be significant at the individual postal sector level in certain areas, but the impact of the errors decreases as these postal sectors are aggregated into larger areas, and is outweighed by the improved overall accuracy by applying the data uplift.

Step 4: BT Service Share Calculation

A6.22 Once the number of customer ends for each operator is determined, BT's share of the retail services is calculated in each postal sector.

A6.23 The service share bands that are used are as follows:

- 0% to 30%
- 30% to 40%
- 40% to 50%
- 50% to 70%
- Above 70%

Wholesale service share analysis

A6.24 Ofcom's intention when it devised the information requests sent to operators was to conduct a service share analysis for each of the relevant wholesale markets similar to that conducted for the relevant markets at the retail level and described above. However, the wholesale data received from providers in the Information Request has been such that it has not been possible to conduct the analysis as originally envisaged.

A6.25 Nevertheless, the data provided BT and OCPs relating to the wholesale markets have been such that it has been possible to conduct an analysis of the wholesale markets. As noted above, in the provision of TI products, there are two types of wholesale input, symmetric broadband origination (TISBO and AISBO) and trunk. Using the data provided by the operators, Ofcom has been able to carry out the analysis as set out below.

A6.26 The market definition boundaries between terminating and trunk markets are based on relevant network topology, and in particular we have used the location of network nodes to inform the break between terminating and trunk segments.

A6.27 For every circuit, based on data supplied by BT, it is possible to match the postcode information on the A and B-end of each circuit to a relevant parented BT Tier 1 node. We combine this Tier 1 parenting information to determine whether the circuit is likely to have a trunk component.

A6.28 Based on the 2003/04 Review definition, trunk segments were defined as providing transmission between two BT Tier 1 nodes. Based on the 2003/04 Review definition, it would therefore be assumed that circuits with ends parented on the same BT Tier 1 node would not include a trunk segment, even though in practice these circuits might be physically routed through trunk nodes. However, as set out in Section 6, instead of defining trunk based on BT's Tier 1 nodes, we have proposed to identify a set of 46 "aggregation nodes". As part of this analysis, we

have mapped BT's 67 Tier 1 nodes to 1 of the 46 proposed "aggregation nodes" based on the criteria we set out in Section 6.

A6.29 Therefore, circuits parented to the same Tier 1 node or to the same "aggregation node" would not include a trunk segment. This could result in some inter-Tier 1 traffic no longer being counted as having a trunk segment (i.e. where both Tier 1 nodes fall under the same aggregation node).

A6.30 We requested the following circuit information from operators:

- Retail circuits – all retail circuits;
- Wholesale Purchase – wholesale circuits purchased from BT and OCPs; and
- Wholesale Provision – wholesale circuits sold/provided to BT and OCPs.

A6.31 However, the information provided by the operators did not include circuits that are self-provided. Therefore the analysis of the TISBO and AISBO wholesale markets used the data available to derive the circuits that are self-provided. The wholesale data was compared against the results of the retail service share analysis by postal sectors to derive the wholesale service share.

A6.32 The wholesale market can be seen as the sum of the following components:

Total Market = BT Self-provide + BT to OCP + OCP Self-provide + OCP to OCP

(A) (B) (C) (D)

A6.33 Components (B) and (D) can be found from the data provided by the operators, while components (A) and (C) are derived by comparing the retail circuit information and circuits purchased from others.

A6.34 Therefore, the wholesale service shares for each operator is found using the following:

Wholesale = Retail – Wholesale Purchase + Wholesale Provision

A6.35 Similar to the Retail Service Share analysis, the methodology to estimate the operator wholesale service shares consists of the following four steps:

i) Data cleansing and normalisation;

ii) Aggregation by postal sector;

iii) Data uplift; and

iv) BT wholesale service share calculation.

Step 1: Data cleansing and normalisation

A6.36 Similar to the retail service share analysis, it was necessary for Ofcom to manipulate the raw data submitted by the operators into a structure suitable for Ofcom's intended analysis.

A6.37 The data cleansing and normalisation were carried out as follows:

- Circuits considered to be outside the leased line markets as defined by Ofcom are removed;
- The circuit bandwidths are checked, to ensure they are all consistently recorded in the same unit. The bandwidths are converted to a common format, expressed in Mbps (mega bits per second);
- Postcode correction was carried out to remedy the common detectable errors made when using automated batch processing techniques to record postcode data;
- Circuit end point analysis was carried out to identify the customer end-points and network end points. Where operators had not provided data on the type of end-point being served, any network end-points were identified using the operator flex point information that was provided;
- Trunk segment determination – the parent Tier 1 and aggregation nodes for the ends of each circuit are determined, using data provided by BT. Each circuit is then identified as having a trunk segment or not; and
- Extraction of the postal sector from the postcode data.

Step 2: Aggregation by Postal Sectors

A6.38 Similar to the step carried out in the retail service share analysis, the postal sectors are extracted from the postcode data for each circuit, as Ofcom opted to aggregate at the postal sector level for the geographic analysis.

Step 3: Data Uplift

A6.39 As with the retail analysis, an uplift factor was applied to the output data from Step 2 to align this with the known number of end points in each market calculated in Step 1 of the analysis:

Step 4: BT Service Share Calculation

A6.40 The wholesale service shares are calculated from the retail, wholesale purchase and wholesale provision information supplied by each operator. Once this is calculated, BT's wholesale service share in each postal sector is then determined.

Trunk service shares

A6.41 To assess circuit counts on trunk routes there are a few instances (on particular trunk routes) where our geographic data might suggest that the total number of trunk circuits that OCPs purchase from BT (or other CPs) is in excess of their total demand for trunk circuits (i.e. the trunk demand arising from the retail markets). This partly reflects the fact that some of our geographic information was not always

complete (as reflected in the need to apply data uplifts where there was missing data). This could also arise, for instance, where CPs rely on a particular trunk route to serve other destinations (e.g. the London Birmingham route might be used in combination with Birmingham Manchester as an alternative to building direct trunk route from London to Manchester).

A6.42 Therefore, when using the trunk data to estimate self-supply, this might suggest that on some routes OCPs are purchasing more circuits than they need to serve the retail demand for that route (e.g. if there were more wholesale sales to third than retail circuits sold with ends in London and Birmingham). In these circumstances, if we were to apply the formula in paragraph A6.32, this might result in a negative value for the implied amount of OCP self-supply (term (C) in the above formula). If we used negative OCP self-supply values in our calculation of service shares, this would tend to overstate BT's service shares on a particular route.

A6.43 To account for this issue, we apply adjustments to circuit counts to ensure that OCPs' total wholesale purchases on a particular route do not exceed their demand for trunk on that route, so that we can avoid negative values for self-supply in the above formula. For the relevant routes where this was an issue, we still have to estimate what proportion of the OCPs' total trunk demand on that route would be met either by BT or OCP wholesale trunk circuit sales. We therefore use the respective proportions of BT and OCP sales of circuits to third parties on that route.

A6.44 For example, if we consider a particular route where BT and OCPs in total sell 100 trunk circuits sales to third parties. If BT sells of these 75 circuits and OCPs the remainder 25 circuits, then the relevant proportions would be 75% and 25% respectively for BT and OCPs. We would then apply these proportions to the total OCP trunk demand to calculate how many circuits would be provided by BT or OCPs on that route. Therefore, if total OCP demand on the relevant trunk route were 80 circuits, we would assume that 60 circuits (75%) would be provided by BT and 20 circuits (25%) by OCPs.

A6.45 Combining this information with circuits that BT self-supplies, we can then calculate a more reasonable estimate wholesale service shares for individual trunk routes.

Network Reach Analysis

A6.46 As set out in Section 6, Ofcom has carried out an analysis of the network operators to assess the extent to which these operators can use their own networks to provide services, either at the retail or at the wholesale level. The network reach analysis can be used to inform an assessment of the extent to which the provision of business connectivity services in different geographic areas is contestable.

Data used

A6.47 To inform the network reach analysis, Ofcom used a variety of data collected following the Information Request from network operators with regards to their network information, and data purchased from Experian for the location of businesses in the UK.

