BT’s response to Ofcom’s consultation document

“Business Connectivity Market Review: Review of competition in the provision of leased lines”

Efficient Network Structure and the Provision of Dark Fibre under Regulation

20 August 2015

Comments on this response should be sent to:
Matt Cherry, BT Group Regulatory Affairs, via email at matt.cherry@bt.com

Section

1. Executive Summary

1.1 This discussion Paper is intended to give some technical explanation to why the mandatory provision of dark fibre at a regulated rate will induce artificial changes to the way that CPs of all types will be incentivised to construct their networks and which will likely not be economically efficient.

1.2 Networks can be designed in radically different ways to offer precisely the same services to end-users. The combination of network topology and technology will jointly define the resulting network structure. The choice of for example whether to have a network which has relatively more electronics than fibre will depend upon the underlying nature of the spatial distribution of demand by end-users and how this is expected to change over time along with innovation in network technologies. A key issue is how to deal with risk mitigation when both of these underlying factors are uncertain.

1.3 Making basic and uncontroversial assumptions about underlying cost structures for different network components, it is possible to illustrate these choices and their inherent trade-offs to underlying uncertainties. The implicit assumption is that the CP is building its own network i.e. it is entirely self-supply and vertically integrated.

1.4 In the alternative framework where some of these 'upstream' network components are available on either a purely commercial basis or from mandated supply from regulation, it is also possible to determine how different upstream pricing structures will affect CPs decisions to build one type of network rather than another type.

1.5 As each CP will likely target somewhat different customers, there is always going to be an element of bespoke requirements for the serving network; these upstream components will likely be best supplied where there is the ability of suppliers to offer a tailored solution to each CP. This is equally true whether there is only one potential supplier in a position of bi-lateral monopoly or it is more accurately characterised as a monopsonistic market of a small number of purchasers and a small number of suppliers.

1.6 The gap between marginal costs of supply and a regulatory price will have a profound effect on network structures when the seller has to maintain a non-discriminatory price for all potential purchasers. In the context of a dark fibre offering, such effects are markedly compounded when the implicit bandwidth gradient of the consequential active downstream service is then lost. This is because the bandwidth gradient itself is a proxy or reflective of the true trade-offs in upstream network structures.

1.7 In general terms, it can be argued that mandating dark fibre such as how Ofcom is proposing will introduce productive inefficiencies by driving CPs to make sub-optimal investment decisions. This will not happen overnight as it takes time to change networks but it nevertheless will be a strong influence on investment decisions.

1.8 This Paper is structured as follows. Section 2 provides some background to the themes and issues which are developed introducing the fundamental concepts under consideration. Section 3 examines in some detail the options and trade-offs open to a vertically integrated

CP in offering services. These are assessed against the conventional economic criteria of allocative, productive and dynamic efficiencies[1].

1.9 Section 4 examines the impact of introducing an intermediate upstream market on the efficient choices open to vertically integrated CPs and examines how the different economic efficiencies are impacted by the introduction of upstream markets. Section 5 summarises the discussion and conclusions drawn.

[1] For the purposes of this paper, we do not make a general distinction between productive efficiency, technical efficiency, and static efficiency. Strictly, the definition of technical efficiency is the optimal use of a firms inputs for its specific output while productive efficiency applies across industry as a whole. A firm is productively efficient when its marginal costs equal its average costs and normally productive efficiency will determine the efficient number of suppliers in a market. Dynamic efficiency reflects the optimal timing of the introduction new technology and best practices. The impact of passive remedies has an important role for dynamic efficiency.

2. Background Concepts

2.1 The way in which CPs build networks and in conjunction with their operational resources then offer services is complex; there are a wide variety of alternative choices on how to design networks and operational processes and even in services.

2.2 An underlying principle of telecoms regulation is that it should espouse 'technology neutrality' i.e. the regulation of a service should not be contingent on the precise nature of the network and process designs to provide an equivalent service.

2.3 However, as this paper will discuss, there are more technical choices open to a CP that just 'technology'; the principle might more properly be called 'network structure neutrality'.

2.4 The principle of neutrality is most easy to appreciate when applied to downstream services i.e. services delivered to customers outside the broad scope of the supplying telecommunications industry[2]. This is the point at which the technical choices of the services in the upstream supplying industry are distinguished from those outside the industry.

2.5 However, the principle can become more difficult to apply for the upstream services themselves as the technical choices of the supplier (to another CP) are often harder to decouple from the specific upstream service being offered. In addition, the downstream CP will also make technical choices which may interact with the technical choices of their upstream supplier.

2.6 In this paper we examine the general nature of the technical choices open to CPs both upstream and downstream from the point of view of overall economic efficiency. In particular we examine the effect of particular upstream regulatory remedies can affect overall economic efficiency.

2.7 Some of the contingent factors of importance here include the following:

- Downstream to upstream linkages. Upstream wholesale supply can be used by a wide variety of downstream applications and downstream CPs draw from a wide range of upstream supplies – there is not a unique mapping of upstream to downstream for any particular end-user service.
- Multiplexing equipment largely removes the dependency of specific technologies (eg TI and AI) from downstream to upstream.
- Network structure is very sensitive to a wide set of parameters which vary significantly between CPs and the structure of upstream pricing can create considerably dependencies between upstream and downstream as the marginal price of an upstream service is a marginal cost to a downstream CP and will have a considerable influence on the downstream CP's choice of network structure.
- When two networks with different network structures offer equivalent services, then the market pricing of the service cannot be directly governed by the costs (average or marginal) of just one network structure or just of the other.

[2] With business services, giving a more precise definition of the scope of the industry is not straight forward and many businesses sit around the edge of the industry including systems integrators and even some large businesses with large internal ICT departments. However, forming a precise definition is not important to this paper, and nor is it, BT believes, important in the wider analysis of BCMR market boundaries and market power.

3. Efficient Network Structure for a Vertically Integrated CP

3.1 In this section we examine the choices open to a self-supply (vertically integrated) CP and the trade-offs they will make in order to achieve an optimal and efficient network structure. While there are many factors which influence strategic decisions of this type, we can broadly assume that, at least over time, they will be driven by a desire to maximise profits and that in order to do this seek to be internally efficient.

3.2 The two main forms of economic efficiency appropriate to the internal operation of the vertically integrated CP are productive efficiency and dynamic efficiency. The CPs desire to produce its outputs with the lowest possible costs using current technology and best practice is reflected as productive efficiency. CPs desire to keep their strategic investments in technology and best practices up to date is reflected as dynamic efficiency.

3.3 The vertically integrated CP is more likely to think in terms of service incremental costs and strategic costs. The service incremental costs are driven directly by sales and are avoidable in the absence of sales. Strategic costs are investments made ahead of sales and are unavoidable in the absence of sales. While they are not the same, there is a correlation with productive efficiency and dynamic efficiency. Once sunk, strategic costs cannot be readily modified in order to contribute to productive efficiency, while it is these same strategic costs which are likely to become obsolete and contribute to dynamic inefficiency.

3.4 At this stage we are not directly concerned with the pricing of the downstream services of the vertically integrated CP. This is a matter of allocative efficiency between the vertically integrated CP and their customers and does not affect the efficiency within the telecommunications industry which is the focus of this paper.

Basic Elements of Network Costs and Drivers

3.5 Decisions a CP must make on optimising their costs including choosing a network structure is largely controlled by the *incremental* costs of the required resources. Table 1 sets out the broad primary domains of cost for a fixed network and the factors that drive these incremental costs. In each case, the identified parameter is the natural unit of volume for the resource.

