Ref. Ares(2016)787773 - 15/02/2016

D 2.8

Dissemination Level:

- **Dissemination level:**

  PU = Public,
  RE = Restricted to a group specified by the consortium (including the Commission Services),
  PP = Restricted to other programme participants (including the Commission Services),
  CO = Confidential, only for members of the consortium (including the Commission Services)

Abstract:

This deliverable focuses on the modelling activities developed within the DISCUS project. Models developed in other work packages are briefly described for completeness while the techno-economic cash flow model developed in WP2 is described in some detail.

COPYRIGHT

The DISCUS Consortium consists of:

Participant Number	Participant organization name	Participant org. short name	Country
Coordinator			
1	Trinity College Dublin	TCD	Ireland
Other Beneficiaries			
2	Alcatel-Lucent Deutschland AG	ALUD	Germany
3	Coriant R&D GMBH	COR	Germany
4	Telefonica Investigacion Y Desarrollo SA	TID	Spain
5	Telecom Italia S.p.A	TI	Italy
6	Aston Universtity	ASTON	United Kingdom
7	Interuniversitair Micro-Electronica Centrum VZW	IMEC	Belgium
8	III V Lab GIE	III-V	France
9	University College Cork, National University of Ireland, Cork	Tyndall & UCC	Ireland
10	Polatis Ltd	POLATIS	United Kingdom
11	atesio GMBH	ATESIO	Germany
12	Kungliga Tekniska Hoegskolan	KTH	Sweden

Name	Affiliation
David Payne	TCD/ASTON
Marco Ruffini	TCD
Christian Raack	atesio
Roland Wessaly	atesio
Andrea Di Giglio	TI
Marco Schiano	TI
Luis Quesada	UCC
Deepak Mehta	UCC
Alejandro Arbelaez	UCC

Internal reviewers:

Name	Affiliation
David Payne	TCD/ASTON

TABLE OF CONTENTS

1 Introduction

The analysis of the DISCUS architecture and comparison with alternative more conventional or business as usual (BAU) architectures requires a range of models for all parts of the network. The overall analysis used for comparisons is economic using cash flow models. Cash flow has the advantage of taking into account operational cost benefits and also the economic advantages of capital expenditure that can be associated immediately with a revenue stream, that is, a just in time (JIT) rather upfront expenditure that has no immediate revenue associated with it and is therefore risk capital.

We also describe the power consumption modelling which is closely linked to the cost modelling within the cash flow model. At the beginning of the project it was thought that this would be a fairly straight forward activity which would use the results from the component and subsystem volume calculations and data from public sources for power consumption however this has proved more difficult than expected due to the inconsistences in the literature and the lack of a physics based model for power consumption projections for future equipment. We have therefore built a model based on VLSI feature size for these future predictions. The physical layer models for the cost models required for the capital expenditure component of the cash flow models has also been more complex and difficult than originally envisaged. Originally it was expected that simple updates to models built before DISUS started would be sufficient but again more detail and a wider range of technical options was required which slowed development and extended the time required for the modelling activity.

To make this deliverable a self-contained report describing modelling activity across the project and not just the direct cash flow modelling activities carried out in work package 2, the first two sections give an overview of supporting modelling activities carried out in other work packages and reported in more detail in other deliverables.

The deliverable describes the core modelling work, the optimisation models and the access and backhaul network models and the cash flow model structures and gives some example results in section n

2 Overview of (previous) DISCUS modelling activities

2.1 Updated core architecture vision and model

This section provides an updated vision of the DISCUS core network architecture and modelling based on the dimensioning results attained in deliverable D7.7.

In D7.7 we have dimensioned the UK reference core network with 73 Metro Core (MC) nodes interconnected with 159 bidirectional fibre links using a challenging traffic forecast provided by work package 2 (57.5 Mbit/s downstream busy hour sustained rate for each PON user). For its topological, geographical, and number of users characteristics this network provides a paradigm of the DISCUS core network applicable to all large European countries and can therefore be considered as a representative network for the larger countries in Europe. It should be noted from D7.7 results that the traffic offered to the core network is very large for this challenging traffic scenario, the total core traffic is 3003 Tbit/s including the over dimensioning that derives from the resilience requirements of a MC node catastrophic failure. The traffic matrix is fully meshed and only 321 traffic demands out of 2628 are smaller than 40 Gbit/s. This value was set in deliverable D6.1 as a threshold for accommodating traffic either on the packet transport layer or on the photonic layer, but the number of demands smaller than 40 Gbit/s being so small, it seems reasonable to accommodate all traffic on the photonic layer, thus avoiding a barely exploited packet transport layer. We can conclude that, at least in the highest traffic scenario considered in DISCUS, the packet transport layer is no longer necessary. It can be possibly used transitorily when core traffic is still moderate, but, in these high traffic scenarios, transport on the photonic layer becomes much cheaper (see next section).

The second architectural revision impacts on the photonic layer node architecture. The high traffic volume makes a conventional ROADM insufficient in many MC nodes even if 1x20 Wavelength Selective Switches (WSS) are used in combination with wideband optical line systems (see deliverable D7.7). Thus, in as many as 58 nodes the adoption of a stacked Optical Cross Connect (OXC) architecture is required. The stacked OXC architecture is shown in Figure 2-1 and described in [1].

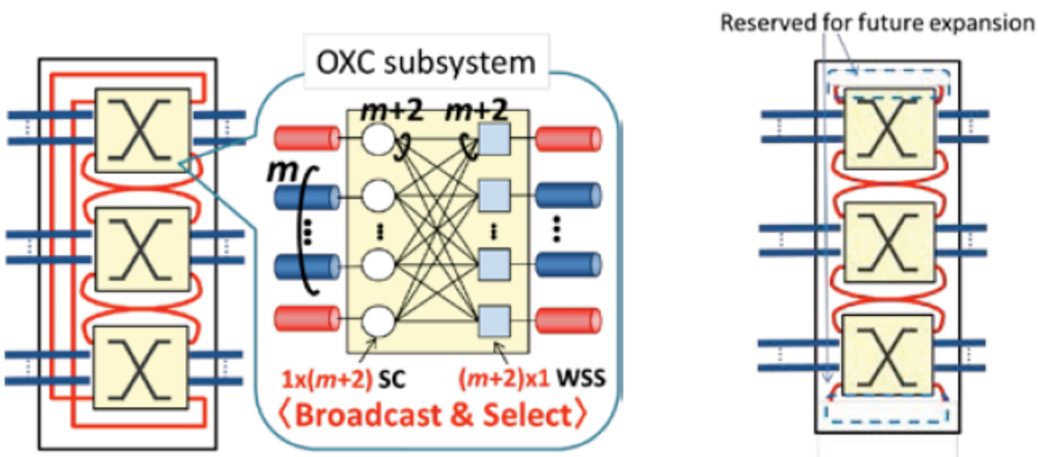

Figure 2-1 OXC stacked architecture (from [1])

Each elementary OXC shown in Figure 2-1 is a ROADM built by 1x20 WSS in *route and select* architecture as explained in deliverables D6.5 and D7.7 (the *broadcast and select* architecture shown in Figure 2-1 comes from the assumptions of reference [1] and is not applicable to DISCUS). The OXC in the stacked architecture are interconnected with each other in a ring or linear architecture to allow Optical Channel (OCh) switching between the elementary OXCs.

As stated in [1], this architecture is not fully non-blocking when an OCh has to be switched between two elementary OXC, but a careful network design can make this blocking issue negligible.

A major transmission technology change has also been described in D7.7. Due to the high number of OCh in many links, if traditional 35 nm optical bandwidth EDFA amplification is used, the number of fibre pairs required in some links becomes very high. To cope with this issue, a 100 nm wideband Raman amplification is proposed. This technology is not yet commercially available, but field transmission experiment have already been performed successfully [2]. Also 100 nm wideband WSS, although not commercially available today, seem to be in the range of present photonic technologies capabilities [3]. Of course, less challenging networks can be built with traditional EDFAs and C band optical components or higher rate optical channels to reduce the total numbers of fibre required.

2.1.1 Migration from Hierarchical to Flat Optical Core Networks

Metro- and core networks have historically been structured as hierarchical networks which consolidate and groom traffic at gateway nodes using packet processing. These hierarchical levels would typically be the backhaul network from the Local Exchanges to the first tier or outer core network nodes. The outer core nodes then connect to an inner tier of larger core nodes, which are often fully meshed. Hierarchical networks enable efficient use of transmission capacity by multiplexing traffic onto more efficient high capacity links by adding and dropping traffic at the inner core gateway nodes.

However, as traffic grows, the amount of traffic passing between any pair of core nodes increases to levels that also efficiently fills optical wavelength channels. When this occurs the alternative architectural option of a flat optical core network performs better than the hierarchical network. As mentioned above and shown in D7.6 and D7.7, cost-optimal core network topologies are flat for larger traffic scenarios. Moreover, recent results see below show that for the UK reference network and customer bandwidth >~7Mb/s (total traffic volumes of order 90 Tbit/s) flat topologies outperform hierarchical architectures.

However, today we still have hierarchical core networks. It follows that an effective and graceful evolution strategy is needed. This strategy should enable the continuation of the hierarchical core, where it is most cost effective, but it should (automatically) transition to a flat core architecture, where significant traffic increase drives up the inter-node traffic.

In this section we give an update on the results from D7.6 and introduce a migration strategy that outperforms the flat core at low bandwidths and the hierarchical topologies at higher bandwidths providing an optimal solution for all traffic demands.

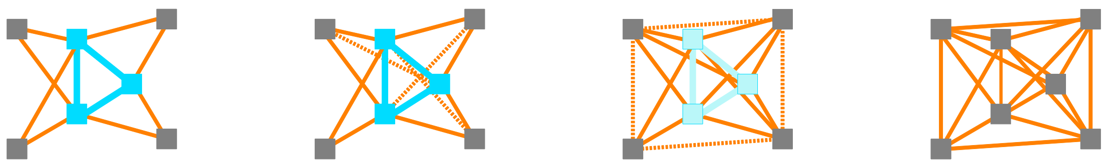

(a) (b) (c) (d)

Figure 2-2: Migration from two-level hierarchy with inner core (a) to flat optical core (d)

We first establish a relationship between user traffic growth and cost evolution for the two architectures (hierarchical and flat). We start from a two-level hierarchical topology (Figure 2-2 (a)), typical of today's networks, consisting of an outer core and a fully meshed inner core. A traffic demand between two core nodes is realized by a transmission path. Paths with more than one hop involve OEO conversion and grooming at an inner core node. Flat core networks ((Figure 2-2 (d)) instead have a full logical mesh of light paths between all nodes in the core network.

As in D7.6 and D7.7 we assume a brown-field scenario using the UK reference network with 73 MC nodes and 159 fibre/cable links. We further use the base traffic scenarios developed in

D2.4., which takes into consideration both residential and business data requirements and generates traffic considering a number of applications that use data centres, Internet peering point and peer-to-peer as data sources. The traffic model also considers the additional capacity required by inter data-centre traffic and leased lines. Instead of a fixed traffic matrix, we use the same traffic distribution and vary the total traffic volume between 20 and 1000 Tbit/s.

For switching we assume MPLS-TP switches with 400G slots and short reach line cards with grey interfaces of capacity 40G, 100G and 400G. Each MPLS-TP link is realized as an optical path using appropriate transponders at both ends that connect to the grey switch interfaces. We assume interface capacities and circuit speeds of 40Gbit/s, 100Gbit/s, 400Gbit/s. Fibres are equipped with WDM terminals at both ends and line amplifiers every 80 km. We provide 1+1 protection to ensure that services are protected against single fibre and single node failures. For cost-modelling we refer to D2.6, D7.7 and following sections in this deliverable. The values for the costs per unit of traffic are assumed to decrease over time (as true for typical price learning curves [4]).

We study a flat architecture (*flat*) among the 73 core nodes and hierarchical networks with 5, 15, and 25 inner core nodes (*twolevel-5, twolevel-15, twolevel-25*), selected among the biggest traffic sources. For all services and in both layers the optimization (described in more detail in D7.6) determines a near-optimal routing among a huge (not complete) set of potential paths using exact methods from Integer Programming. We note that in all scenarios the designated topology (flat or hierarchical) strongly restricts the solution space of possible routings and hardware realizations. In fact, the routing in the MPLS-TP domain is mainly fixed for each pair of active nodes by the given grooming gateway nodes. The optimization decides which grooming location is used for which service, how channels should be realized in the fibre topology for 1+1 protection, and which circuit speeds are used between which pair of nodes.

Results for a moderate traffic of around 30 Tbit/s core traffic volume (2 Mbit/s sustained bandwidth per user in the busy hour) are shown in Figure 2-3. Costs are categorized as switches plus interface line cards (yellow), WDM equipment (red) and optical transponders (blue). For this moderate traffic level the two main observations we make are: (i) Costs are mainly driven by transponders and switching equipment required to terminate the wavelength channels and forward the packets. (ii) The hierarchical structure with 5 inner nodes is most cost effective.

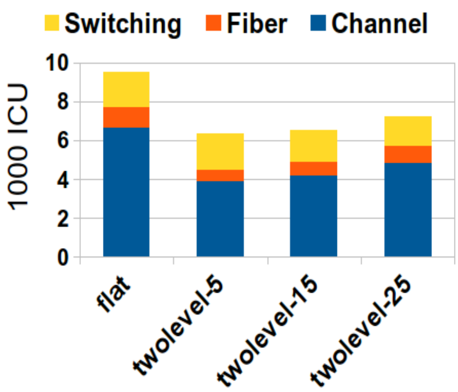

Figure 2-3: Brown field network cost in 1000 ICU, busy hour traffic per customer 2 Mb/s

However, the overall network cost is strongly dependent on user traffic growth and we thus considered a range of average busy hour customer traffic patterns up to 35Mbit/s. As traffic grows we completely re-optimize all topologies. Figure 2-4 shows the accumulated upgrade cost (applying price learning curves) with user bandwidth growth. Clearly, as traffic increases flatter architectures become more cost effective. The flat core has the largest upfront cost but also the smallest slope and outperforms all other studied network topologies when user traffic exceeds around 10-20 Mbit/s for the model parameters used.

Our results show that as user traffic increases the most cost-effective network solution moves from a hierarchical structure to one based on fully meshed optical islands. In order to efficiently migrate today's hierarchical networks we suggest the following strategy: Consider a hierarchical core network with a subset of the nodes forming a fully meshed inner core as in Figure 2-2 (a). The remaining outer core nodes are connected to inner core nodes by (at least) two (disjoint) connections. The inner core nodes are consolidating traffic from a number of outer core nodes and will be the largest traffic nodes. As network traffic grows the consolidated traffic between a pair of nodes (traffic that needs OEO conversion at both nodes) will also grow. If that traffic exceeds a threshold of, e.g., 40Gb/s a direct optical channel connection is provided over which all the consolidated traffic is sent. This threshold is in line with using 100 Gb/s light paths and is the value defined in deliverable D6.1 for accommodating traffic directly on the photonic layer but other values could also be considered which could enable the use of lower capacity light paths to be compared, this was

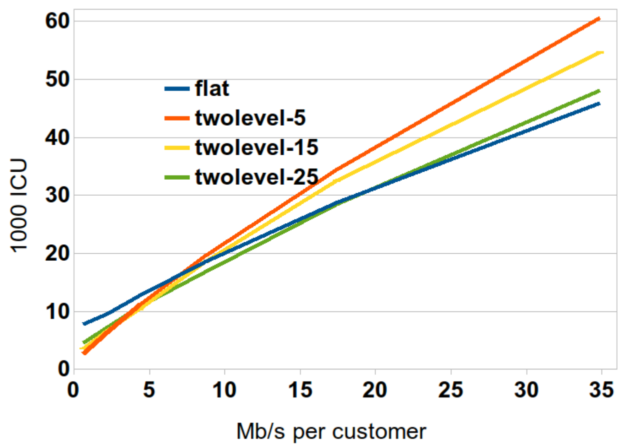

Figure 2-4: Cumulative upgrade costs for different network topologies in 1000 ICU

not done in the current modelling due to time and resource limitations.

As a result, in the first phase outer core nodes will get connected to most of the inner core nodes. Later, with traffic growth more light paths can be directly connected between pairs of nodes without the OEO conversions and eventually even smaller outer core nodes will have a direct channel connection. This process is illustrated in Figure 2-2.

To demonstrate the efficiency of this particular migration strategy, we started with an inner core of the 5 largest nodes in the UK and optimized the network for a relatively low user capacity utilization i.e., 400kb/s sustained rate per user~4x todays sustained rates) leading to a sparsely connected two level hierarchical topology as in Figure 2-2. We then increased the user bandwidth and checked the consolidated traffic for all non-connected pairs of MC nodes. If the consolidated bandwidths exceeded the threshold of 40Gbit/s new direct channels were implemented between the node pairs and the routing changed as described above.

Figure 2-5 (*migration-5-40G*) shows that this transition strategy outperforms the flat core at low bandwidths and the two level hierarchical topology at higher bandwidths, providing an optimal solution for all traffic demands. We believe that this migration strategy should now be implemented to ensure networks remain at lowest cost and lowest power into the future. It perfectly suits the transition from today's network topologies to architectures that follow the DISCUS principles leading to the lowest cost and most efficient future network.

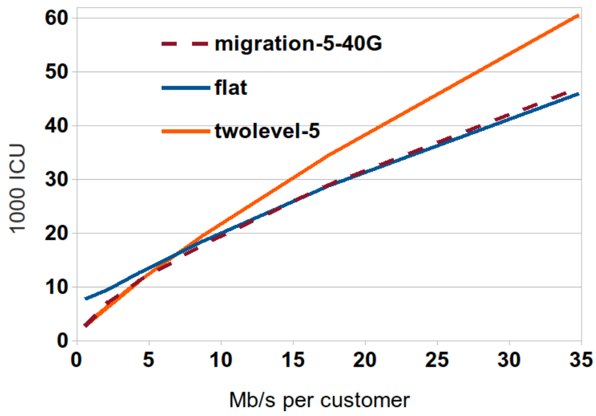

Figure 2-5: Cumulative upgrade costs in 1000 ICU over time

2.1.2 Optical equipment cost and power consumption models

The DISCUS core network optical equipment cost parameters have been mostly derived from the models developed in the EU IDEALIST project [3] (that is in turn an enhancement of the work done in the EU STRONGEST project[6]), with some more forecasts on future, innovative components introduced in the DISCUS network. Cost parameters are expressed in the IDEALIST Cost Unit (ICU) that corresponds to the cost of a 100 Gbit/s fixed DP-QPSK transponder. In general, for the costs of flexgrid component a 20% increment has been assessed w.r.t. the fixed grid components used in the STRONGEST project.

Starting from the general architecture of ROADM and OLA developed for DISCUS [7], [8], the cost and energy consumption of each functional block, namely line interfaces, add/drop blocks and line amplifiers have been obtained by adding the costs and powers required for all the cards that compose the subsystems. The results are summarized in Table 1.

In particular, the calculation on line interfaces includes the data on cost and power consumption of two 1x20 WB WSS, one WB Raman booster, one WB Raman pre-amplifier, one Control board, one Optical Performance Monitoring (OPM) card and one Optical Supervisory channel (OSC) card. The WB WSS cost has been calculated summing up the cost of three conventional WSS (each covering 1/3 of the 100 nm bandwidth), reduced by 25% as an estimation of the economy of scale arising from the component integration in one single package. The power consumption has been estimated to have a 30% increase due to control circuits increased complexity w.r.t. the present device.

Item	Cost (ICU)	Power consumption (W)	Notes

Line interfaces	3.1	260	1 dual WB WSS and 2 WB Raman amp.
Line amplifier	1.6	200	2 WB Raman amp.
Dual Multicast Switch 8x32[1]	1.3	45	1 for each group of 32 transponders

Table 1 - Power and cost model of the photonic switching functional blocks

The cost and energy consumption of the WB Raman booster and pre-amplifier used for the line interfaces have been calculated as the 75% of the cost of the WB Raman used for the line amplifiers. The OLA WB Raman price has been quoted to 0,8 ICU, about 10% less than 3 times the cost of a common EDFA plus Raman amplifier, with 35 nm optical bandwidth. The power consumption of a WB Raman amplifiers has been estimated 100 W.

The cost and energy consumption of the Colorless, Directionless and Contentionless add/drop functional block, namely the dual 8x32 Multicast switch (MS), includes 8 amplifiers whose cost and energy can be considered equal to simple, single-stage EDFA amplifiers, while the values for the 8x32 MS has been derived from the IDEALIST cost of a 16x16 MS switch.

Table 2 summarizes the cost and power consumption values for the Bandwidth Variable Transponders (BVTs) and the Sliceable Bandwidth Variable Transponders (S-BVTs), that are connected to the tributary ports of the MS switch.

The cost of the single carrier BVT has been quoted substantially equal to the ICU unit, while the power consumption has been reduced of about 25% w.r.t. to the IDEALIST cost due to power saving expected from photonic components integration.

The DISCUS forecast on cost and power consumption of the quadruple carrier S-BVT are just three times the corresponding values of single carrier BVT. This can be considered a reasonable figure considering components integration and consequent cost and energy reduction.

Table 2 - Power and cost model of S-BVTs for the DISCUS core network

1. **Single-carrier BVT**

N. of optical carriers	**N. of 37.5 GHz slots**	**Configurations (services)**	**Client signals**	**Modulation format**	**Cost (ICU)**	**Power consumption (W)**
1	1	Single 40 GE	1x40GE	DP-BPSK	0.95	140
		Single 100 GE	1x100GE	DP-QPSK	1	150
		Dual 100 GE	2x100GE	DP-16QAM	1.1	160

2. **Quadruple-carrier S-BVT**

N. of optical carriers	**N.of3 7.5 GHz slots**	**Configurations (services)**	**Client signals**	**Modulation format**	**Cost (ICU)**	**Power consumption (W)**
4	4	Single 400GE	1x400GE	DP-QPSK	3	440

[1] Estimated by TI.

		Twin 400GE	2x400GE	DP-16QAM	3.2	450
		Twin 100GE	2x100GE	DP-BPSK	3	440

2.2 Summary of optimisation models and results

As discussed in Deliverable D2.6 [12], we the following are the three high-level optimization activities within Discus (see):

- In the optimization of the optical distribution network (ODN) after assigning customers to local exchanges (LE) and fibre access network using a passive optical network (PON) architecture, is designed optimizing both the cable tree routing and the location of passive optical splitters.
- In the optimization of the Backhaul we determine the number and location of metro-core (MC) nodes so that all local exchanges can be connected to two different such MCs via disjoint fibre routes (directly or within a chain or partial ring topology).
- In the optimization of the core network we study the connection of MC nodes among each other including core fibre routing, the embedding of light-paths, wavelength (or spectrum) assignment, the location of (Raman/EDFA) amplifiers, and optimal MC node core network hardware.

Clearly as each sub-problem is solved independently, decisions made in one network may impact decisions taken in another part of the network. Therefore we took care to investigate different criteria and parametrizations leading to different sub-solutions and hence a huge set of possible combinations that could be evaluated with respect to different criteria such as

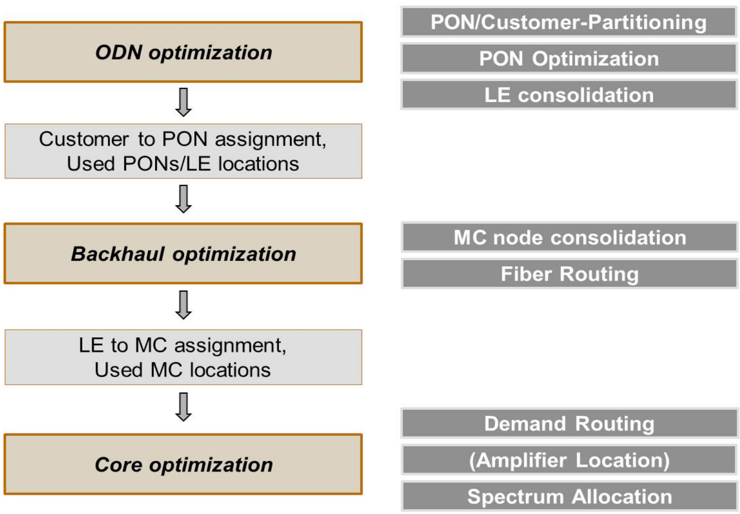

Figure 2-6 Discus end-to-end optimization process

cost, power consumption, or resiliency.

As a basis for all optimization activities we used various reference networks for different European countries. These reference networks have been developed as a combination of operator specific data such as LE locations in Italy, Spain, the UK, and Ireland and public available data such as street network structures. The networks are used to decide the location of nodes (splitters, MC nodes) and the routing of fibres in the ODN, backhaul, and core, see Figure 2-7 as an example of results.

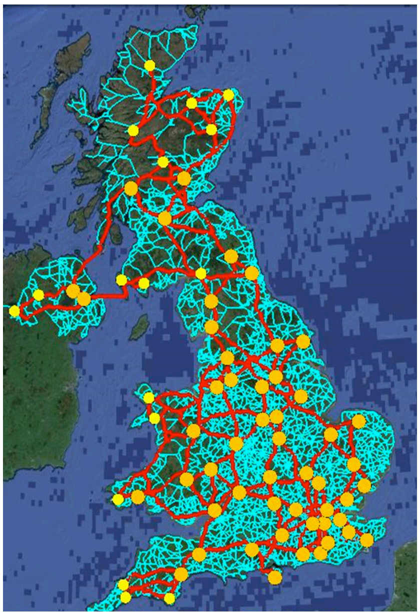

Figure 2-7: Reference networks for the UK: (i) metro fiber network

2.2.1 Optimizing the ODN

In the optical distribution network (ODN) optical fibres are routed from a local-exchange site to a set of customers forming the distribution network.

The optical signal attenuation in passive optical networks (PONs) is due to the splitter loss determined by the number of optical splits required for connecting a given number of end-users to the PON and the length of the fibre between the local-exchange site and the customer. The total path loss is one of the constraints to be observed within the optimisation of the ODN. Placement and distribution of the splitters also affects path length and of course costs. For example, placing all the splitters in a local-exchange site and then using point to point fibre to directly connect customer to the to the local-exchange leads to shorter path connections. However, the drawback is the increased cost of the total amount of fibre cable used. The optimisation therefore places the splitters to minimise the total fibre length and cost.

It is important to notice that the problem has a multi-criteria objective function: minimising the number of PONs and the total cable fibre cost. These two objectives might be in conflict, i.e., the solution providing the minimum number of PONs might not give the overall minimum cable fibre cost and vice-versa. Generally speaking, we start by first minimising the number of PONs for a local-exchange site to maximise average PON utilisation and afterwards we minimize the cable cost, and we focus our attention to three PON splitter configurations with a maximum split of 512, 256, and 128 customers located at maximum distances of 10 km, 20 km, or 30 km from the local-exchange site, respectively. Note the split of the PON is being traded for distance to meet the power budget requirements.

Ireland case study.

Ireland is the least urbanised country in Europe and is therefore an interesting case study. We have data on the Irish population of 2,189,120 customer sites connected to 1120 local-exchange sites distributed throughout the country.

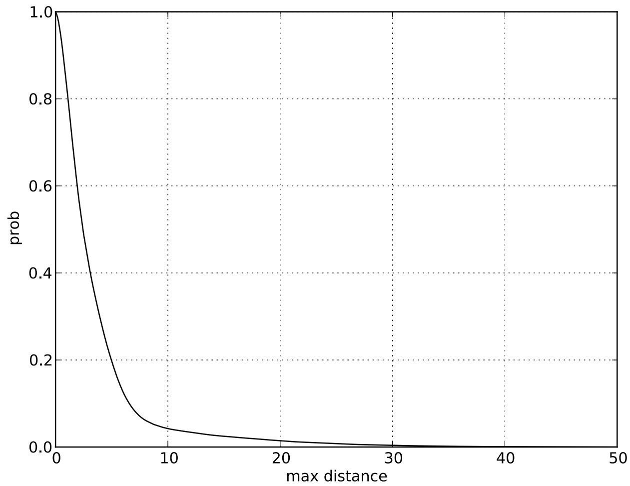

Figure 2-8: Distance from customer to closest LE

As expected the majority of the customers are located very close to the local-exchange. Indeed, as observed around 95% of the customers in Ireland are within 10kms from their closest LE see Figure 2-8: Distance from customer to closest LE. Therefore, the 512 PON type will cover about 97% of the population, and PON types 256 and 128 are covering about 1.96% and 0.74% respectively.

Additionally, for the back haul network connecting the LE sites to their corresponding MC-nodes we further optimize the cable costs by using the cable chain model. Unlike the cable tree model where each PON should be dual parented by explicit fibre connections from the PON to the primary and secondary metro-core nodes, in the chain model a set of PONs and their associated Les are chained through a fibre cable and the metro-core nodes terminate the two ends of the chain. However, in order to build chains, a set of constraints, including distance between PON OLTs and customers ONT must be fulfilled. We have observed that the chain model helps to reduce the cable link distance by 65 to 67% compared to direct connections to metro-core nodes.

2.2.2 Optimizing the Backhaul Network

The DISCUS LR_PON is based on a dual-homing architecture (see D2.1 [10] and D2.3 [11]), which assigns each LE site to two different MC nodes by disjoint fibre-paths with total LR-PON distances (ONU to OLT) up to 125 km.

In the deliverables D2.6 [12], D 4.5 [13], and D4.10 [14] we showed how to optimise the number of MC locations given a country-wide distribution of local exchanges and the task of assigning LE to MC nodes under distance and dual-homing constraints.