A6.48 The Experian Business Database provided information on the location of businesses with 250 or more employees within the business. The 250 employee cut-off point was used since, based on discussions with industry, Ofcom is of the view that this is a reasonable proxy for the size of business that could be a potential

customer of leased line services. It is also the case that the cost of a leased line is less likely to be justifiable in the case of smaller businesses.

A6.49 Data on network reach were collected from each of the Other Communications Providers (OCPs) and BT. We requested the most recent data on their network infrastructure, geographically located by means of their postcodes / addresses / coordinates. This included the following information:

- Fibre network maps;
- Flex points or aggregation points: where existing fibre can be added to in order to connect to end-users. Flexibility points may well be buildings where fibre terminates on an Optical Distribution Frame or underground chambers where the fibre can be accessed, where ducts meet at a junction (etc). The fibre in the ground/duct would have to be added to by fibre-splicing and duct dug in order to connect an end-user premise to the fibre optic cabling;
- Points of interconnect with BT; and
- Points of interconnect with OCPs.

Overview of analysis

A6.50 The network reach analysis consists of the following parts:

- Flex Points analysis – count of OCPs' flex points for each postal sector;
- Contestability analysis – an assessment of the number of operators to which each large business location could seek supply, taking into account different economic build distance assumptions; and
- Interconnect analysis – the ability for operators to interconnect.

Flex Points Analysis

A6.51 The flex points analysis used geo-analysis software to plot each of the OCP's flex points information to calculate the number of flex points in each postal sector. This gives some sense of geographic variations in competitive conditions that could exist, as the areas of greater concentration of operator flex points are likely to generate a stronger competitive constraint than those areas where there is less concentration or no flex points present.

A6.52 However, to provide a more comprehensible picture of the geographic variations in competitive conditions, further analysis is required as the operator's flex points are not limited to providing services to end users located within a single postal sector. Depending on the size of the postal sector, the distribution of large businesses around flex points and assumptions on the economic build distance, an operator could provide services to end users in a number of different neighbouring postal sectors.

Contestability Analysis

A6.53 This analysis seeks to find the average number of operators that are able to provide services to end users in each postal sector. The main assumptions for this analysis relate to:

- the number of operators required to provide a sufficiently greater level of competitive constraint as compared with those areas where there is no choice of operator; and

- the economic build distance, which is the distance that an operator would build out from their network in order to provide services to end users.

A6.54 In order to illustrate the differences in competitive constraint that may exist, Ofcom's base case assumption is that in the circumstances of leased lines markets there would need to be at least two additional operators (i.e. at least 3 operators) in an area in order to provide a sufficiently different competitive constraint.

A6.55 The build distance assumption made in this analysis is that operators would be willing to extend their network by a distance of 200m[117] to serve a business customer. Ofcom recognises that this distance would vary on a case by case basis. However, an assumption on build distance has to be made for the purpose of this analysis.

A6.56 As the base case analysis uses information on flex points, which are located deeper into an operator's network, rather than on points of presence, a shorter economic build distance can be assumed.

A6.57 The analysis involves plotting geographically the location of all large business sites in the UK with more than 250 employees across the business, and comparing this with the geographic location of OCP's flex point information. The average number of operators per business location in each postal sector is calculated, assuming the build distance.

A6.58 Figure A6.2 below shows the results of the contestability analysis for Central London Zone (CLZ). This assumes a build distance of 200m and a constraint of at least two OCPs.

[117] In the July 2008 consultation we set out our reasons for using an assumed economic build distance of 200m as opposed to 250m as used in the January 2008 consultation. In Section 6 above we address comments received in response to the two consultations on this assumption.

Figure A6.2: Number of operators in the CLZ, assuming 200m build distance

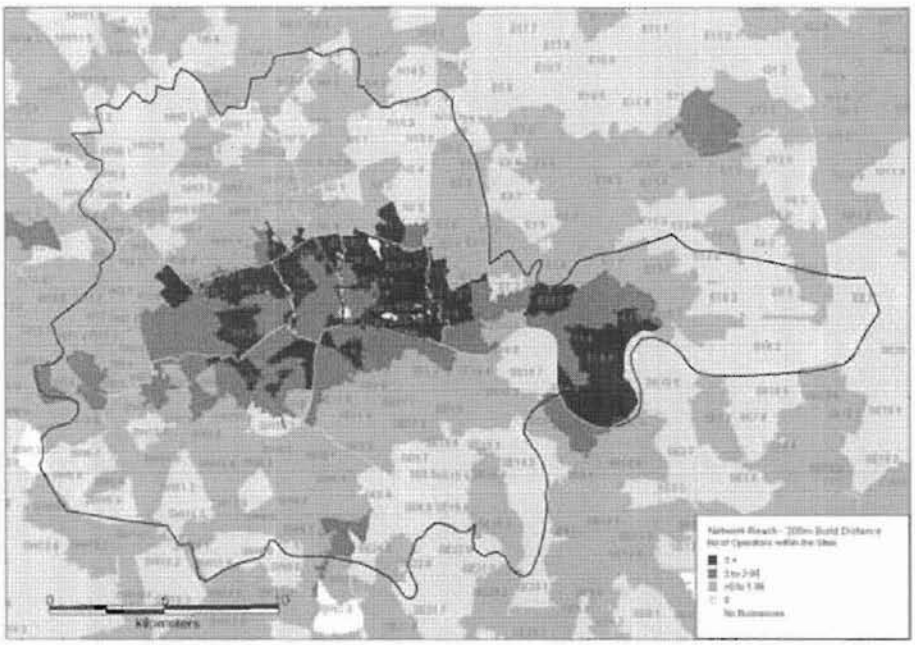

A6.59 From this figure it can be seen that a significant number of postal sectors within the City of London area appear to show that on average, at least three OCPs are able to provide leased lines services to businesses located within these postal sectors.

A6.60 This analysis was also conducted for other cities in the UK and the results show a similar pattern to that observable in London, as there are greater concentrations of operator flex points in the centres of these cities. However, the geographic coverage of any greater constraints that may exist is more limited, covering a fewer number of postal sectors.

A6.61 Figure A6.3 to A6.8 below show the results of the contestability analysis, assuming a build distance of 200m for major cities in the UK, based on population. These major cities include Birmingham, Glasgow, Liverpool, Leeds, Sheffield, Edinburgh and Manchester.

