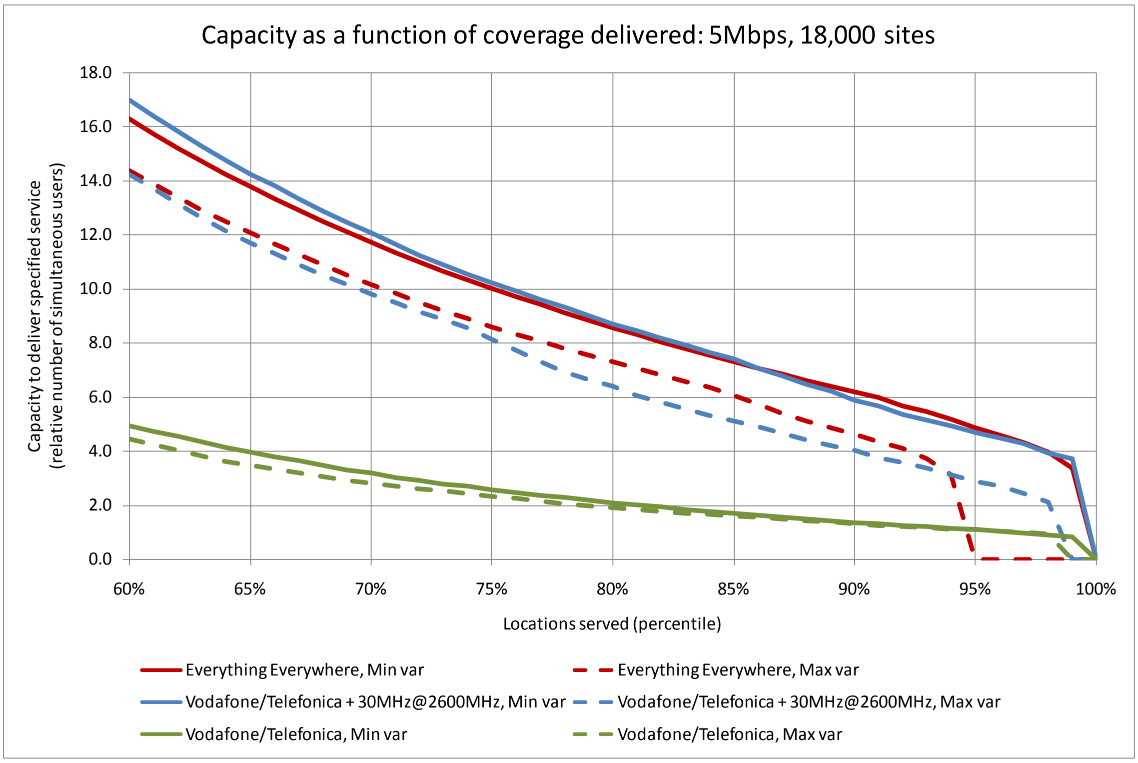

Figure 70: Impact on Vodafone/ Telefónica's network relative capacity of 30 MHz of 2600 MHz spectrum: 5Mbps, 12,000 sites