The core of this problem is a facility location problem that we solved using methods from mixed integer and constraint programming. The 'facilities' in this context are the potential MC node locations, and the 'clients' are the LE nodes that want to connect to the MCs.

Assignments are only allowed if they fulfil the desired conditions on distance and disjointedness. It turns out that already parametrizing the facility location models is a hard problem as it requires computation of distance-limited and disjoint fibre routes for all potential LE-MC assignments on a nation-wide scale, see D 4.5 [13] for details.

Nevertheless, we were able to present a series of optimized MC node distributions for Italy, Spain, the UK, and Ireland using the previously mentioned reference networks. We also showed how numbers and locations of MC nodes depend on the different side-constraints:

- Distance limits
- Resiliency level: single homing, dual homing, edge/node disjointedness
- Maximum/Minimum size of MC nodes in terms of connected households
- Regional versus Non-Regional LE-to-MC assignments

Table 3 provides an overview about available MC node distributions with different characteristics. The solutions in this table already stated in D4.10 [14] are the basis for all backhaul and core optimization studies within DISCUS. Very often we used variations of these solutions with different number of MC nodes for sensitivity analysis.

Country	LE-MC-km	MC size	Assignments	MCs
UK	115 km	no limit	Dual homing	45
UK	115 km	[100K,2Mio]	Disjoint Dual homing,	53
UK	90 km	no limit	Dual homing	65
UK	115 km	[100K,1Mio]	Disjoint Dual homing,	73
UK	115 km	no limit	Disjoint Dual homing	75
UK	90 km	no limit	Disjoint Dual homing	95
UK	115 km	[0,2Mio]	Disjoint Dual homing	86
UK	90 km	[0,2Mio]	Disjoint Dual homing	98
UK	115 km	[0,1Mio]	Disjoint Dual homing	106
UK	90 km	[0,1Mio]	Disjoint Dual homing	114
UK	115 km	[0,200K]	Disjoint Dual homing	308
UK	90 km	[0,200K]	Disjoint Dual homing	324
Spain	115 km	no limit	Dual homing, Regional	95
Spain	115 km	[0,1Mio]	Disjoint Dual homing	110
Spain	115 km	[0,2Mio]	Disjoint Dual homing	115
Spain	115 km	[0,1Mio]	Disjoint Dual homing	116
Spain	90 km	no limit	Dual homing, Regional	155
Spain	115 km	no limit	Disjoint Dual homing,	179
Spain	90 km	no limit	Disjoint Dual homing,	257
Italy	115 km	no limit	Dual homing, Regional	80
Italy	115 km	[100K,1Mio]	Disjoint Dual homing	111

Italy	115 km	[100K,2Mio]	Disjoint Dual homing	111
Italy	115 km	[0,1Mio]	Disjoint Dual homing	111
Italy	115 km	[0,2Mio]	Disjoint Dual homing	111
Italy	90 km	no limit	Dual homing, Regional	120
Italy	115 km	no limit	Disjoint Dual homing,	116
Italy	90 km	no limit	Disjoint Dual homing,	171
Ireland	115 km	no limit	Dual homing	12
Ireland	90 km	no limit	Dual homing	17

Table 3: MC active nodes necessary to connect customers to LR-PONs using the DISCUS architecture. We report on maximal LE-to-MC distances in kilometre, constraints on the MC size (stated as an interval of households), the resiliency level, the allowed assignments and eventually the number of MCs. In the Spain (continent) instance we removed the 6 MCs for the Balearic Islands.

In the following we review some more detailed studies that go beyond the aspects from Table 3.

2.2.3 Fibre routing and cost

In D4.10 [14] we elaborated on the aspect of fibre, cable, and duct cost in the backhaul section (between LE and MC). Based on the UK reference network and MC node distributions between 75 and 324 MCs we studied two different fibre routing principles: (i) fibre routing to minimize the total fibre kilometres and (ii) fibre routing to minimize the commonly used trail kilometres. These principles can be seen as two conflicting extreme solutions. While the motivation for (i) is to have short fibre routes, the motivation for (ii) is to share cable and duct resources.

We made three major observations, which are crucial for end-to-end evaluations of cost including the core network.

- There is significant cost offset in the backhaul for providing the cable and duct system. We estimated the cost to be at least 500 Mio EUR. Surprisingly, the cost differences between the different MC node distributions (75 nodes are even 324 nodes) are relatively small. generally the backhaul costs increases with decreasing MC numbers. However, for the same cost-model and fibre routing principle we observed a maximum cost difference of at most 150 Mio EUR. As the cost for the ODN is independent of the number of MC nodes and the cost in the core drastically decreases with smaller MC numbers there is a clear indicator to have as few MC nodes as possible.
- The difference in cost between the two extreme routing principles of minimising trails or minimising fibre distance (or any optimized solution in between) is insignificant, in particular for MC node distributions with few MC nodes (below 20 Mio EUR). In fact, keeping the number of MCs small and increasing the LE-to-MC distance means less path alternatives which together with the requirement to have disjoint fibre routes leads to a situation where the solution space is actually limited. Optimizing fibre/cable/duct routes does not pay off independently of the cost model.
- Which of the two principles (i) or (ii) is most cost-efficient strongly depends on the cost-models used. With a simple cable and duct model (cable size independent of the number

of active fibres and duct availability probability independent of the number of cables) solutions that share resources (principle (ii)) have lower cost (savings however below 100 Mio EUR). However, if a more detailed cost-model was assumed (cables installed based on the number of fibres, duct probability increases with the number of cables) then principle (ii) surprisingly led to more expensive solutions. The gain from the decrease in the cost-per-fibre was not enough to accommodate the increase in fibre kilometres.

2.2.4 Node and edge disjointedness

In the access network we have considered dual-homed solutions with two protection mechanisms. Edge-disjoint solution protects links in the connectivity between the metro-core network and the local-exchange sites, that is, the solution allows switching to an alternative path whenever a single link fails in the network. Node-disjoint solution allows switching to an alternative path whenever a node fails. Certainly, node-disjoint solutions are stronger than edge-disjoint solutions; however, a node-disjoint solution is usually more expensive due to a longer secondary path to the protection MC-node.

In this project we have observed that the actual cost of the solution varies according to multiple factors. For instance, providing only dual-homed solutions to 80% of the population and the remaining 20% with edge disjointedness (with respect to node- disjointedness) only increases the cost up to 0.3% w.r.t. to fully dual-homed solutions. Alternatively, another important factor in the cost is the population coverage; we have observed a total solution reduction of 166% when limiting coverage to up to 80% of the customers for the Italian network.

2.2.5 Trading Distance for reducing Capacity of MC nodes

Given a fixed MC node distribution and without any special constraint it is generally optimal to assign LEs to the two closest MCs to minimise the corresponding fibre/cable resources. Only when adding additional constraints such as the disjointedness of the fibre routes, minimal/ maximal customer constraints at the MC nodes or assigning LE sites to optimal backhaul cable chains, does this not necessarily hold anymore.

In this section we elaborate on the effect of relaxing the parenting constraint between local exchanges and MC nodes compared to forcing that the two MC nodes associated with a local exchange should be the closest and the second closest. In our relaxation approach we allow the two MC nodes to be further away as long as they can both cover the local exchange.

We evaluate three solutions computed using the over-provision approach presented in [13]. In this approach the extra capacity needed at each node to support single failures is minimised. Previous experiments have shown that this approach tends to balance the primary load of the MC nodes.

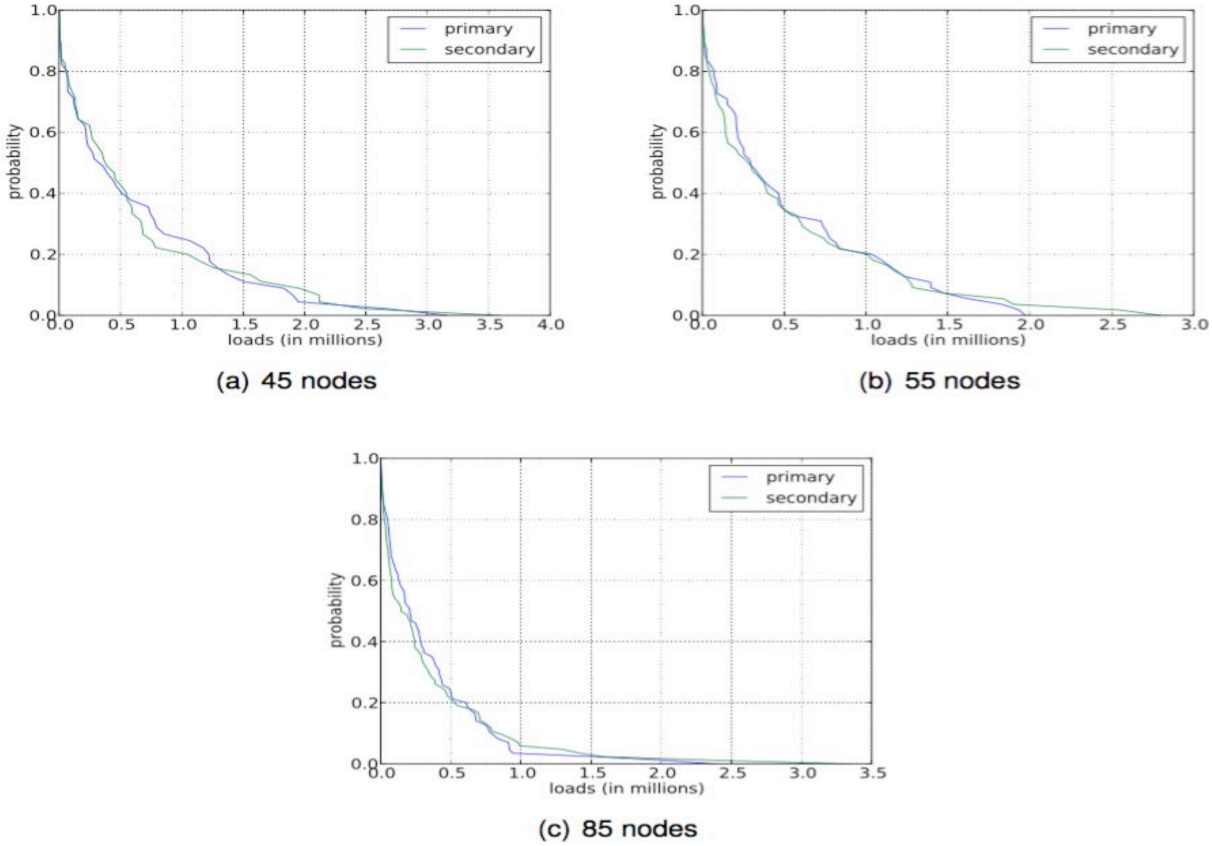

Figure 2-9 Load distribution for three MC node placements for UK using the over provision approach

In Figure 2-9 we show the load distributions obtained for three different numbers of MC nodes (45, 55, and 85) considering a maximum PON reach of 125km. Even though it is possible to double cover UK with 45 MC nodes with 125km reach, the maximum primary load is significantly above the suggested maximum target value of ~1 million. In Figure 2-9 the *x* axis is the number of customers (in millions) and the y axis is the probability of having a MC node covering at least that many customers. That is a point *(x, y)* should be read as "the probability of having at least x million customers is y". The figure shows the distribution for both the primary load and the secondary load.

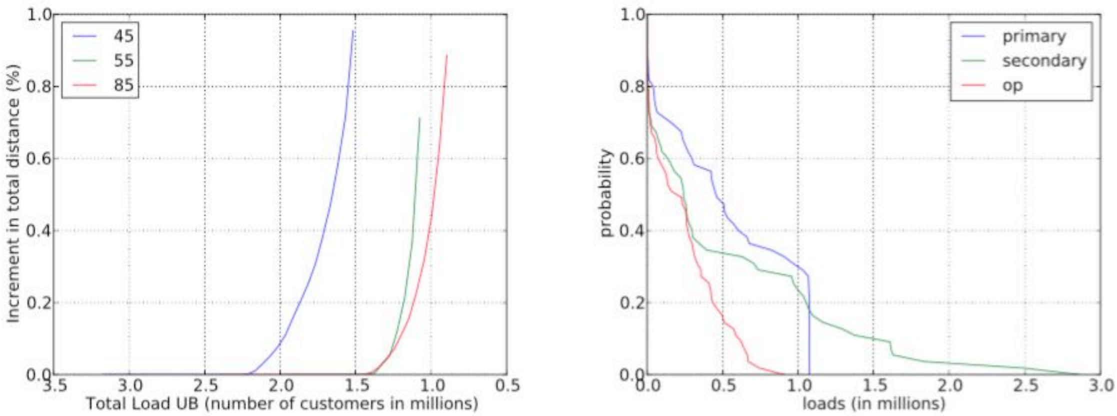

(a) Total distance vs max load after parenting relaxation

(b) 55-node solution after relaxing the parenting constraint

Figure 2-10 Results obtained after the relaxation of the parenting constraint

Figure 2-10shows the results after the parenting relaxation. As we see in Figure 2-10(a) we can substantially reduce the load without increasing much the total distance from local exchanges to their two MC nodes. A point *(x,y)* in this plot means that to enforce an upper

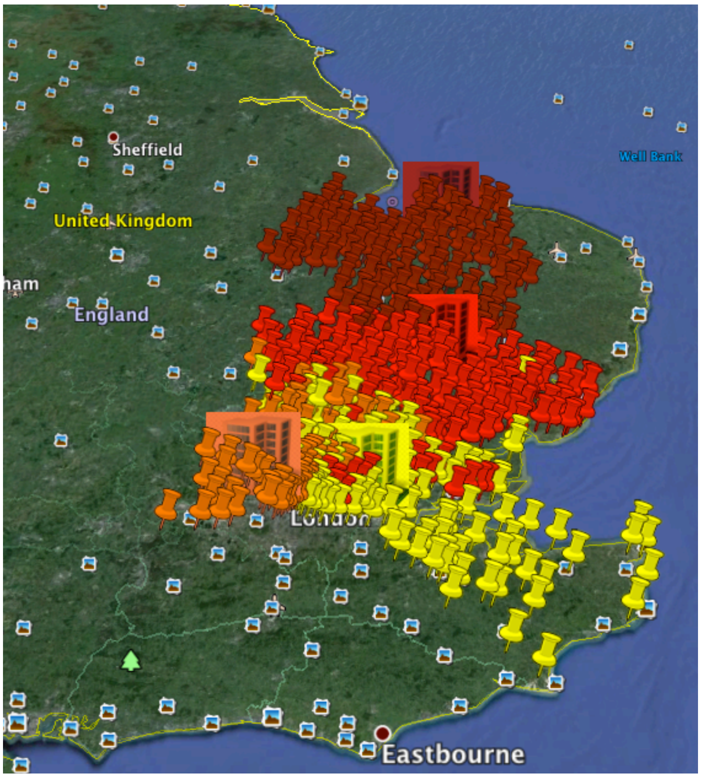

Figure 2-11 Showing how restricting the load reduces the reach of the MC node

bound on the primary load of *x*, an increment of *y%* on the total distance is required. For instance, in order to achieve the 1, 075, 137 upper bound on the primary load using the 55 nodes solution, an increment of 0.7% (less than 1%) is required. Figure 2-10(b) shows the new load distributions. This time we are also showing the over-provision distribution (op). As it can be observed in the figure, the over-provision is always below the primary load threshold. This is expected since in the worst case a secondary node will take the full primary load of one its neighbours, which is below the threshold. We can also observe that for more than 80% of the nodes, the over-provision is below 500, 000.

Figure 2-11 shows four MC nodes (represented by buildings of the corresponding colour) in the 55-node solution and the corresponding primary local exchanges. In a way we can say that, by relaxing the parenting constraint, we are adapting the reach of the MC node depending on the density of its neighbourhood. Indeed the area of coverage of the orange MC node is significantly smaller than the other ones because it is covering bigger local exchanges. Certainly one consequence of relaxing the parenting constraint is that the secondary MC node can be closer than the primary node (helps with balancing the primary load), but as both need to be connected to the local exchange that would not increase the total cable distance.

2.2.6 Optimizing the Core network

The DISCUS architecture envisages a transparent flat optical core network where all MC node pairs have a direct optical channel connection. The advantage of transparent networks is the absence of OEO conversion, which is one of the major sources of cost and power consumption in optical core networks. The disadvantage is that the number of physical connections required to interconnect the MC nodes grows with the square of their number.

Already in deliverable D7.2 [15] we learned that for typical European countries a single transparent optical island suffices to construct core networks in terms of signal reaches and flex-grid design. Moreover, in D7.6 we demonstrated that as user traffic grows beyond a threshold value the transparent flat optical core is the most cost-efficient among different architecture alternatives that are based on hierarchical structures or several connected optical islands. More precisely, to save cost it turns out that it is best to directly connect a pair of MC nodes with a light-path connection once the traffic between these nodes exceeds the capability of the existing interface capacity (e.g. 40Gb/s or 100 Gb/s) and new capacity would need to be added. see previous discussion in section2.1In the same section, we also introduce an upgrade strategy that shows how to safely migrate from today's hierarchical networks to flat optical core networks when the traffic increases.

A first challenge in the optimization of the DISCUS core is to design a fibre topology that minimises the total length of the fibre links while guaranteeing that the distance between any pair of MC nodes is within a given threshold in order to ensure transparency and guaranteeing that there exist light-path alternatives in order to ensure disjoint path protection. The second challenge is to dimension the network optimally in terms of fibres and transponders. The cost of a nation-wide core network is typically dominated by the cost of connections between pairs of MC nodes (which includes the cost of fibre deployment, placement of optical amplifiers at regular intervals etc.), and the cost of transponders placed at each MC node. We demonstrate the impact of different network designs on the cost of dimensioning such networks by considering national networks of Ireland, the UK, Italy, and Spain with different numbers of MC nodes.

The most important elements for dimensioning a transparent optical core network topology for a given a traffic matrix are transponders and optical fibres. To determine how many transponders of what types are required at each MC node one needs to decide what type of optical channels should be used for each traffic demand. To determine the number of fibres in each physical link, one needs to solve the routing and spectrum allocation problem for a given set of optical channels. Dimensioning a given optical core network topology involves finding the right balance between the cost of the transponders and the cost of the optical fibre cable. One option to design a transparent optical core network is to connect all pairs of MC nodes using shortest paths in the input graph. This topology is referred as the "Initial Size Network" (ISN). The advantage of ISN is that the average path length between all pairs of MC nodes would be minimum. This would allow us to dimension the network with transponders in the least expensive way. However, the disadvantage is that the size of the network would be bigger as the total length of links is going to be very high. In the worst-case each pair of MC nodes is could be connected directly with its own cable route which would increase the cost of optical fibre cable and installation significantly.

Another possibility is to design a network by selecting only a subset of the links L such that the total length of the sum of the link connections is minimum and the distance between any pair of MC nodes does not exceed the maximum optical signal reach. We call this network the "Minimum Size Network" (MSN). The MSN can be obtained by solving the Diameter and Degree Constrained Network Design problem (DDCND). The advantage of MSN is that it would allow sharing more optical fibre and might reduce the cost of fibre. However, the disadvantage is that the average path length would increase which might not allow the use of transponders of high capacity and might increase the number of transponders and ultimately increase the transponders cost.

Although the size of ISN is significantly more than that of the MSN the average distance between a pair of MC nodes in the former is less than that of the latter. Consequently, a wider applicability of different optical signals is feasible in ISN as shown in Table 4. We remark that there are 2 optical signals associated with each distance limit of 2430, 1170 and 500KMs respectively.

			Diameter		Network Size (KMs)				Average Path Length (KMs)		2430KMs		1170KMs		500KMs	
Network	#MC	#MC Pairs	ISN	MSN	LB	ISN	MSN	Time (sec.)	ISN	MSN	ISN	MSN	ISN	MSN	ISN	MSN
Ireland	18	306	509	749	1559	35636	1621	1.7	233	378	100%	100%	100%	100%	99%	68%
Ireland	20	380	558	830	1684	45792	1754	1.9	241	417	100%	100%	100%	100%	98%	63%
Ireland	22	462	549	968	1726	53725	1810	2.0	233	427	100%	100%	100%	100%	99%	62%
Ireland	24	552	576	1253	2004	67020	2044	2.5	243	519	100%	100%	100%	97%	98%	50%
UK	74	5402	1747	2398	6389	1577035	6940	20.2	584	951	100%	100%	94%	67%	48%	23%
UK	79	6162	1751	2397	6923	1799987	7429	20.0	584	956	100%	100%	94%	66%	47%	20%
UK	84	6972	1761	2412	6823	1990870	7256	34.8	571	901	100%	100%	94%	70%	49%	22%
UK	89	7832	1751	2429	7165	2217170	7798	30.6	566	903	100%	100%	95%	70%	50%	25%
UK	94	8742	1754	2428	6893	2401359	7523	54.5	549	904	100%	100%	95%	70%	51%	26%
UK	99	9702	1748	2422	7334	2777481	8002	44.0	573	888	100%	100%	94%	70%	49%	25%
Spain	179	31862	1524	2416	12118	8765621	13458	130.1	550	1068	100%	100%	98%	56%	45%	15%
Italy	132	17292	1779	2423	7611	5899528	9297	203.1	682	979	100%	100%	88%	62%	39%	22%
Italy	189	35532	1844	2426	12746	12078528	15699	1263.9	680	984	100%	100%	88%	66%	39%	17%

Table 4: Signal Type Distributions for Initial and Minimum Size Networks

On one hand the average distance between a pair of MC nodes is less in ISN (as shown in Table 4), which means that more optical channels of higher capacity can be used to efficiently route the traffic. Consequently the number of optical channels required by ISN could be less than that required by MSN. Therefore, the cost of transponders would be less for ISN and more for MSN. On the other hand the total length of physical connections between pairs of MC nodes is significantly less for MSN which might help in consolidating the traffic in fewer fibres. Consequently the total length of fibres could be less for MSN than that of ISN. Therefore, the cost of fibre deployment could be less for MSN and more for ISN. Certainly, there is a trade-off between the total length of the fibres and the number of optical channels. This trade-off not only depends on the total length of physical connection and the applicability of different optical signals but also on the traffic matrix and the inherent characteristics of the reference network itself.

We therefore analyse this trade-off and present the dimensioning results for UK and Italy instances with 74 and 132 nodes respectively. Our analysis is based on iteratively increasing

MSN by adding some links and demonstrates the sensitivity of the following performance measures to the size of the network: (1) percentage of pairs of MC nodes that can be served with the same signal type in ISN (or optimal reachability), (2) normalised cost of transponders and (3) normalised cost of fibres.

Reachability of UK-74 and Italy-132 networks follows the same pattern as shown in Figure 2-12 and Figure 2-13 and reveals that at most 7% of the size of the ISN is enough to serve all pairs with their optimal signal type for both. In other words, the minimum transponder cost can be achieved with at most 7% of the ISN for these two networks.

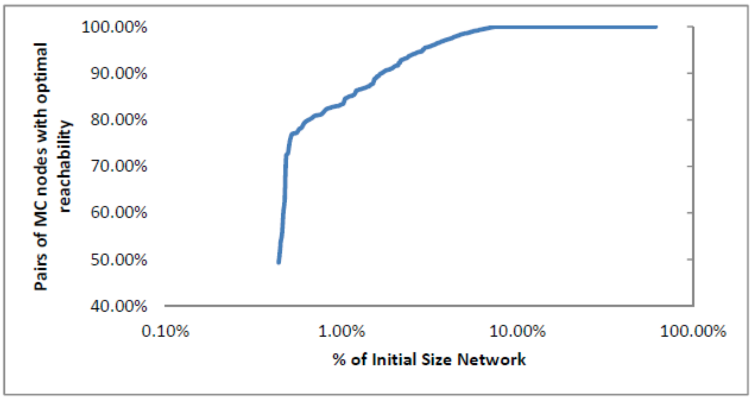

Figure 2-12 Reachability for UK network with 74 nodes

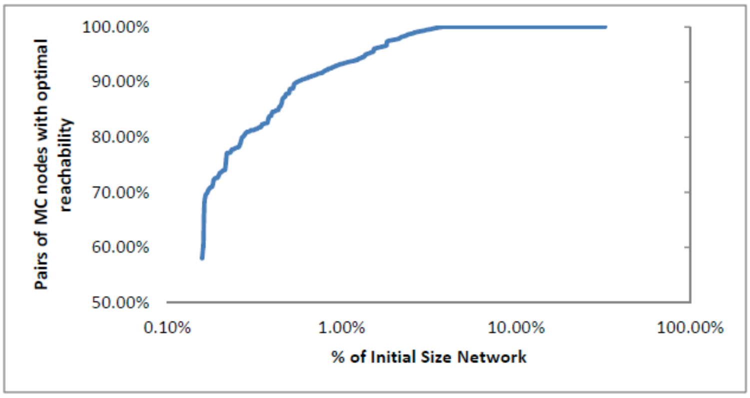

Figure 2-13 Reachability for Italy network with 132 nodes

Similar to the reachability analysis, transponder cost of each network depicted in Figure 2-14 and Figure 2-15 which are normalised with respect to the cost in the corresponding MSN design, follows the same monotonic decreasing trend and show that minimum transponder cost can be obtained with only 7% and 4% of the ISN. Furthermore, both figures show that the total cost of transponders is highly sensitive to the size of the network. Even increasing MSN by 1% reduces the total transponder cost around 10%.

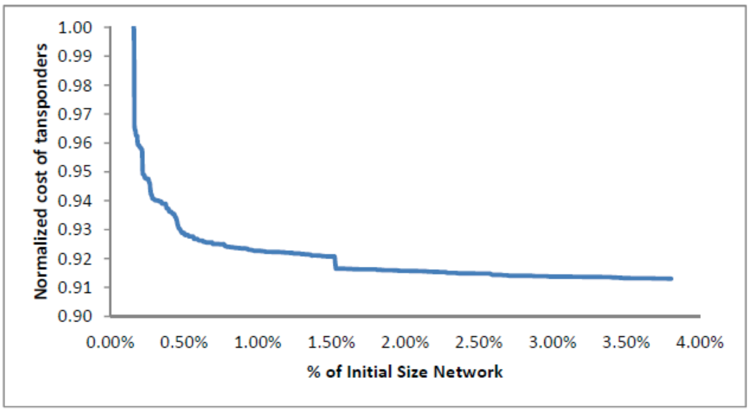

Figure 2-14 Transponder cost for Italy network with 132 nodes

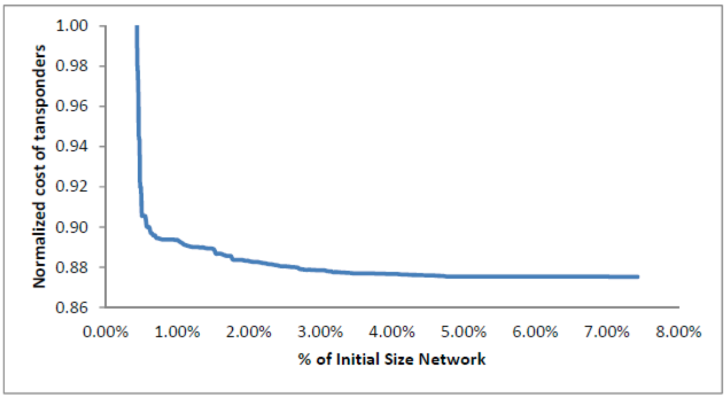

Figure 2-15 Transponder cost for UK network with 74 nodes

The normalised cost of fibres follows the same trend for both UK 74 and Italy 132 nodes as shown in Figure 2-16 and Figure 2-17. We notice that increasing MSN slightly helps to reduce the total fibre cost up to a certain point (i.e., 1% of ISN) in both networks. When MSN is increased slightly, it reduces the length of the paths for some pairs of MC nodes and eventually makes it possible to use the optical channels in ISN. Using ISN channels also reduces the number of slots required and consequently, fibre consumption too. But after increasing MSN to 1% of ISN, the number of pairs of MC nodes that can be served with ISN channels does not change with the same ratio of increase in the length of the fibres. Furthermore, sharing of fibres by different pairs of MC nodes is diminishing. Therefore, total cost of fibres starts to increase.

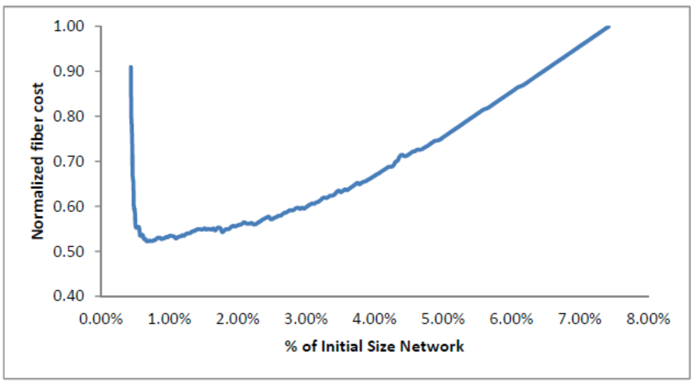

Figure 2-16: Fibre cost for UK network with 74 nodes

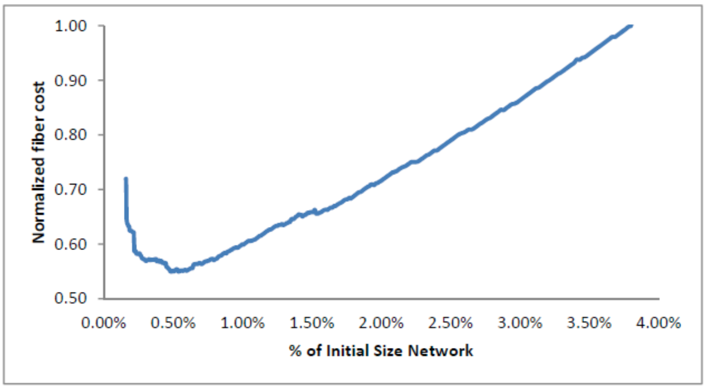

Figure 2-17: Fibre cost for Italy network with 132 nodes

In this section, we provided a dimensioning analysis to demonstrate the effect of the size of the network on the cost of transponders and fibres. Our empirical results show that the network of minimum size is providing a very good quality base solution to be extended for minimum cost of transponders and fibres. Furthermore, for networks such as UK and Italy, at most only 8% of the initial size network is required to serve all traffic with minimum cost of transponders, which is the most dominant cost component if the network is brown (i.e. existing fibres can be re-used for a significant portion of the new network). However, if networks are green (i.e. requiring new fibre build) the fibre cost can be also an important factor, but less than 2% of the initial size network results the minimum fibre deployment cost.