Figure A6.3: Number of operators in Birmingham, assuming 200m build distance

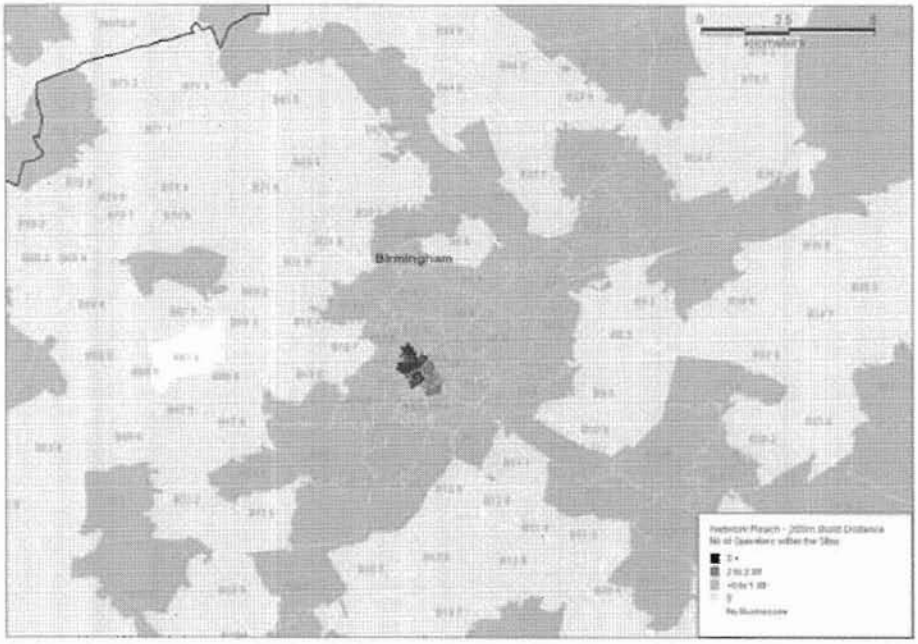

Figure A6.4 – Number of operators in Glasgow, assuming 200m build distance

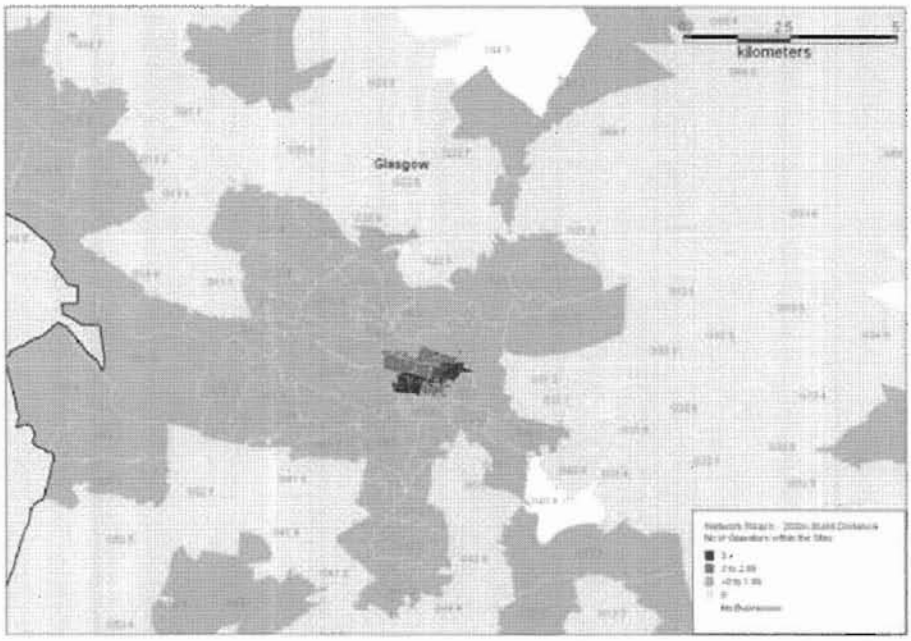

Figure A6.5: Number of operators in Liverpool, assuming 200m build distance

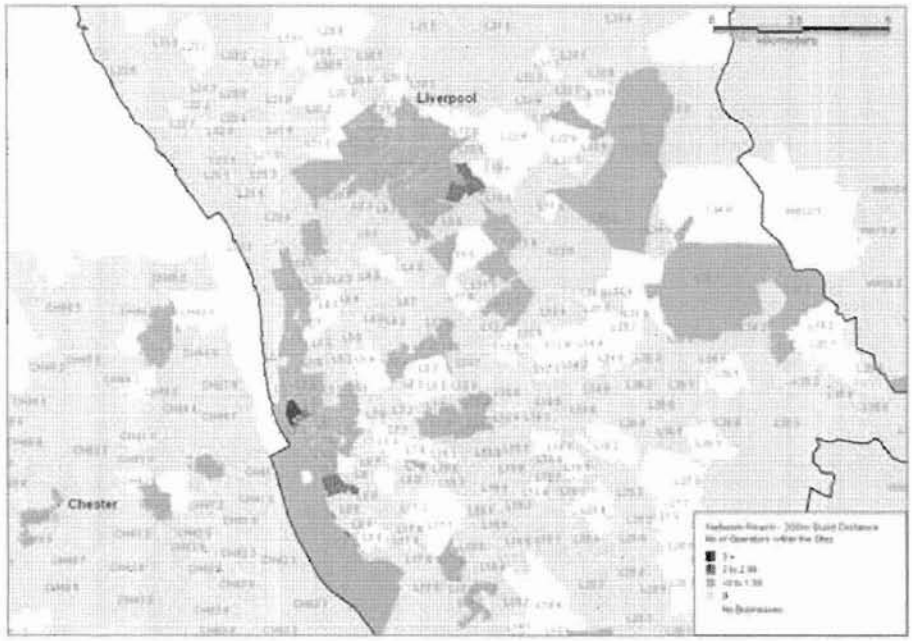

Figure A6.6: Number of operators in Leeds, assuming 200m build distance

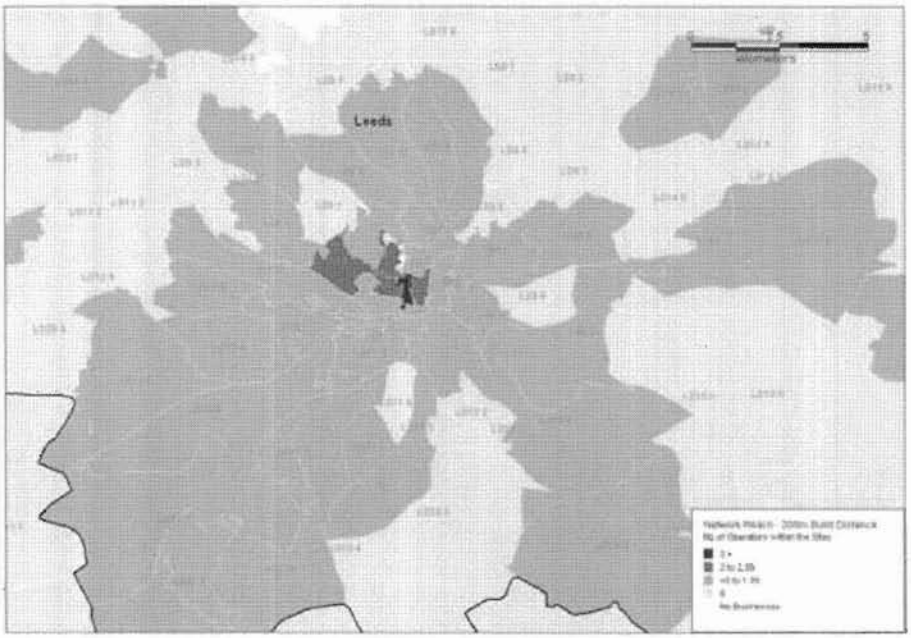

Figure A6.7: Number of operators in Sheffield, assuming 200m build distance

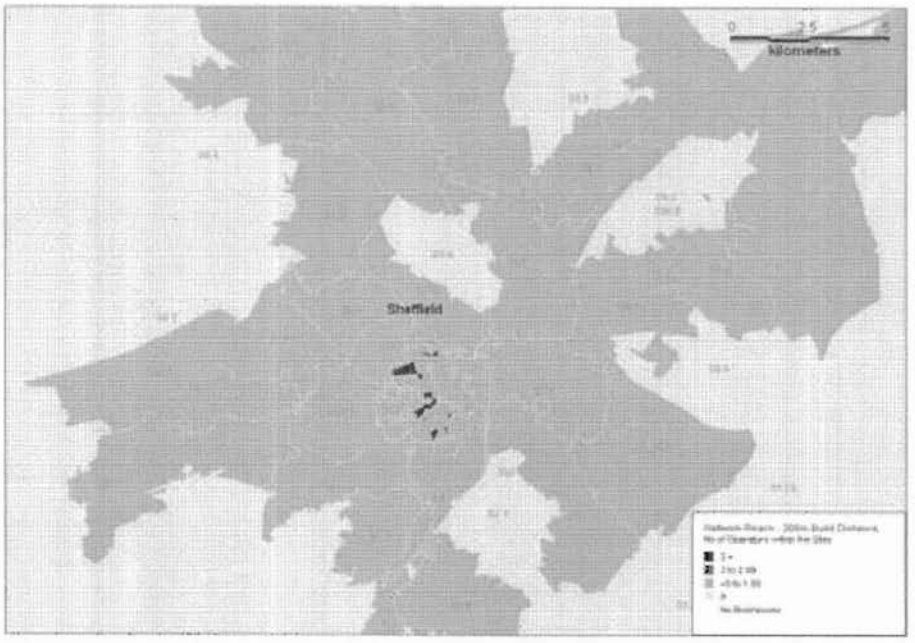

Figure A6.8: Number of operators in Edinburgh, assuming 200m build distance

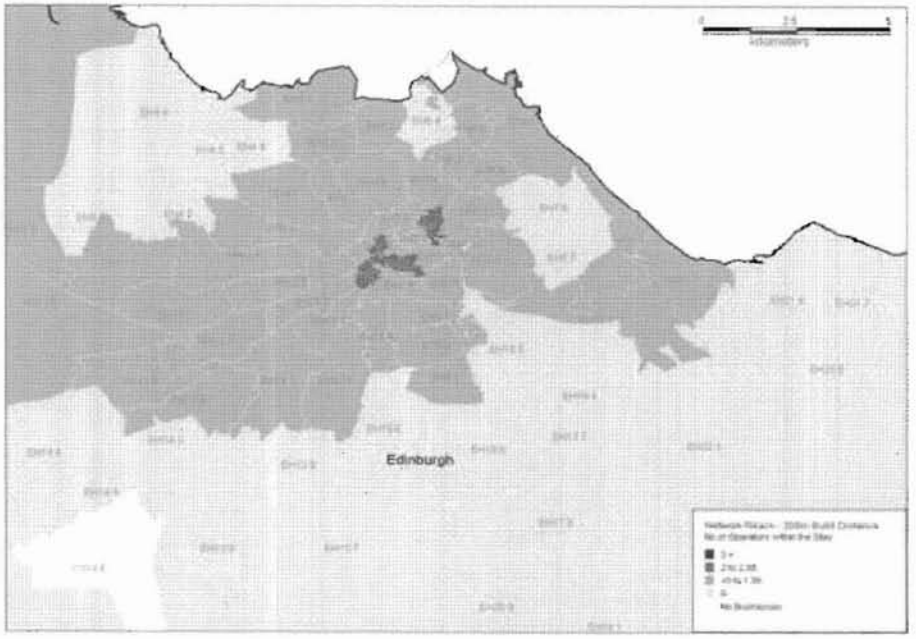

Annex 7

Aggregation nodes and geographic trunk analysis

Introduction

A7.1 In this Annex, we set out further analysis we have undertaken to finalise our views on the scope of the trunk market. This further analysis is in light of responses to the January 2008 consultation and considers two main areas:

- **London TI aggregation nodes:** in the January 2008 consultation, we proposed a single aggregation node covering the Greater London area. We have re-assessed this proposal and, based on the analysis below, we now conclude that we should identify eight aggregation nodes for TI trunk in the London area; and

- **The scope of the AI market:** in order to inform the scope of the AI market, we have also identified relevant AI aggregation nodes. Our assessment is that the aggregation nodes identified for the TI market provide a useful starting point (as these reflect the key areas where businesses are concentrated). There are some differences, however, that are likely to increase aggregation opportunities for the AI market. Therefore, we have identified an additional ten AI aggregation nodes based on further analysis.

A7.2 Prior to the detailed discussion of the above issues, we provide a brief reminder of the "aggregation nodes" concept and how we went about using this approach to define the break between trunk and terminating markets. We already covered many of these issues in Section 6, but we have repeated the discussion in this Annex. This is because it helps introduce some of the issues associated with identifying the aggregation nodes for the AI market and for the TI market in the London area.

A7.3 We therefore highlight the key concepts behind the aggregation nodes approach below. Following this discussion, we then discuss each of the steps we have followed to finalise our view on TI and AI aggregation nodes.

Summary of the aggregation nodes approach

A7.4 In the 2003/04 Review, we identified the breakpoint between symmetric broadband origination (terminating segments) and trunk segments based on BT's Tier 1 nodes. This definition also meant that "equivalent" nodes on other communications providers' networks identified the relevant breakpoint between origination and trunk markets services on their networks.

A7.5 In the January 2008 consultation, we considered, however, that we should revisit the 2003/04 Review trunk market definition, as we considered that it no longer captured sufficiently the key competitive differences between trunk and terminating markets. The evidence we looked at in the January 2008 consultation suggested that OCPs have not interconnected at all of BT's Tier 1 nodes. This was reflective of the fact that, in many cases and for many urban or business centres, there was often more than one Tier 1 node serving that location. Our assessment of the available aggregation opportunities in each geographic area suggests that OCPs are only likely to interconnect at a single point serving a particular area. Hence,

although OCPs typically interconnect at some of BT's Tier 1 nodes, the total interconnection coverage by an individual OCP is not likely to extend to all of BT's 67 Tier 1 nodes. As such, when OCPs purchase PPC terminating segment from BT this often included a short element between Tier 1 nodes to an OCP's relevant Point of Handover.