Table 1

Cost domains and drivers

Cost Domain	Approx. Life	Service incremental	Primary Drive of Incremental Spend	
Duct	20-30 years	Fixed	Extending geographic reach Renewals	**Distance**
Cable	10-20 years	Fixed	Extending geographic reach Renewals	
Equipment	3-10 years	Fixed	Increasing network bandwidth capacity Increasing network functionality	**Bandwidth**
Service OSS	1-5 years	Fixed	New services	**№ of services** **Service complexity**
Service development and marketing	1-5 years	Fixed	New services	
Service sales and operations	1-5 years	Variable	Service provisions, rearrangements, cessations	**№ of contracts** **№ of service instances** **Service complexity**
Network OSS	1-5 years	Fixed	New vendor/equipment types	**№ of equipment** **Range of equipment**
Network development	1-5 years	Fixed	New vendor/equipment types Increasing network bandwidth capacity Increasing network functionality	
Network operations	Current account	Fixed	Network failures	**№ of equipment** **Route length**
Other overheads	Current account	Fixed		**Fixed**

3.6 There is little that could be considered controversial in this table which is more or less accepted industry wisdom. As can be seen, the two parameters of bandwidth and distance that are often assumed to feature in transacted services do indeed drive some costs. In particular, the costs of duct and fibre are driven by distance and the cost of electronics is driven by bandwidth**.** It is also important to note that the converse is not generally true and that electronic costs are not driven by distance (in a UK network, the cost of optical amplification is a small proportion of total costs of supplying business connectivity) and duct and fibre costs are not driven by bandwidth, at least up to moderate bandwidths[3].

3.7 It is also clear that most costs are strategic involving a scope and timescale beyond that of any one service contract; therefore the vertically integrated CP is concerned about recovery of common strategic costs. Closely coupled to this, the vertically integrated CP is likely to be

[3] 'Moderate bandwidths' will normally be set by the economics of DWDM. A DWDM system can carry today, up to around 1Tbit/s on a single fibre pair. Bandwidth capacity beyond this will require additional fibre strands. However, the economics of DWDM versus using additional fibre strands will vary with distance. On longer links, DWDM is likely to be more economic and on shorter links additional fibre is more likely to be economic.

concerned about the pace of innovation in their strategic investments and hence in their own ability to be dynamically efficient.

3.8 The costs of duct, cable, and equipment are often taken as the obvious costs of a vertically integrated CP, however, in addition, there are many costs which are not driven by distance or bandwidth at all. Some of these are sensitive to the scale and scope of individual contracts such as contract complexity, marketing costs and the like. Generally speaking, fewer, longer term contracts will likely be more profitable than many short term contracts of the same overall size with all other things being equal.

3.9 This suggests that in a competitive marketplace, pricing will reflect this cost variation and large businesses tendering for consolidated longer term contracts will be offered significant price reductions compared to small businesses taking short term contracts as both the direct costs will be relatively lower and there will be significant risk mitigation of utilisation of the network in the future.

3.10 In summary, competitive pricing of network services will typically not follow the parameters of bandwidth and distance in a simple linear fashion. It is shown below that as communications services share infrastructure with other communications services through service aggregation, this factor will also profoundly affect both the level and structure of overall incremental costs for a CP.

Technology Choices for Aggregation

3.11 The essence of a telecommunications network is that a good proportion of costs can be shared across many customers, notably duct and fibre costs, but also all strategic costs. Indeed, it is this sharing of the strategic costs which is how telecommunications manages to achieve general economies of scale and scope.

3.12 For example, without sharing, it would be necessary to dig a duct from every customer end point to every other customer end point and the costs of duct infrastructure would grow roughly as the square of the number of the number of customer end points. With sharing, costs grow roughly linearly with the number of end points[4]. The concept of aggregation to share infrastructure is profound to telecoms networks but, it can be easily taken for granted.

3.13 In this section, the primary focus in on the sharing of network infrastructure across different service instances, that is duct, fibre, and equipment; however there are points where the sharing of other strategic costs play an important part in determining the efficient network structure of a CP.

3.14 While aggregation is fundamental, there are many different ways of achieving aggregation and it is important to clearly establish the extent to which different methods of aggregation are equivalent and are, in economic terms, substitutable alternatives. This needs to be taken in the context of an overall network structure.

3.15 We start by considering the very broad alternatives for aggregating telecoms traffic. In order to share the duct and fibre costs, communications from different customers must be aggregated together at aggregation points which itself can arise in a number of ways:

- **Physical Aggregation** – aggregation of fibre strands into a duct route. The current generation of optical fibre cables can carry large numbers of fibre strands (100-200) in a small cable size and it is possible to arrange for a very large number of fibre strands in a standard 4 or 12 bore duct-way. The aggregation and technology is either cable joint fibre

[4] Each end user will require dedicated network to some degree.

splicing or a fibre distribution frame and both are done manually. Cable joint fibre splicing makes permanent connections and carries a modest cost compared to the cost of the cable itself. Fibre frames are semi-permanent connections but are at a significantly higher cost than cable joint fibre splicing. (Note, the bandwidth capacity of a fibre strand is not essential to the fibre strand, it is set by the technology used to 'light' the fibre[5].)

- **WDM/TDM Aggregation** – aggregation of fixed bit rate, high capacity bitstream (e.g. 10G Ethernet) onto a fibre. TDM technologies like SDH and more recently OTN as well as WDM technology multiplex fixed bandwidth pipes in a way that is software controlled, but is still semi-permanent. While TDM and WDM are technologically different, in terms of what and how they aggregate, they are very similar[6]. Current systems can achieve an aggregated capacity in excess of 1Tbit/s of bandwidth on a single fibre pair. However, for some network scenarios, it is the fixed bandwidth, semi-permanent nature of the connectivity which is attractive.
- **Packet Aggregation** – aggregation of packet data flows onto high capacity bitstreams. This aggregation uses packet switching technology (normally either Ethernet or IP/MPLS technology) and can support a high level of connectivity flexibility switching individual packets to a destination on a network of very large numbers of nodes. Broadly speaking, Ethernet offers a lower cost of bandwidth with lower connectivity flexibility while IP/MPLS offers greater connectivity flexibility but at a higher bandwidth cost. In many cases a CP will optimise by deploying both technologies.

3.16 In most cases, when used in conjunction with each other, these aggregation technologies will be used in the order described, so a packet is carried on WDM/TDM which is then carried on fibre. This sets a general layering[7] and so it is normal to discuss: (a) the fibre (or physical) layer; (b) WDM/TDM layers (often referred to as circuit[8] layers); and (c) packet layers. In many cases, there may be more than one WDM/TDM layer (e.g. WDM and SDH) and/or more than one packet layer (e.g. Ethernet and IP/MPLS).

3.17 This ordered layering will generally be the case. However, in some cases, a layer may have no effective aggregation/switching. For example, a 10Gbit/s Ethernet signal may be carried directly on a fibre and be multiplexed together with other signals. However, the 10Gbit/s signal is still a valid TDM signal with TDM features and could be multiplexed at a future time if desired. In this case the WDM/TDM layer is called a redundant layer.

3.18 In addition, it is possible to swap around the layer order with TDM traffic carried across a packet network which is especially useful in evolutionary interworking scenarios. This is often called circuit emulation. While the emulation may not be ideal in every regard (it tends to create extra one way delay and early systems did not correctly handle network synchronisation), for most applications circuit emulation gives a direct equivalent to TDM.

3.19 These are the basic technological options available for aggregation and the general way in which they interrelate. These different layers of aggregation are illustrated in Figure 1 below.

---

[5] The exact optical specification of the fibre, normally to an ITU-T standard, can be relevant.

[6] In fact, in strict terms, 'pure' WDM does not generally exist as a technology today and most commercial systems are a hybrid of TDM and WDM technology. 'Alien wavelength' technology gets closer to 'pure' WDM.

[7] Layering as discussed here should not be confused with 'tiering' of a network into, for example, access, backhaul, and core

[8] While there is a common origin in the terminology, these circuit layers should not be confused with circuits in the PSTN.