3 Traffic model

Forecasting traffic growth is a notoriously difficult problem. Short-term predictions can be made with some accuracy, considering previous trends, and the CISCO Visual Networking Index (VNI) has been in the past quite successful estimating traffic growth. Predictions however tend to become highly inaccurate in the medium to long term, as the correlation of past and future trends become weaker and weaker. Indeed simple extrapolation of past trends over many years in the future has proven dangerous as it is believed to have triggered the telecommunications bubble around the year 2000. The issue with extrapolating aggregated past traffic matrices is that it does not rationalise the traffic estimation with the user applications that could generate such traffic. Without such rationalisation it is easy to fall into the temptation of extrapolating exponential traffic growths too far into the future, generating scenarios that are highly unlikely to occur.

Operators have typically solved the issue by planning for growth on an annual basis using short-term extrapolations and typically a medium-term 5 year rolling plan for technology upgrades and investment planning. Planning rules for a level of overprovisioning in the network are used to counter uncertainties in the short to medium term forecasts, with good probabilities that traffic demand would not exceed the available capacity for the period considered. While metro and core networks can be gradually upgraded and increased in capacity so that the operators can react to increase in traffic as demand arises, the access network is more problematic often requiring a new generation of technology to be installed for significant upgrades. Traffic prediction constitutes an important tool for access networks, as many network operators will need to choose the next access technology to replace outdated copper pairs using xDSL lines. Upgrading the access has a high cost per user and under the present climate of ever-decreasing operating margins, operators do not want to overestimate the traffic demand and upgrade prematurely. This high cost per user has, at least in Europe, been one of the reasons that has delayed FTTH deployment.

In DISCUS we have propose a method to rationalize traffic demand predictions by generating aggregated demands in access networks from estimated statistical behaviour of individual users. While we do not pretend to be able to accurately predict what the traffic will be in the timeline considered, the tool relates traffic growth to application usage thus allowing informed discussion over the growth predictions obtained.

The traffic modelling tool developed by DISCUS has been thoroughly described in Deliverable D2.4. In this section we report the scenarios that have been investigated and potential timelines for identifying traffic growth.

In order to carry out a comparison of different access technologies over different levels of traffic demands, we use three service usage scenarios: the first for a conservative scenario for the short-term, the second a moderate scenario for the medium term and the third for the long term which is more bullish and assumes FTTH will drive significant growth in high capacity services.

The traffic model we have created is made up of three individual components. The first component considers traffic generated by users (we differentiate between moderate residential user, high residential user and small business user); the second component considers the inter-datacentre traffic; the third component considers traffic associated with leased lines (thus including medium and larger businesses).

3.1 Traffic generate by applications

Application-generated traffic is generated by a traffic modelling tool we have implemented in Matlab and recently released as open source software. This considers the traffic produced by end-user applications. For this tool, the short, medium and long-term scenarios differ in terms of applications available (each with a different data rate) and user behaviour (i.e., probability of using a given application during the day, most likely time of usage, number of application instances used during the day and their duration). As mentioned above we also define three types of user behaviour: moderate residential user, high residential user and small business user. For each application, within each scenario and user type we associated to a particular set of beta distributions. While it would be tedious to report here a full matrix with all the parameters used to define the different statistical distributions of such scenarios, we report in Table 5 the statistical means for the distributions used for the high residential user type for the third scenario. Looking at this table the reader can understand the types of applications and usage that we have for example envisaged for the longer term.

It should also be noticed that the tool differentiates applications depending on data sourcing, i.e., data centre sources, Internet exchange sources or peer-to-peer sources. Each application can be associated with any one of these sources or with a mix of them.

Table 5 Statistical distribution parameters for high residential user for scenario 3

Service	Mean Duration [h]	Mean Sessions	Avg. downstream capacity [Mb/s]	Avg. upstream capacity [Mb/s]
e-life	0.25	1.14	10	10

e-commerce	0.25	3.43	10	1
e-learning	1.9	1.9	10	4
e-social	0.54	6.67	2	1
VoD 4K	1.67	4	120	1
VoD 1080p	0.2	6	30	1
VideoConf 720p	0.5	4	30	30
Online gaming	1.67	2	20	20
VoIP	0.05	9.14	0.2	0.2
File sharing	0.83	1.14	16	4

The results in terms of average daily download and average busy hour download rate are reported in Table 6.

Table 6 Summary of scenarios parameters

Scenario	Average daily download [GB]	Average busy hour download rate [Mb/s]
Scenario1	4	1.3
Scenario2	35	7.5
Scenario3	250	58

Figure 3-1 shows an example of a daily downstream traffic pattern generated by assuming scenario 3 traffic levels on 100 PONs with a possible 512 way split but with the number of connected customers per PON selected randomly from a uniform distribution spanning the range 100 to 500 customers. the number of active users each curve in the figure represents the traffic pattern for one of the PONs, the lower curves being the smaller take up or customers per PON and higher traffic levels being high numbers of customers per PON.

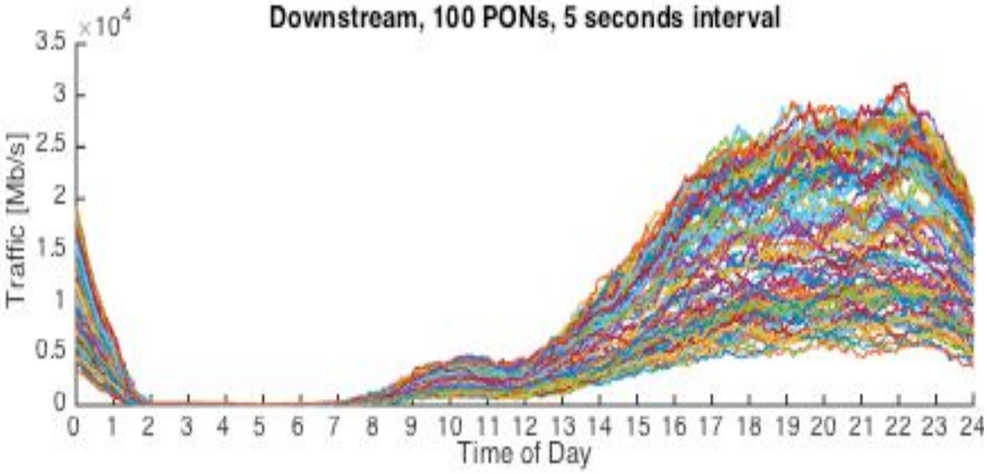

Figure 3-1 Daily downstream traffic pattern

For the purpose of the cash flow model, it is also important to try and define a timeline for such traffic growth. For this purpose we have defined three traffic growth scenarios: a pessimistic scenario, a medium scenario and an optimistic scenario, all of which we believe are potentially realistic. The pessimistic scenario puts the Compound Annual Growth Rate (CAGR) at 20%, which is lower than that predicted in the short term by the CISCO VNI (at 29% for the USA for peak times); the medium scenario considers a CAGR of 40%, which we believe can be realistic with moderate FTTH deployment; the optimistic scenario considers a CAGR of 60%. This is quite high for developed countries, but some operators, such as BT for example, have experienced such traffic growth in the past few years, this is driven by the rapid uptake of FTTCab in the UK and we believe that a wide scale FTTH deployment would stimulate at least similar if not even greater growths if services become available to exploit the additional capacity and high bandwidths supported by FTTH networks. We do not consider rates higher than this because although they have been measured in the past, they tend to be associated with short periods of unsustainable growth (e.g., during the years of the millennium telecomm bubble) and thus are not practicable as long-term forecast.

Table 7 shows the potential years when the CAGR traffic growths discussed above are associated with the service Scenarios 2 and 3 could occur if those growths are maintained. Scenario1 is not represented as that corresponds to the CISCO VNI prediction for 2018[2](averaging between Europe and the US).

Table 7 Traffic scenario timelines in relation to CAGR

Growth rate	Scenario 2 timeline	Scenario 3 timeline
20%	2028	2039
40%	2023	2029
60%	2021	2026

[2] The VNI index is only available up to 2018. While it is not our scope to define precise timelines, extrapolating the current 20% annual growth rate, would put scenario 2 about a decade away, while scenario 3 two decades away.

3.2 Inter-datacentre traffic modelling

Inter-datacentre traffic is calculated as a percentage (45%, from and extrapolation of the CISCO GCI index study ref 79[16]) of the total data centre traffic delivered to the end users. The value obtained is then added to the overall core traffic matrix that determines the traffic demand between Metro-Core nodes.

3.3 Leased lines traffic modelling

For the estimation of leased lines we have followed an approach where we consider the number of lines to be a function of the number of businesses in a given country (data for the U.K. for example was taken from refs [17] & [18] and growth in numbers proportional to overall traffic growth. The leased line capacities are then also related to the size of the business. We have used a number of points to identify curves for different MC node sizes, and then used extrapolation to fit the scenario under examination.

The points used for scenario 3 were calculated using a model within the Metro-core dimensioning and cost model based on the UK data in refs [17] & [18] and are reported in Table 8.

Table 8 Values used for extrapolating leased line services for scenario 3

MC node size [thousand of lines]	Number of 1G lines	Number of 10G lines	Number of 40G lines	Number of 100G lines	Number of wavelength services
50	1104	0	7	0	181
300	6666	0	44	0	1108
800	17785	0	117	0	2957

From the values obtained, we assumed that an average 70% of leased lines will be terminated within the same node, while the remaining 30% are directed towards other MC nodes.

These values generated by the leased line component are then added to the core traffic demand matrix to generate the matrix that is used for modelling the architecture.

4 Cash flow modelling

Cumulative discounted cash flow modelling rather than simple cost modelling was chosen as the methodology for economic comparison of the DISCUS solution, and its variants, with the alternative, more conventional, solutions for FTTH using GPON, point to point fibre or hybrid fibre copper solutions such as FTTCab. Cash flow models can give results that are better economic comparators because it produces risk indicators and time to a positive return on investment that cost modelling alone cannot do.

Capitals expenditure has two important components the balance of which can affect the business case choice despite costs being equivalent or even when the better business case has higher costs. These two expenditure components are upfront expenditure and just in time (JIT) expenditure. Upfront expenditure can be defined as capital expenditure that is risk capital with no guaranteed link to revenue return. The just in time expenditure is expenditure that is only incurred when enabling a revenue stream.

In the network build situation, the upfront expenditure is the capital required to build network infrastructure to "pass" potential customers and is usually up to the point where further infrastructure build is not shared across other customers. The JIT expenditure is then the expenditure required to "connect" the customer to the upfront provisioned network, because a request for service has been received and a revenue stream comes on line.

Upfront expenditure has a significant effect on the time to cash flow positive and the level of risk investment. While Just in time expenditure has relatively little effect on those same parameters. It is therefore economically expedient to design networks that minimise upfront expenditure even at the expense of increased JIT expenditure. To show these economic effects on business cases it is necessary to go beyond cost modelling and into cash flow modelling. Cumulative discounted cash flow models require models for; capital expenditure, operational

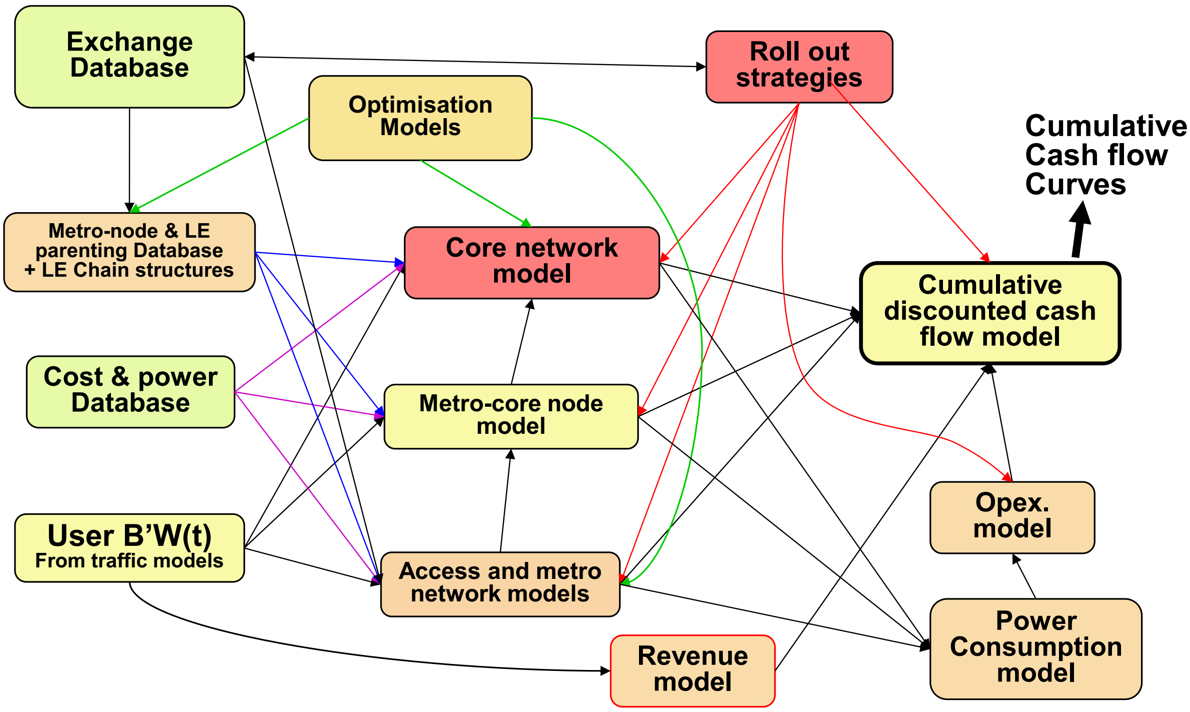

Figure 4-1 Cash flow modelling structure and process

expenditure, revenue, network rollout time lines and deployment strategies which ideally need to be included into an integrated model. The model structure developed within DISCUS is outlined in Figure 4-1, the various model components are described in the following subsections.

The model components are built as relatively independent modules that pass results and parameters between them. The models have been constructed in Microsoft Excel using worksheet functions where ever possible and VBA coded modules where additional functionality is necessary. The modules are being continuously developed and additional modules can be added in the future if new projects can be found to extend the work. For DISCUS the main focus has been to compare the DISCUS architecture consisting of LR-PON plus flat optical core with a small set of metro core nodes against a more conventional GPON, point to point fibre based FTTH and FTTCab access networks with local exchanges left in place and a metro access back haul network to connect these LEs to a set of outer core nodes.

4.1 Exchange Database

A key part of the model input data is the exchange database, network operators have detailed knowledge of their network statistics, infrastructure locations and dimensioning but this data is usually kept confidential and is not available for outside parties to use for modelling activities. Therefore for projects such as DISCUS which needs a detailed data base for the internal models these "Exchange Databases" need to be built from much sparser source data. The most basic data that can be available from public domain sources is population density data, for example from [21], road networks such as open street map and number of potential address sites (sites where a telecommunications service could be delivered to). This data is often open source and is available free of charge. There are also sources from commercial surveys, reports and studies but these can be expensive. In DISCUS we have been able to get limited data from our partners and for the UK there is a public domain data source from [22] which gives dithered locations of exchange buildings and the business and residential lines served by those exchanges. For DISCUS therefore the starting data set is typically:

Dithered Exchange site location

Number of residential sites connected

Number of business sites connected (these latter two figure may be combined)

The dithered locations only produce small errors in cable distances in general and can be ignored a bigger problem is if the LE physical cable route degree is derived from the placement of the Le onto the street maps. The dithered location may well place the Le in a road (degree 2 rather than a road junction (degree 3 or greater) or vica-versa. However in the current model Exchange degrees were simply allocated to obtain national average figures

An additional parameter that is essential is the physical serving area of each local exchange. This is not generally given in public sources and needs to be derived indirectly, one approach is to use Voronoi polygons which although they do not give true exchange boundaries can give a good approximation to the total area of a local exchange and can then be used for geotype classification. Voronoi polygons are polygons surrounding a point in a set of points (in our case the points are the LE site locations) where for a point within the boundaries of the polygon the nearest LE will be the one surrounded by the polygon. At the edge of the network the physical boundary of the country forms some of the polygon edges. (See appendix n for a methodology we use for computing Voronoi polygons)

Given the exchange area and the number of customer sites within the area a geotype classification can be applied.

4.1.1 Geotypes

Geotypes are a commonly used classification of exchanges used by operators for various modelling and analysis activities. There is no strict definition of what constitutes a particular geotype but it is usually associated with customer site density or more typically, historically, telephone line density. This latter definition is of limited value today and in DISCUS we use a site density classification system.

Once the geotypes are defined and the exchanges classified some general statistics about the exchange areas can be generated that aid other parameter calculations. For the UK the Samknows database can be used as a starting point and the geotype classification shown in Table 9 can then be defined.

Geotype	Site Density (lower bound)	Exchange Count	No. sites	Ave sites/exch	Ave site density
Sparse	0.02105500	1383.0	563,329	407	5.3
Rural 3	15.0	1405.0	1,428,985	1,017	26.4
Rural 2	45.0	1010.0	2,970,893	2,941	77.7
Rural 1	135.0	666.0	4,484,930	6,734	220.4
Urban 3	365.0	414.0	5,639,119	13,621	508.4
Urban 2	730.0	350.0	6,162,699	17,607	992.0
Urban 1	1460.0	250.0	4,458,268	17,833	1,899.7
Metro	2920.0	77.0	1,398,025	18,156	3,731.8
City	5840.0	24	384,453	16,018	7,022.5

Table 9 Geotypes definitions for UK network

We have defined nine geotypes from sparse rural to dense city. To classify the exchanges the site density is looked up and fitted into the bounds defined by the values in the second column of the table, "Site Density (lower bound)". Adjustment of these bounds effectively determines the distribution of geotypes for the exchange data set and determines the site distribution across the geotypes. The distribution of sites across the geotypes using the classification defined in table 3 for the UK network is shown in Figure 4-2.

The number of geotypes is fairly arbitrary and is chosen so that a reasonable statistical variation between geotypes exists without the edge effects of misclassifying exchanges being too important to the overall cost model calculations. Also the chosen classification is solely used within DISCUS and any similarity with classifications used by other organisations is purely coincidental.

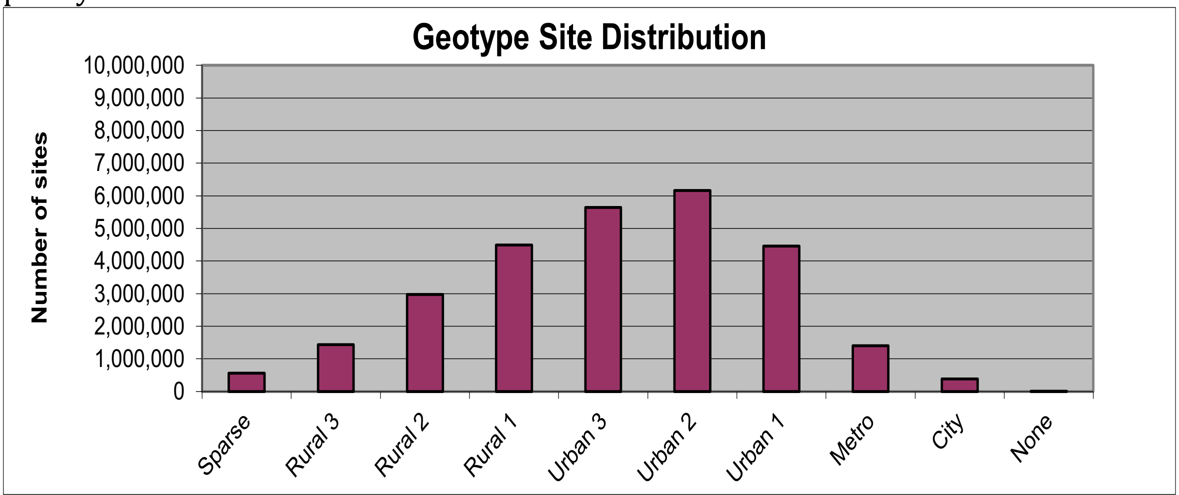

Figure 4-2 Distribution of customer sites across the geotypes using the definitions shown in Table 9.

The reader may have noticed that the lower bound for sparse rural geotype is non zero, whereas it would be expected to be zero as it is the lowest density classification, the reason it is not zero for this density is that the Samknows data set has a few anomalous entries with a few exchange having zero sites connected (possible a core exchange with no directly connected customers or an error in the data set, or possible sites that they are not local exchanges but simply a remote multiplexer parented off a local exchange site). These are ignored and removed from this classification. In the model they are captured for completeness but infrastructure is not built in those areas. There are only 7 such exchanges in the UK data set out of a total of 5586 exchanges so the omission is negligible and effectively neighbouring exchanges subsume the area of the few LEs when computing the exchange areas.

This geotype classification is then used to calculate a number of parameters that are used within the cost model. Where possible these parameters are consistent with any available national network statistics, if they are known, such as proportion of overhead plant to underground plant, total number of DPs and cabinets etc. The results in Table 10show some of the derived parameters for the UK network geotype classifications.

Geotype	PCP Degree	DP splitter size	% drop cable OH	% drop cable UG	Target PCP size (sites)	% D-side cable OH	% D-side cable UG	% D-side cable DB	Ave LE ODN degree	% E-side cable OH	% E-side cable UG	% E-side cable DB
Sparse	2	8	100%	0%	128	95%	0%	5%	2	90%	5%	5%
Rural 3	2	8	100%	0%	256	90%	0%	10%	3	80%	10%	10%
Rural 2	2	16	100%	0%	384	80%	5%	15%	3	70%	15%	15%
Rural 1	3	16	90%	10%	384	65%	5%	30%	4	55%	25%	20%
Urban 3	3	32	60%	40%	512	50%	10%	40%	4	40%	50%	10%
Urban 2	3	32	40%	60%	512	20%	55%	25%	4	20%	80%	0%
Urban 1	4	32	20%	80%	512	0%	90%	10%	4	0%	100%	0%
Metro	4	32	0%	100%	640	0%	100%	0%	4	0%	100%	0%
City	4	32	0%	100%	768	0%	100%	0%	4	0%	100%	0%

Table 10 A selection of derived parameters for geotype classification

The LE and PCP degree is an arbitrary but hopefully pragmatic allocation of the average physical cable route degree from the LE/PCP by geotype. they are parameters needed in the access cost models and is used as part of the calculation for E-side and D-side cable sizes.

The DP splitter size is the average sized splitter assigned to the geotype classification. In the real network there would be a distribution of splitter sizes to meet local variations these figures are used as an "average" size for the geotype based on site density and the need to limit average drop lengths.

A separate study at Trinity College Dublin, outside of the DISCUS project, is applying optimisation techniques to fibre cable and splitter layouts in sparse rural areas in Ireland using road layout and site location data, when this work is complete it could be applied to other countries and used to generate better statistics for sparse rural areas.

The DISCUS model computes the costs of overhead (OH) cable and drop separately from underground (UG) cable and drop and so the relative percentages of these technology deployments are needed for the different geotypes. The percentages shown in Table 10 meet nationally known statistics for the UK but the figures shown for the distributions across the geotypes are assumed.

The target PCP size is the probably the parameter that will be most different from real world statistics, this is because the figures in the table are derived to fit binary splitter sizes. To accommodate the real world, non-binary, cabinet sizes spare splitter ports from smaller cabinets would be routed over the cable chain to larger cabinets on the cable chain. It is assumed that cabinet splitters are fully utilised (or at least very highly utilised) splitter port spare capacity for growth is provided at the DP sites not the cabinet sites. The cabinet size is the cabinet splitter multiplied by the subsequent DP splitter sizes, so for example a sparse rural geotype with DP splitter size od 8 would have a cabinet splitter size of 16 to accommodate the target cabinet size of 128. The additional infrastructure column with "DB" as an infrastructure type is for direct buried cables. This is a cabling practice used in some countries and was used in the UK in the 1970s and 1980s it is cable that is directly trenched into the ground without first installing duct ways. This is a lower cost installation practice but has major implications for upgrade costs and operational costs when cables get damaged as these cables cannot easily be recovered and replaced.

There have been developments for removing the copper pairs from these cables and leaving the sheath in the ground. This sheath can then be used for insertion of a blown cable or blown fibre tube which can then have optical cables/fibre installed within them. This is still an expensive process compared with properly ducted routes but is claimed to be cheaper than installing new duct. However in the cost model we cost these regions as full duct rebuild as it is not known how well these alternative techniques work in practice or if operators will adopt them.

The average LE ODN degree is the average physical degree from the LE site, which is the number of separate physical cable routes into and out of the LE. The assumptions are that in rural areas villages tend to have a main road running through them and exchanges are near that road and cables run from the exchange in both directions along that road before branching to side roads. In denser area there is greater probability that exchanges are near junctions that increase the potential cable route degree. As with the PCP degree it is used to determine E-side cable sizes.

4.1.2 Exchange data base parameters

The geotype classification is just an aid to populating the Exchange database and ensuring that overall national statistics are adhered to. The data base has a large number of parameters that need to be calculated, many of them are intermediate parameters used to calculate the parameters used within the cost models

4.2 Database build methodologies – ODN, Cable chain and infrastructure dimensioning models

The ODN is the optical distribution network from the local exchange site to the customer premises, this is the case also for LR-PONs where the LE site is by-passed and replaced with an optical amplifier node. The ODN structure has evolved from the original copper network infrastructure and consists of an E-side which is the cable network from the LE site to the Cabinet or PCP (Primary Cross Connect) sites, the D-side which is the cable network from the cabinets to the distribution points (DPs) and the drop section which provides the connections from individual customer premises to the DPs.

Telecommunication cables in the E-side and D-side generally follow the road layouts this is the generally the case even in rural areas, but occasionally overhead cable routes will cross fields rather than follow roads to save cable length, however these are relatively rare as access to infrastructure is much easier if placed along public roads. The road layout will produce a tendency for cables to pass through successive cabinets or DPS in chain structures. With the copper network this chain structure was implemented as a tree and branch structure with larger cables near the exchange which would branch to smaller cables as cabinets and DPs were passed. A similar structure could be used for optical cables but each branching point requires a splicing joint to connect the larger cables to the smaller cables an alternative that is available for optical cables particularly with blown cable and blown fibre technology is to limit these joint housing to the Cabinet and DP sites and provide cables between these site without additional branching points. This would mean that cabinets and DPs would be provisioned and connected by dedicated cable chains.

It would require a detailed study on real network areas to determine which of the two approaches is actually the lowest cost and only operators or network owners have the detailed data for such an analysis although our optimisation work is attempting to compere optimised ODN cable layouts (see previous sections on ODN optimisation). For the DISCUS cash flow model we will assume a cable chain model rather than a cable branching model. A hypothetical example of a cable chain layout for the E-side and D-side parts of the ODN is shown in Figure 4-3 and Figure 4-4 respectively.

From these structures a number of parameters need to be derived or be given as source data for cost modelling purposes such as:

For the E-side:

- Exchange degree, this is given from the geotype definitions previously described it is used primarily to ensure that there is a least one cable chain for each LE spatial degree or cable route out of the LE.
- Average size of cabinets
- Total number of cabinets in LE area
- Number of cable chains for E side cable network
- Average spacing between cabinets, this would be derived as a radial (straight line) spacing and then a routing factor applied to represent actual cable lengths.
- Number of fibres required in each cable chain section – this is then converted to

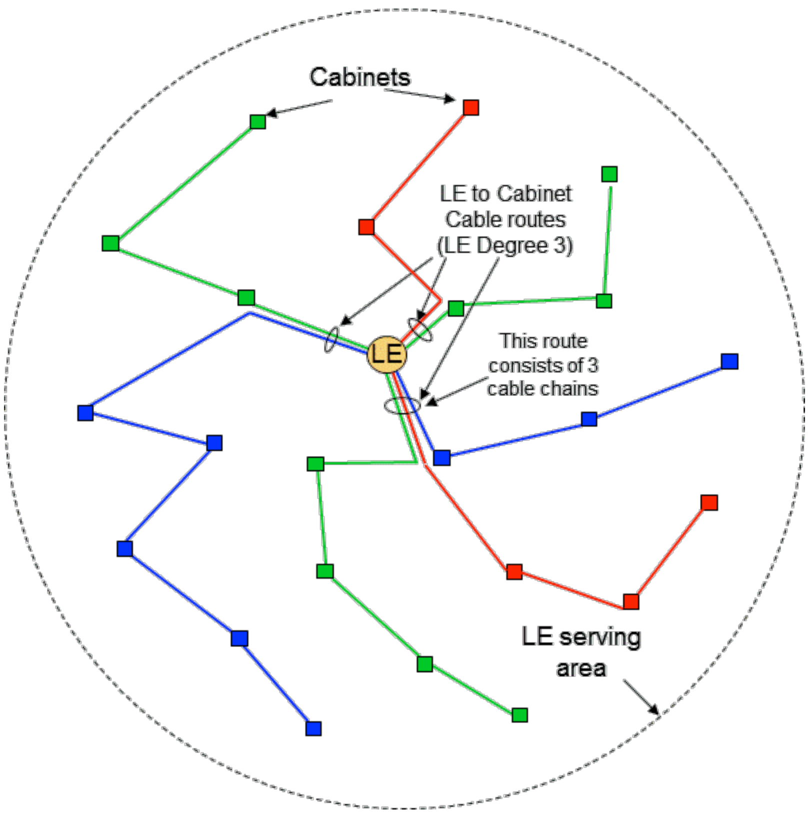

Figure 4-3 Hypothetical E-side cable chain structure

commercially available cable sizes and will include spare fibre for growth.