Why CPs are unlikely to interconnect at all Tier 1 nodes

A7.6 To see why a CP is unlikely to interconnect at all Tier 1 nodes, we consider an example below, which looks at the possible wholesale inputs used to supply a retail circuit from London to Reading. We have referred to this example (which was also used in the January 2008 consultation) to help explain the factors driving CPs' interconnection decisions and, having identified these factors, we then go on to explain how we used this information to analyse the break between trunk and terminating segments.

A7.7 As shown in Figure A7.1 below, in our example, we assume that an OCP has its own trunk capacity from its point of handover (PoH) in London to the Tier 1 node in Reading, so it can self-provide the trunk across this route. We have assumed that this OCP would require a terminating segment and would need to purchase this from BT. If the CP wanted to provide a retail circuit to the end-user shown in Figure A7.1, this would require a terminating segment to connect the London premises to the parent Tier 1 node (Tier 1 (b) in the Figure below). And as the CP in our example is interconnected (Tier 1(a)), it would also need to purchase a circuit connecting the two London Tier 1 nodes Tier 1(a) and Tier 1(b).

Figure A7.1: CP interconnection at Tier 1 nodes serving the same urban centres

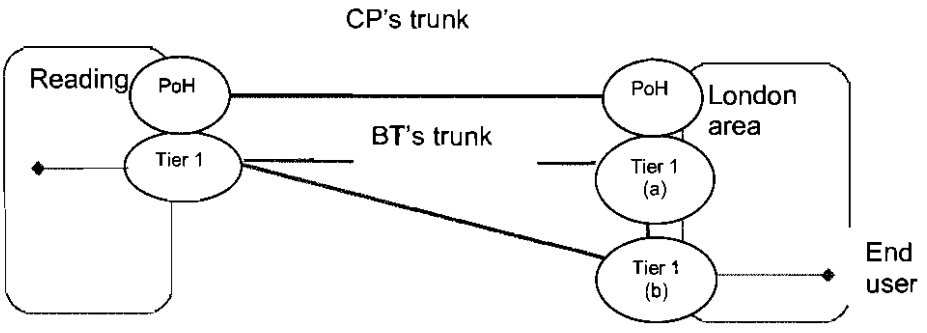

Source: Ofcom, November 2008

A7.8 In the above example, the CP could potentially compete to supply retail circuits between London and Reading using its own trunk capacity (from its PoH at Tier 1 (a)). Indeed, in order to exploit (and achieve) sufficient economies of scale on its trunk route, an OCP may well require that all of its traffic between urban centres is conveyed over a single high capacity trunk circuit if possible. This is because the benefits of using its own trunk capacity might only be realised if it can aggregate together sufficient volumes of circuits from all of its customers London area.

A7.9 In the January 2008 consultation, reflecting the above logic, we considered that we should define circuits linking major urban centres as trunk. Following the same logic, the competitive conditions in links between Tier 1 nodes in the same urban area are more likely to resemble those associated with backhaul markets. This reflects the fact that the opportunities to exploit aggregation opportunities are much

smaller on intra-urban links. In our example, the node at Tier 1(b) might serve insufficient end-users to further invest in an additional interconnection point. In addition, in order for a CP to able to provide competitive trunk (between London and Reading) it might also be necessary that the CP aggregate all of its London traffic over a single trunk link (i.e. trunk will only be competitive where the CP achieves sufficient scale).

A7.10 This means that CPs will only interconnect where sufficient aggregation opportunities are present to exploit the benefits of trunk circuits (i.e. a single PoH at Tier 1(a)). In our example above, the OCP would need to purchase wholesale circuits from BT to provide links from customer premises to its PoH. The links purchased from BT would comprise a circuit for the end-user to Tier 1(a) and also to a circuit to Tier 1(b).

A7.11 As stated above, the evidence we looked at in the January 2008 consultation suggested that OCPs were not interconnected at all Tier 1 nodes serving the same urban area. OCPs therefore often purchased a short trunk element between Tier 1 nodes to their relevant Point of Handover. These short trunk elements had characteristics closer to backhaul than trunk. Therefore, we thought that it would be more appropriate to treat circuits between Tier 1 nodes in the same area as TISBO rather than trunk.