Figure 1

Technology Layers of Aggregation

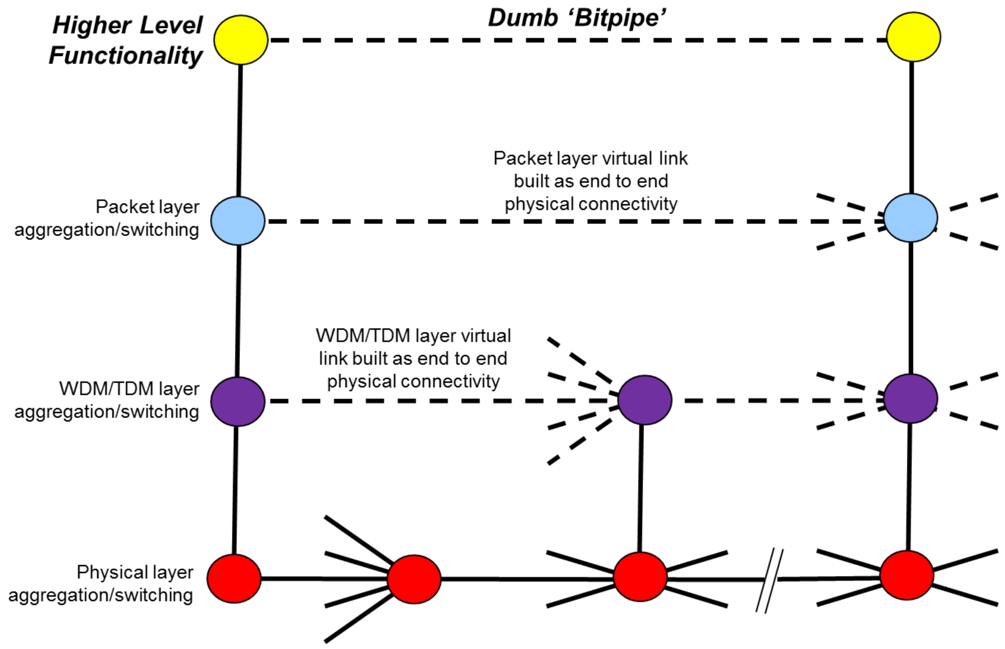

3.20 The layer, or set of layers of aggregation that is optimal for a CP to use will depend on a number of factors:

- The bandwidth and destinations of the traffic being aggregated.
- The number of customer communications being aggregated.
- The statistical characteristics over time of volume and destination of the traffic being aggregated.
- The incremental cost of bandwidth between aggregation points.
- The incremental cost of bandwidth at aggregation points.

3.21 We can say with approximate and broad generality, that:

- The higher the bandwidth of each service instance being aggregated, and/or
- The smaller the number of customer communications being aggregated, and/or
- The more certain the statistics on volumetrics,

then the greater the drive is to lower layer aggregation and vice versa.

3.22 In summary, aggregation is fundamental to any telecommunications network. There are three basic means by which aggregation can take place – physical aggregation, WDM/SDH aggregation, and packet aggregation and all of these achieve the basic objective of aggregation. Vertically integrated CPs may come to very different ways of achieving aggregation in their network structures but all will be able to offer equivalent services.

Optimal Network Structure[9]

3.23 In order to examine the underlying choices more closely, we need to consider the way the choice of aggregation technology couples with the choice of network topology – the siting of aggregation/switching and the logical way in which they are interconnected. In this section, we consider what factors drive optimal solutions for a CP's network structure – the network structure being the combination of the choice of aggregation/switching technology and the network topology.

3.24 In order to develop the analysis on optimal network structure more clearly, we start by considering aggregation/switching at any one layer of aggregation. This discussion considers only one aggregation technology; however, the exact nature and in particular the cost structure of the links - is left open. This will allow the analysis to be generalised in the following discussion taking into account multiple layers.

a) Meshed versus Flat and Consolidation versus Distributed Topology

3.25 Given this restricted scenario of only one type of aggregation/switching, it is possible to identify three extreme general topology solutions:

- **Fully meshed topology** – with a fully meshed topology, there is a direct link from every end node to every other end node (note the number of links will grow as the square of the number of end points).
- **Flat Consolidated Topology** – with a consolidated topology, the switching/aggregation is done centrally and each end point has a single long link to the central node.
- **Flat Distributed Topology** – with a distributed topology, the topology will attempt to 'turn around' local traffic in order to reduce the distance traffic has to travel across the network[10].

3.26 These three extremes are illustrated in Figure 2 below and there is a full spectrum of possible topologies in between these extremes.

[9] In order to improve the clarity of the paper, this analysis and discussion of topology does not include the requirement for network resilience which does not affect any of the overall results drawn.

[10] Mathematically, an optimal flat distributed topology is a minimum spanning tree.

Figure 2

Choices of network topologies

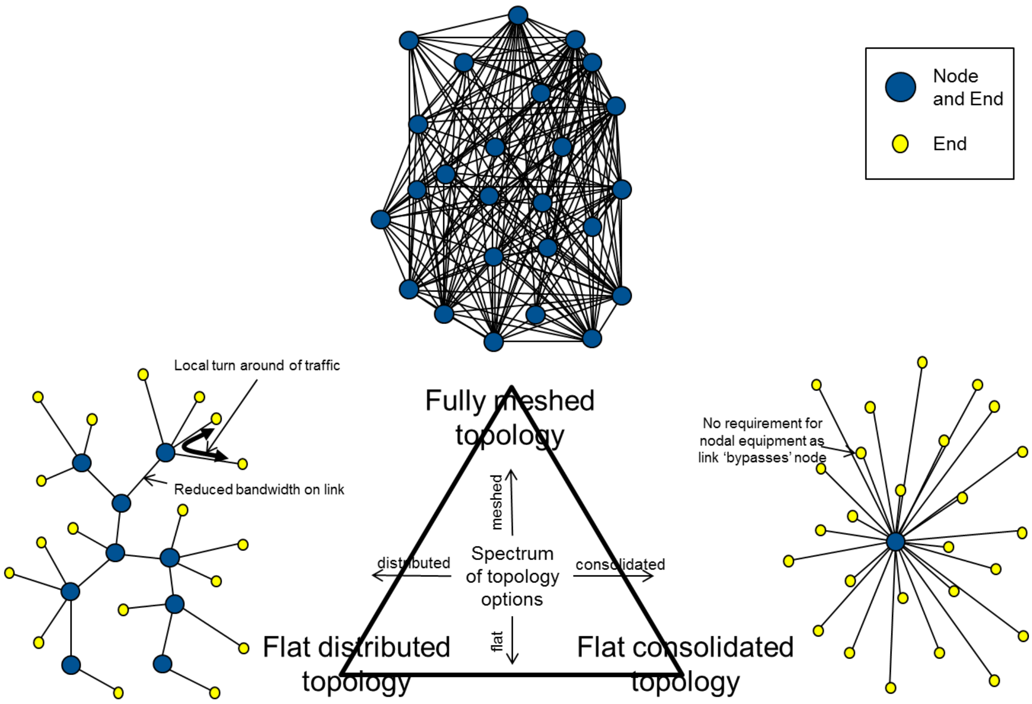

3.27 The scenarios which drive towards one of the extremes can be broadly summarised by considering two axes – 'meshed versus flat' (covering both flat consolidated and flat distributed[11]), and 'consolidated versus distributed'.

3.28 The **'meshed versus flat'** axis optimises the number of links compared to the number of end points. This is likely to be controlled by the non-linearity of the cost of links. If the cost of links is linearly controlled by bandwidth, there is little gain in reducing the number of links and increasing the number of links will reduce the nodal costs by reducing the number of nodal equipment points traffic must pass through[12].