For the D-side:

- Average cabinet degree
- Average size of DPs, this is determined from the geotype and the binary number size assumed for DP splitters within the geotype
- Total number of DPs
- Number of D-side cable chains
- Average spacing between DPs
- Number of fibres in the DP cable chain sections

And for the Drop:

- The proportion of UG and OH drops
- The average lengths of the drops

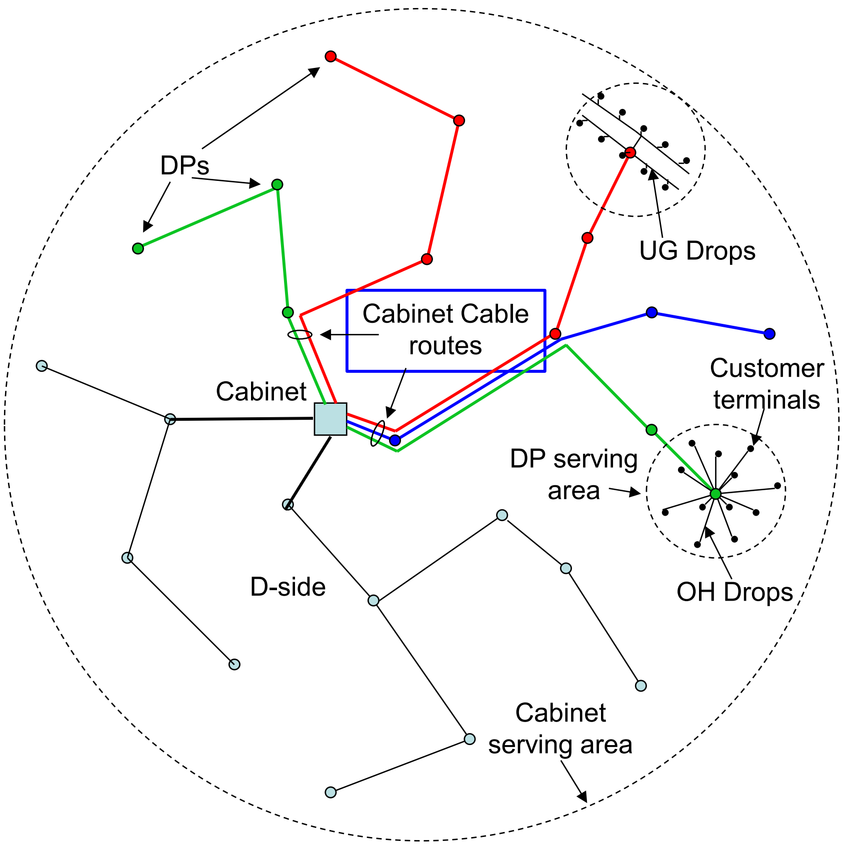

Figure 4-4 D-side cable chain structure

A full list of modelling parameters in the exchange data base is given in Appendix 8.3 note that many of these parameters are intermediate parameters and not directly used in the cost models but are required to ensure that the model parameters are internally and self-consistent.

To simplify the analysis and to meet the assumption that only minimal initial data will be available, a number of simplifying assumptions are made to aid calculation of the critical network parameters:

- Uniform customer site density within the LE area, D customers/km2,
- Circularly symmetric LE, Cabinet and DP coverage areas
- Nominal optimal triangular packing of infrastructure nodes within the coverage areas, *essentially this is hexagonal packing with nodes at all vertices*

These assumptions are illustrated in Figure 4-3, Figure 4-4 and Figure 4-5.

Note that although the cabinet and DP locations are assumed to be similar to those used for the copper network, when considering an existing infrastructure situation, this does not need to be the case for passive optical network installations particularly the LR-PON case with lean fibre infrastructure. Efficient fill or utilisation of the PON splitters is an important consideration in the design of the physical layout for the optical network and therefore the optimal location of the splitter points may not coincide with the copper network infrastructure locations. Further work to analyse and compare the strict use of legacy locations and more optimised locations for cabinet and DP particularly in the LR-PON scenario would be necessary to determine and quantify the advantages, however this analysis requires detailed knowledge of existing infrastructure so would probably need to be carried out by the network operator and is beyond the scope of the DISCUS project. But note that moving the splitter locations away from copper cabinet and DP locations is not a difficult operation as the splitter housing can fit in UG footway boxes and do not need the street furniture associated with the copper infrastructure, so that the surface cabinets used as copper pair cross connect points are no longer required and can be removed when full FTTH deployment has occurred and the copper legacy has been removed.

The drops to each customer are the cable infrastructure between the DP and the customer premises network termination equipment (NTE) (note the NTE can be the ONT or ONU if the operator owns this equipment. Otherwise it will be a wall box where the external fibre is terminated and connected to internal fibre for connection to an internal ONU (note that internal cable materials are different from external cables due to environmental and fire regulation requirements and external cable should not be used internally). Generally the drop and NTE are provided only when the customer takes services. To “pass” a customer the operator only needs to build network infrastructure to the DP. The drop and NTE installation can be referred to as just in time (JIT) provisioning whereas the infrastructure to pass customers is upfront or pre–provisioned and as discussed previously this has an impact on the cash flow and economic viability of specific network designs.

There are two main drop cable systems in use within communication networks; underground drop cabling (UG drops) and overhead drop cabling (OH drops). A common cabling infrastructure for UG cabling is called “frontage tee” where a cable or cables are run from the DP along the street or frontage of premises and tee junctions are used at each of the premise passed to route drops to the individual premises. This is illustrated in the UG Drops diagram in Figure 4-4. The OH drops use a pole with a radial point to point drop cable system to all the premises within some maximum distance typically around 70 metres from the pole. The drops can be longer than this if multiple pole hops are used but rarely exceed three additional pole hops (except in sparse rural areas where very long drop pole routes can be found. The DP can be fed by either aerial cable or UG cable the latter being typical in built up areas and the former being more common in rural areas. This system is illustrated in the OH drops diagram in Figure 4-4.

From the above it can be seen that there are three infrastructure coverage areas:

- The DP coverage area
- The cabinet coverage area
- The LE site coverage area

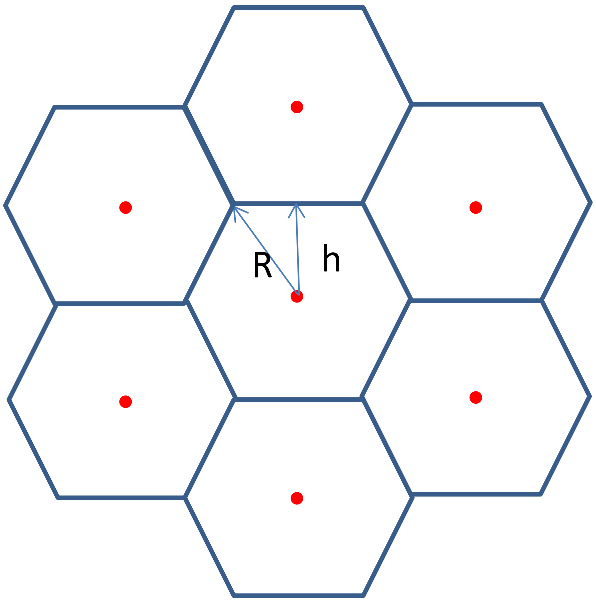

Figure 4-5 hexagonal optimal packing assumption of infrastructure nodes

Assuming the base source data has the number of customers in the exchange area and we also have an estimate of the physical coverage area (either given or estimated e.g. using Voronoi polygons) then assuming a uniform distribution of those customers within the exchange area estimates for the infrastructure node separations and sizes of the nodes can be made using geometrical arguments:

Consider a set of infrastructure nodes optimally packed such that they form a triangular (or hexagonal) optimal packing structure as shown in Figure 4-5.

The node capture area is $A = \frac{3\sqrt{3}}{2} R^2$

The separation between nodes is $2h = \sqrt{3}R$

The number of infrastructure nodes or customer sites captured by the infrastructure node coverage area $N_P = AD_{sites}$

Where D_{sites} = the customer site or infrastructure node density. If it is site density then it is given by $A_{LE}/N_{LEsites}$ where A_{LE} = the LE serving area and $N_{LEsites}$ = the total sites within the LE area.

Using the geotype classifications and taking into account the binary numbers for the splitter size at the infrastructure node locations the various areas of the infrastructure serving areas can be estimated and then the values for the separations calculated.

The pattern of node locations in Figure 4-5 looks very idealistic and indeed it is, so the question therefore arises; does it bear any relation to real networks for cost evaluation purposes? Fortunately it does as long as cable costs are linear functions of length which in practice is usually the case. In appendix 1 it is shown that even with an arbitrary statistical distribution of cable lengths within an infrastructure node serving area (varying cable lengths is equivalent to a distribution of node spacing's which more in line with real world situations), if the cost of the cable is a linear function of the form a+bl.

Where: a is the length independent cost such as splice housings and splicing etc. b is the length dependent component of the cable cost and l is the length so that the the total cost is given by $C_T = aN_p + bN_p\mu_l$.

Where: Np is the total number of sites served in the serving area and μ_l is the mean of the length distribution. So although the separation of the nodes 2h is a random variable the only parameter required is the mean and for a uniform distribution with tight packing as assumed then $\mu_l = 2h = \sqrt{3}R$.

The density of points in the various serving areas can be considered as a hierarchy starting at the customer site level then the DP nodes, then the cabinet nodes and finally the LE sites for the backhaul network to the MC-nodes. The density of the DP and cabinet nodes will therefore depend on the average size of these nodes.

Using this hierarchical concept we can build a table of densities and distances as a function of geotype statistics. In the actual cost model each exchange is dealt with in turn and the actual densities for each exchange is used for these calculations.

4.2.1 DP serving area:

There are a number of parameters in the exchange database for the DP serving area, three of which feed directly to the cash flow cost model these are: The relative percentage of OH and UG drops and the average length of the drop, the other parameters aid derivation of the length parameters for OH and UG drops.

From the base data the exchange area and number of customer sites is known and from the geotype classification the DP splitter size is defined. Applying a fill factor for future growth an average number of sites per DP is obtained. Using the total number of sites, the number of DPs for the exchange area can be calculated from which the average size A_{DP} of the DP serving area is obtained. Given the DP serving area the radius R_{DP} assuming a circularly symmetric serving area is calculated.

That is the number of customers per DP = $N_{DPCust} = S_{DP} * f_f$

Where S_{DP} = DP splitter size and f_f = the design fill factor for future growth

Therefore the DP serving area radius is $R_{DP} = \sqrt{\frac{A_{DP}}{\pi}} = \sqrt{\frac{N_{DPCust}}{\pi D_{Sites}}}$

Mean DP to customer radial distance $\overline{R}_{DP} = \frac{2}{3} R_{DP} = \frac{2}{3}\sqrt{\frac{N_{DPCust}}{\pi D_{Sites}}}$

The factor 2/3 comes from the mean distance from a point to all other points in a circular segment of randomly distributed points about the point, see appendix 8.2.

Once the mean drop radius is known then the drop cabling cost can be derived. OH drops can use the radial figure directly as OH drop are radial distances from the DP, for UG drops a routing factor of $\sqrt{2}$ is used for the average drop length.

Note that OH drops in rural areas can be longer than the allowable single span length from a pole (about 70 metres) and more than one pole is used. However for simplicity in the model it is assumed that the pole infrastructure is already in place for the existing copper network and is reused for the fibre drop installation and therefore only the fibre drop cable is included in the cost model.

4.2.2 Cabinet serving area:

The same approach can be used for the cabinet serving area which is also often called the D-side network for "distribution side". The Cabinet is a flexible cross-connect point required for the original copper network. It is not a required flexibility point for PON networks but is a convenient splitter point as D-side cable routes come together at these points in the physical duct and cable network.

The D-side cable network can be considered to be equivalent to the drop network for the DP serving areas but the points in the network to be served by the cabinet locations are now DPs rather than customer sites. The density of DPs is the customer site density divided by the average number of customers per DP i.e. $D_{DP} = \frac{D_{Sites}}{N_{DP}}$

The Cabinet size in terms of customer sites is given in the geotype table and is an integer multiple of 128 to facilitate a 100% fill target of cabinet based optical splitters, that is each D-side splitter port at the cabinet will connect to a DP in the cabinet serving area. It needs to be recognised that these cabinet sizes may deviate from the original copper cabinet sizes (we are assuming that this original data is not available, which is generally the case except for the incumbent operator). This potential difference in cabinet size can be accommodate to some extent by passing spare ports from one cabinet in the E-side cable chain (discussed later) on to other cabinets, in the same chain, that may be short of ports. Also for small split PONs in the more dense LE areas there may be no cabinet splitters, all the PON split being accommodated at the DP locations (this can be typical for 32 way split GPON systems).

Let the cabinet size = N_C; where N_C is the number of customers (optical terminations) in the cabinet serving area.

The cabinet coverage area is therefore N_C/D_{Sites} km^2.

And the radius of cabinet serving area $R_C = \sqrt{\frac{N_C}{\pi D_{Sites}}}$.

The mean cabinet to DP radial distance can be calculated in a similar manner used for the average DP to customer radial distance calculation.

Let the number of DPs in the cabinet area = $N_{DP} = N_C/N_{DPCust}$

Then the average radial distance cabinet to DP $\overline{R}_C = \frac{2}{3}R_C$

However for the D-side network cables will not be provide directly from the cabinet to the DP but rather there will be cable chains from the cabinet, that serve DPs as the cable chain route progresses from the cabinet location to the edge of the cabinet serving area. This requires two additional parameters to be found: The number of DPs in a cable chain and the average separation between DPs. With these additional parameters cable chain lengths and cable section sizes can be derived and then costs determined.

The average separation (l_{DP}) of DPs can be determined using the method described above in the introductory sub-section and $l_{DP} = \sqrt{3}R_{DP}$

Determining a value for the number of DPs in a cable chain is more problematic as there are no rigorous rules or approaches for doing this.

The method used in the model for calculating the number of DPs in a cable chain is to select the lower of three bounds:

- The first bound is calculated for the specific LE by taking the number of DPs per cabinet and dividing by the cabinet degree (number of physical cable routes out of the cabinet). This ensures that at least one cable chain per degree will exist but will also produce the longest chains as the DPs are shared across the minimum number of cable chains for the cabinet serving area. In general these chains would be longer than lengths derived by the other methods and will be rarely selected within the model
- A second bound uses the average cable route length from cabinet to DPs at the edge of the cabinet area divided by the average separation of DPs, this gives a more pragmatic value for the number of DPs expected along a typical cable chain route.
- A third bound is an arbitrary maximum bound of 8 DPs per cable chain a limit not expected to be seen but a bound that aids implementation of the code of the Excel model.

The minimum of these bounds is selected for the model value. In practice the calculations in the model nearly always selects the value from the average cable route length divided by the DP separation.

The above methodology derives a working value for the number of DPs in a D-side cable chain. This is generally a fractional value and of course only integer numbers of DPs can physically exist in a cable chain. This is handled in the model by using two lengths of cable chain with one chain length being the integer part of the average number of DPs per chain and the longer chain being this length plus 1. The ratio of the two DPs/chain numbers is calculated to fit the average value of DPs per chain previously derived.

The cable installation techniques in the D-side network are Overhead, Underground and Direct Buried (DB). Direct buried is an alternative underground cable installation technology that was used in the 1970s and 80s and was a low cost installation by directly burying the cable into the ground without pre-installing duct. Although this has lower installation cost and therefore reduces upfront capital investment it has serious operational problems in that it is difficult to maintain and very expensive to replace. Because of these limitations it wasn't used to a great extent after these early years when it was fashionable but unfortunately it is scattered across the network and there are generally not good records kept of where it was installed. It is also the case that not all operators adopted the practice and the extent of deployment varied greatly between different countries.

To accommodate DB deployment of cable a percentage figure for different geotypes is incorporated into the geotype table and these percentage figures are fed into the cost model together with the OH and normal ducted UG percentages so that all cable routes in the exchange areas of geotypes where it is determined to have been deployed will have a mix of these cable installation types. The percentage allocations are consistent with the national statistics for this cable type deployment.

Using the derived results for the D-side statics within the exchange data base the costs of the DP splitter installations and the cable installed in the cable chains are derived as part of the access network model

4.2.3 Local Exchange serving Area

The E-side network is that part of the ODN from the local exchange site to the cabinet locations. The modelling requires very similar parameter values to those for the D-side network and the majority can be calculated via the same processes/methods. In this case the nodes are the cabinets and the cable chains are from the local exchange site outwards to the local exchange area edge picking up and serving cabinets along the E-side cable chain route.

One parameter that is different for the E-side is the occurrence of direct exchange lines (DELs). These are a hangover from the copper network and are copper lines connected directly to customers without going through a cabinet but instead go directly from the LE to a DP and then to the customers premises. For the long reach PON solution these can easily be accommodated by placing a cabinet splitter at the LE site with a cable directly to the associated DPs (it could well be that this cable could be a shared cable with other cabinets, that is the first cabinet in this particular E-side cable chain is the LE site. The problem for DEL connected customers is for FTTCab because VDSL signals cannot be fed over copper cables from the LE while ADSL is present. This stops those customers getting the high speed broadband services offered by FTTCab even though it is available within their exchange area and they are the closest customers to the exchange.

Infrastructure costs can be calculated in a similar manner to the D-side infrastructure costs as previously describes but now using the E-side data parameters within the Exchange data base.

4.2.4 Exchange data base summary

The Exchange data base is constructed assuming that very sparse real data is available as detailed infrastructure data is usually kept strictly confidential by the network operator/owner. However some basic data can often be found and our starting assumption is that we get a minimum of dithered exchange site locations, and an indicative value of the total number of customer sites connected to the local exchanges. We also assume there are some global statistics available about the network that can be used to help proportion infrastructure types across the exchanges and also help classify the exchanges into geotypes. The most critical other parameter required is the Exchange serving area size. If this is not given then it must be derived in some way, for DISCUS we use Voronoi polygons as a reasonable methodology for getting pragmatic exchange area sizes.

All other infrastructure parameters are then derived as described in the previous sub-sections for a full set of the exchange data base parameters with descriptions see appendix 8.3

4.3 Access and metro-network

This section outlines the main features of the physical layer model used for cost modelling purposes within the cash flow model. For a fuller description of the current modelling methodology see appendix 4 note this is a snap shot of the modelling process as this is continuously being updated and extended and will continue after DISCUS finishes. The description starts at the customer termination options and progresses up the network towards the LE site.

4.3.1 Operator owned ONU and Drop model

The diagram in Figure 4-6 shows the model scenario we are assuming when the Operator owns the ONU and NTE. In this case the operator usually installs the ONU on the outside of the premises and also provides battery backup, power is usually supplied from the customer's premises. The ONU connects via a copper cable to a copper technology NTE inside the customer premises which the customer can plug his equipment into. This copper NTE is the physical and legal demarcation point between the Operator's network and the customer's network. The ONU forms the boundary of the optical network and therefore there is no separate optical NTE. The ONU is connected to the copper NTE inside the customer premises by a copper cable e.g. Coax, UTP, Cat5 etc. An electronic router/hub, which could be customer owned, would plug into this copper NTE (the router function could also be part of the ONU

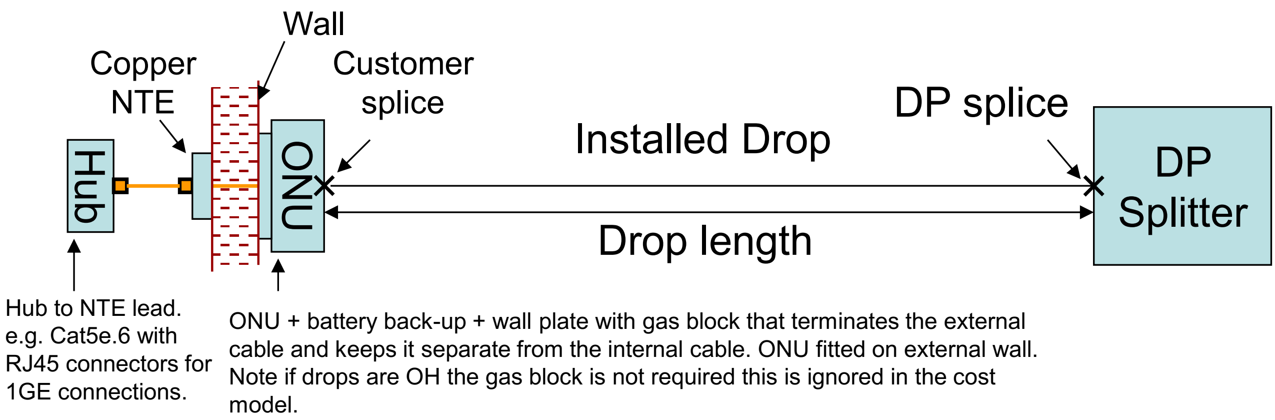

Figure 4-6 Operator owned ONU and NTE model structure

functionality if desired).

The model includes the cost of the above installed equipment with the exception of the customer hub, this can be optionally included to get a complete system cost (service cannot be received without it). In the current version of the model it is not included.

Note: all the above costs are incurred on a "just in time" (JIT) basis in the cash flow model and are therefore associated with a revenue flow not upfront risk capital.

The NTE and ONT are included within the CPE components in the model. The LR-PON ONU price = GPON ONU - GPON optics + drive electronics + 2.5 * GPON optics + drive electronics. Currently in the model the GPON optics + drive electronics is assumed to be 25% of the GPON ONU cost.

4.3.2 Customer owned ONU and Drop model

In this version of the model, illustrated in Figure 4-7, we are assuming the customer will own the ONU and connects it into an optical NTE owned by the operator which forms the physical and legal demarcation between the operator's network and the customer's network. Note: in this scenario it is assumed that battery back-up is the customer's responsibility and they bear the cost if they choose to implement it e.g. the ONU could be supported by a customer owned UPS device or customer supplied batteries if they so desired. Because the internal ONU does not have to be as rugged or operate in such a harsh environment as the external ONU the cost of the External ONU is set 30% higher than the internal ONU. Note that the wall plate is always fitted to separate the external cable from internal cable but the gas block is only required for UG drop cable feeds this small cost difference is ignored in the model. The choices for ONU ownership strategy can be toggled within the model to enable comparison of the economic impacts of the two approaches.

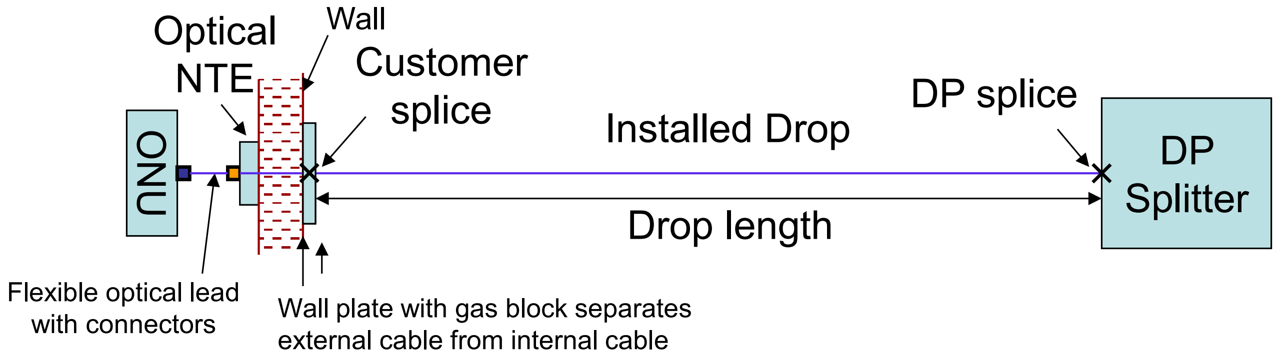

Figure 4-7 Customer Owned ONU model

4.3.3 Cable chain models: dimensioning and growth

Using cable chains assumes that cables will run along streets and drop off fibre to service DPs and Cabinets as required. As discussed above the number of cable chains and the number of DPs or cabinets in a cable chain are derived within the Exchange data base used within the cash flow model. These values are based on geometrical assumptions and uniform distributions of customers within the serving areas of street furniture and building nodes. For simplicity the first version of the model will use separate cables for each chain. This could be modified later to assume branching points along physical duct routes that share cable chain routes. In these situations there may be advantages in combining smaller cables into larger cables which could reduce per fibre costs further if the cables have not already exceeded the maximum cable size used in the model (276 fibre cables). However these savings are offset by the need to have a cable joint at the cable branching points it these do not occur at the cabinet and DP locations so we expect relatively small costs difference to occur in practice.

With the exchange data base populated as discussed in previous sub-sections the cable sizes and lengths for all the cable sections in the exchange area are be computed. The drop section is computed first then the D-side cable chains of DPs then the E-side cable chains of cabinets. The backhaul chains of local exchanges that span a pair of metro-nodes are computed using optimisation modelling described in section 4.3.7.1. The cable model needs to be generic enough to accommodate LR-PON, GPON (and other PON variants that could be added later) and also point to point fibre cable solutions, For the FTTCab model a point to point fibre E-side model will be computed. G.Fast solution could be added at a later stage but this is beyond the resources available in DISCUS and would need to be added in follow on projects, however the modular structure of the model should readily enable additional technology modules to be added after the end of the DISCUS project.

The main complexity is deciding on the splitter housing structure which in addition to the splitter and associated splices must accommodate the spare fibre handling which includes bespoke fibre and fibre for future PON growth. The fundamental assumption here is that there will be an allowance for growth in the day one build and dimensioning. This is provided for by including a fill factor across the model. Typically this is set at 80% allowing 20% spare capacity for growth. However in PON solutions the number of spare fibre needs to reduce as the cables run from the customer premise through the D-side, then the E-side and through the backhaul network to the core. This is because most growth will be randomly distributed and not all growth provided near the customer needs to be translated up through the network. Also for PON solutions additional splitter nodes would be added which reduces upper network fibre counts. Making this sparing scalable and internally self-consistent adds further complexity to the model and it is necessary to get the sparing strategy and methodology clearly defined upfront.

To make spare resource effective, access to spare fibre needs to be available at the access edge i.e. at the DP position. However it is neither necessary nor economic to carry all the spare fibre at the edge deeper than necessary into the network. We therefore will have spare fibres unterminated and not spliced through at intermediate street furniture nodes in the D-side and E-side cable networks, accommodation of these non-spliced fibres in housings also needs to be included.

The basic design rules used within the model are therefore as follows:

1. The number of spare fibres entering the LE site should be working fibres *(1 - fill factor).
2. The D-side cable chain needs to carry four categories of fibre:
   - Fibres feeding the PON splitters (includes fibre for additional network-side splitter ports if NxM splitters are deployed).
   - Fibres for the bespoke networks that use dedicated fibre - usually only very large or specialist customers would want this capability.
   - Spare fibre for growth (additional PONS)
   - Spare fibre arising from fitting actual cable sizes to the required fibre count (the fitted cable size will invariably be of greater size than the minimum required fibre count which includes fibres for growth).

As mentioned above, not all the spare fibre needs to be spliced through at intermediate DP and cabinet locations in the cable chains and that cables on the input and output sides of a splitter position will generally not have equal numbers of spare fibres. Splicing through spare fibre adds cost at the initial "day one" installation and could be omitted until required in the future, however to minimise future disturbance, all spare fibre in a D-side cable chain, that can be spliced through will be spliced, through to the cabinet location, regardless of statistical demand. This is to minimise the number of future network visits required to implement growth or bespoke fibre links and to minimise intervention faults. This strategy maximises a "hands off" operational network strategy which can reduce operational costs and increase network reliability, however quantifying these additional benefits was beyond the resource of the DISCUS project.

If NxM splitters are used where M>1, the additional upstream splitter port fibres can be used for a number of operations, a minimum of one port is of course required for the working PON, another port could be a test access port for field engineers to aid fault finding and diagnostics, other ports could be used for split by-pass, for example to accommodate different wavelength bands or to accommodate systems that cannot tolerate the longer roundtrip time of the access network of LR-PONs. Optical filters may also be applied to these port fibres to spatially isolate or separate different wavelength bands. The model will enable comparison of different splitter installation strategies, with different numbers of upstream splitter ports, which is set be a variable parameter in the model, and the effects on upfront cost and pay back times.

4.3.4 D-side Cable Chain model

This sub section and the following sections briefly describes the infrastructure models, that is cable dimensioning and splitter nodes designs etc. For fuller descriptions see appendix 8.4 (note all the figures in these sections are also reproduced in the appendix for the convenience of the reader).