Why the aggregation nodes better captured the break between trunk and terminating

A7.12 In the above discussion, we highlighted why, in general, CPs might only seek to interconnect at one node to serve a particular urban centre. In the January 2008 consultation, we thought that the simplest way to reflect this would be to group together Tier 1 nodes in similar locations. We referred to these groupings of Tier 1 nodes as "aggregation nodes". An OCP serving an urban area would need only one major point of handover within an aggregation node and so it would choose only one of the Tier 1 nodes that fell within a particular aggregation node.

A7.13 However, the process to identify of an appropriate set of aggregation nodes needed to be more involved than simply grouping Tier 1 nodes that happen to be located in the same urban area. Where traffic is highly concentrated, it may still make commercial sense for an OCP to connect at more than one node (even if those nodes are relatively close to each other). We therefore needed to find an appropriate methodology to group together relevant Tier 1 nodes to form our consolidated list of aggregation nodes.

We used proximity analysis to identify relevant Tier 1 groupings

A7.14 In the January 2008 consultation, we proposed to use "proximity analysis" to group together nodes within a certain "proximity" to another node.

A7.15 Our logic for grouping nodes based on proximity was that it reflected the relationship between circuit volumes and distance, which we identified as key determinants of CPs' interconnection decisions. Volumes are important, as the benefits of grooming traffic sooner onto the CPs own network will be higher the greater are the opportunities for the CP to aggregate that traffic to exploit the economies of scale associated with trunk. At the same time, the potential benefits of additional interconnection would also depend on the distances involved (between an existing interconnection point and a potential new interconnection location). If a CP only has to backhaul its traffic relatively short distances, the savings (in grooming its traffic onto its trunk network sooner) would be smaller.

A7.16 Therefore, the volumes/distances relationship is the key determinants of the potential benefits of interconnection, which a CP would compare to the costs of additional interconnection. For each Tier 1 node in the UK, the proximity analysis first considered the location of that node relative to other potential Tier 1 node interconnection points nearby and the volumes of circuits served by each Tier 1 node. Using information on the volumes of circuits sold at each Tier 1 nodes, we then determined an appropriate "proximity assumption" for each node.

A7.17 The proximity assumptions we used in the January 2008 consultation are described in detail in Section 6 of this document and in paragraphs 6.120-6.133 of the January 2008 consultation. In summary, we proposed three broad categories (i.e. based on the trunk traffic volumes originating or terminating at that node). For example for nodes serving a high-volumes of circuits (in excess of 1,500 circuits), we assumed that unless the CP was already interconnected at another node within a radius of up to 10 to 15km, it would be worthwhile locating at that node. For nodes serving a smaller volume of customers the "proximity assumption" would increase to 20 to 25km.

A7.18 For each node, we assessed whether there was another Tier 1 node nearer than the distance implied by the relevant "proximity assumption" for that node. If there was a Tier 1 node within reach of the node in question, we would group those Tier 1 nodes together as part of the same aggregation nodes. Hence, using volumes of circuits sold at each node and the relative distances between nodes, we could use our proximity assumptions to establish which Tier 1 nodes we might group together.

We used other information to verify our proximity assumptions

A7.19 The "proximity assumptions" we used in the January 2008 consultation were not based on precise calculation of operator costs and build decisions. Nevertheless, we sought to verify the proposed range of proximity assumptions (i.e. 10-15km for high volumes through to 20-25km for low volumes) against information that a CP provided on its build decisions for trunk (see paragraph 6.126 of the January 2008 consultation). The information the CP provided us allowed us to determine, for example, its view of likely circumstances where it would be economic to build additional trunk circuits (particularly the volumes and distances where trunk investment would be feasible).

A7.20 In parallel to this analysis, we also looked at where OCPs typically located their network and points of interconnection. Had the wider analysis of OCPs' interconnection suggested that they were present at nearly all Tier 1 nodes, then this could have suggested much shorter proximity assumptions. If OCPs were located at limited number of core nodes, this would have tended to suggest much longer proximity assumptions. In the event, our analysis broadly supported the aggregation node proposals generated using our proximity analysis.

We identified 40 aggregation nodes in our January 2008 consultation

A7.21 The proximity analysis and the analysis we conducted in parallel that looked at actual on CPs' interconnection decisions resulted in us identifying 40 aggregation nodes. In general, the aggregation nodes we identified grouped together BT's Tier 1 nodes in particular regions or urban centres. For example, for the London area, the proximity analysis resulted in grouping of 19 Tier 1 nodes into a single London aggregation node. This is shown in Table A7.1 below.

Table A7.1: Summary of Tier 1 nodes falling with London aggregation node (using 15km radius)

BT TIER 1 NODES WITHIN LONDON AGGREGATION NODE	
FARADAY	MAIN NETWORK ELTHAM RS
SOUTHBANK	WOOLWICH
BISHOPGATE	MILE END
COVENT GARDEN	POPLAR
MAIDA VALE	POTTERS BAR
MUSEUM	WOODGREEN SSC
COLINDALE SSC	CROYDON SSC
EALING SSC	KINGSTON SSC
HARLESDEN	WATFORD HERTS
ILFORD SSC	

Source: Ofcom, 2008

Further assessment of the London aggregation nodes

A7.22 In Section 6, we highlighted why, in light of the respondents views to our January 2008 consultation, there was a case to support additional aggregation nodes in the London area. In particular, there is a far larger density and volumes of traffic, which suggests that OCPs have been able to achieve the necessary economies of scale (and scope)[118] to justify further network build and interconnection at a number of locations within London.

A7.23 To assess this further, we have revisited the proximity analysis we undertook for the London area. We have also considered more detailed evidence on OCPs' interconnection.

Revised proximity assumptions for the London area

A7.24 As discussed above, we did not base the "proximity assumptions" in the January 2008 consultation on detailed explicit modeling of operator costs and build decisions. However, the proximity assumptions we used did appear to correlate with

[118] Economies of scope (in the context of the trunk market) refers, for instance, to the ability to aggregate together different bandwidth TISBO services onto a single high capacity trunk link.

information an operator provided on its trunk network build decisions. In parallel, we also sought to look at evidence on OCP interconnection to verify our results.

A7.25 However, the responses to the January 2008 consultation suggested that our proximity analysis did not capture in sufficient detail the particular aggregation opportunities available in the London area. To address this concern we have sought to model more explicitly the type of cost / benefit assessment that a CP might undertake when considering possible interconnection. We have then used this model to determine the relevant "proximity assumptions" for each of the 19 London nodes and for this information the relevant aggregation nodes.

A7.26 In the January 2008 consultation, we applied a single (10-15km) proximity assumption for "high volume" Tier 1 nodes (and we applied the same assumption for any nodes in the London area). Our revised approach provides a more granular view of appropriate proximity assumptions in the London area, with an individual proximity assumption calculated for each Tier 1 node serving the London area. We think that this more detailed view is warranted because the London area accounts for approximately one third of all trunk circuits sold. On further examination of the traffic at different Tier 1 nodes in London, we have also found that there is also quite a large difference in traffic concentration across different London nodes. As such a uniform "proximity assumption" applied to each of the London Tier 1 nodes would not capture the different aggregation opportunities that may exist at different nodes.

A7.27 To ensure that our results are not reliant solely on our revised proximity assumptions for the London nodes, we have also sought to verify our results by assessing observed levels of interconnection in the London area. We have also compared our results to BT's own future roll-out decisions on its 21CN network, which it will use (in part) to support leased lines traffic alongside other traffic streams such as broadband and voice. This latter information is relevant, as the design of BT's new network has provided it with an opportunity to re-assess the optimal location of nodes.

Determining revised "proximity assumptions" for the London area

A7.28 As stated above, the aggregation nodes we identified seek to capture the key factors differentiate trunk from terminating segments. In particular, the trunk market reflects the bulk transfer of traffic on high capacity routes (typically between major urban centres) to exploit economies of scale and scope. By contrast, we have identified terminating segments in the traditional interface markets are likely to be less competitive due to more limited aggregation opportunities.