3.29 There are two clear examples of such non-linearity of the link costs:

- If there is a fixed cost to providing a link, this will create non-linearity and reducing the number of links will reduce the instances of this fixed cost.
- If there is uncertainty in the exact bandwidth required between any two end points, resulting in a planning policy of over-dimensioning links to allow for a high probability of meeting peak traffic demand, and further if the uncertainty between different end to end routes is not fully correlated - this will create a non-linearity in dimensioning. The proportion of over-dimensioning needed will reduce if the meshing is reduced and traffic aggregated onto common links.

---

[11] Increasing the meshing from either a flat distributed or a flat consolidated topology will move towards the same fully meshed topology, hence the triangular spectrum of option, not rectangular.

[12] If resilience were to be included in the discussion, the flat distributed topology rather than being a tree structure would show rings while the flat consolidated topology would have more than one hub node (there is no need to interconnect these resilient hub nodes).

3.30 The **'consolidated versus distributed'** axis optimises the 'local turn around' in the network. This is likely to be controlled by the distance element of link costs. If there is a significant distance component to the cost of links along the a bandwidth component compared to the cost of nodal equipment, it becomes worthwhile to introduce local nodal equipment in order to turn around local traffic thus reducing the distance the traffic needs to travel. Again uncertainty in the traffic is an important factor. The success of local turn around depends on a knowledge of the amount of traffic from an end point which is destined for local end points versus traffic for more distant end points. If it turns out that most of the traffic is not local and the dimensioning of the links needs to be increased accordingly, then it can be that the distributed mesh eventually built may be more expensive than a consolidated mesh.

3.31 The way in which the optimisation across the consolidated versus distributed axis works is complex, but a practical observation is that it is highly sensitive to assumptions. Generally, if it is assumed that the incremental costs of nodes and links are linear, then the optimal solution tends to go directly to one extreme or the other as illustrated in Figure 3 below. The practical optimal solution will be determined by the non-linearities in the link and node costs.

Figure 3

Trade-offs of topologies to costs

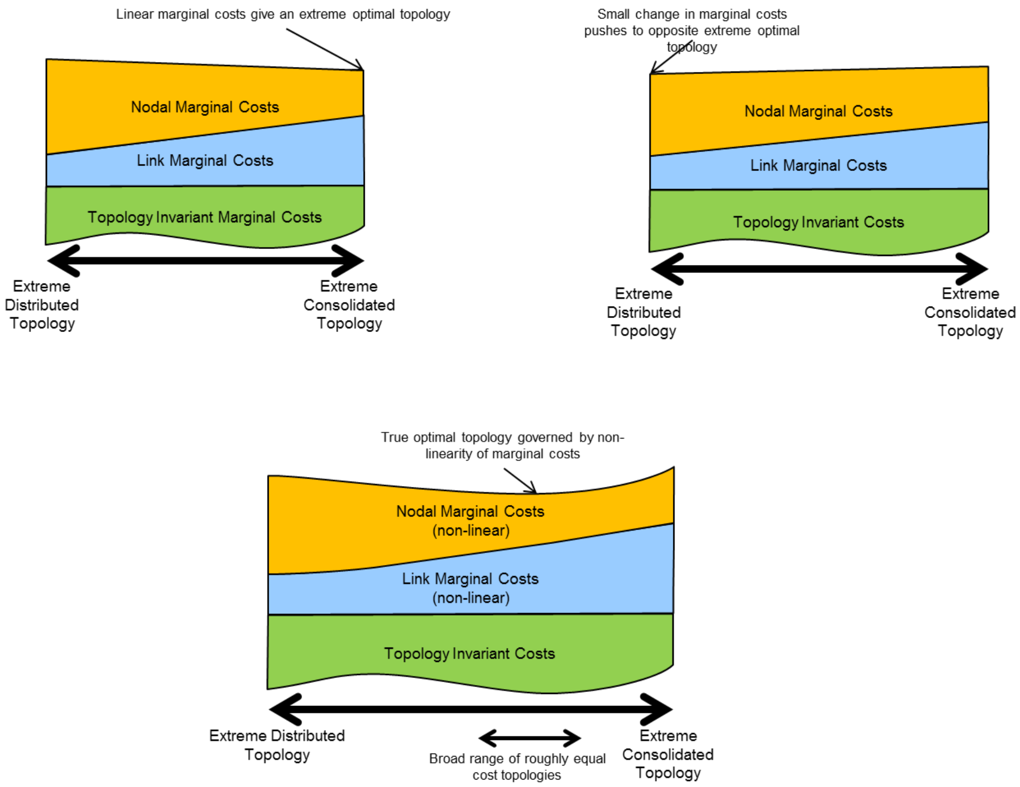

3.32 On both axes, meshed versus flat and distributed versus consolidated, we see that the optimal topology is controlled by the non-linearity of costs. This makes the optimal topology a very sensitive parameter; however, we also note from Figure 3, that there may be a wide range of topology solutions which are all equivalent optimality. This means that different CPs can apparently have quite different topologies and yet still be of equivalent cost optimality. However, it should also be noted that it is the nature of topology that apparently small departures from a broadly optimal topology solution can be anything but optimal.

3.33 This leads to three conclusions on network structure:

- The optimal topology is very sensitive to key parameters which control the incremental costs of links and nodes, which depending on the precise layer of aggregation/switching will include the incremental costs of duct, fibre, and nodal equipment.
- An overriding consideration may be uncertainty at the time of investment as to where traffic will go from and to and this gives rise to an over-dimensioning factor which may dominate the non-linearity of incremental costs.
- Many apparently wildly different topologies may be broadly 'optimal' given the circumstances of the CP taking into account all the uncertainties it faces; however, apparently small external impositions on the topology, induced for example, by regulation, may drive topology far from optimal solutions.

b) Multi-Layer Topology

3.34 The second aspect of aggregation reviewed is the choice of aggregation/switching layer for a node. The discussion above considered only the optimal topology using one layer of aggregation/switching. However, there are a number of possible layers to choose from, and most significantly, they can be used in conjunction with each other. Indeed, most practical solutions do involve a combination of aggregation/switching technologies. In order to keep terminology clear, we use 'network structure' to describe the overall selection of aggregation/switching technologies together with the topology of the way they are interconnected –both physically and logically.

3.35 The distinction between logical and physical topology is highly significant when connecting a multi-layer network structure. In particular, if aggregation is carried out at a node at a lower layer, say physical aggregation, then traffic cannot be turned around at a higher layer, say packet aggregation/switching – the lower layer aggregation is essentially invisible to the higher layer.

3.36 For example, suppose a separate fibre pair is used to connect each of many customer sites to a central node and that each fibre pair is connected to a separate interface on an Ethernet switch in the central node, then there is a logical consolidated topology as far as the Ethernet switch is concerned. However, it is very likely that the access fibre is aggregated together from many different customers in footway boxes in the street into one large cable and that the physical layer aggregation is a flat distributed topology. But this physical layer topology will be completely invisible to the Ethernet layer. Moreover, all the traffic which will be Ethernet switched must come all the way to the central node and cannot be turned around in a footway box. BT's network structure is broadly of this form.

3.37 Different CPs may choose quite different solutions to this mix of aggregation to create a network structure. For example, faced with the same situation of customer sites as in the example above, a different CP might choose to site Ethernet equipment at each customer site and then 'daisy chain' these together with the physical fibre. This does achieve a much greater degree of local turn around (even if at a customer site rather than in the footway box outside) and would reduce the total length of fibre required but would require more Ethernet switch capacity and would be less flexible to the exact end-to-end traffic requirement. This situation can be further optimised by using hybrid WDM/Ethernet technology. BT understands that some city fibre networks are more akin to this form of network structure than that of BT's.

3.38 As noted before, each of the three basic layers of aggregation has different technical characteristics and different costs. This means that not merely each may suit different forms of underlying communication demand, they can be used optimally in combination in the same network.