As previously mention the fibre network up to the metro-core node can be considered as a series of infrastructure sections from CPE and customer drop, DPs and D-side cable network, Cabinets and E-side cable network LE site equipment, Backhaul (or metro-access) network.

The D-side network consists of the D-side cables and the DPs which house splitters for PON solutions or are just splice housings for point to point solutions. For the model we have implemented it is assumed that optical cables run from cabinets to the cabinet serving area edge passing and serving DPs on route as previously described. This effectively breaks the cable chain into a number of sections with reducing fibre counts as the cable passes DPs from the cabinet to the cabinet serving area edge this is illustrated in Figure 4-8.

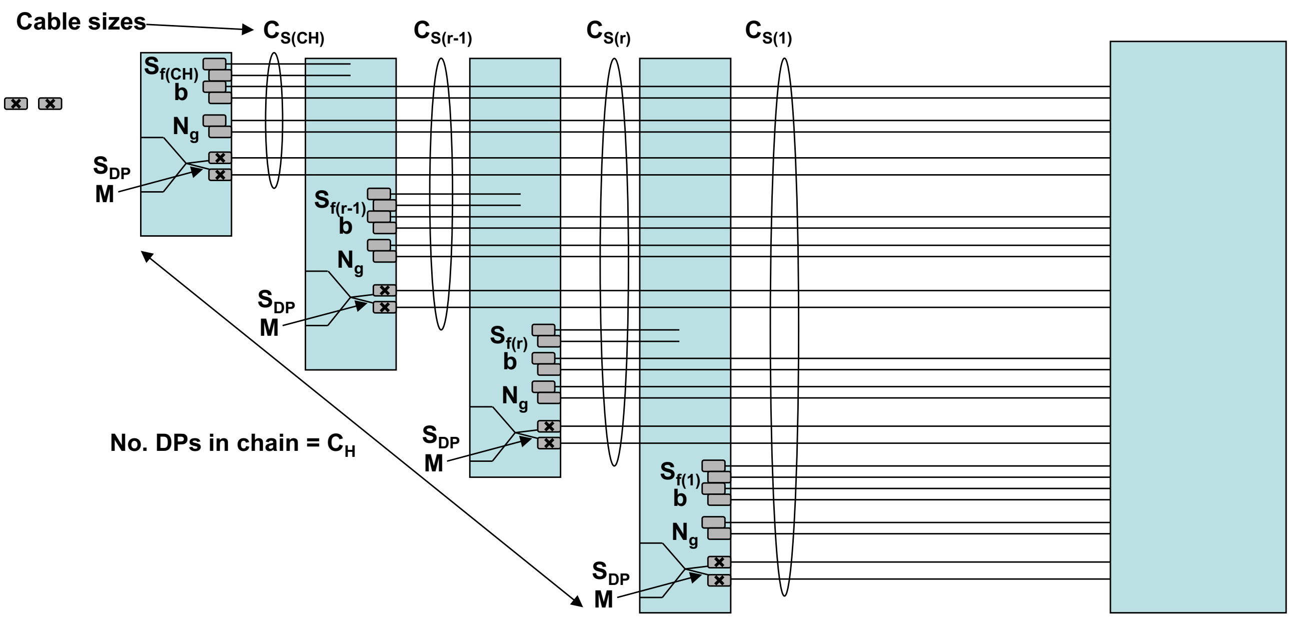

Figure 4-8 D-Side cable chain model structure

As fibres from the cabinet location pass the DPs fewer fibres need to be carried through to serve successive DPs. The model calculates the number of fibres for the various categories describes in section 4.3.3 and it can be seen that some fibre are spliced through and others are stored on splice trays in the DP housing without splicing. The variables in Figure 4-8 are:

S_{DP} = DP splitter size

M = number of D-side splitter ports $M = 2^n$; where n= 0, 1, 2.

C_H = number of DPs in the cable chain.

b = number of fibres per DP reserved for bespoke fibre services.

Ng = fibres reserved for future additional PONs for customer numbers growth

Note there is also a variable f_f which is the fill factor parameter that reserves spare DP splitter ports for customer growth on the existing PON/s, the Ng spare fibres are for growth that exceeds this factor and enables additional PONs to be added to provide additional fibre connections.

$S_{F(r)}$ = The spare fibres from the mismatch between actual cable size and required number of fibre in a cable section. The model computes the required fibre count in a cable section, including the allowance for growth, and then looks up the smallest cable size that is greater or equal to this value as the section cable size.

$C_{S(r)}$ = The rth section cable size.

Deciding how many of the spare fibres that link the input cable to the output cable to splice through at a DP site is quite a complex process. The maximum number of spare fibres that span the whole chain due to cable sizes exceeding the required number of fibres either side of a DP is the minimum of the set of absolute values of the spare fibre count difference about the set of DPs in the chain and these spare fibres should be distributed across the DPs in the chain. However this is a complex planning rule to practically implement for field staff in terms of fibre splicing of these spare fibres at the DPs. Therefor, although there is a small cost penalty, a simpler planning rule is assumed for the model which is to splice all spare fibre through at each DP, this number of splices at the r^{th} DP is:

Spiced through spare fibres = $\mathbf{ABS(C_{S(r)} - C_{S(r+1)} - (2C_H - 2r + 1)(b + (1+f)(N_g + M_{DP})))}$.

This does mean more splicing at the DPs but is a simple rule for field staff to implement and although this increases the number of spare fibres spliced through to the cabinet they do not all need to be spliced through to the local exchange over the E-side cables.

The DP splitter will be installed with splice trays for the incoming fibre cable/fibre unit and the drop fibre units. Drop fibre splices however are not included in the splitter node cost as they are included in the customer drop costs as a "Just in Time" cost (JIT).

Note: Often the drop fibre unit will be two or four fibres. For this model it is assumed that only one fibre is spliced for single fibre working (typical) or two fibres for two fibre working. Any spare fibres in the drop fibre unit are assumed to be stored on the splice tray for that drop cable, but without splices, with up to two splices per tray. For the urban areas typical splitter sizes at the DP location will be 32 way with some 16 way splitters and in the sparser rural areas splitters as small as 8 way. In the current version of the cash flow model the DP splitter size is an input variable and the cabinet splitter size is calculated using the DP splitter size value, it also assumes all DP splitters in the exchange area will be the same size. In sparse rural areas the DP splitters could be split into additional tiers with smaller end point splitters, as small as 4-way and maybe occasionally 2-way. Optimisation techniques to fit cables and splitters to these areas are being developed outside of the DISCUS project but they are not included in the cash flow model at present and sparse rural is accommodated by assuming longer drop lengths dependent on mean customer density.

The model for the DP splitter node is shown in Figure 4-9 and shows the equations for the input and output cable sizes for the r[th] node in a D-side cable chain. The cost for the installed DP splitter housing = cost of housing + cost of splitters + cost of splice trays + cost of splicing + cost of installation + cost of site visit.

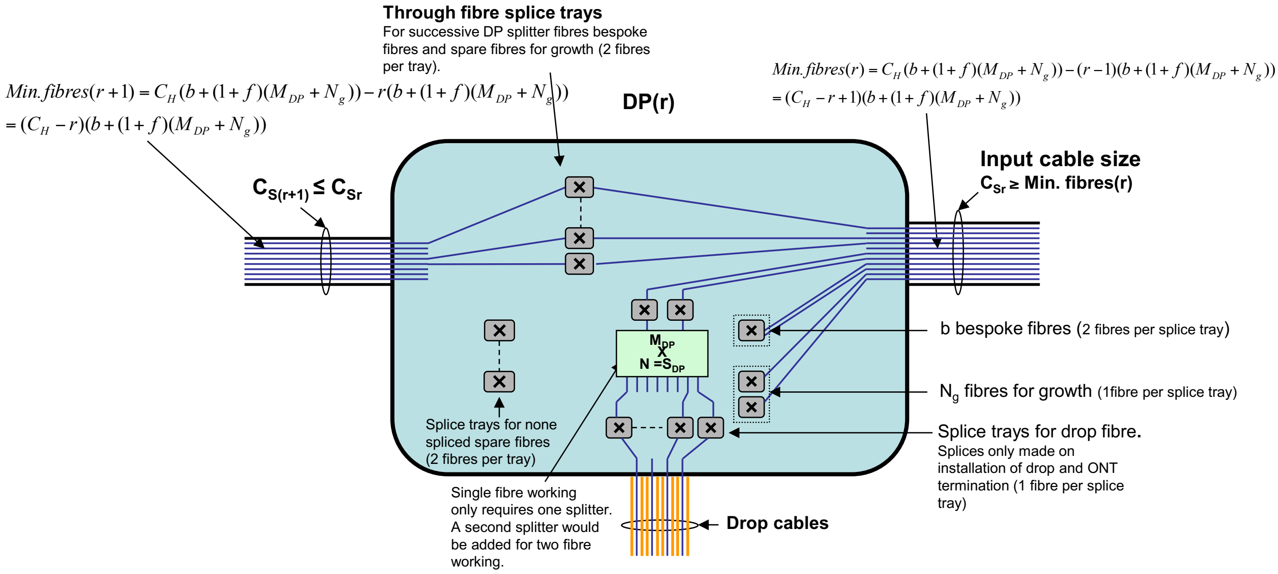

Figure 4-9 DP splitter node model

For point to point fibre systems there is of course no splitter at the DP and the DP housing is simply a splice joint housing. The model for the point to point fibre DP node in a point to point D-side cable chain is shown in Figure 4-10 and also shows the input and output cable sizes for the rth DP node in a D-side cable chain. Again in the DP point to point fibre cable chain the cable sizes on the input and output size will generally be larger than the calculated required number of fibre and spare fibre will be available. The same simple rule of splicing all spare fibres that can be spliced through at each DP node can be applied and any additional spare fibre in the input or output cable are stored on splice trays without splices.

4.3.5 Number of fibres from cable chains entering the cabinet node:

Figure 4-11shows a D-side cable chain entering the cabinet node. This cabinet node will house a splitter for LR-PON and other larger split PON solutions or just a splicing joint housing for point to point fibre solutions or smaller split PONS where all the split is placed at the DP position.

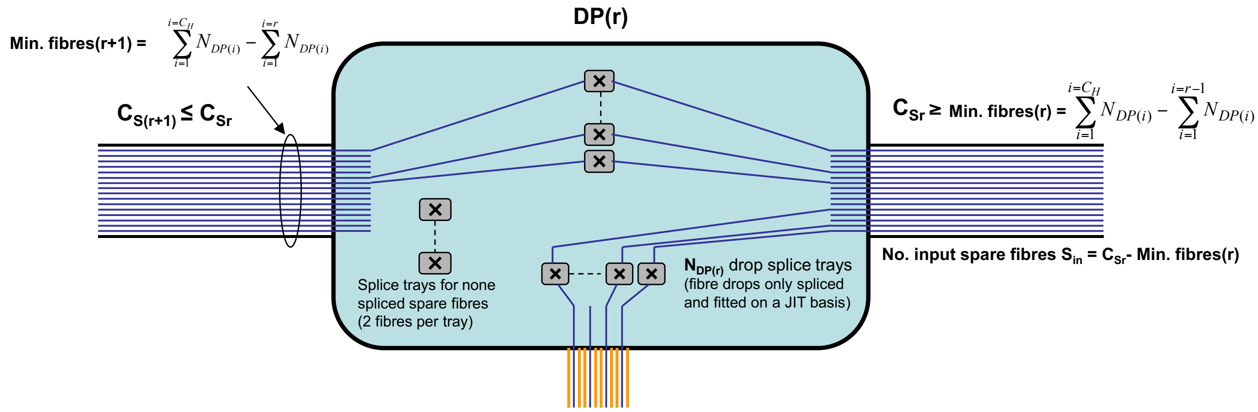

Figure 4-10 DP node model for point to point fibre systems

Fibre entering the cabinet from the chain are:

Fibres from DP splitters = $C_H M_{DP}(1+f)$ the factor f = 0 for single fibre working and f=1 for two fibre working (the reason the flag f is 0 or 1 rather than set as 1 or 2 is because the flag state in the Excel implementation is sometimes tested via functions that return a logical value of true or false which needed 1 to be added to get the appropriate multiplier)

Bespoke fibres = $C_H b$; (Minimum b = 2)

Growth fibres = $C_H N_g(1+f)$

Spare fibre from cable = C_{S1} – Min.fibres(1); r= 1 for the first cable section from the cabinet

At the cabinet $C_H(1+f)$ splitter fibres will be connected to cabinet splitters and the number of E-side fibres for servicing these splitters will be $C_H(1+f)/S_{C.}$ per D-side cable chain. Where S_C = cabinet split size).

An option is provided in the model to either pass the remaining splitter fibres up to the LE site or just store them in the cabinet for future use as required (this is set with a flag f_2).

A weighting factor W_b is also used to reduce the number of bespoke fibres carried up to the LE site to reflect the relatively low probability that any one DP will use bespoke fibres and to minimise E-side cable sizes. This parameter is assigned by geotype as there will be higher probability in city areas and lower probability in sparse rural areas that such fibre will be required. In the model bespoke fibres are not spliced through to the E-side cable at the cabinet until required but are stored on splice trays (two fibres per tray).

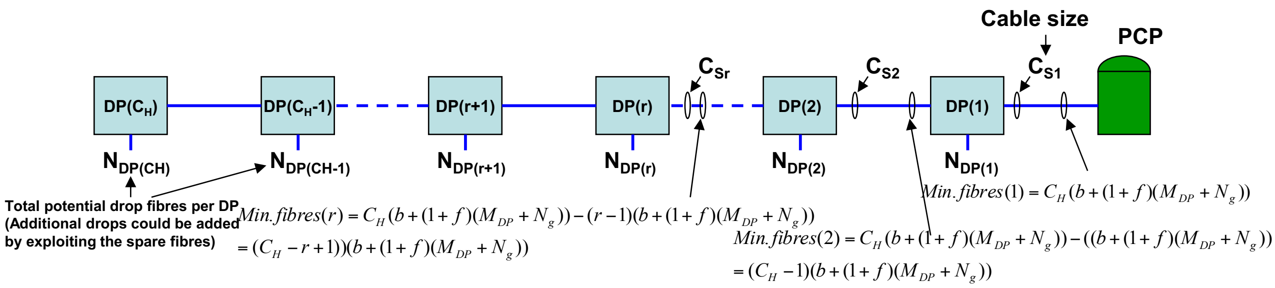

Figure 4-11 D-side cable chain showing fibre count entering the Cabinet

Growth fibre for future additional PONs would be connected to further cabinet splitters in the future and therefore can be divided by S_C, the cabinet splitter size, for the E-side fibre count.

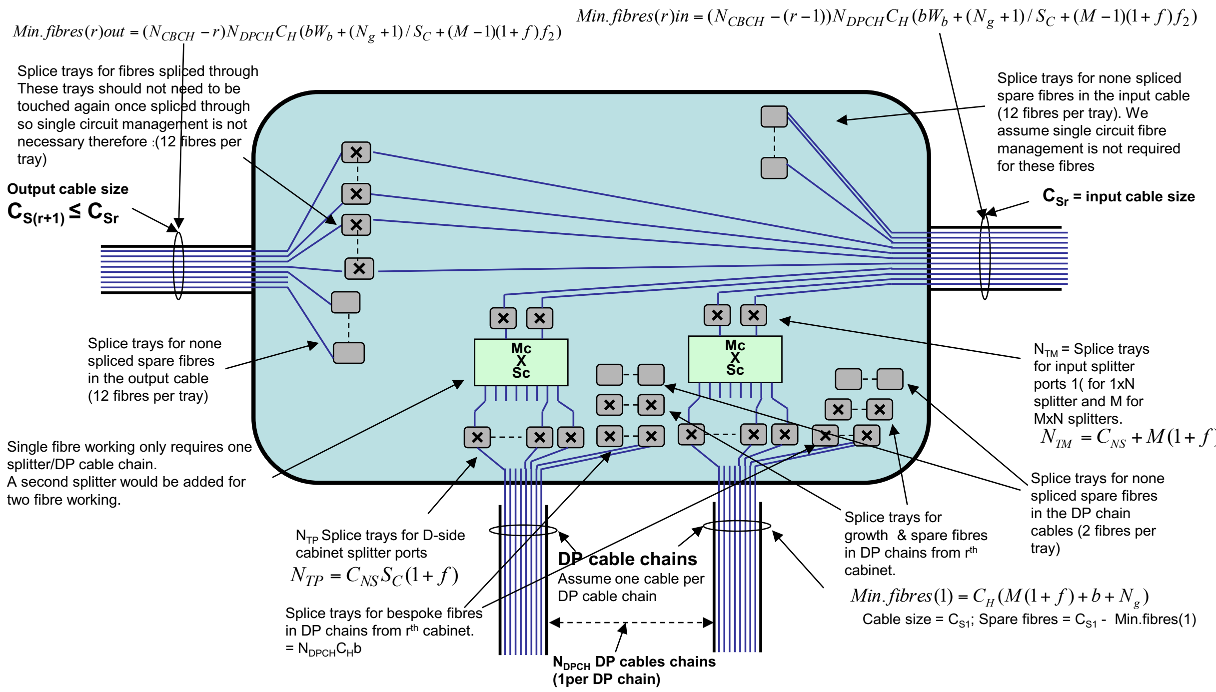

Figure 4-12 The model structure for a cabinet splitter node

Within the model it is assumed that the DP cable chain is installed as one job and site visit travelling time and installation time are averaged over all DPs and therefore a single value for

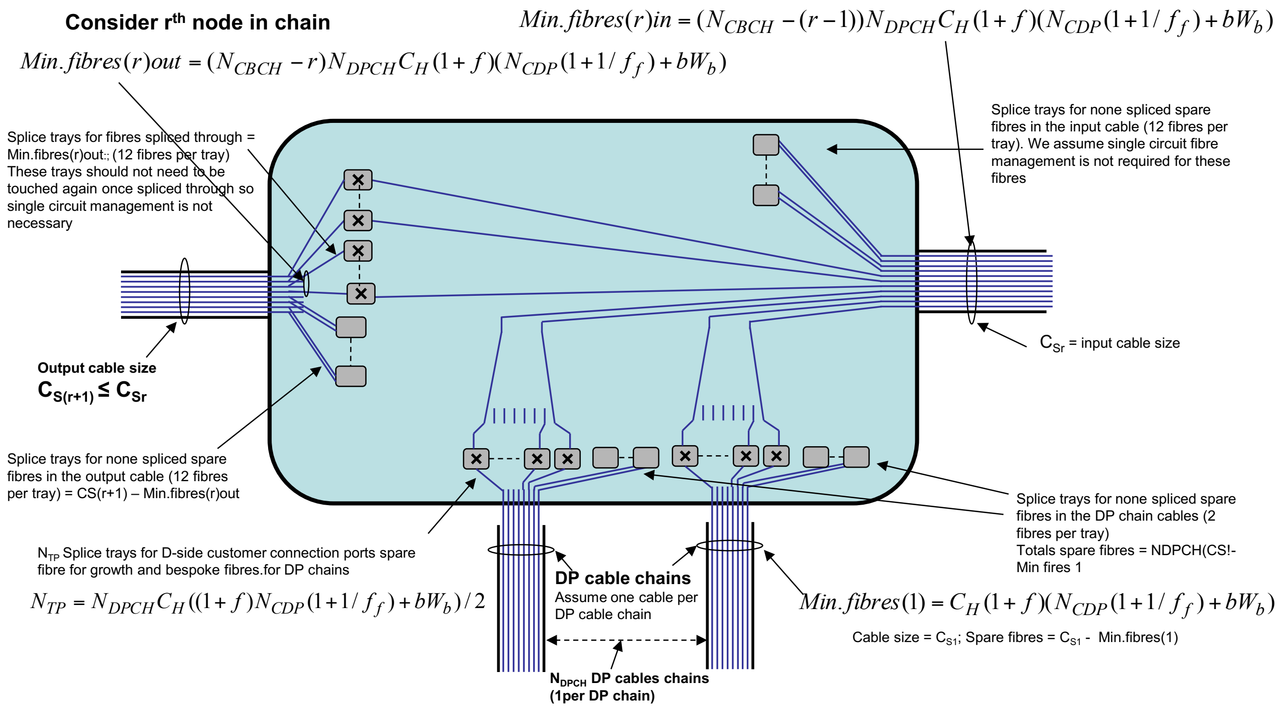

Figure 4-13 Model structure for point to point fibre cabinet node

the installation time per DP is used. Splicing labour time is included within the splice cost figure.

The model structures for the splitter cabinet node and the point to point fibre cabinet node are shown in Figure 4-12and Figure 4-13 respectively.

The cabinets and cables in the E-side network are also in cable chains in a similar manner to the D-side chains and the Figure 4-12 and Figure 4-13 show the cable sizes at the rth node in such an E-side cable chain, r=1 is the section and node closest to the local exchange.

4.3.6 Backhaul network cable chains

The fibre network between the local exchanges and the metro-core nodes for the LR-PON network solution proposed in DISCUS uses dual parenting protection to two separated metro-core nodes. The cable route from the LE site to the two metro core nodes will pass other LE sites that are also amplifier nodes for the LR-PONs serving the customer in those exchange areas. Therefore it would make economic sense to combine the fibres required for these other LE sites into common cables effectively forming cable chains between metro-core node pairs with a number of LE sites per chain see Figure 4-14. So the backhaul network now also becomes a set of cable chain routes. This is also compatible with historic SDH/SONET ring structures where the cable chains can be sections of the old ring structures, it is also

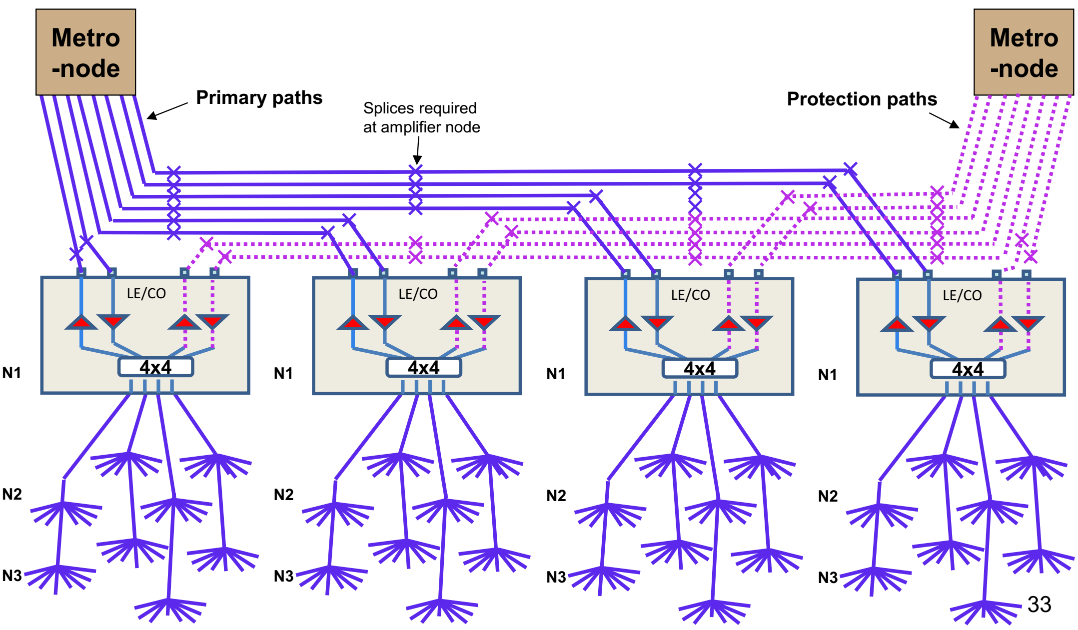

Figure 4-14 Backhaul cable chain model structure

compatible with the newer networks such as the BT 21cn network which also introduced WDM network chain structures in the back haul network. The backhaul cable structure for LR-PON, as shown in Figure 4-14, requires four fibres for each LR-PON at an LE site for dual parenting protection but because these fibre are in opposite directions only two fibres per LR-PON are required within the backhaul cable. The total fibre in the back haul cable for the LR-PONs in the chain is therefore $F_T = \sum_{i=1}^{i=n} 2N_{LE}(i)$ where F_T is the total LR-PON fibre and $N_{LE}(i)$ is the number of LR-PONS at the ith local exchange node in the chain. At each LE there are also spare fibres for growth in each ODN and also an allowance for bespoke fibre. These also need to be supported by additional fibre in the backhaul cable chain. Again there is a design decision whether to provide all the backhaul bespoke and spare fibre entering the LE sites from the ODNs or just a percentage based on statistical estimates of the total bespoke and spare fibre required for the LEs in each cable chain. The latter is selected in the cash flow model but a toggle can be added to compute the impact of providing all that fibre at day one. The actual number of these fibres is derived from the size of the cables from the E-side cable chains entering the LE node sites as previously discussed. Any statistical factor to reduce these fibre counts within the cable chain is a simple multiplier of the sum of all these fibres entering the LE nodes from the E-side cables.

The spare and bespoke fibre will be terminated on splice trays in the housing, this could be the amplifier node housing, if space permits, or a separate housing for cable splicing can be employed. In the model we assume these splices are in the amplifier node housing although until development of such a node is carried out the practicality of that approach is uncertain. Once all the fibres are known the cable size and costs can be determined.

To generate these backhaul dual parented chain structures optimisation models were used as described below, an example of results for a 73 metro-core node solution for the UK network is shown in Figure 4-15. The optimisation modelling also looked at tree and branch cable topologies for the backhaul network however for the cash flow model only the chain

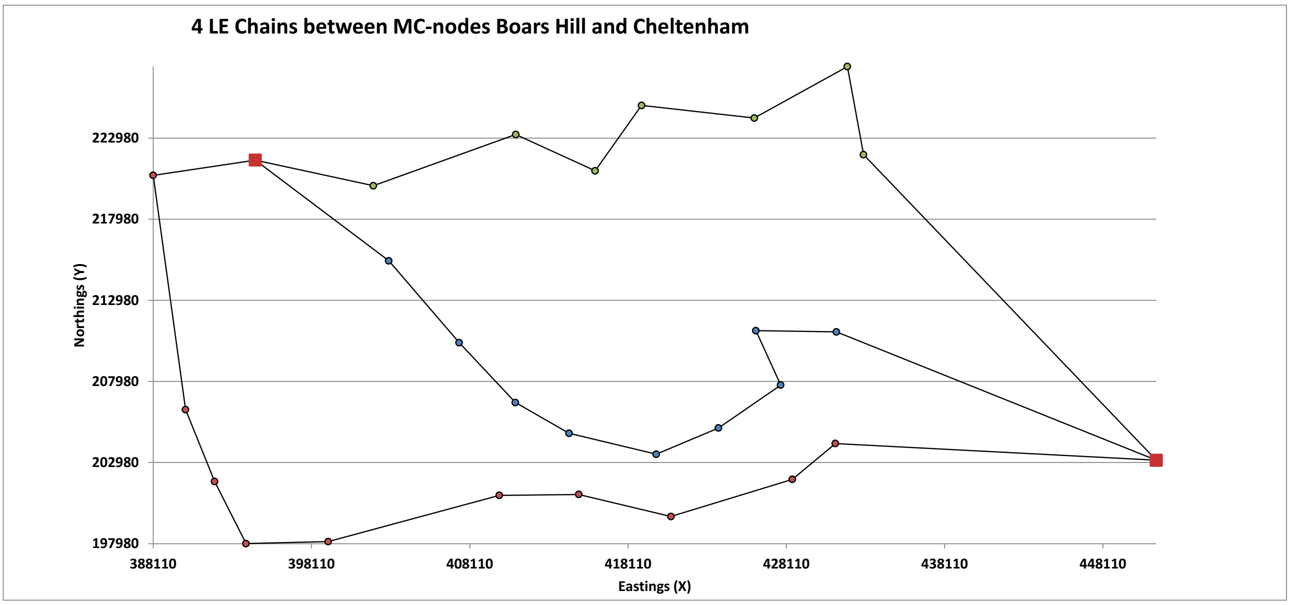

Figure 4-15 Example of backhaul cable chain structure for UK network

topology was implemented

For compatibility when comparing LR-PON solutions with more conventional solutions that do not bypass the LE sites the same cable chain structures are assumed for all solutions although the cables sizes can change between solutions.

The cash flow model computes each LE area sequentially to determine the costs of the various elements making up the solution. From these costs the cost per customer passed and connected can be determined. These costs vary with user bandwidth and also depend on the roll-out rate which gives a time line for expenditure depending on the rollout strategy and rate. For the backhaul network it is required that a cost per customer as a function of the same user bandwidth parameters is also calculated for each exchange the difference now however is that the backhaul cable chain cost is shared over other LE areas on the same cable chain. To dimension the chain, the number of fibres required for each LE on the chain needs to be calculated. This complicates rollout strategies as it implies that all the LE areas on a cable chain should be rolled out together. However once all the fibre requirements are known the cable size can be determined and the costs can be calculated for the total cable chain. This cost is then divided by all the customers on all the LEs on the cable chain to get an average backhaul cost per customer. This simple strategy implies rollout on a chain by chain basis which can restrict some roll-out options e.g. tackling rural exchange areas first to address the digital divide, as an urban LE might share a common chain with the rural exchanges and would therefore also need to be included. One way round this problem is to compute the backhaul cable requirements for 100% roll out and derive the average backhaul cost per customer for each LE. Then when applying roll out strategies these cost per customer are used even if the other LEs in the chains would not be part of the build being considered. There are practical, economic, and subsidy argument issues with this approach but it does enable simple rollout strategies to be compared

4.3.7 Optimizing the backhaul network

This section describes the optimisation methodology used for generating the backhaul network cable chains and also compares the chain topology with tree and branch topologies. The objective is to design the optimum cable fibre network between a set of MC nodes to connect their sets of LE sites such that there are two-bounded node disjoint paths from each LE site to its two MC nodes Figure 4-16 shows an example for the two alternatives, in this example we observe two MC nodes m_1, m_2, and six LE sites. Let us assume that each LE site in the set {*a*, *d*, *e*, *f*} requires 1 PON, and each LE site in the set {*b*, *c*} requires 2 PONs. For each PON we need four fibres to connect to their two MC nodes (i.e., two fibres per PON per MC one for upstream and one for downstream).