A7.29 In the most part, we suggested that an OCP would use a single node to serve a particular urban area. But for the London aggregation node, however, there are reasons why the concentration of traffic makes it worthwhile to have more than one aggregation node. Take for example an OCP that has initially built out its trunk network to a location in Central London (shown by Node A in Figure A7.2 below).

Figure A7.2: Relative costs of trunk and backhaul

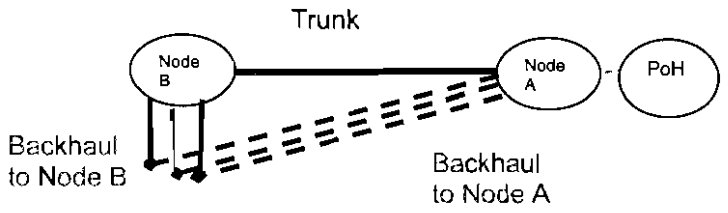

Source: Ofcom, November 2008

A7.30 If we assume that a CP is located at Node A, it may consider, for instance, whether to invest at an additional interconnection point (Node B) in London. If we find that it is not worthwhile to invest at both Node A and Node B, then both are regarded as part of a single aggregation node (hence the circuit between those two nodes would be a terminating segment). On the other hand, if aggregation opportunities suggest interconnection at both nodes is worthwhile, this is indicative that the circuit between Nodes A and B is trunk.

A7.31 When deciding whether to interconnect at more than one node a CP would need to consider the potential lower costs of trunk (particularly whether sufficient aggregation opportunities exist) compared to backhauling a circuit over longer distances (i.e. backhauling to Node A).

A7.32 The backhaul costs saved (i.e. backhauling from Node B to Node A) depend on the relative distances involved. If Node A is quite close to Node B, a CP would not have to backhaul those circuits much further from its retail customer sites to get to its existing point of interconnection at Node A (than if it had interconnected at Node B). The cost and benefits would also depend on whether it is cheaper to use a trunk circuit rather than backhauling the additional distance to Node A.[119] It will only be cheaper to use trunk capacity if there are sufficient volumes of traffic, given the larger required economies of scale associated with trunk.

A7.33 The discussion suggests that the benefits of interconnection depend on there being sufficient volume of traffic served at the possible new interconnection point (at Node B) and sufficient distance between the nodes to make any saving in backhaul costs worthwhile.

Expressing interconnection decision as a break-even relationship

A7.34 An alternative way to express the relative costs and benefits implied by the above interconnection decision would be in terms of a break-even relationship. This would entail identifying, for a given distance between two potential nodes (i.e. between Node A and B), the required volume of circuits (at Node B) to make any investment in an additional interconnection (at Node B) worthwhile (assuming that a CP is already located at Node A).

[119] We are seeking define the possible scope of the trunk and terminating markets (i.e. whether a terminating segment should be defined as potentially being longer than the nearest parented Tier 1 node and includes circuits between Tier 1 nodes in the London area). It would not be appropriate to assume that the regulatory requirement exists to provide trunk or terminating segment to particular locations, such a regulatory requirement would arise from a finding of SMP finding, and as this finding could potentially be affected by where we identify the boundary between trunk and terminating segments, we should assume the absence of regulation (for the purposes of market definition).

Figure A7.3: Break-even for additional trunk interconnection

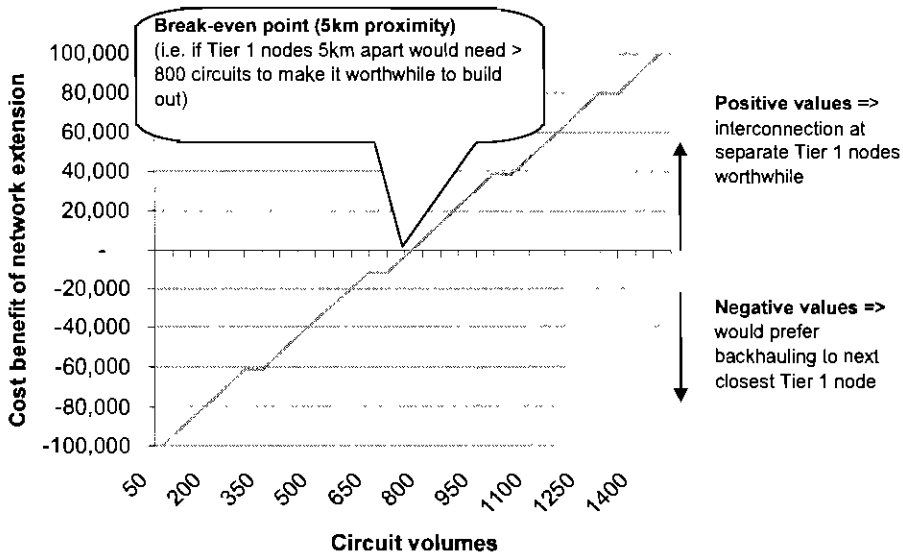

Source: Ofcom, November 2008

A7.35 We have calculated the above break-even relationship using estimates of the relative costs of additional interconnection and trunk (which are largely fixed) and the ongoing rental costs of terminating circuits (for those interested in the precise details, this is discussed in Box A.1 below). Figure A7.3 shows that a CP considering interconnection at two nodes within 5km of each other would require it to sell a sufficient number of circuits (at least 800 circuits) at the new interconnection point (Node B) for this investment to break-even. In Figure A7.3 above, the point on the x-axis where the line crosses the y-axis would represent the "break-even" volume. Hence, an additional interconnection investment would only break-even for volumes of around 800 circuits (and above).

A7.36 We can also use the above cost-benefit comparisons for other interconnection scenarios (i.e. where there are greater or smaller distances between Nodes A and B) to calculate the relevant break-even volumes. We can then plot this as a distance/volume "break-even" relationship, as shown in Figure A7.4 below.

Figure A7.4: Break-even for additional trunk interconnection

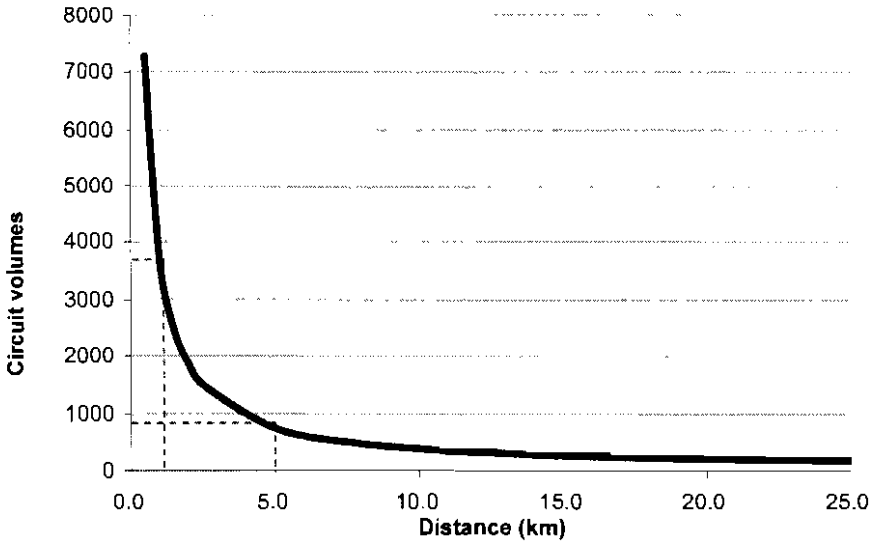

Source: Ofcom, November 2008

A7.37 Figure A7.4 shows that the closer (farther) network nodes are to each other the greater (smaller) the volume of circuits required to make it worthwhile to interconnect at both nodes. The key point to this analysis is that we can use this volume/distance relationship to determine "proximity assumption" for each of BT's Tier 1 nodes in London. Hence, for a given volume of circuits at a particular node (Node B), using the above break-even analysis, we can determine whether it would be worthwhile to interconnect at that node if an OCP already had an existing interconnection point nearby. For example, as shown by the red dotted line in Figure A7.4 above, if an OCP sold 800 circuits at Node B then it would only be worth interconnecting at Node B, where its existing interconnection point (Node A)) is farther than 5km this node (i.e. at a point to the right of purple line). On the other hand, if Node A is less than 5km to Node B, it would not be worth interconnecting both at Node A and at Node B.