3.39 In particular:

- Physical aggregation has a relatively low marginal cost of bandwidth but once a connection between fibre strands has been made, this connection is fixed and inflexible.
- Packet aggregation has a significant cost of bandwidth but is very flexible – each and every packet can be directed to a different destination on a packet by packet basis.
- TDM/WDM is between these two, both as regards the marginal cost of bandwidth and as regards flexibility.

3.40 This suggests an optimal topology in which nodes where aggregation can be fixed and permanent, is best implemented using physical aggregation/switching, while the residual nodes which must offer flexibility are best implemented using packet aggregation. There may well also be nodes, and given particular communications requirements, where the implementation using TDM/WDM aggregation/switching is optimal. Such a scenario is illustrated in Figure 4 below.

Figure 4

Optimisation of networks

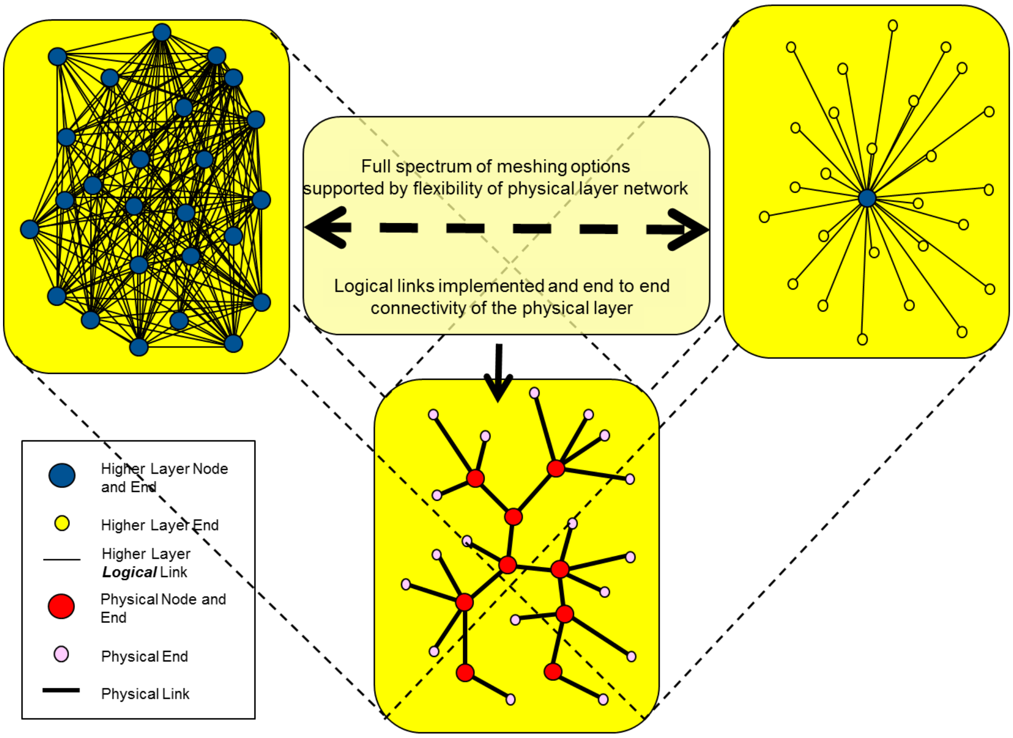

3.41 A hybrid network structure where there is a mix of aggregation types makes the calculation of an optimal network structure more complex. Assuming the issues of flexibility and uncertainty are taken into account, we can draw the following immediate conclusions:

- From the perspective of the uppermost layer (normally a packet layer), the incremental cost of the logical links of this network now depend on the incremental cost of nodes and the incremental cost of links of the lower layers as well as the topology of the lower layers.

- The optimisation problem is framed as a problem in the node incremental costs of all the different types of aggregation, the link incremental costs at only lowest layer, and solving, simultaneously, for the optimal network structure which is now a combination of the logical topology at each layer.
- Once again, any scenario with constant incremental costs will drive directly to extreme optimal structures.
- Actual practical optimal network structures, including the relative proportions of each type of aggregation/switching technology will again be governed by the non-linearity of incremental costs.

3.42 The final issue examined here is that of risk mitigation. As has been already noted in the foregoing discussion, uncertainties can have a profound influence of the choice of network structure.

3.43 As well as being as source of risk, some network structural choices can themselves have a mitigating effect on risk. For example, aggregation will combine demand volume from different customers and assuming some non-systematic risk between customers, the uncertainty in the aggregate demand should be less than the sum of the uncertainties in the individual customer demands. A network structure that maximises aggregation will therefore mitigate volume risk.

3.44 Similar structural opportunities exist for both spatial and technology risk. Spatial risk (capacity in the wrong location) can be mitigated by limiting the investment in aggregate infrastructure which is committed to a geographic location and increasing the amount of investment that is specific to a particular customer demand. This would result in less density of aggregation nodes and longer 'spokes' to each customer site. However, it can be seen by comparing with above, that this may well be in conflict with mitigating volume risk if there is a significant incremental bandwidth cost to the 'spoke' links.

3.45 With aggregation technology risk, the risk associated with the fit of different aggregation types to different customer demand, some technologies are more broadly applicable to meet customer requirements than others. For example, packet aggregation/switching technologies will be capable of handling a very broad spectrum of communications requirements and so a risk mitigation strategy might be to minimise, or even eliminate, the use of WDM/TDM technology and use packet aggregation as a 'universal' solution. However, this might then not be optimal if there is considerable demand for fixed high capacity bit pipes which are more optimally handled with WDM/TDM technology.

3.46 In summary, the factors which drive a vertically integrated CP's selection of network structure are very sensitive to underlying assumptions and the strategy of risk mitigation with different outcomes quite feasible. Moreover, the choice for the vertically integrated CP is essentially a strategic choice as the primary elements of cost are not service incremental.

Volume Growth and Generations of Technology

3.47 A network structure tends to be a coherent decision for the whole of a CP's network. Altering a network structure is highly disruptive and normally very expensive so such changes will manifest over a longer period. A network structure is often optimal for a certain volume of traffic and even where the traffic originates and terminates within the network. If the volume of traffic grows and/or occurs in different places to forecasts, the network structure may need to be changed.

3.48 In addition, the network structure is an integrated decision covering both technology and network topology, that is the siting and interconnection of equipment. This means that any change to the technology may also be disruptive to overall network structure.

3.49 Telecommunications is an industry of rapid technological advancement and so any vertically integrated CP is likely to consider future evolution of their network in their network structure as a way of ensuring ongoing dynamic efficiency. This means that network structure cannot be simply optimised at one point in time with one mix of technology.

3.50 These two dynamic factors, volume growth and changes in technology - are often major constraints on a CPs decision on network structure at one point in time.

3.51 For example, a CP may choose to design a network with a given technology and which is optimal with a specific volume of traffic. As this network fills up towards its optimal capacity, the CP may then decide to introduce a new ‘overlay’ network with a newer generation of technology and with a much higher optimal volume of traffic and it may also target new services at the new network. This strategy would minimise the need for interworking between the generations of network.

3.52 A different CP may choose a very different strategy and choose a topology which is reasonably tolerant to growth in volume and changes in technology[13]. With such a network structure, this CP can grow and adapt their network in line with volume growth. However, it is likely that at any particular demand volume and placement, such a structure may be less productively efficient that one specifically designed for that volume and placement of traffic.

3.53 This illustrates that when growth and technology innovation are taken into account, any CP will face a balance between productive efficiency and dynamic efficiency.

3.54 In general, the vertically integrated CP has a number of opportunities to minimise any dynamic inefficiency. A vertically integrated CP can:

- Select a network structure consistent with the CP’s growth forecasts and confidence in those forecasts;
- Make use of the timing of major service innovations to coincide within an major technology introduction and/or change to network structure;
- Make use of the timing of major equipment obsolescence to coincide within an major technology introduction and/or change to network structure;
- Design a network structure which is more flexible and adaptive to future technology change.