In the tree topology fibres are distributed from MC nodes to LE sites through cables consisting of fibres that form a tree distribution network. The tree distribution network must be resilient to edge and node failures, that is, the two paths to the two MC need to be disjoint so that the PON can switch to the alternative path whenever an edge or node in the distribution network fails. Figure 4-16(a) depicts an illustrative example, black arrows shows the cable distribution network for MC node m_1 and grey arrows for MC node m_2. The weights associated with the edges indicate the length of the link (left) and the minimum number of fibres required (right). For instance, (*e*, *f*) requires 2 fibres as *e* and *f* only use one PON each and the distance between e and *f* is 1 KM; the link between (*a*, *b*) requires eight fibres: four fibres for LE *a* (with two PONs) and four fibres for *c* (with two PONs) and the distance between *a* and *b* is 3 KMs. Complete details of the tree topology are available in Deliverables 2.6 and 4.10.

In the chain topology fibres are distributed from the MC nodes to LE sites through cable chains. In the chain LE sites are connected one after another with no branches in the sequence, the first and last LE sites in the chain are directly connected to the MC nodes. Unlike the tree topology where we need fibres from two different cables for a PON to design a resilient network, in the chain topology we only need one cable to connect to two MC nodes for any PON. Figure 4-16 (b) depicts a chain solution for our example, in the example we observe two chains, one highlighted with black arrows and another one with grey arrows. Similarly to the tree topology the weights in the edges indicates the length and the number of fibres for each link. In this toy example we observe that the tree topology requires 306 km of fibre and the chain topology requires 280 km of fibre.

4.3.7.1 Optimization methodology

Considering the combinatorial nature of the problem, we use a local search (LS) approach to compute near-optimal solutions by minimizing the cost of the solution. Broadly speaking, the LS algorithm starts with an initial solution and iteratively improves the solution, little by little, by performing small changes. In order to move from one solution to another we use the following 4-step procedure:

1. Select a random LE *e* from the current chain *c*.
2. Remove *e* from *c*.
3. Identify the best location (i.e., a new chain *c'*) for *e* satisfying all constraints.
4. Insert *e* in *c'*.

Additionally, the LS algorithm is equipped with incremental way of maintaining information related to checking constraints and computing the cost of the solution.

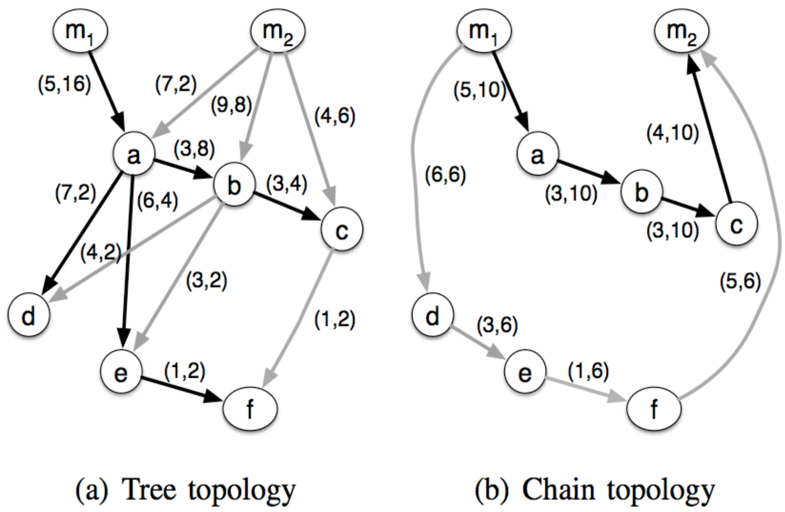

Figure 4-16 Backhaul network example for two MC nodes and 6 LE sites. The weight of the edges indicate the distance in kms of the link (left) and the total fibers required for the link (right), assuming that LEs a, b, c, e, and f are associated with 1 PON and LEs b

Table 11 Optimisation modelling results for tree and chain topologies summaries the total trail distance (i.e., sum of the distance of all edges in the solution), cable, and fibre required for the tree and chain topologies. We recall that connecting two LE sites might require more than a single cable if the required number of fibres exceeds 276 (max. number of fibres in a cable used in the model).

In the table we observe that the tree topology tends to use fewer cable kms than the chain at the cost of increasing the total number of fibres. We recall that the chain topology requires two fibres per PON while the tree topology requires four. In fact, in all the experimented scenarios the tree topology uses less cable (about 12%) than the chain topology, however, at the same time, the tree topology uses more fiber (about 15%). The reason is that the tree topology aggregates more fibres in a trail (or link between two LE sites). On average in the UK scenario the tree and chain topologies deploy 244 and 144 fibres per km, roughly speaking, a km of the distribution network needs a 144-fiber cable for the chain topology and a 276-fiber cable for the tree topology.

Country	Number of MCs	Topology	Trail	Cable	Fibre
			Thousands of kms		
UK LEs in model = 5393	75	Tree	110.2	164.9	23,865
		Chain	171.9	197.5	21,570
	80	Tree	108.3	165.3	24,324
		Chain	167.4	192.0	21,533
	85	Tree	106.3	161.2	23,544
		Chain	160.4	185.2	20,694
	90	Tree	103.7	158.4	23,318
		Chain	150.9	175.3	20,057
Italy LEs in model =10708	100	Tree	168.8	236.1	32,516
		Chain	267.3	287.3	29,143
	120	Tree	160.6	227.0	31,793
		Chain	236.3	254.7	27,090
	140	Tree	156.0	221.2	30,764
		Chain	218.8	240.2	26,213
	160	Tree	151.5	211.8	29,148
		Chain	212.4	231.4	25,072

Table 11 Optimisation modelling results for tree and chain topologies for UK and Italian networks

4.3.8 LR-PON amplifier nodes

For the LR-PON solution the optical amplifiers required for increasing the power budget for the larger split reach and 10Gb/s symmetrical line speed are placed at LE exchange site. The structure for the amplifier node is shown in Figure 4-17.

The amplifier node contains the m x 4 splitter (typically m = 4, so the splitter is typically a 4x4), the LR-PON amplifiers plus tap splitter for PON monitoring and amplifier node management and any additional splitters required for the rural amplifier chain structure in the backhaul cable network. The current cost model only has the amplifier node for the "lollipop" version of the LR-PON design with only the LR-PON 4x4 splitter installed. The structure shown in Figure 4-17 e shows a configuration for 4 LR-PONs. The main module assembly including splitters is factory assembled so that the only field operation is splicing the E-Side and back haul fibres and possibly installing cards for the system capacity required.

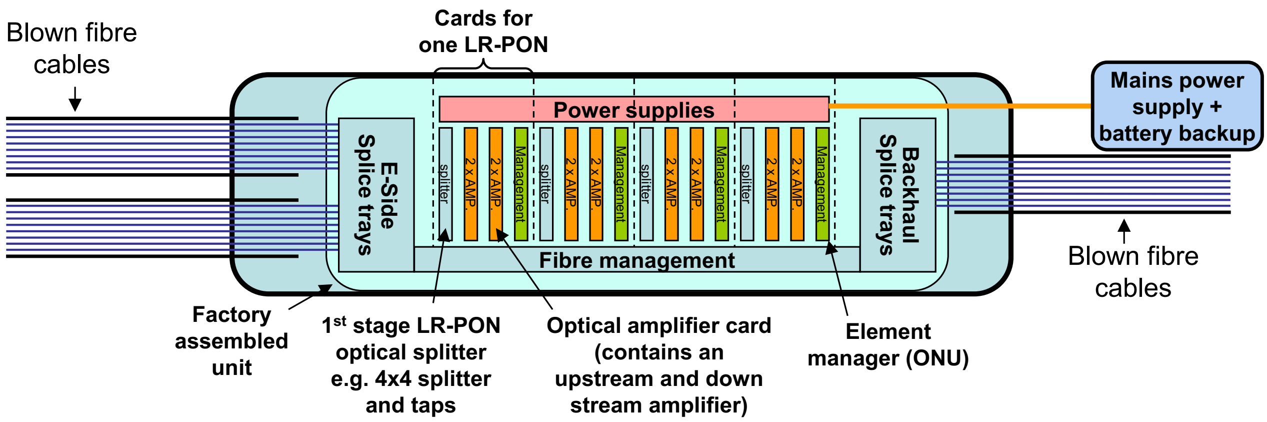

Figure 4-17 LR-PON amplifier node model structure - four LR-PON ODNs can be supported

The structure includes four EDFA amplifiers per LR-PON, an element manager which includes an ONU for communication back to the metro-core node OLTs. There is also the power supply and batteries pack for standby power when the mains supply fails. This power unit is assumed to be outside the main amplifier module housing. There are also separate splice tray assemblies for the E-side (to the ODN) and the backhauls fibre cables which will contain trays for splices for through fibre and fibre terminating on the node and also trays for storing spare fibres and bespoke fibres if they are not used at initial installation, which is assumed to be the case for the model.

Note although the housing in Figure 4-17 is shown dual ended for convenience and clarity the actual housing may well be single ended with ODN and backhaul cable entering from the same end. In either case the fibre management would allow interconnect and splicing between E-side (ODN) fibre and amplifier node splitters or backhaul fibres and backhaul through fibre etc, in a protected and managed fibre environment. From a cost model perspective the total number of splices and splice trays are calculated and the fibre management and installation costs determined from those values.

4.4 Metro-core node model structures

The metro core node model was described in some detail in D6.1[23], it focused on the dimensioning aspects of the cost model which are of course the core calculations that need to be performed for the cost modelling. In this sub section only the updates to the model will be described.

The models structures for the optical switch now includes the two sided non blocking Clos switch, the non-blocking single sided Clos switch and the single sided partitioned switch. These options are selectable in the model and for the partitioned structure the degree of partitioning can also be preselected. The single sided clos switch with the first stages using 2-sided matrices has also been included for completeness but is not an option that would be used. The design using both stages as single sided matrices is the preferred solution for both the full single sided Clos switch and the single sided partitioned switch.

The electronic layer2 and layer3 packet switching and router structures have been revised in the light of the power consumption modelling which required more detailed structures than originally included in the cost model. The various structures have however been kept generic and are shown in Figure 4-18 which also shows the configuration of the packet switches and

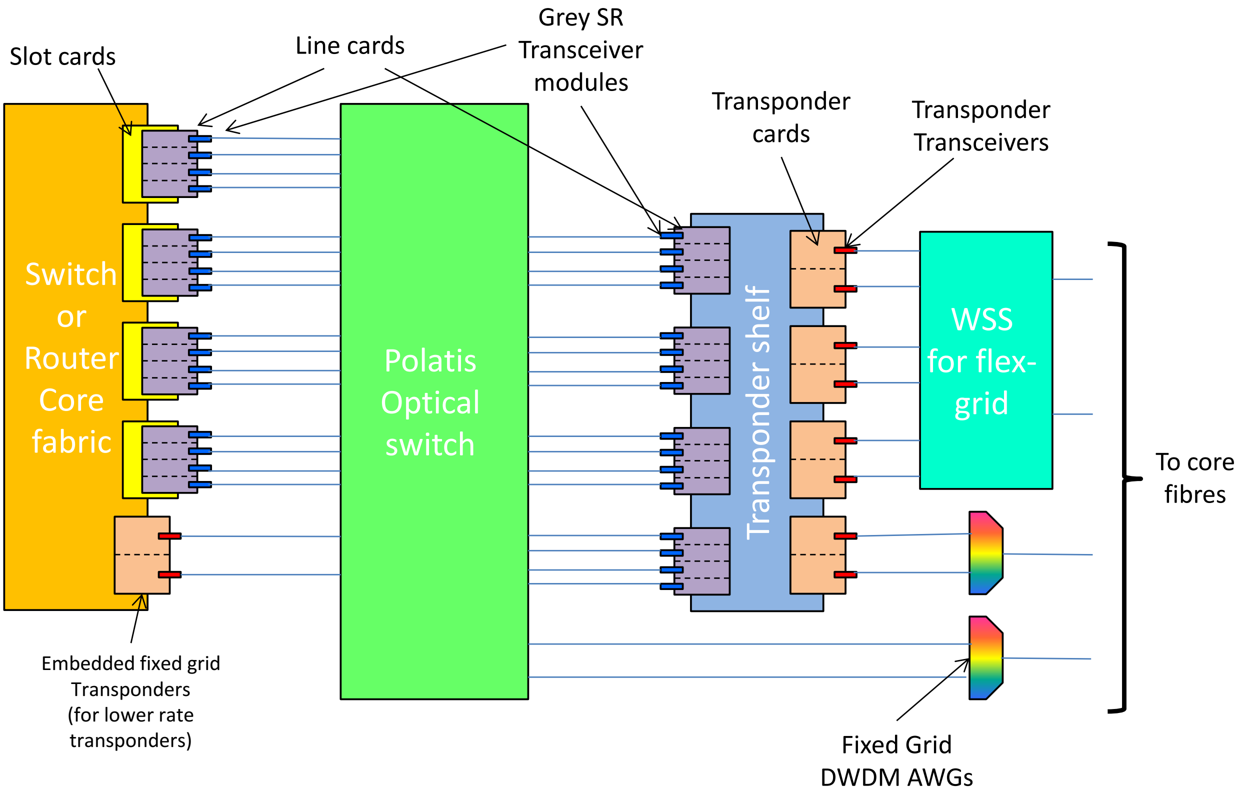

Figure 4-18 Configuration of electronic subsystems in the metro-core node

routers, the transponder shelves for both fixed and flex-grid systems and the optical multiplexing to core fibres, the optical switching layer is shown with ports on two sides, this is not to imply a two sided switch structure but is drawn in this way for clarity of the figure.

The routers and layer 2 switches are generic with an assumed common scalable switch fabric that have the router or switch slot cards which contains the forwarding engine electronics and into which plug the line cards. We also have an option for embedded fixed grid transponders in the layer 3 routers avoiding the interconnecting grey optical ports and the need for transponder shelves as separate equipment. The line cards are shown generically to represent both the access side and the core side of the packet switch layers. The port/line cards for both are calculated separately to match the architecture being modelled, the size of the metro-core node and customer bandwidth demand it needs to service.

The difference between the layer 2 switch and layer 3 router functionality is assumed to be handled in the slot card forwarding engine with the router functionality being greater than that of the layer 2 switch. In the model this is reflected in the relative cost and power consumption of router and layer 2 switch slot cards. In the version of the model described in

D6.1 the maximum port rate considered was 100Gb/s that has been increased to 400Gb/s and would be easy to increase further to say 1Tb/s ports capacities. The model also assumes generic shelf and rack structures so that interconnect, cooling and power supply costs can be taken into account. For example the switch fabric cards for switches and routers can, for smaller switches, be include in the shelves with the switch blades containing the slot cards and line interface or port cards. These switch cards are in turn assumed to be a programmable fabric that enables them to function as any stage in a Benes re-arrangeable non-blocking switch structure. If more capacity than a single shelf is required the switch fabric cards have capacity for inter-shelf connectivity allowing up to four shelves to be configured as a single switch. If further capacity is required switch fabric shelves are introduced which only contain switch fabric cards. These switch fabric shelves are in turn interconnected using the programmable switch fabric cards with additional capacity for interconnect purposes so that the switch can be expanded gracefully to very large sizes. The general concept is similar to the CISCO CSR3 router fabric structure.

4.4.1 Integration of the metro-core node and the access & backhaul models

It was originally intended to fully integrate the metro-core node model with the access and backhaul network models. However the unforeseen additional complexity and work required on the detailed structures within the model plus the requirements of the power consumption model has meant that full integration has not been possible within the project resources. However to use results for end to end analysis we can run the metro-core node model separately from the access and backhaul model and generate sets of results in terms of cost per customer as a function of metro-core node size and user bandwidth requirements. These results can then be used as look-up values for the cash flow model.

Effectively because the two major input variables for the metro-core node model is the size of the metro-core node and the required user sustained bandwidth. The cost per customer can be evaluated for the required range of these variables and these values used within the capex elements of the cash flow model. In this way the metro-core model is preconfigured with a chosen configuration, that is all design options are pre-selected and then the model is run for the bandwidth and metro-core node size ranges required. Full integration is still planned and would be an advantage as a wider range of scenarios could quickly be examined and compared but this will probably be completed after the formal end of the DISCUS project. If results of the integration are completed before the project review then results will be presented at that review.

The major disadvantage of not having a fully integrated model is that the impact of changing configuration options for given access and backhaul parameters is more difficult to investigate and quickly generate results.

4.5 Opex and Revenue models

4.5.1 Opex model

A cash flow model requires capital expenditure, revenue and operational expenditure as a function of time. The capital expenditure is derived from the cost models described above coupled with a rollout or deployment time line, see section4.6. Operational expenditure is a complex subject in its own right and a full model is well beyond the scope and resources of a project like DISCUS. However some meaningful operational expenditure parameter needs to be included in the cash flow model if the results are to have any real value. To cut through the complexity of operational expenditure modelling a simplified approach has been adopted in the Cash Flow model based on the assumption that operational costs ultimately come down to the manpower cost for running the business. There are of course recurring materials costs, accommodation costs, energy costs etc. over and above just manpower direct costs and these are accommodated by an overhead component applied to the manpower cost figure, so that manpower cost is a multiple of the average salary costs. The overhead factor used in the model is 2.5 with an assumed average salary of £30,000 per annum. This assumption aligns well with the 2015 BT Report and Accounts figures for OpenReach, which is the closest organisation in BT for physical network operating cost estimations that we require for the cash flow model. Note equipment wear out which is usually accounted for by a depreciation element in the business accounts are included in the capital expenditure part of the cash flow model not the operational cost.

The major operational cost benefit of end to end optical networks with FTTH/P, which minimizes electronic nodes and packet processing in the network and is the objective of the DISCUS solution, is the manpower savings that can potentially be achieved. However determining those savings is problematic as there are very few organisations with fibre only networks and even less public financial data available. However there have been occasional snippets of information from small operators with FTTH only networks that can be used for estimates of potential manpower savings. Estimates can also be made using failure and fault rates of network equipment, plant plus deployment rates, time to repair etc. estimates of the number of back office support staff and customer handling staff plus management overhead to give estimates of overall manpower requirements for FTTH networks. These figures when compared to the sparse information from companies operating fibre only networks can give estimates for the average manpower per customer connection for 100% deployment and take up. These figures in turn can then be used for the manpower based operational cost estimates within the model. This work has not been carried out within DISCUS but has been done by one of the authors in previous projects and these earlier estimates have been used for calculations of manpower reduction as a function of the number of sites connected, within the DISCUS

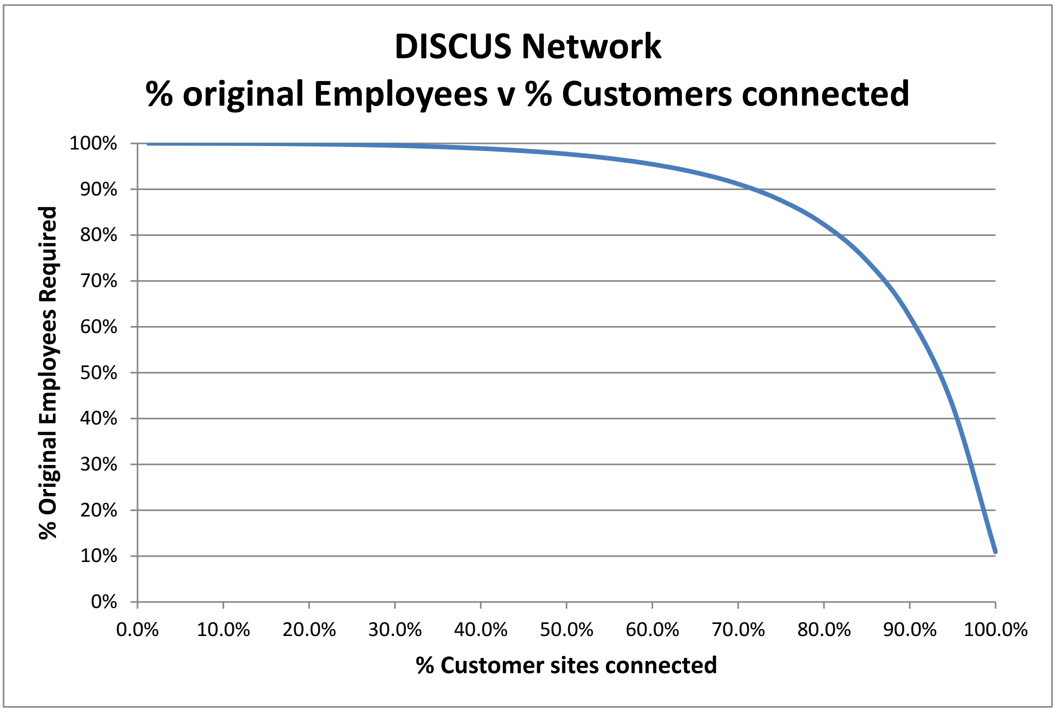

Figure 4-19 The percentage of employees retained v percentage of customer sites connected to the DISCUS network

cash flow model.

The curve in Figure 4-19 shows the assumed relationship between the percentage of employees required and the percentage of customers connected to the LR-PON based DISCUS network for 100% take rate. At lower service take rates the manpower saving will be reduced pro-rata. Also because of the dependencies on business models and the competitive environment etc. the actual manpower reduction achievable could vary considerably from the model parameters. Therefore a selectable toggle is built into the model that turns off the manpower saving parameter in the opex cost saving, so that networks can be compared without assuming manpower savings in the opex costs. Also all the parameters affecting the potential manpower savings are accessible variables in the model and if better published empirical data from use cases of FTTH deployments become available then the model parameters can be simply adjusted to reflect the new data. Other architectures such as point to point fibre, FTTCab and GPON have also had manpower saving estimates made for comparison purposes.

4.5.2 Revenue model

The third parameter required for cash flow modelling is a revenue time line. Revenue will be directly related to the number of customer sites that are connected to the network being evaluated. It is also dependent on the economic climate, the distribution of disposable income, regulation and the competitive environment amongst other market factors. However for the DISCUS cash flow model we consider three revenue parameters to encompass all revenue scenarios, these are:

- Fixed upfront connection charge
- Flat rate monthly charge
- Network usage charge

The upfront connection charge, if used, is usually applied to offset the one-off cost of physically connecting the customer to the network, that is; providing the customer drop, the customer premises termination and the CPE necessary to connect to service providers (if this is owned by the network provider). This is often waived in today's competitive environment and recouped as part of the monthly rental particularly if the contract is for a minimum term e.g. 1 year, as is often the case. But upfront charging has historically been extensively used and is often used in the mobile market where a mobile handset is part of the mobile service contract.

The monthly charge is, in today's market, the usual method for operators to receive the majority of their customer revenue. It is usually a flat rate charge based on a service bundle. The bundling element in the charge levied is often the closest operators get to a usage related charge. Historically this was not the case and before the broadband data era most revenues were received via the usage charge. The dependence on a monthly charge has been attractive to the customer base as their service cost is easily predictable and there is no financial barrier to using the services. But it does produce distortions and a level of unfairness to the market, for example heavy users of the services can get much better value than light users. Flat rate charging, which the monthly charge model is a version of, also effectively decouples revenues from service usage and the subsequent demand on network resources; this in turn reduces the incentive for operators to invest in the network capacity as there is no direct link to that network investment and any increased revenue stream to pay for it.

A usage charge, the third parameter in the model was historically very important and is a charge levied for some measure of the usage the customer makes of the network resources, usually related to the information transmitted to or from the customer site, e.g. Gigabytes per month of data usage. However this charge is very often waived and in effect there is no usage charge and little relationship between customer revenue and their usage of the network. It does however get used for bandwidth capping if customers pass a usage threshold or even when on "unlimited" contracts the customer is deemed to be an unfair user of the network resource (and limiting access for other customers). As the bandwidth capability of networks increases in the future capping may become untenable or at least impracticable in a truly superfast broadband future where customers could be using bandwidths orders of magnitude greater than today's average rates. This might require a return to some form of regular usage charging in order to couple network investment to revenue streams to recoup the network build cost required to support the high bandwidth demand. Although the charge might be more service related than direct bandwidth usage charging it would need to be linked to the sustained used bandwidth in some way and in the model this parameter is coupled directly to user bandwidth for simplicity and to avoid the complexity of service usage models.

In the model the revenue parameters are pre-set to initial "day one" values per customer and then a revenue growth parameter is applied to reflect anticipated revenue gains as new broadband services are introduced over time and as user bandwidths increase.

A complexity in setting initial values is defining what constitutes legitimate revenues that can be used for new network build and growth. For an incumbent operator should existing service revenue be allowed to be incorporated as part of the revenue streams of the new network? These operators have the problem that these existing services are already provided on the current network and the revenues earned are required to support that legacy network. The alternative is to only allow incremental revenue growth due to the enhanced service capability the new network enables to be used as the revenue stream for the new network build. This is a much more difficult business case to achieve and of course forces service packages on the new network to be at a higher price than the service packages on the existing network, which could limit take up. New build operators however can use all revenue streams against their network build as they have no legacy network or legacy services to support however they may have greater infrastructure build costs as they may have limited access to existing network infrastructure. To cut through these complexities in the model a compromise revenue stream of ~50% of today's super-fast (FTTCab) broadband published revenue rates are used as initial values for the model. These values can of course be changed to any values in order to see the sensitivity and challenges produced by varying revenues.

4.6 Rollout strategies

In order to generate the time lines for capital expenditure, opex expenditure and revenues a rollout of network build and customer take-up over time is required. In the earlier version of this model the rollout strategy was a simple total sites per annum figure and either a representative exchange was used for the cash flow calculations against that volume rollout or each exchange was calculated in turn, after being ordered in some way, and the costs and revenue accumulated until the per annum number of customers was captured. This was then repeated for each successive year. It should also be noted that in this earlier model the installation costs only included the access and backhaul transmission network up to the metro-core node. The switching and routing functions in the metro-core node were not included.

This basic process has been continued in the current version of the model with metro-core node costs (including core transmission transponder equipment) included. However the use of backhaul cable chains has complicated the rollout. To implement the chain and derive the costs equitably the whole chain needs to be included within the calculation which by default implies they are included in the rollout. The cable chain cable size could be pre-computed and spare fibres reserved for all the exchanges that will ultimately be in the cable chain but, depending on the roll-out scenario, all the exchanges may not be included at initial deployment.

The problem with that approach is the derivation of the cost per customer for the initial deployment. In this approach there will be capital spend for example on the backhaul cable that is shared over relatively few customers. This would artificially put up the cost per customer for those early deployments making them more expensive. We have therefore decided that for all rollout strategies the full chain and all exchange on the chain will be computed and an average cost per customer for the back haul for each LE is derived. This derived cost per customer will then be used for the rollout regardless of whether all the exchanges in the backhaul cable chain are included. It recognises that this causes some economic and possibly subsidisation issues that would need addressing with a real world rollout but it is a pragmatic way of not burdening areas with costs that would ultimately be shared.

The rollout strategies considered are:

- Rollout by geotype,
- Rollout by geographical area
- Rollout by customer demographics
- Rollout by exchange size

The strategies can be chosen to address various political and economic scenarios, for example rolling out by geotype could be used for early addressing of the problem of the digital divide whereby rural areas could be targeted first on the grounds that they are the most deprived areas when it comes to broadband service provision and the town and city areas will already be served by DSL and FTTCab. Indeed such a strategy could invert the digital divide as the rural areas having an ultra-fast FTTH solution deployed, would have better service capability than the dense areas. This could help to attract inward investment to these areas by companies wishing to exploit the advantages of these networks and offer attractive areas for their employees to live. Deployment via demographics could be an economic strategy to target areas with particular customer types that might take-up specific superfast broadband services. Rollout by geographic region or area could be a more pragmatic approach for an operator enabling more efficient use of manpower and equipment resources etc.

Whatever strategy is adopted they all are effectively a particular ordering or selection of the exchanges within the exchange database, the software routine sequentially selects the exchanges from the chosen ordering and for each exchange selected finds the associated exchanges in the same backhaul cable chain. After the necessary cost and revenue calculations these exchanges are then removed from the search and the next exchange and cable chain is found until the required number for the assumed rollout rate is obtained. The cash flow results are calculated and stored and then the next tranche of exchanges selected until all the exchanges in the data base or the search rollout range of interest have been selected.

4.7 Cumulative cash-flow models

While implementation of the cumulative discounted cash flow model is quite complex the basic concepts are fairly straight forward. Cash flow is cash in minus cash out, if this is positive then the business is in a state of positive cash flow and cash in exceeds cash out. Cash flow however itself has no history and the financial state of the business needs to take into account the cash flow history and produce cumulative cash flow figures. This is simply summing over time the total cash in and subtracting, for the same period, the total cash out, when the figure goes positive the business is at a breakeven point. Doing this periodically or even continuously a cumulative cash flow curve can be produced. The income and expenditure timelines are usually different curves and to further refine the breakeven point future expenditure and income should be discounted back to a nett present value (NPV). In this way discounted cash flow simply takes into account that future money is worth less than present money.