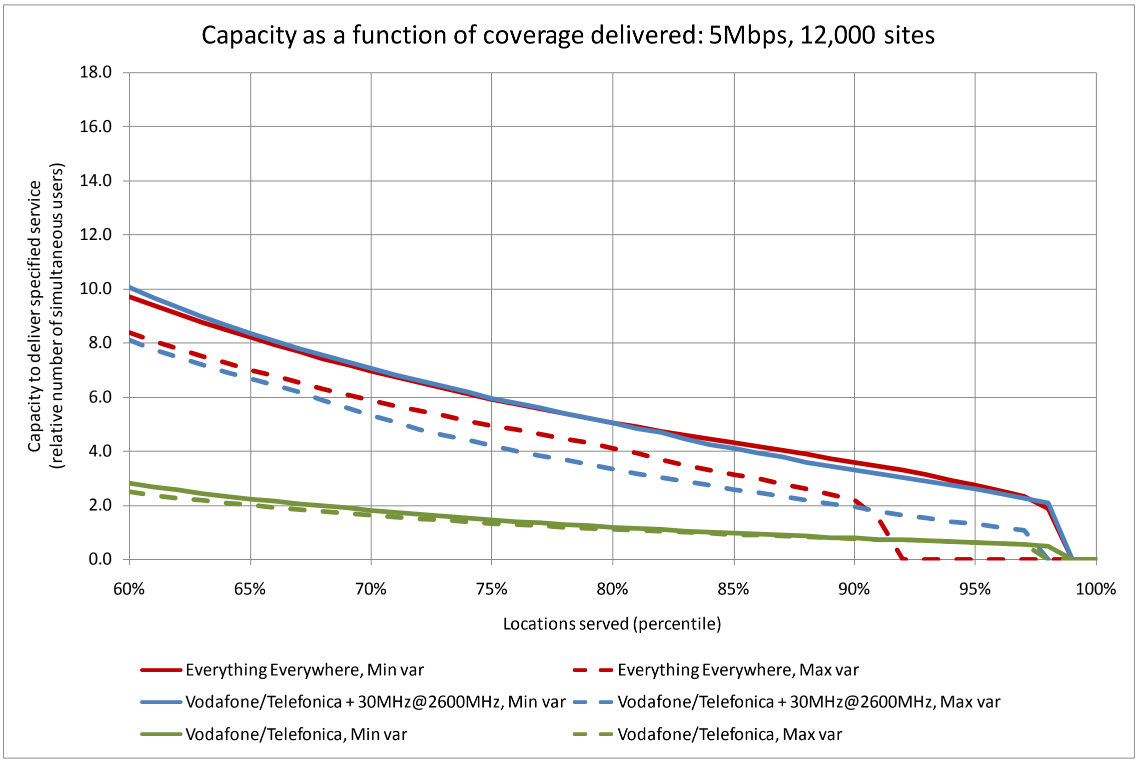

A7.131 In Figure 69 and Figure 70 we see that the capacity of a network that combined the existing spectrum holding of either Vodafone or Telefónica with 30 MHz of 2600 MHz spectrum is predicted to almost exactly match the predicted capacity of Everything Everywhere's existing spectrum portfolio in the long term, for all coverages up to about 90% of locations.

A7.132 If the additional spectrum were to be in lower frequencies then slightly less bandwidth might be required, for example Figure 71 and Figure 72 shows how only 25 MHz of additional spectrum would achieve much same result if 10 MHz of that were at 800 MHz.

Figure 71: Impact on Vodafone/ Telefónica's network relative capacity of 10 MHz of 800 MHz plus 15 MHz of 2600 MHz spectrum: 5Mbps, 12,000 sites

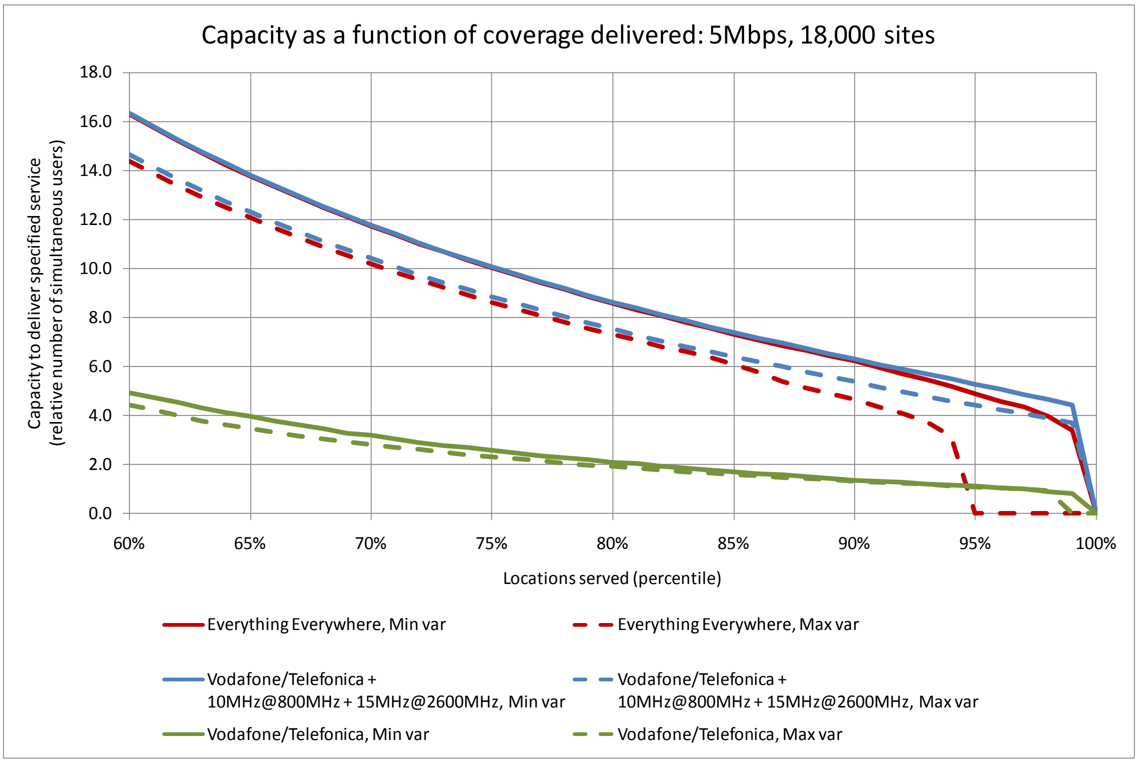

Figure 72: Impact on Vodafone/ Telefónica's network relative capacity of 10 MHz of 800 MHz plus 15 MHz of 2600 MHz spectrum: 5Mbps, 12,000 sites