3.55 In practice, both volume growth and new technology can be a cause of major cost disruption and/or productive cost inefficiency. While it might be reasonably expected that new technology would cause a major disruption, it is often the case that volume growth causes the greater levels of productive inefficiency.

3.56 As an example, it have been the experience of many CPs that capacity planning for SDH networks became highly complex and inefficient as volume growth exceeded the forecast capacity of a chosen topology, notably ring topologies. Over time, SDH networks became more consolidated with greater use of SDH cross-connects rather than add-drop multiplexers which

[13] An example would a topology which is variously called a ‘fat tree’ or ‘Clos core’ which is insensitive to the placement of traffic and can be incrementally increased in size as needed. However benign topologies such as this do have scaling limits which can only be overcome by newer generations of equipment with higher capacity interfaces or by radical change to the network topology.

were used much more widely early on. The consolidated topology is much more tolerant to growth in volume.

3.57 In summary, the vertically integrated CP will likely have a trade-off between a network structure which optimises the productive efficiency of current known traffic volumes and current technology with the dynamic efficiency for less certain future growth and technology innovations. The impact of volume growth on efficient network structure can be as large if not larger than that of new technology.

4. Network structure with Upstream Services

4.1 In the previous section we considered the case of a vertically integrated CP and the factors which govern their choice of network structure. We saw that the choice of network structure can be highly sensitive to particular incremental costs.

4.2 In this section we consider the situation where a CP does not have complete physical infrastructure and seeks upstream supply from another CP to complete their required physical connectivity. This is predominantly an issue of access connections, assuming we take a broad definition of access. The downstream CP is looking for an upstream service to connect an end node of some description to a network mode of their network.

4.3 The end node may vary and certainly includes business customer sites, mobile base station sites (for the mobile operators), and LLU local exchanges (for the LLU operators). In each case, the downstream CP is looking for a service which will connect these sites to a consolidated node which is within the footprint of their own physical network.

4.4 Much less commonly, a CP may look for connectively to interconnect two of their own consolidated network nodes. This is less common as, generally, the economics for the downstream CP to invest in their own infrastructure between their own consolidated nodes is more compelling.

4.5 This upstream market has been one of the earliest targets of regulatory intervention (dating back to at least the Open Network Provisioning (ONP) initiative of the 1990s). Regulation has therefore been a feature of upstream supply and it therefore not easy by looking at historic services to disentangle genuine free market solutions from solutions directly or indirectly influenced by regulation i.e. on modified Greenfield site assumptions.

4.6 In order to avoid drawing inappropriate conclusions from historic solutions which derived from regulation (rather than lead to regulation), in this section we look at the general economic factors which will shape pricing in a competitive free market absent regulation.

4.7 A key economic tool in carrying out this analysis is to consider the impact of upstream pricing on three forms of economic efficiency, that is allocative efficiency, productive efficiency, and dynamic efficiency.

4.8 The first subsection looks at the general factors which will influence the units by which upstream pricing are likely to be offered. The second subsection the looks at the consequences of this on allocative efficiency. The third subsection looks at the way upstream pricing can drive the choice of network structure of the downstream CP while the fourth subsection considers the specific case of the proposed dark fibre remedy. The fifth subsection looks at how demand from downstream CPs based on their chosen downstream network structure will drive the choice of upstream network structure of the upstream CP while the sixth subsection considers the specific case of the proposed dark fibre remedy. The final subsections examines the effect of these upstream and downstream network structures on dynamic efficiency and again considers the specific case of the proposed dark fibre remedy.

Pricing of Upstream Services

4.9 We noted above that different CPs for very good reasons may choose quite different network structures. A critical consequence of this is that different CPs in direct competition with each other could have quite different incremental costs in order to supply the same downstream service. In this section, we look at how incremental costs are likely to drive upstream pricing structures.

4.10 If we assume that pricing structures will broadly follow the largest element of cost, we see that different broad classes of services have different largest elements of cost. Table 2 below sets this out by taking each of the primary elements of cost in turn and identifies the broad classes of service where that primary cost element is dominant, or is sufficiently important to affect pricing structures.

Table 2

Scenarios where cost drivers affect pricing

Overall Cost Domain	Specific Cost Domain	Service Parameters Potentially Relevant for Pricing	Scenarios when Important
Physical infrastructure	Physical infrastructure dedicated/lightly shared by a customer end(s)	• Customer density at each circuit end • Contract term	Important when • non-dense circuit ends
	Physical infrastructure heavily shared by customer ends	• Circuit length	Important when • dense ends • longer lengths • very high bandwidths
Equipment	CPE	• Bandwidth at each customer end • Contract term	Important when • dense ends • short lengths • very high bandwidths
	Network equipment	• Circuit bandwidth	Important when • dense ends • longer lengths • moderate and low bandwidths
		• Total number of circuit ends	Large numbers drive • pt-mpt (eg E-tree) • mpt-mpt (eg E-LAN, IPVPN)
Sales, marketing, customer ops		• Total number of circuit ends • Contract term	Economies of scale
Fixed costs and overheads		• Many options possible	Most scenarios

4.11 This table might suggest that different components forming services will likely have quite different units of pricing. For example, it might be expected that services which are predominantly based principally on physical infrastructure are likely to show a dependency on distance but have little bandwidth dependency while a service which makes extensive use of equipment is likely to price bandwidth and not price by distance.

4.12 However, we noted above from the discussion on network structure, that the choice between physical layer aggregation and packet layer aggregation is often very sensitive and different CPs may well make different choices. This means that CPs may be offering identical service using network solutions with quite different cost characteristics.

4.13 This has a profound impact on pricing. For example, a CP which has chosen physical aggregation may think that pricing is optimal based on distance and insensitive to bandwidth. At the same time a CP which has chosen packet aggregation may think the exact reverse. But they will likely offer services in competition with each other even potentially at the upstream layer.

4.14 This will drive a market price which will involve both distance and bandwidth as each CP sees the competitive impact of the other. A very important general conclusion on the pricing of connectivity is that competitive pricing is not just based on the supplying CP marginal costs, but is also based on the technological alternatives, indeed, even if the alternatives have not yet been deployed.

Upstream Pricing and Allocative Efficiency

4.15 The underlying economic principle at work with upstream pricing is that the marginal price of the upstream service is a marginal cost to the downstream CP. The downstream CP will now base its decisions on this marginal price of its upstream service.

4.16 If the marginal price of the upstream service is the same as the marginal cost of the upstream CP, then, at least in principle, the downstream CP should be equally efficient compared to a vertically integrated CP. However, even in the most simple scenario this may not be the case as typically the upstream CP will have to price above marginal cost and then the customers of the downstream CP will be subject to the well-known problem of 'double marginalisation'[14.]

4.17 There are also well known solutions to this upstream pricing problem and resulting allocative inefficiency, most notably a two-part tariff. However, as we noted above, the situation for the upstream supply of connectivity is more complex and the marginal costs of the upstream CP will depend on the choices it has made on its network structure which affect not just the level of the marginal costs, but also the units of marginal cost, for example, distance, access bandwidth, circuit bandwidth, number of contracts, etc.

4.18 This means that double marginalisation is likely to occur even with two-part tariffs. Unless the upstream CP has exactly the network structure that the downstream CP would have chosen were it vertically integrated, then it is very likely that the units of the two part tariff will pass the incorrect marginal cost information from upstream to downstream.

4.19 It is plausible to construct multi-part tariffs as a solution but they may be other mismatches in the network structure which still result in allocative inefficiency. Indeed, any form of imposed pricing structure is highly likely to create double marginalisation and allocative inefficiencies and as discussed below this will also include a regulatory price for dark fibre.