In the model a periodic accounting period of 1 year is used and a curve is fitted to these discrete points to generate the cumulative discounted cash flow results.

Major parameters in the cash flow model are:

- Take rate
- Interest rate
- Learning curves applied to plant and equipment costs
- Discount rate
- Manpower overhead factor
- Installation rate
- per customer values for the capex, opex and revenue
- Network dimensioning parameters such as total network size, number of customers, number of employees at year one (for incumbent model)

The take rate is simply the proportion of customers passed that actually are connected and generate revenue, it is used for computing cost per customer connected and cost per customer passed.

The interest rate is the interest charge on the negative cash flow values prior to the breakeven point. It is assumed for consistency that negative cash flow is supported via loans which will have an associated interest charged. There is a toggle in the model that allows this to be turned off to see the effect of interest rate on time to breakeven. Interest charges, as would be expected, will delay the time to breakeven.

The learning curve is the technique used in the model to take into account volume related price declines of network equipment and infrastructure. Most of the equipment is electronics based and is assumed to have the industry norm of an 80% learning curve that is as the volume of the equipment doubles the price will fall to 80% of the previous price. Describing learning curves as a percentage is usual practice. The price per unit of equipment or unit of bandwidth as a function of volume is given by:

$$P_2 = P_1 L^{\log_2(\frac{V_2}{V_1})};$$

Where: P_2 = price per unit at volume V_2

P_1 = price at initial volume V_1

L = the learning curve (usually expressed as a percentage).

Learning curves are a very useful way of enabling price decline as equipment is deployed with the cumulative volume increase producing price decline per unit or the product. Although always an approximation to reality they have proved to be amazingly resilient in the telecoms equipment market. An 80% learning curve produces the old quoted rule of thumb for telecoms transmission equipment of "four times the bandwidth for two and a half times the cost" this rule of thumb was valid from 140 MB/s systems through to 10GB/s systems. In recent years with 40Gb/s and 100Gb/s systems the price decline has not kept up with this historical value which is bad news for the telecoms sector which requires even faster price decline not slower to stay economically viable over the next decade! This is because of the decoupling of bandwidth usage and revenues where bandwidths can grow dramatically and the cost of providing the equipment to service that bandwidth grows much faster than revenues (which have been in decline over the years since the financial crash).

The discount rate is used for the discounting of future money to a present value or nett present value (NPV), in the model all cumulative cash flow comparisons use NPV values. For an example of the parameters and calculation of discount rate see for example [24], for the model a value of 11% is used for the discount rate which has been a fairly common rate in the industry for the UK and, it is believed, elsewhere.

As mentioned previously the manpower rate overhead factor is used for opex calculations and is a multiplier of the average direct costs for manpower. It is chosen to be 2.5 in the model which is a typical ratio of average wage cost to operating cost in the network business see [25].

The installation rate is a variable defining the total customers passed in a deployment scenario. It is used as the maximum rate over the duration of the period being considered. It is assumed that there is a ramp up to the deployment rate over the first two years to facilitate training and recruitment that may be required. A very simple ramp up is used; in the first year it is assumed that one quarter of the installation rate is deployed, in the second year one half and in the third and successive years the full rate is achieved. This is assuming that a major deployment rate is being targeted so that national rollout could be achieved in a reasonable time scale. However it is recognised that for a country the size of the UK a full FTTH roll out could take 10 years or more, particularly as intermediate technologies such as FTTCab and G.Fast will tend to delay full FTTH deployment.

The other parameters capex, opex and revenues come from the model sections previously discussed and the network dimensioning parameters come from publicly available sources such as the office of national statistics [26] and the national regulator and OECD reports.

5 Power consumption modelling

The cost model within the cumulative cash flow models are closely linked to the power consumption modelling and the additional detail that was required for the power consumption model for switches and routers was fed back into more detailed cost models for these items of equipment in the access and metro-core node models. The dominant network power consumption (neglecting the CPE) is the access network termination equipment (OLTs shelves and racks) and the layer 2 switches and layer 3 routers in the metro-core nodes.

The problem when trying to build a power consumption model, that was both consistent with historically performance data and also capable of projecting future power consumption performance, was that there are considerable inconsistencies in the published literature. Although this was recognised in some publications the basic problem appears to be the lack of a model based on the physical properties of VLSI technology that links performance improvements to technology improvements. It was therefore decided to build our own model based on linking switch and router energy efficiency to the improvements in VLSI feature size.

5.1 ASIC power consumption as a function of feature geometry size and the relationship to Router and Switch power consumption

In addition to the inconsistencies in publically available power consumption data for network equipment and in particular the routers and switches, there are also uncertainties about the definition of router and switch capacity, both bi-directional and unidirectional capacities (unidirectional is used for transmission system capacity) are used and it is often unclear which definition is being adopted. For example the Cisco CRS 1 data sheet uses unidirectional capacity while datasheets for the CRS 3 uses bidirectional capacity. It now appears that Cisco and other vendors are using the bi-directional definition of capacity but it is not certain that is always the case.

To address the inconsistency problems, which has made power consumption modelling and projections much more difficult than originally anticipated it was necessary to build a power consumption model that has its foundations firmly based on VLSI physical design parameters. It is the performance improvements of the underlying VLSI technology that has enabled the performance improvements and the power consumption reductions obtained during the evolution of switches and routers over the past couple of decades and power consumption calculations need to be based on the intrinsic properties of the VLSI technology used for the ASICs within such equipment,

The model therefore takes the underlying energy performance trends of the VLSI-ASIC technologies lying at the heart of routers and switches and fits them to the historical performance data that appears to be the most reliable in the published literature. The following sections describe the analysis and the proposed solution for the power consumption modelling work in DISCUS.

5.1.1 Power consumption evolution of VLSI

The power consumption of integrated circuits has two parts one is the power used moving charge into the gate regions of the transistors to switch the transistor state from on to off or off to on. The amount of charge to be moved is proportional to the capacitance which in turn depends of feature geometry size W. How fast the charge has to be moved also affects the power consumption such that power required is proportional to the operating frequency. The other part is leakage current which is independent of operating frequency but does depend on the thicknesses of the gate oxide insulating layers which are also proportional to the feature geometry size W.

The equations relating these parameters [1] are:

$$P = ACV^2 f + VI_{leak}$$

$$f = k_0 (V - V_{th})^{\alpha} / V$$

$$I_{leak} = I_{sub} + I_{ox}$$

$$I_{sub} = k_1 W e^{-\frac{V_{th}}{nV_\theta}} (1 - e^{-\frac{V}{V_\theta}})$$

$$I_{ox} = k_2 W (\frac{V}{T_{ox}})^2 e^{-aT_{ox}/V}$$

Where: P is the power dissipation

A, k_0, k_1, k_2, α, a and n are experimentally determined constants.

Vth = threshold voltage

V = operating voltage

C = capacitance of device which determines the charged to be moved to effect the required change in the switch state of the transistors.

I_{sub} = sub threshold leakage current

I_{ox} = Gate-oxide leakage (both leakage current terms are proportional to device feature size (W).

V_θ = Thermal voltage = ~25mV at room temperature and increases linearly with temperature.

T_{ox} = oxide thickness and is proportional to W. ∴ $T_{ox} = k_3 W$

$$P = ACV^2 f + VI_{leak} = ACV^2 f + V[k_1 W e^{-\frac{V_{th}}{nV_\vartheta}} (1 - e^{-\frac{V}{V_\vartheta}}) + k_2 W (\frac{V}{T_{ox}})^2 e^{-aT_{ox}/V}]$$

The parameter we are interested for power consumption comparisons is P/f which is a performance parameter which can be measured in W/Gb/s or nJ/bit, (these two sets of units are numerically equal).

Power consumption improvements of electronic systems come from improved design and better technology which when remaining in the silicon material technology comes from reducing device feature size (W). So we need to get expressions in terms of the device geometry feature size and for this exercise we assume the parameters deemed as constants in the expressions above remain constant as device geometries shrink and improved insulation materials are used (probably not true in practice). Reducing supply voltage obviously reduces power consumption and has been used as technology has developed over time, but it also reduces performance. Because we are also trying to maximise performance in terms of switch and router throughput we will assume that this will not be the option used as device geometry shrink further in the future. Over the last decade or so there have been many design techniques used for maximising performance while minimising power consumption and these techniques have been known and understood for many years. We therefore assume the best techniques have already been implemented and for the period of interest the major parameter affecting improvement in performance and power consumption is the reduction in feature size and any additional further improvements need to come from architectural design and introduction of new technologies (e.g. optical interconnect at chip and board level etc.).

With the above assumptions we can simplify the above equations:

The capacitance $C = \varepsilon_0 \frac{Ar}{d}$; The area parameter Ar will be proportional to W^2 and the distance parameter d will be proportional to W, so that $C = \varepsilon W$ where ε is a constant. Therefore:

$$P = A\varepsilon WV^2 f + V[k_1 We^{-\frac{V_{th}}{nV_\vartheta}}(1 - e^{-\frac{V}{V_\vartheta}}) + k_2 W(\frac{V}{T_{ox}})^2 e^{-aT_{ox}/V}]$$

Substituting for Tox = k_3W, then:

$$P = A\varepsilon WV^2 f + VWk_1 e^{-\frac{V_{th}}{nV_\theta}}(1 - e^{-\frac{V}{V_\theta}}) + \frac{V^3 k_2}{Wk_3^2} e^{-ak_3 W/V}$$

The first two terms are proportional to W and therefore as the device geometry size W shrinks the power consumption reduces, see Figure 5-1 for the evolution of VLSI feature size.

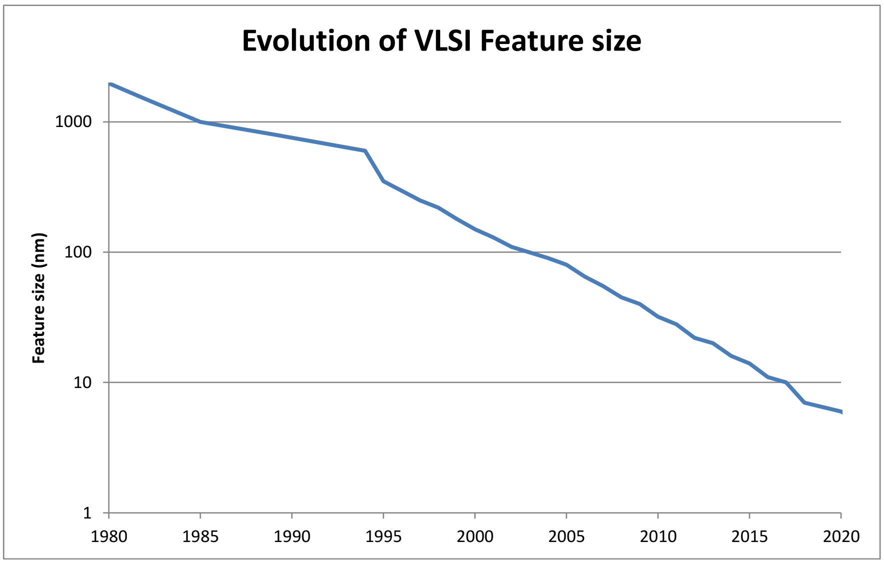

Figure 5-1 Reduction in VLSI feature size with time

The third term which is the leakage current associated with the oxide thickness in the junctions and the insulating properties of additional insulating layers is inversely proportional to W (this is modified by the exponential term which decrease with W so the overall effect is not as bad as an inverse relation). It is this term however that will stop silicon electronics continuing to improve as Moore's law.

Assuming all constants have remained constant over the time of interest except supply voltage and threshold voltage the equation becomes:

$$P = A'WV^2 f + k_1 VWe^{-\frac{V_{th}}{nV_\theta}}(1 - e^{-\frac{V}{V_\theta}}) + \frac{V^3 k'}{W} e^{-DW/V}$$

$$= A'WV^2 f + k_1 VW(e^{-\frac{V_{th}}{nV_\theta}} - e^{-\frac{V_{th}}{nV_\theta}} e^{-\frac{V}{V_\theta}}) + \frac{V^3 k'}{W} e^{-DW/V}$$

$$= A'WV^2 f + k_1 VW(e^{-\frac{V_{th}}{nV_\theta}} - e^{-\frac{V_{th}}{nV_\theta}\frac{V_\theta}{V}}) + \frac{V^3 k'}{W} e^{-DW/V}$$

$$= A'WV^2 f + k_1 VW(e^{-\frac{V_{th}}{nV_\theta}} - e^{-\frac{V_{th}}{nV}}) + \frac{V^3 k'}{W} e^{-DW/V}$$

$$= A'WV^2 f + k_1 VWe^{-\frac{V_{th}}{nV_\theta}} - k_1 VWe^{-\frac{V_{th}}{nV}} + \frac{V^3 k'}{W} e^{-DW/V}$$

Where A’ = Aε; $k' = \frac{k_2}{k_3^2}$ and $D = ak_3$;

The first term is the dynamic power consumption without leakage current effects the last three terms are components of the leakage current, of these the last term is the leakage due to oxide layer thickness or high k insulation layer thickness. In the early days with relatively large device feature sizes the leakage current terms would be small compared to the dynamic power consumption term and the equation could be approximated to:

$$P = A'WV^2 f$$

The performance parameter P/f = A’WV^2 reduces in proportion to W and V^2 and is a lower bound for the performance improvement relationship, any leakage terms will always be to the detriment of this bound. To apply this lower bound equation to router performance the constant A’ can be approximately estimated by fitting to early examples of router technology with 5-volt supply voltages and families of technologies using known feature size. A value of 0.0391 J/nm/V^2 was selected by this empirical fitting process

The core supply voltage V and the transistor threshold voltage Vth have also reduced as geometries have shrunk however the threshold voltage has reduced at a slower rate than the supply voltage and there is probably little prospect of significant future reductions in silicon technology see Figure 5-2 and [27]. The core supply voltage also needs to be significantly greater than the threshold voltage to ensure high performance and good switching characteristics. The expected power consumption improvement for dynamic power performance can be derived using the above equations and an assumption for the technology feature size as a function of time.

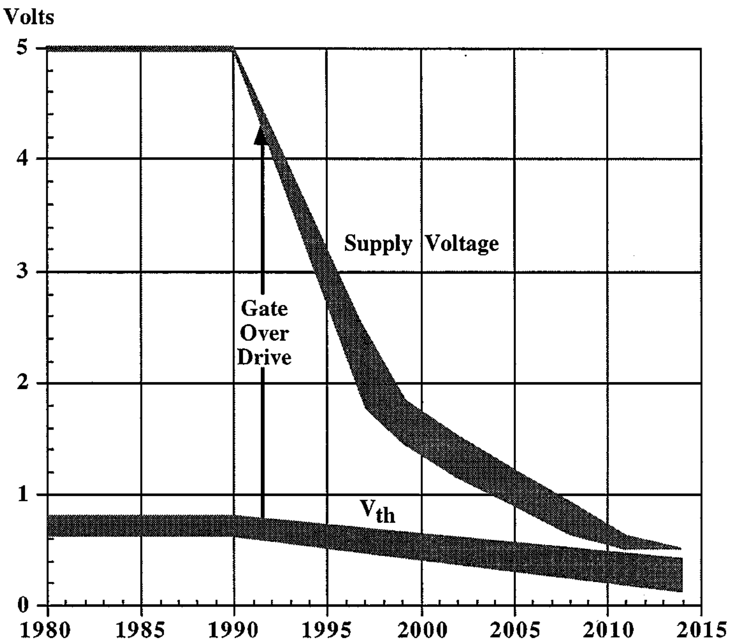

Figure 5-2 Evolution of VLSI supply and threshold voltages

In order to perform this comparison an assumption for the technology feature size as a function of time must be assumed, this is based on the emergence dates of reduced feature sizes in VLSI. It may be optimistic to assume router and switch technology adopted the feature size reductions immediately as there could be a time lag between chip technology capability becoming available and router design implementations using that technology. The time line used is as shown in Figure 5-1 for source of data see [28]

A public domain source of router energy performance with some identified products is a publication by Nielson, the data are shown in Figure 5-3 [29] is shown in Tucker's Greentouch presentation which references Nielson's results.

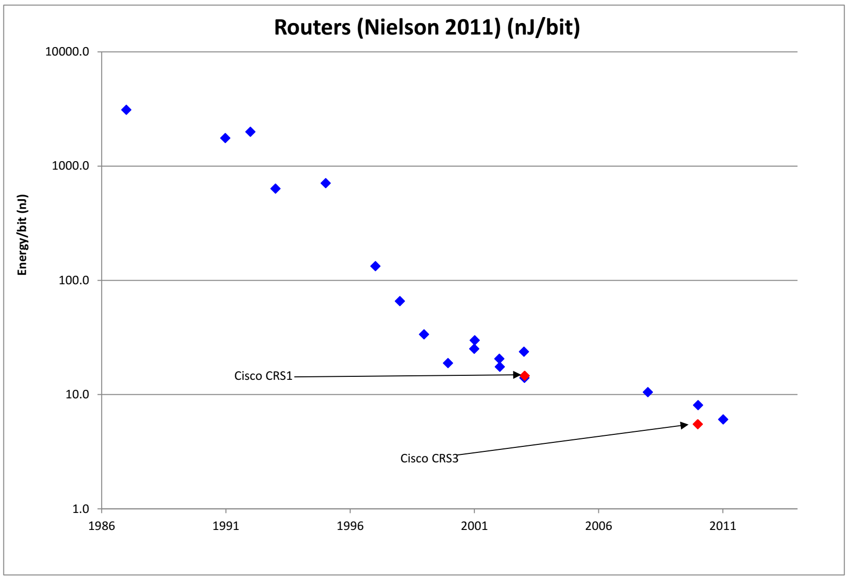

Figure 5-3 Router energy performance (Nielson 2011)

By using the above power equations with the data in Figure 5-1to Figure 5-3 potential energy efficiency evolution curves can be fitted to the Nielson data to show expected improvement trends. The Neilson results seem to be based on the unidirectional router/switch capacity definition.

The curves derived from the technology power equations above will tend to be lower bounds on energy efficiency improvements which may not be met because manufacturers are not just designing for energy efficiency but also need to design for performance and these two requirements tend to have conflicting demands on the technology designs.

If we take the power performance equation above and divide by f to arrive at the expression P/f which is typically used for power performance comparison (nJ/bit or W/Gb/s; they are numerically the same) then:

$$P = A'WV^2 f + k_1 VWe^{-\frac{V_{th}}{nV_\theta}} - k_1 VWe^{-\frac{V_{th}}{nV}} + \frac{V^3 k'}{W} e^{-DW/V}$$

$$P/f = A'WV^2 + (k_1 VWe^{-\frac{V_{th}}{nV_\theta}} - k_1 VWe^{-\frac{V_{th}}{nV}} + \frac{V^3 k'}{W} e^{-DW/V})/f$$

We now have the leakage term inversely proportional to the operating frequency f but we can substitute f for the expression $f = k_0(1 - \frac{V_{th}}{V})^\alpha$ where ko and α are experimentally derived constants so that:

$$P/f = A'WV^2 + (k_1 VWe^{-\frac{V_{th}}{nV_\theta}} - k_1 VWe^{-\frac{V_{th}}{nV}} + \frac{V^3 k'}{W} e^{-DW/V})/k_0(V - V_{th})^\alpha / V$$

$$= A'WV^2 + (k_1 V^2 We^{-\frac{V_{th}}{nV_\theta}} - k_1 V^2 We^{-\frac{V_{th}}{nV}} + \frac{V^4 k'}{W} e^{-DW/V})/k_0(V - V_{th})^\alpha$$

The above equation shows the functional form of the power performance relationship the first term being the dynamic power performance and the second term being the result of leakage currents. However there are now a number of unknown constants that require values to be assigned in order to get a curve fit to Nielson's router power performance figures. Usually the constants are experimentally determined but we can only fit values to match the available data as far as possible. The problem is that the curve generated from the equations show the power performance that could be achieved if power performance was the driving goal of the router design however this is probably not the case and performance is an equally likely if not more important driver for manufactures particularly in the earlier years when power consumption was less of an issue. For example the potential supply voltage decline shown in Figure 5-2 may not be adopted and higher voltages for increased performance might be used. Supply voltage is a major parameter affecting both performance and power consumption and although the trend shown in Figure 5-2 is reducing supply voltage manufactures may not be able to take full advantage of the potential reduction if high performance is also an important design consideration, this will be discussed again later.

There are some indicative values published for some of the parameters; in a 2003 paper by Nam Sung Kim et al "Leakage current: Moore's Law meets static power" from IEEE Explore, α is quoted as ~1.3, V_θ is quoted to be ~25mV at room temperature but increases linearly with temperature, the temperature co-efficient is not given but this value can be used initially. The value for A' was previously assigned to fit the dynamic power curve through the CRS1 performance point. The remaining unknowns are therefore k_0, k_1, k', n and D. The parameter k_0 is related to operating frequency which is taken to be the operating frequency of the majority of transistors in the VLSI chips. Because of the degree of parallelism that will be used within the chip design the transistor operating frequency will be a fraction of the device throughput speed and will be of the order of 100 to 200 Mb/s, a compromise figure of 150Mb/s was chosen for high speed devices e.g. devices with ~3Gb/s clock frequency for processor chips, the frequency parameter f will then scale with V and V_{th} once a value for k_0 is derived. The chosen point for the voltages to determine k0 are those assumed for the technology of the CRS1 router which gives a figure for k0 = 200. The other constants should make the leakage terms relatively small compared to the dynamic power component in earlier years, e.g. early 90s, and then modify the other terms to approach the CRS3 performance for the 2010/2011 time scale assuming these routers were designed for reasonable balance between good energy efficiency and performance (which may not have been the case). We can then extrapolate the curve to future years to estimate probable performance figures and compare with recently quoted but uncertain values.

The results shown in figure 4 follow empirical fitting of parameter values to get a close match to the most energy efficient of the routers shown in Figure 5-3. Those values for the constants after empirical fitting are shown in Table 12.

Constant	**Value**
A'	0.0391
k_0	200
k_1	1.148
k'	144.5
V_θ	0.025
n	43.52
a	1.3
D	0.00150

Table 12 empirically fitted values for the constants in the power efficiency equations

The thinner pale green line is the dynamic power performance curve alone and the brighter green and thicker line is the power performance curve with leakage current terms included. This latter curve is fitted to the most energy efficient route which is the Avici TRS and then parameters are adjusted to get close to the Cisco CSR3 router without impacting future performance too seriously and also keeping the impact of leakage in early years relatively small compared to dynamic power performance which was the situation at that time. The dashed curves show the leakage current terms. Of these the third term is the one that damages future performance and is assumed to be having an impact from 2008 onwards and does affect the energy efficiency of later routers such as the CRS3.

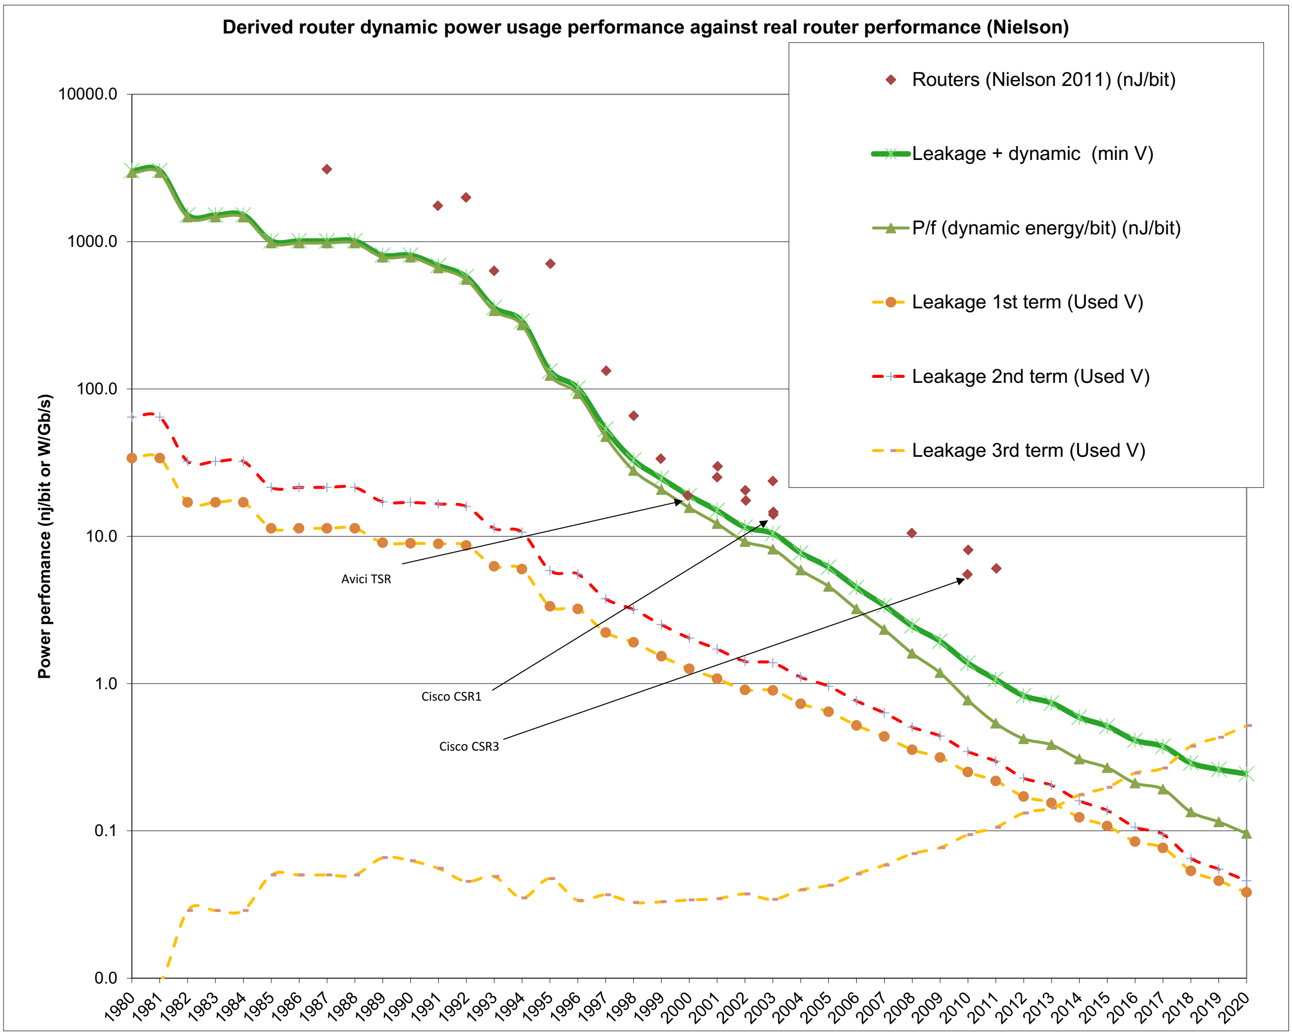

Figure 5-4 Comparison of dynamic VLSI chip power performance and router power performance

It can be seen from Figure 5-4 that earlier routers were not as efficient as they could have been but this is reasonable as they were almost certainly designed for performance rather than energy efficiency. It can also be seen that the later generation routers including the CRS3 do not follow the power performance curve even when the leakage current affects are included. This possibly indicates the continuing need to target performance and trade off power efficiency. One obvious way of trading performance against power efficiency is not to fully follow the voltage decline curve and operate at higher than the minimum supply voltages. To illustrate this the supply voltage curve in Figure 5-2 is approximated by a logistics curve with the lower asymptote set higher than the minimum voltage this was chosen to be 1.5 Volts rather than 0.7 volts. The curve selected to be used is shown in Figure 5-6 the minimum voltage of 1.5 v is a compromise between trying to match the power consumption of the CRS3 router and not having too great an impact on the leakage currents in the later years. The third term which is the gate leakage term has a V3 relationship and is therefore sensitive to high values of V.