A7.38 We have therefore used the above model to derive relevant "proximity assumptions" using the volume of circuits served by different Tier 1 nodes in London. Given the information we have on other London Tier 1 nodes in close proximity to each Tier 1 node, we can then compare this with our "proximity assumption" for each node to consider whether a CP would be likely to have a point of interconnection at a particular Tier 1 node (given it may already have an interconnection point nearby).

Assumptions used in our model

A7.39 The above model cannot provide a definitive view of the relevant cost and benefits as it depends on various assumptions regarding the relative costs of trunk versus terminating segments. We have explained in more detail our key assumptions in Box A.1 below. We think that the assumptions we have used as inputs to our model are sufficiently robust for it to generate an accurate view of the key likely interconnection points (given information on the proximity of nodes and the volumes of traffic).

A7.40 However, the view we come to on potential aggregation nodes has not relied solely on this modelling exercise. To ensure that our outputs are sufficiently robust, we have also followed the approach we adopted in the January 2008 consultation by examining, in parallel, wider evidence on actual CP interconnection and operator build decisions in the London area. We have also conducted sensitivity analysis to understand the effect of changes to our assumptions of the relative costs of trunk and terminating segments on our view of relevant London aggregation nodes, which we discuss further in paragraphs A7.59 to A7.64 below.

A7.41 Given that the results generated by our "bottom-up" modelling appear to correlate quite well with a "top-down" assessment of actual interconnection, we conclude below that it is appropriate to identify a number of aggregation nodes rather than a single node (as proposed in our January 2008 consultation).

Box A.1: Key assumptions used

This Box sets out some of the key assumptions underlying our break-even model. As discussed above the break-even model is based on a relative simple network build model. This relies on a comparison of the costs of backhaul to the nearest node (i.e. Node B in Figure A7.2 above) plus the cost of trunk and any additional Point of Handover equipment (i.e. interconnection equipment) versus the costs of backhaul to the next nearest node (i.e. Node A in Figure A7.2 above).

The final output of the break-even model is a relationship, which shows for a given volume of circuits at a particular node whether it would be economic to interconnect at that node. In deriving this model, we have had to rely on estimates of the relative costs of trunk and terminating segments and made some simplifying assumptions regarding the nature of demand for traditional interface services.

Modelling trunk costs:

- We used estimates of trunk costs using BT's published WDM costs. This includes the relevant split between relevant fixed elements and distance related elements.

Terminating costs:

- In modelling terminating costs, we have assumed that a number of terminating costs would be the same no matter which London node a CP would backhaul its circuits to (e.g. access related costs such as fixed local end costs)

Distance related terminating costs:

- The break-even calculation assumes that if a CP does not interconnect at a nearer node (Node B), it would not have the opportunity (in the absence of regulation) to purchase regulated trunk. Therefore, one of the costs of not interconnecting at the nearer node (Node A) is that it would have to purchase longer distance terminating segments to the next nearest Tier 1 node.

Additional interconnection costs:

- The key costs of locating at Node B would relate to the costs of an additional interconnection point. We have therefore used relevant published BT prices for its Customer Sited Handover charges to reflect the costs of interconnection.

Retail and wholesale circuit demand

- As a simplifying assumption, we have assumed that all customers demand 2Mbit/s circuits. This may understate the total bandwidth demand since circuit volumes are based on all relevant bandwidths. However, as low bandwidth circuits tend to dominate, we think that this assumption is reasonable.

Using the above break-even model to identify aggregation nodes

A7.42 As stated above, the above break-even model provides an appropriate "proximity assumption" for each of BT's Tier 1 nodes. As we have available data on the volumes of circuits currently sold at a particular Tier 1 nodes and the proximity of Tier 1 nodes to each other, it is possible to use the above break-even information to determine the possible Tier 1 nodes we should group together as aggregation nodes.

Table A7.2: Proximity assumptions for each node

TIER 1 NODE	Estimated circuit sales	Proximity Assumption for node (km)	Distance to nearest node (km)	Group with another node?	Nodes within "proximity"
FARADAY	1395	2.5	1.5	Yes	Bishopsgate, Covent Garden, Southbank
BISHOPSGATE	1467	2.5	2.0	Yes	Faraday
COLINDALE	225	20	5.5	Yes	All other nodes (except Croydon, Eltham, Ilford, Woolwich)
COVENT GARDEN	937	5	1.5	Yes	Faraday, Bishopsgate, Museum, Southbank
CROYDON	336	10	13.9	No	-
EALING	562	7.5	5.3	Yes	Harlesden
HARLESDEN	218	20	4.1	Yes	All nodes (except Croydon, Ilford, Eltham, Woolwich)
ILFORD	441	10	7.5	Yes	Mile End, Poplar, Woolwich
KINGSTON	385	10	11.5	No	-
MAIDA VALE	173	25	3.8	Yes	All other nodes
MAIN NETWORK ELTHAM	305	15	4.4	Yes	Faraday, Bishopsgate, Covent Garden, Croydon, Ilford, Mile End, Poplar, Woolwich
MILE END	121	25	1.6	Yes	All other nodes (except Watford)
MUSEUM	1101	5	1.6	Yes	Faraday, Bishopsgate, Covent Garden, Maida Vale
POPLAR	2301	1.5	1.6	No	No other nodes within proximity
POTTERS BAR	170	25	12.3	Yes	All nodes (except Croydon, Kingston, Eltham, Woolwich)
SOUTHBANK	950	5	1.9	Yes	Faraday, Bishopsgate, Covent Garden, Museum
WATFORD HERTS	364	10	13.3	No	No other nodes within proximity
WOODGREEN	415	10	8.7	Yes	Faraday, Bishopsgate, Colindale, Covent Garden, Maida Vale, Mile End, Museum
WOOLWICH	154	25	4.4	Yes	All nodes (except Ealing, Kingston, Potters Bar, Watford)

Source: Ofcom, November 2008

A7.43 Table A7.2 shows data on the 19 Tier 1 nodes in the London area. It shows the potential circuit sales for each of the 19 London nodes for a "typical" CP (based on estimated service shares). From these volumes, we have determined an appropriate "proximity assumption" for each node using our break-even model. Finally, based on the location of each node relative to other Tier 1 nodes, we have then determined whether or not to combine Tier 1 nodes. For each node, we have considered whether the distance to other nodes is less than the "proximity assumption" for the node in question. If the distance to the nearest node is less than our proximity assumption then this suggests (in principle) that we could group these nodes together. Each of these steps is discussed in turn below.

We adjust volumes at each node to account for potential CP service share

A7.44 It is important to note that for each Tier 1 node, we are seeking to determine whether an average competitor to BT might interconnect at more than one location in the London area. Clearly, on average, a potential competitor to BT would not achieve 100% share of circuit sales at particular Tier 1 nodes. Therefore, the volumes we have used to assess an interconnection decision for an individual CP is based on the CP in question achieving a 10% share of any particular Tier 1 catchment area.

A7.45 For example, using data on the Woolwich node, approximately 1,537 wholesale circuits are sold at this node (i.e. from to an end-point in the Woolwich catchment area to an end-point elsewhere in the UK). So we would assume that on average a competitor to BT might only sell 150 circuits in this catchment area.

We use these adjusted volumes to derive proximity assumptions for each node

A7.46 Based on the estimate volume of circuits a typical CP might sell at each node, we can determine an appropriate proximity assumption for each node. For example, give that the Woolwich node might only serve around 150 circuits, we estimate that it would only be worthwhile interconnecting at the Woolwich node if the distance to the next nearest node was above 25km. Therefore, this result would not suggest Woolwich would be a major interconnection point in the London area.

We compare proximity assumptions for each node to other Tier 1 node locations

A7.47 The next step in the analysis is to compare the proximity analysis to the information we have on other potential interconnection points nearby. For example, if we compared the Woolwich node with the location of other Tier 1 nodes, there are four other nodes nearby to this node, namely: Eltham, Poplar, Ilford and Mile End (between 4.2 km to 8.2km). Given there are four nodes between 4.2 to 8.2km, this analysis would not suggest identifying (in the first instance) Woolwich as a separate node if a CP were already located at another of these nodes. This is because these other nodes (as well as other London nodes) are within the 25km proximity assumption for Woolwich. In other words, the relatively low volumes of circuits and the proximity of other nodes does not suggest it is worthwhile investing in additional interconnection just to serve customers located at the Woolwich node.