Upstream Pricing and Downstream Network Structure

4.20 Units of marginal price are likely to lead to some degree of allocative inefficiency, however, the consequences go considerably further. This marginal price of upstream service is now the marginal cost for the downstream CP in its choice of network structure. As was noted in the discussion of a vertically integrated CP, the choice of network structure can be very sensitive to marginal costs and so the difference between the vertically integrated marginal costs and the marginal price of the upstream service may be sufficient to drive a radically different optimal network structure for the downstream CP.

[14] The term arises because the upstream CP and the downstream CP are both taking a profit margin above the marginal cost. In contrast, the vertically integrated CP will only take one profit margin and will offer a lower price to the downstream customer. This arises even when all three CPs are profit maximising.

4.21 For example, if the marginal price of distance on a per circuit basis, is above their vertically integrated marginal costs, then it is likely to drive towards distributed topology when the vertically integrated marginal costs might have driven towards a consolidated topology. The downstream CP has an incentive to minimise the number of circuits that are co-routed on the same length of physical infrastructure (often called 'tromboning' as traffic routes up the network and back again in the shape of a trombone) by inserting their own multiplexing/switching equipment. However, the true marginal costs of co-routing may be very low and the cost of the downstream CP equipment could have been completely avoided by a vertically integrated CP.

4.22 Moreover, if the price of bandwidth is reasonably linear when added across circuits and there is little or no fixed cost to a circuit, then this is likely to drive towards a meshed topology while the non-linearities of the true vertically integrated marginal cost might well have driven towards a flat topology.

4.23 This means that where the upstream marginal price does not reflect upstream marginal costs, a downstream CP is likely to have a radically different network structure compared to the same downstream CP vertically integrated with their upstream supplier. By definition, these structures will be less efficient than that of the vertically integrated CP.

4.24 Again in a competitive free market, bespoke contracting can provide a way of avoiding this situation as it gives both the downstream CP and the upstream CP the opportunity to bilaterally negotiate a complex technical and commercial contract which can allow them both to have mutually efficient network structures. Where there are regulatory obligations of non-discrimination, especially when coupled with transparent pricing, the drive to inefficient downstream network structures is much more likely.

4.25 In summary, even in a free competitive market, achieving allocative efficiency and efficient downstream network structures when an upstream CP supplies a downstream CP is highly complex and is very unlikely to be achieved outside bespoke contracting[15]. Regulatory obligations of non-discrimination with a determined price for upstream services will not satisfy the needs of all downstream network operators depending on their precise circumstances.

Dark Fibre and Link to Downstream Network Structure

4.26 Given that outside careful bespoke contracting, the marginal price of the upstream service is not likely be the same as the upstream marginal costs, the optimisation of network structure of the downstream CP is likely to be very different to that of a vertically integrated CP.

4.27 Importantly, the current proposal for a dark fibre remedy is likely to cause a profound change to optimal downstream network structures as it undermines and in places completely removes any marginal price for bandwidth., However, there will most likely be a marginal price for each dark fibre circuit (a requirement for non-discrimination would suggest that the marginal price per circuit would be linear). This is likely to drive an optimal downstream topology which minimises both the number of dark fibre circuits and their length, i.e. a flat distributed topology[16].

[15] It should be noted that the entire discussion in the subsection has effectively assumed that there are no additional costs arising from the upstream supply which will undoubtedly be the case. This discussion has focussed entirely on allocative efficiency on the assumption that productive efficiency is unchanged.

[16] In effect, a minimal spanning tree.

4.28 When there is bandwidth gradient, there is an incentive on the downstream CP to use aggregation provided by the upstream CP especially at the physical layer and the downstream CP will have a more meshed topology. This has the important effect of creating a trade-off in the per circuit costs which is much more reflective of the true upstream marginal costs and especially when the possibility of different upstream network structures is taken into account.

4.29 In practice, the more meshed downstream topology with significant upstream aggregation (especially at the physical layer) is likely to be much closer to the network structure of a vertically integrated CP and much more productively efficient. In practice, dark fibre is likely to drive away from cheaper physical layer aggregation by the upstream CP and drive towards more expense WDM aggregation by the downstream CP.

4.30 In summary, dark fibre, by reducing or even eliminating the bandwidth gradient is likely to create a significant change in optimal downstream network structure in a way that is overall significantly less productively efficient than an upstream service with a bandwidth gradient.

Upstream Pricing and Upstream Network Structure

4.31 The impact of regulated upstream services on network structures does not stop however, with the downstream CP. The change in the network structure of the downstream CP will change the nature of demand for upstream services by the downstream CP and this demand profile is then likely a major influence on the upstream CP's choice of network structure. There is a profound and very strong feedback mechanism whereby the marginal price of the upstream service will change demand in a way that is much more complex than a simple demand curve showing simple elasticity.

4.32 As a downstream CP adjusts (maybe radically) its choice of network structure in response to a change to the marginal price of the upstream service, this will cause a consequential profound and radical change to the volume, location and technical form of the demand for the upstream service.

4.33 This feedback mechanism to the network structure choice of the upstream CP is likely to be just as profound as the initial impact of upstream pricing was on the network structure of the downstream CP. Two examples can illustrate this.

4.34 Consider first the case discussed above where the introduction of upstream pricing meant that the marginal price of distance, on a per circuit basis, is above its upstream marginal costs. This, as discussed above, lead the downstream CP to develop and distributed topology when a vertically integrated CP would have built a consolidated topology. In this case, instead of requiring a smaller number of longer circuits as would have been the case had they been vertically integrated, the upstream CP is now required to supply a larger number of shorter circuits. This means that the circuit termination costs for the upstream CP have now considerably increased.

4.35 In this example, we can see that as a result of the introduction of upstream pricing, compared to the efficient network structure of the vertically integrated CP the downstream CP has increased costs of multiplexing/switching equipment in distributed locations and also the upstream CP has increased costs of terminating equipment at these distributed locations.

4.36 If this solution were to be reviewed *ex post* and the marginal cost of distance assessed, there would indeed be a high overall cost of distance because of all the distributed equipment. This would then give the false impression that the original pricing assumption on the marginal price of bandwidth was correct. In fact, the feedback mechanism has made the original pricing assumption self-fulfilling through altering the optimal network structures themselves from consequential sunk investments.

4.37 In terms of regulation, this example illustrated two very significant points:

- Upstream pricing which does not align the upstream marginal price with upstream marginal costs, will affect the choice of network structure of both the downstream CP and the upstream CP;
- Examining the cost structure of the overall solution of an upstream CP and a downstream CP will not normally indicate what the marginal costs of a vertically integrated CP would be.

4.38 This latter point is especially important when considering regulated upstream services. It is tempting to use the observed marginal costs of the upstream CP as indicating the marginal costs of an efficient vertically integrated solution. This may well not be the case and the observed upstream costs may well be a simple self-fulfilment of historic regulation of upstream services.

Dark Fibre and Link to Upstream Network Structure

4.39 The second example is specifically that of dark fibre. As discussed above, the effectively linear marginal pricing on a per circuit basis coupled with the absence of a bandwidth gradient is likely to drive downstream CPs to a flat distributed topology minimising the number and length of dark fibre circuits. At the same time, the downstream CPs are likely to have the incentive to buy additional multiplexing equipment (for example WDM) instead of purchasing addition fibre strands from the upstream supplier.

4.40 If this is indeed the case, this will feedback to the upstream CP in the form of demand for fewer and shorter dark fibre circuits compared to the provision of active services across multiple fibre strands. The fibre assets required and the positioning of fibre management are profoundly altered in this scenario. As a result of the introduction of dark fibre, the upstream supplier may well be left with stranded fibre assets.

4.41 When compared to a vertically integrated CP which can understand the true marginal costs of fibre strands in different part of its network and will undoubtedly follow a 'fibre-rich' policy, the introduction of regulated dark fibre is likely to drive to a 'fibre-lean' overall infrastructure.

4.42 It is evident from the examples of vertically integrated CPs in the UK, that they generally follow a fibre-rich policy and therefore this gives reasonable demonstration that this is the overall efficient solution. If, as seems likely, dark fibre drives towards an overall fibre lean solution, this is clearly profoundly different to the optimal network structure and clearly less efficient.

4.43 It should also be emphasized, that it is not dark fibre per se that causes the inefficiency, it is the regulatory imposition of dark fibre which likely denies the opportunity for bespoke bilateral contracting that is the principal source of economic inefficiency.

Upstream Pricing and Dynamic Efficiency

4.44 The above discussion principally considered efficiency at a particular point in time with technology current at that time. We see that considerable inefficiencies which may well arise through the introduction of an upstream service, especially when imposed under regulation, and that these inefficiencies are likely to be made worse by the imposition of certain regulatory remedies, notably a requirement for non-discrimination. In short, the introduction of upstream pricing is likely to introduce both allocative and productive inefficiencies.

4.45 However, it is also important to consider the impact of upstream pricing on dynamic efficiency and whether compared to a vertically integrated CP, the introduction of upstream pricing, especially when imposed by regulation, between a separate upstream CP and downstream CP is likely to impact this also.

4.46 Clearly, the vertically integrated CP has full freedom to decide its own strategy for the evolution of its network structure covering changes arising from both increases in volume and changes in generations of technology. If the vertically integrated CP is subject to effective competition, i.e. there are a number of such CPs each with their own networks, it will have full incentive to maintain a competitive network structure over time. It would therefore follow that the introduction of upstream pricing will inevitably create some level of dynamic inefficiency compared to this scenario.

4.47 The argument for mandatory provision of upstream network services is that, absent regulation, competition would not be adequate to create full effective competitive between vertically integrated CPs; the introduction of regulated upstream inputs would then allow competition downstream of these inputs and create dynamic efficiency at this downstream level.

4.48 However, this general policy needs to fully appreciated and give consideration to the impact of network structure and the consequential impact of upstream regulation with respect to dynamic efficiency.

4.49 We have observed when examining the vertically integrated CP above, that network structure is normally deeply embedded in a CP's investment strategy and that any change to network structure is normally both costly and slow to achieve. For any CP to maintain its dynamic efficiency, periodic changes to its network structure are essential and these normally require careful strategically focussed long term planning. A CP will very often seek to tie in changes to network structure with other strategic initiatives for example a major new service launch.

4.50 When considering the relationship between upstream and downstream CPs, we have seen that the pricing of the upstream service, and especially the structure of the pricing - have a profound impact of both CPs' choice of network structure. Any change in the upstream pricing structure is therefore a major strategic issue and will have a profound effect on dynamic efficiency.

4.51 Any significant unplanned changes in the upstream service and associated pricing are very likely to cause three dynamic effects:

- Cause short term productive and allocative inefficiency (over and above those described above) as the upstream supply mismatches with existing network structures, either downstream or upstream or both;
- Require significant new investment in non-strategic assets to adapt to a new optimal network structure resulting from the changed upstream service;
- Result in stranded assets where investments were made anticipating the previous upstream service and pricing would have continued.

4.52 Most significantly, all three of these will create a major distraction from a general strategy to promote dynamic efficiency by absorbing significant investment funds and management time adjusting to the new unanticipated and changed upstream service. This will therefore reduce funds and management focus to true strategic investment promoting dynamic efficiency.

Dark Fibre and Dynamic Efficiency

4.53 The introduction of a dark fibre is firmly in this category of a change to an upstream service with a changed upstream pricing structure and which has no meaningful positive impact or effect on dynamic efficiency. As we have discussed in detail in our main response, no meaningful service innovation is predicated on the availability of dark fibre. On the other hand, dark fibre introduces a completely different pricing structure for upstream supply by removing the bandwidth gradient.

4.54 In the case of dark fibre, we noted above that this will likely lead to a fibre lean overall network structure with downstream CPs investing in WDM equipment. In terms of the three effects described above, the introduction of dark fibre will:

- In the short term, by removing the bandwidth gradient changes the way the upstream CPs can recover their common costs away from the differential cost recovery which has always been considered efficient and no evidence has been offered to the contrary[17;]
- Downstream CPs are likely to make additional investment in WDM equipment to replace multiple active products;
- The upstream CPs are likely to be left with stranded assets of both fibre and terminal equipment.

4.55 None of these have any dynamic efficiency benefit but will take the industry several years to adjust to and absorb significant capital investment. During this time, the strategic focus of both management and investment is deflected away from investment which is truly promoting dynamic efficiency.

4.56 Finally, a further cost is imposed on the industry: increased regulatory risk. That Ofcom can act to propose such a profound change in regulation at such short notice, will create a climate of increased regulatory uncertainty. As with all business risk, this new climate of increased regulatory risk will come with significant cost to the industry.

[17] Indeed the opposite is the case. Ofcom's proposal to use a 1G minus pricing formula was intended to and minimise the adverse impact of removing the bandwidth gradient.

5. Conclusions

5.1 In conclusion, we see that network structure is a critical part of a CP's investment and once committed, it takes considerable time, cost, and planning to change a network structure. The choice of network structure can be very sensitive to particular features of traffic demand and particular marginal costs. Different CPs are likely to select different network structures and while different CPs can have very different network structures, is it also possible for small changes in anticipated traffic or marginal costs to make a chosen network structure significantly productively inefficient.

5.2 A competitive downstream market of vertically integrated CPs serves as a reference point of efficiency as competition will maintain productive and dynamic efficiency. Issues of allocative efficiency within the industry can be assumed to be solved from sufficient numbers of vertically integrated CPs.

5.3 The introduction of a break in the production chain between an upstream CP and a downstream CP with an upstream service traded between the two, introduces a number mechanisms for potential economic inefficiency. While the issue of double marginalisation and allocative inefficiency can easily arise, the more significant problem is the way upstream pricing can drive radical and inefficient changes in the network structure of both the downstream CP and also the upstream CP. This can create significant productive and dynamic inefficiencies in addition to the allocative inefficiency from double marginalisation.

5.4 In a free competitive market, technical and commercial bargaining with complex contracts under conditions similar to those of bi-lateral monopoly likely go a significant way to avoiding these inefficiencies. However, if the upstream service is mandated by regulation and especially if it requires non-discrimination with limited flexibility in pricing structures, the likelihood of inefficiency is very real and significant.

5.5 In this latter scenario, there is a considerable cost as the time taken for CPs to adapt their network structures means that in the short term, CPs will be given incorrect price signals leading to inefficiency and in the medium to long term will end up with both additional investments and stranded assets which do not contribute to dynamic efficiency. Indeed, the management and investment resources required to adjust to the upstream service change is deflected from proper investment to promote dynamic efficiency.

5.6 Dark fibre fits squarely in this category of a change to an upstream regulated service. It is clear from the analysis in this paper that the introduction of a mandated dark fibre remedy will offer no improvement to any allocative, productive, or dynamic inefficiencies that may be present in the current active remedies; however, it would seem certain to introduce a completely new set of economic inefficiencies.

5.7 It seems likely that dark fibre will drive towards an overall fibre lean architecture when it is clear that any efficient vertically integrated CP would choose a fibre rich strategy and assuming mandated dark fibre does indeed drive to lean fibre networks, this must be productively inefficient compared to the current active remedies which give the correct outcome of fibre richness.

5.8 However, probably the largest inefficiency arises from making such a large change in such a short timeframe and without CPs being able to plan the change into their ongoing strategy. This change will cause considerable cost with no identifiable efficiency benefit, be it productive, allocative, or dynamic. Importantly, the change will deflect management and investment resources which should be focussed on evolving network structure to promote dynamic efficiency, and as a result risks substantial dynamic inefficiency.