Determine the optimal Tier 1 node groupings to determine aggregation nodes

A7.48 The above proximity analysis provides a view for each Tier 1 node of the other nodes that could in principle be grouped with this node. From the analysis in Table A7.2 above, it was possible to come to a firm view on an initial set of aggregation nodes. This is because there are a number of nodes, such as Watford, Croydon, Poplar and Kingston where our proximity analysis suggests that a CP would locate, irrespective of whether it is located at another nearby node. For example, Table A7.2 shows that the proximity assumption for Watford is 10km and no other Tier 1 node falls within this radial distance. We show this diagrammatically in Figure A7.5 below.

Figure A7.5: Comparison of "proximity assumptions" against Tier 1 node locations

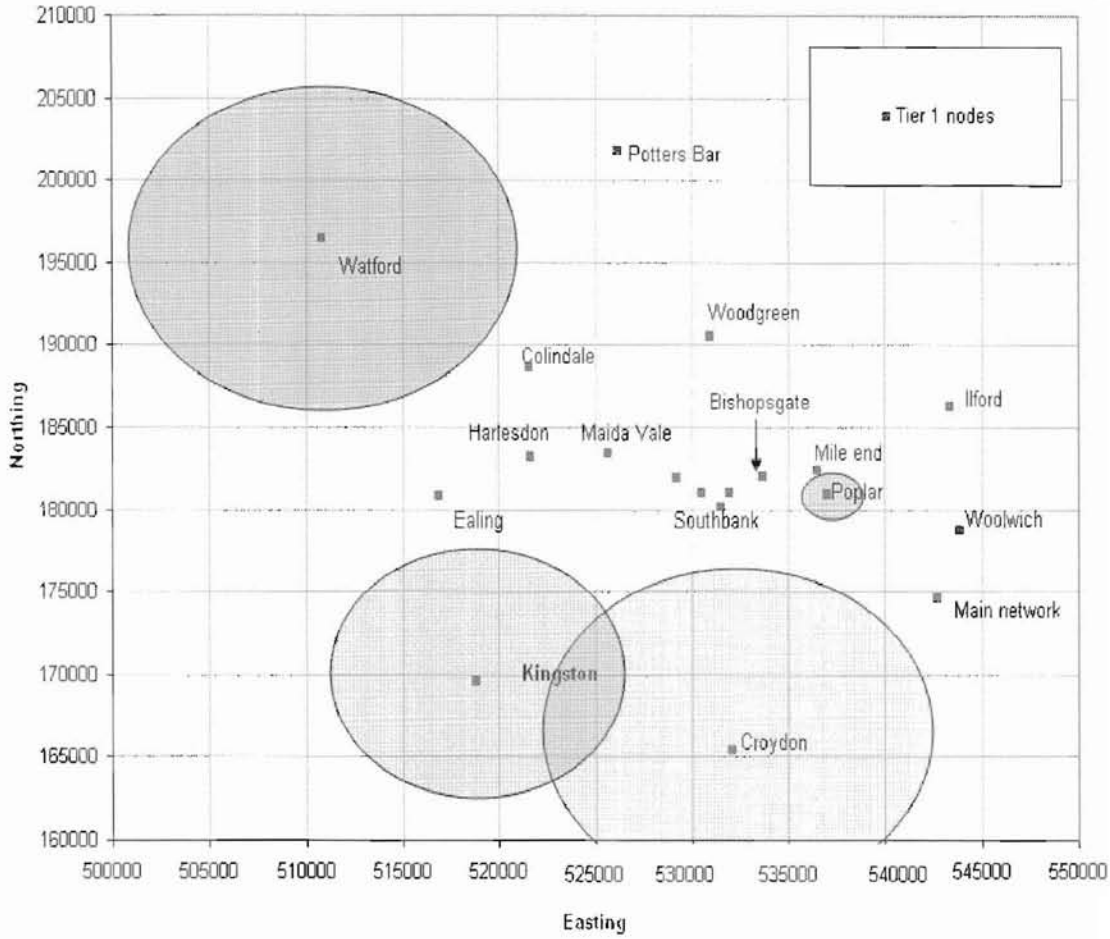

Source: Ofcom, November 2008

A7.49 In Figure A7.5, we show the nodes that do not have another Tier 1 node within "proximity" to that node. In Figure A7.5 we have shown the (approximate) radial distance around each node, reflecting the proximity assumptions for certain BT Tier 1 nodes. The different radius of each circle reflects the fact that we have different proximity assumptions for each Tier 1 node. Therefore, the above Figure suggests, as a starting point our revised analysis suggests a more disaggregated view of the London market than proposed by our January 2008 consultation. In other words, more than one aggregation node should be identified for the London area (represented, at the very least, by the above nodes).

A7.50 On the other hand, relative to the existing 2003/04 Review definition our analysis points to a consolidated list of aggregation nodes than the list of 19 Tier 1 nodes. In most cases, the above analysis suggests that there would be at least one other Tier 1 node within "proximity" to that node (see column 6 of TableA7.2 above). This suggests that we should group most Tier 1 nodes with at least one other node (and potentially more than one).

A7.51 We highlighted above that we can come to a firm view on an initial list of aggregation nodes (Croydon, Kingston, Poplar and Watford). But, even if a CP

located at each of these points, it would not necessarily follow that it would backhaul all other traffic served in the London area to one of those interconnection points. In other words, it would also be optimal for a CP to locate at another interconnection point (in addition to Croydon, Kingston, Poplar and Watford).

A7.52 Therefore, in addition, to our initial list of potential aggregation nodes, we had to consider a number of alternative additional points, where it could also be optimal for a CP to locate. Deciding which additional Tier 1 nodes formed aggregation points was a potentially complex task, as some groupings of Tier 1 nodes would be more efficient than others. We therefore needed a way to find the lowest-cost grouping from all potential groupings.

A7.53 For example, the proximity analysis would suggest that we might group Harlesden and Faraday, where a CP chose to interconnect at this latter node. There are however numerous other nodes that might be within proximity to the Harlesden node (in fact only Croydon, Ilford, Eltham, Woolwich were not within proximity of this node). A CP might choose to interconnect at one of these other Tier 1 nodes within proximity to the Harlesden node. If this node were nearer than Faraday, then it is likely to be optimal for the CP to backhaul to this new interconnection point rather than groom its traffic via the Faraday node. For example, in the case of Harlesden, if a CP also chose to interconnect at the Ealing node, it could be optimal to backhaul circuits served by the Harlesden area to this new point.

A7.54 The decision over which nodes to group together was therefore a complex task. Many Tier 1 nodes could potentially fall within a number of other nodes to form aggregation nodes, so there were a number of permutations to consider. This additional layer of complexity meant that we had to go through a number of iterations to identify appropriate number of aggregation points, and to find a final list that resulted in the "optimal" grouping of nodes. We discuss below the process we used to identify this final list.

Identification of the final list of aggregation nodes

A7.55 We have explained below each stage of this analysis we used to find the final list of aggregation nodes. For each Stage, we have explained the purpose of the analysis; an example of how we applied each stage; and the final aggregation nodes implied by each stage.

- **Stage 1:** we took as our starting point the nodes where the proximity analysis definitively suggested a CP would interconnect. As shown in Figure A7.5, these were the nodes where there are no other Tier 1 nodes within "proximity", so that no matter where else in London a CP might choose to interconnect, a CP would still seek to have an aggregation point at these initial nodes.

Output from Stage 1: Croydon, Kingston, Poplar, Watford

- **Stage 2:** based on the nodes identified under stage 1, we then identified any nodes that we thought we could group with the initial nodes. At this stage, however, we only identified one node (Mile End) that we would unambiguously group with one of the initial nodes. For other nodes there was the possibility that there could be another node nearer than one of the "initial nodes" under Stage 1. Therefore, at Stage 2 in our analysis we only grouped together nodes that where we were certain this would not be affected by the identification of other