The impact of limiting the minimum supply voltage using the logistic curve in Figure 5-6 is shown in Figure 5-5. The results show a significant increase in power consumption compared

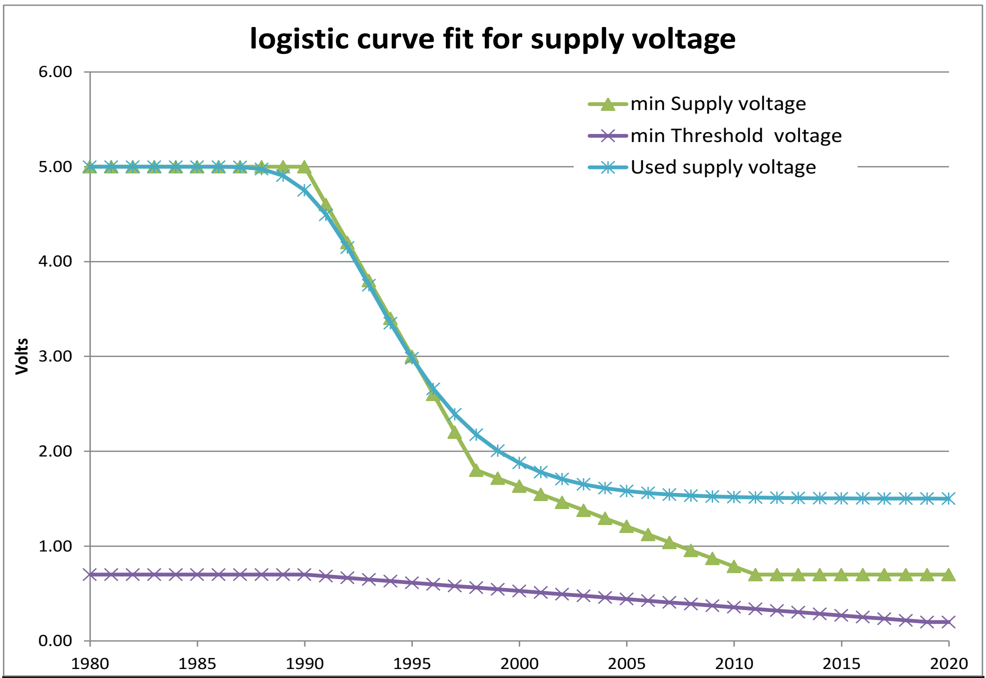

Figure 5-6 Logistic curve fit to supply voltage used to limit minimum

to the minimum voltage curves and are now closer to the trend of the 2008 to 2011 router performance figures from Nielson. If this trend is assumed into the following few years then it

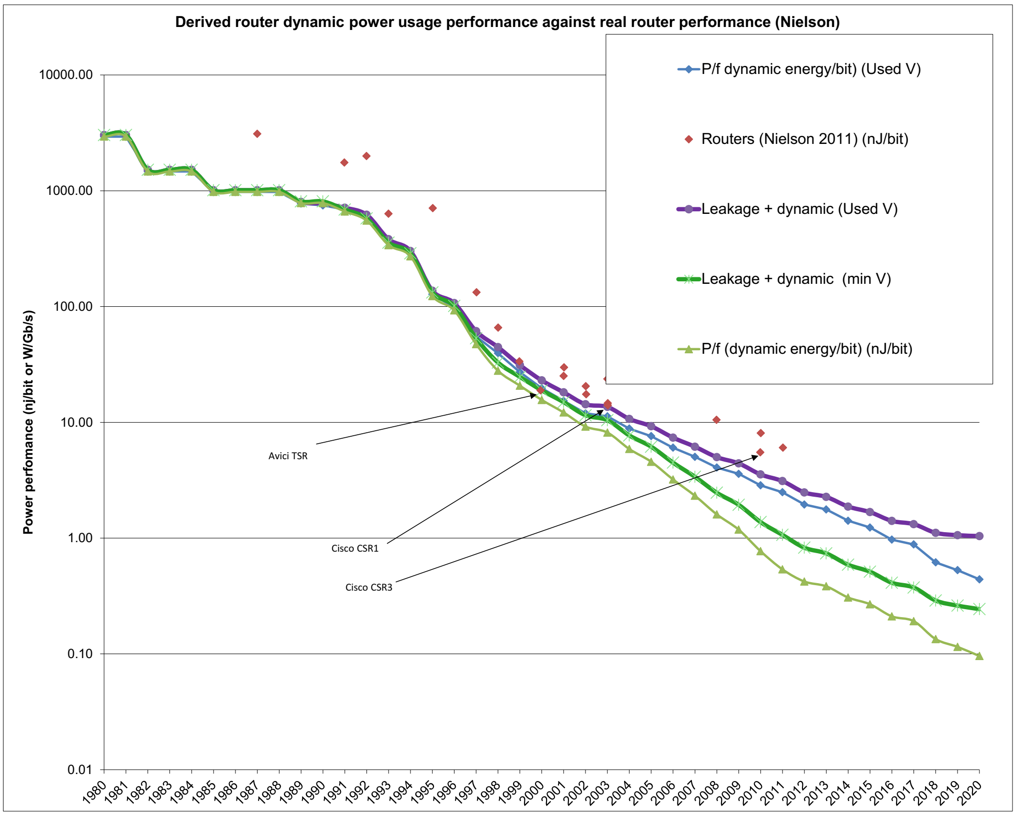

Figure 5-5 Impact of limiting minimum supply voltage.

can be used as indicative of potential router power consumption performance.

We can now compare the purple (leakage +dynamic (Used V) curve with other projected performance figures and also with the layer 2 switch energy efficiency which has power consumption performance of order 50% lower than router power performance figures. These comparisons are shown in Figure 5-7 Derived results compared to some published values for router and layer 2 switches power performance they also illustrate the discrepancies that have appeared in published sources.

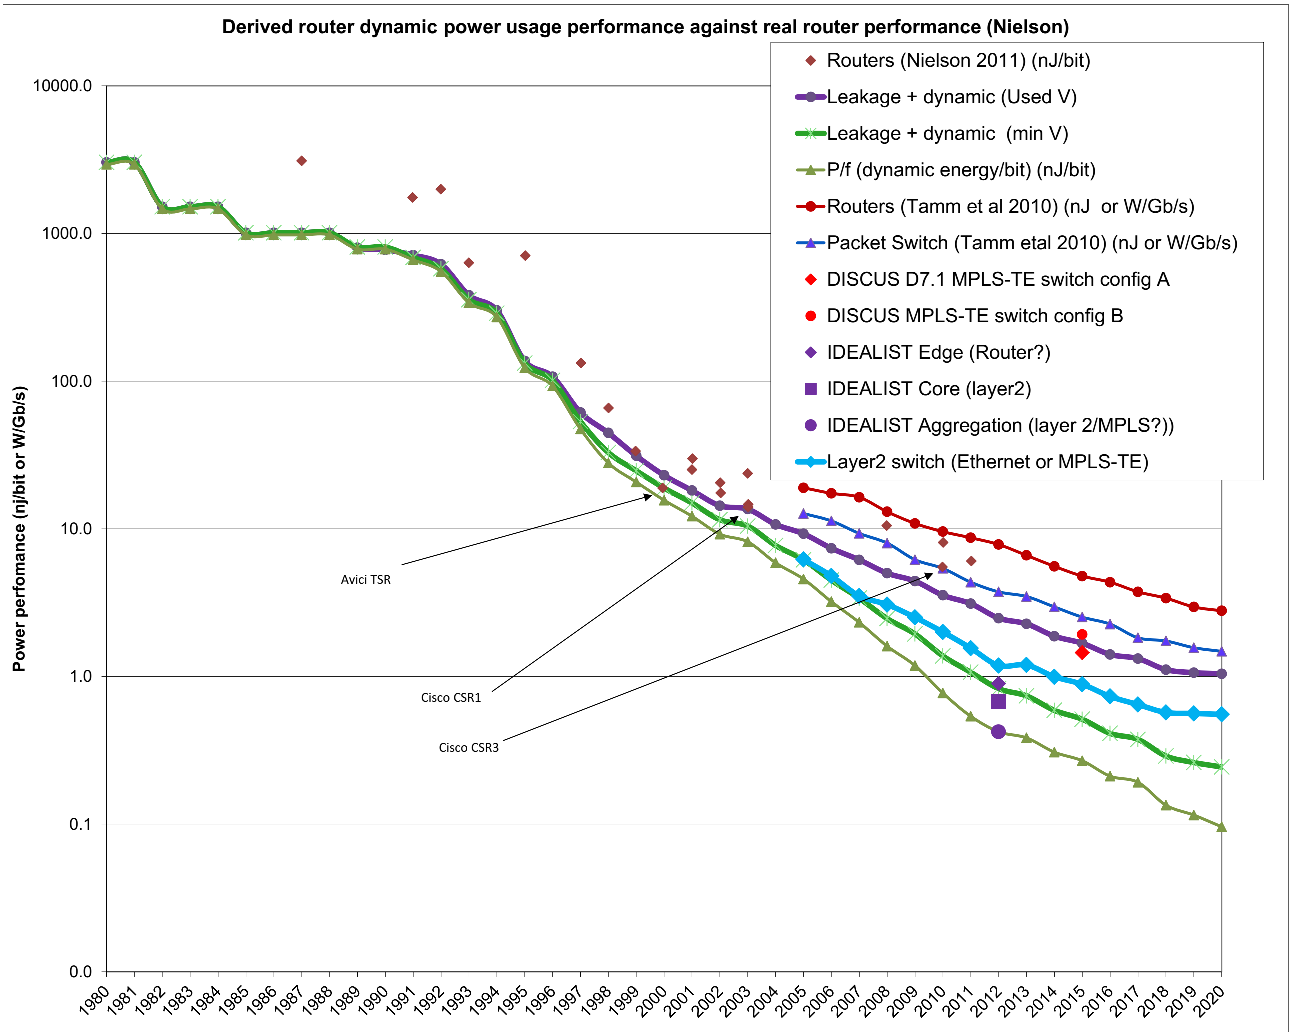

Figure 5-7 Derived results compared to some published values for router and layer 2 switches power performance

The thick lighter blue line shows the layer 2 switch performance curve and can be compared to the thick purple router power performance curve. This was derived by assuming the same ratios apply for routers to layer 2 switches as the results from the Tamm et al paper indicates ref**Error! Reference source not found.**However it would appear from these results that the Tamm paper results are an over estimate of router and layer 2 switch power consumption although they continue to be used and references in presentation and papers (e.g. GreenTouch ref[29]

The 2 red points are the values for the MPLS-TE switch described in DISCUS D7.1, when compared to the blue thick line for layer energy efficient layer 2 switches these values now seem reasonable if 2016 technology is assumed for their implementation. It is quite reasonable that the lower bound curves for routers and switches are not met in practice as switches will need to be designed for performance as well as energy efficiency and therefore some energy efficiencies will need to be sacrificed for performance gains.

The 3 purple points are for the Huawei “cloud switches” reported in the IDEALIST deliverable D1.5 ref [32]. These are below the blue and purple curves and appear to be far too low, these points are actually doubled compared to the figures actually quoted in D1.5 ref [32] to normalise them to the unidirectional definition for switch and router capacity used in the above graphs. Although not explicitly stated it is believed the capacity figure quoted in the IDEALIST report are using the bidirectional capacity definition, but even with this adjustment the figures seem to be too low and require further investigation.

5.2 Conclusions

Modelling fundamental VLSI chip performance and fitting the equations to published router values are a method of indicating minimum expected energy efficiency performance for routers and switches. By projecting the results and making some assumption about energy efficiency and performance trade-offs some pragmatic lower bound performance curves can be produced. These are the “leakage + dynamic (used V)” (dark purple curve in Figure 5-7) and the “layer2 switch “Ethernet or MPLS-TE” (thick blue curves in Figure 5-7). In practice it is reasonable to expect router and layer 2 switch power consumption to be above these lines, depending on the relative importance of energy efficiency compared to router/switch performance. The simple logistic curve used to implement the energy efficiency to performance trade-offs in the results of Figure 5-6 and Figure 5-7 could be modified if there is better data available to fit the curves to that would aid better power consumption forecast estimates to be made..

6 Summary

The modelling activities within the DISCUS project have been very broadly based and are a crucial part of the evaluation process of the architectural proposals. It has been a cross project activity involving many of the partners on various parts of the modelling problems. The overall objectives of the modelling activity have been largely achieved despite the total modelling task being much more complex and difficult than originally anticipated at the beginning of the project. One area where progress has not been as far as originally hoped for, is the full integration of the cash flow network model components. Current status is that the core model, the metro-node model and the access network models still need to be run independently and results incorporated into the cash flow model (which is embedded into the access and metro-network model). There is also further work incorporating optimisation results into the access and backhaul network models is required so that the latest version of the cash flow model currently does not produce meaningful cash flow results. We have earlier, simpler and more limited cash flow models working.

The lack of integration makes it more difficult to compare the range of effects on the end to end network economic performance of the individual parameter variables in specific areas of the network. For example the Metro-core node model has a lot of configuration options which currently have to be preselected and then the MC-node performance pre-calculated for that configuration for a range of node sizes and user bandwidth requirements. Those results are then passed to the cash flow model to be combined with a corresponding set of results from the core network model and results from the access and metro-network model. It is still intended to complete the model integration but that will now have to be carried out after DISCUS formally ends.

The results of modelling within the project have however confirmed the validity of many of the proposed benefits arising from adoption of the DISCUS architecture and which formed part of the original project objectives such as:

- Reduction in buildings housing electronic switching, routing and transmission equipment by closure of the majority of LE/COs and simplification of the core network
  - This was achieved by the optimisation modelling that showed LE site assignments of MC-nodes coupled with the flat core network modelling could enable removal of traffic processing routers switches line cards etc. from least 98% of LE sites in the smallest and largest countries in Europe
- Reduction in customer network ports (cf. xDSL) (by increasing LR-PON split to ~500 ways) ~99.8%.
  - This was achieved by power budget modelling and experimental demonstration (not described in this deliverable see D 4.6 and D8.5) and showed the 512 way split required for the objective could be achieved both for the “lollipop” LR-PON design for dense areas and the amplifier chain design for sparse rural areas.
- Reduction in network ports (by elimination of metro network and using flat optical core) ~70%
  - This was achieved by showing that the long reach of 100km to 125 km can be achieved, again via power budget modelling and technology demonstration showing that the reach and split can be simultaneously be achieved. Then by the optimisation and core modelling for the MC-node placement which enables the small number of MC nodes for the Flat core requirements, reduces the number of backhaul links required and eliminates the separate metro-network transmission systems. and associated port cards
- Reduction in network power consumption (neglecting CPE) with respect to Business as Usual (BAU) (from elimination of LE/COs, backhaul transmission systems and flat circuit switched optical core to minimise packet processing) ~95%.
  - The much improved power consumption modelling has shown that the power consumption for the DISCUS architecture is even better than originally proposed in the project proposal see figure.

All these objectives were important to establish in order to demonstrate that the DISCUS architecture is the most promising and viable future, superfast broadband network.

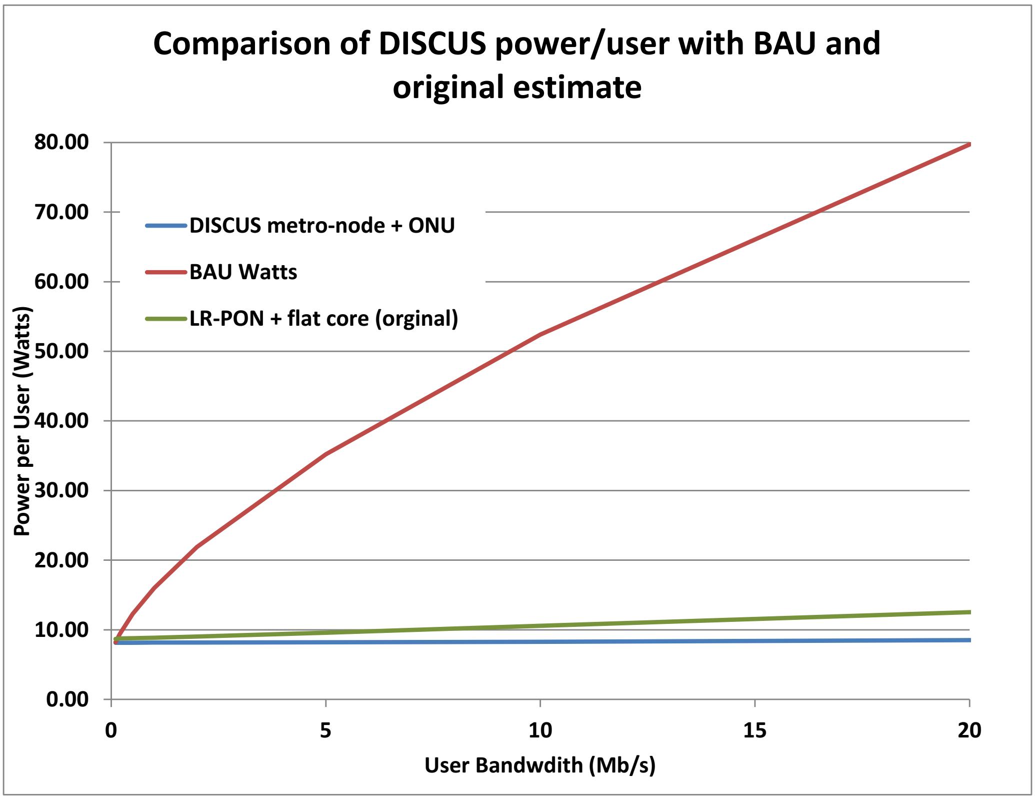

Figure 6-1 Improved power consumption for DISCUS architecture compared to BAU and original proposals

- Scalable to **>1000** times today's (2012) ADSL broadband capacity.
  - **The models scale beyond 200Mb/s sustained bandwidth per customer this is over 1000 times todays bandwidth which is in the region 100 to 200kb/s.**

.

7 References

[1] H. Hasegawa, "Large Scale Optical Cross-connect: architecture, performance analysis, and feasibility demonstration," in Optical Fiber Communication Conference 2015, OSA Technical Digest, paper W3J.1

[2] Puc, A.; Grosso, G.; Gavrilovic, P.; Fevrier, H.; Kaminski, A.; Burtsev, S.; Chang, D.; Foster, M.; Pelouch, W.; Perrier, P., "Ultra-wideband 10.7 Gb/s NRZ terrestrial transmission beyond 3000 km using all-Raman amplifiers," 31st European Conference on Optical Communication, 2005, ECOC 2005., pp.17,18 vol.1, 25-29 Sept. 2005

[3] Do-il Chang, Wayne Pelouch, Philippe Perrier, Herve Fevrier, Sergey Ten, Christopher Towery, and Sergejs Makovejs, "150 x 120 Gb/s unrepeatered transmission over 409.6 km of large effective area fiber with commercial Raman DWDM system," Opt. Express 22, 31057-31062 (2014)

[4] Ralf Huelsermann and Andreas Gladisch and Roman Klahne and Roland Wessa}lyn,"Applying a Discounted Cost Model and a Time-Dependent Demand Model to Access Network Cost Evaluation", Proceedings of the CTTE, pages 1--10, 2011, Berlin, Germany

[5] IDEALIST project, http://www.ict-idealist.eu/

[6] Rambach F., Konrad B., Dembeck L., Gebhard U., Gunkel M., Quagliotti M., Serra L., Lopez V., "A Multilayer Cost Model for Metro/Core Networks" in Journal of Optical Communication Networks, Vol. 5, No. 3, March 2013

[7] DISCUS Deliverable D6.5, "Final report on the specification of the Metro Core node architecture", July 31st 2015

[8] DISCUS Deliverable D7.7, "Consolidated/integrated optical core design", July 31st 2015

[9] Van Heddeghem W., Idzikowski F., Vereecken W., Colle D., Pickavet M., Demeester P., "Power consumption modeling in optical multilayer networks", in Photonic Networks, January 2012

[10] DISCUS, Deliverable 2.1 Initial reference architecture, The DISCUS Project (FP7 Grant 318137), 2014.

[11] DISCUS, Deliverable 2.3, Updated reference architectur, The DISCUS Project (FP7 Grant 318137), 2014.

[12] DISCUS, Deliverable 2.6, Architectural optimization for different geo-type, The DISCUS Project (FP7 Grant 318137), 2014.

[13] DISCUS, Deliverable 4.5, Optimization models for long-reach access networks . Technical report, The DISCUS Project (FP7 Grant 318137), 2014.

[14] DISCUS, Deliverable 4.10, Updated optimization models and methods for wireless/wireline integration, The DISCUS Project (FP7 Grant 318137), 2014.

[15] DISCUS, Deliverable 7.2, Optical island description, The DISCUS Project (FP7 Grant 318137), 2014.

[16] http://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/Cloud_Index_White_Paper.pdf

[17] http://stakeholders.ofcom.org.uk/binaries/consultations/bcmr/statement/research.pdf

[18] https://www.gov.uk/government/statistics/business-population-estimates-2015

[19] DISCUS, Deliverable 7.4, Preliminary quantitative results for flat optical network, The DISCUS Project (FP7 Grant 318137), 2014.

[20] DISCUS, Deliverable 7.6 Core network optimization and resilience strategies, The DISCUS Project (FP7 Grant 318137), 2014.

[21] http://sedac.ciesin.columbia.edu/data/collection/grump-v1

[22] https://www.samknows.com/

[23] DISCUS D6.1 metro core node

[24] http://www.investopedia.com/university/dcf/dcf3.asp

[25] BT OpenReach report and accounts 2015

[26] http://www.ons.gov.uk/ons/index.htm

[27] Plummer and Griffin: Material and process limits in silicon VLSI technology, Proceedings of the IEEE, Vol. 89,

[28] Die size reduction source data:. https://en.wikipedia.org/wiki/Semiconductor_device_fabrication

[29] Nielson, Router power dissipation historical datas ECOC 2011 presentation.

[30] Oliver Tamm and Christian Hermsmeyer and Allen M.Rush", "Eco-sustainable system and network architectures for future transport networks", BLTJ V14(4) 2010

[31] www.greentouch.org/uploads/documents/Rod_Tucker - TIA Charting a Path to Sustainable and Scalable ICT Networks.pdf

[32] http://cordis.europa.eu/docs/projects/cnect/9/317999/080/deliverables/001-IDEALISTD15Ares20144048842.pdf

8 Appendices

8.1 Appendix 1 - Linear cost parameters - “average” cost models

Consider an area containing a uniform distribution of N_P points and a reference point P. Assume the distribution of distances (cable lengths) from P to any other point P_r is as shown in Figure 8-1 with a mean μ_ℓ and a standard deviation σ_ℓ we also assume the distribution is

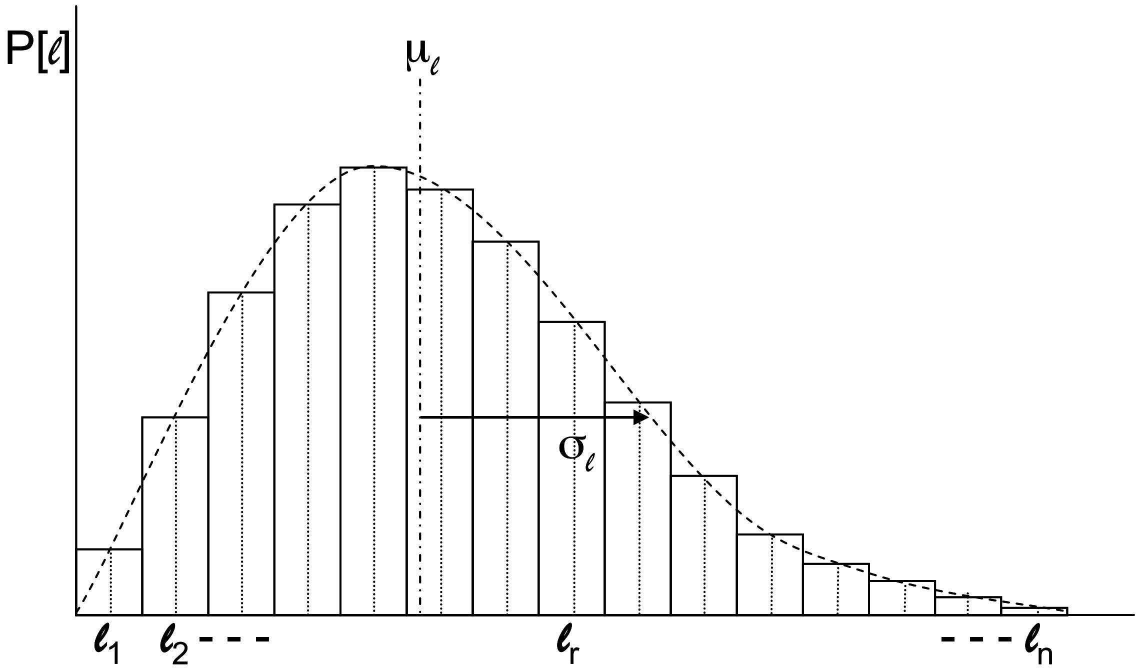

Figure 8-1 Assumed distribution of distances (cable lengths)

truncated at $l_{\max}$

The distribution is divided into n sample groups of mean lengths ℓ_1 to ℓ_n with spacing $\Delta\ell = \ell_{\max}/n$.

Therefore $\ell_1 = \Delta\ell/2$; $\ell_2 = \Delta\ell+\Delta\ell/2$; $\ell_3 = 2\Delta\ell+\Delta\ell/2$

$\ell_r = (r-1)\Delta\ell+\Delta\ell/2$ & $l_n = (n-1)\Delta\ell+\Delta\ell/2$

The number of cable lengths in the r^{th} sample group = $P[\ell_r]N_p$

The cost of an installed cable length = $a+b\ell$

Where a = fixed, length independent, cost; it includes such things as housings, travel time, splicing etc.

b = the length dependent cost.

ℓ = length of cable

Therefore the cost of installing the r^{th} sample group $C_r = P[l_r]N_p(a+bl_r)$

The total cost $C_T = \sum_{r=1}^{r=n} C_r = \sum_{r=1}^{r=n} P[l_r] N_p (a + bl_r)$ - - - - - (A1)

$$= \sum_{r=1}^{r=n} aP[l_r]N_p + bl_r P[l_r]N_p$$

$$= \sum_{r=1}^{r=n} aP[l_r]N_p + \sum_{r=1}^{r=n} bl_r P[l_r]N_p$$

$$= aN_p \sum_{r=1}^{r=n} P[l_r] + bN_p \sum_{r=1}^{r=n} l_r P[l_r]$$

However $\sum_{r=1}^{r=n} P[l_r] = 1$ and $\mu_l = \sum_{r=1}^{r=n} l_r P[l_r]$

And $C_T = aN_p + bN_p \mu_l$

Therefore we do not need to work with actual distributions to calculate the cable cost in an area only the mean lengths of cable routes are required so an "average" model suffices for the cost modelling purposes as long as cost has a linear relationship to length.

Optimisation techniques can be applied to optimise cable layouts within example areas/geo-types so that representative means can be determined but optimised layouts for whole countries are not required for country wide cost model analyses.

8.2 Appendix 2 - Average distance for uniform point distributions

Consider a point P about which there is a circularly symmetric segment of radius R and angular width θ (for a full circle θ = 2π). The segment contains N points uniformly distributed so that the average density is D points/unit area; we wish to find the mean distance to these points from point P.

Consider an annular segment (or annulus for a full circle) of radius r and width δr so that the inner radius is r-δr/2 and the outer radius is r+δr/2 see figure A2.1

The area of the annular segment $A_a \approx r\theta\delta r$; the number of points falling within the annular segment $N_A \approx Dr\theta\delta r$

Let L_{TA} = the sum of the lengths from point P to the points falling within the annular segment. Then $L_{TA} \approx rD\theta r\delta r$

Limit $\delta r \rightarrow 0$ then $L_{TA} = rDr\theta dr = D\theta r^2 dr$

Let L_T = the total length to all points from point P in the segment area. Therefore $L_T = \int_o^R D\theta r^2 dr = \frac{D\theta}{3}R^3$

But the total points in the circular segment $N = D\frac{\theta}{2}R^2$

The mean distance $\mu_l = \frac{L_T}{N} = \frac{D\theta}{3}R^3 / D\frac{\theta}{2}R^2 = \frac{2}{3}R$

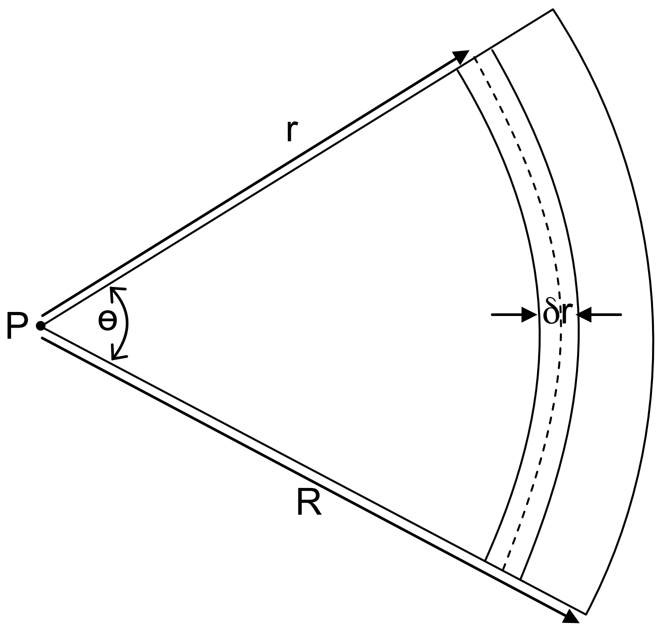

Figure A2.1 Circular segment radius r angle θ with uniform distribution of network locations with density D points/km²

8.3 Appendix 3 Exchange data base parameters

The table in this appendix lists the parameters stored in the exchange data base most of the values are not give and have to be estimated or calculated using simple geometric assumption about the exchange, cabinet and DP serving area.

Parameter	Type	Notes
exchange name	**Given**	Name from Sam Knows data base
exchCode	**Given**	ID from Sam Knows data base
pcode	**Given**	Post code from Sam Knows data base
Eastings	**Given**	OS 6 fig. Eastings from Sam Knows data base
Northings	**Given**	OS 6 fig. nothings from Sam Knows data base
serve Res	**Given**	Number of residential sites in LE area
serve NonRes	**Given**	Number of business sites in LE area
Total sites	Calculated	Sum of Res plus Bus sites in LE area
Estimated population	Estimated	Given Res sites x average household size
Area km^2	Calculated/given	Calculated from Voronoi polygons or given
Site Density	Calculated	Calculated from Total sites and LE area
Geo-type	Defined/ Calculated	From Geotype table. Defined by site density
Edge or middle	Defined/ Calculated	Defined by Voronoi polygon boundary , if at least one edge is an area boundary then exchange is Edge
Number of DELs	Estimated	A problem for FTTCab. Defined by the number of sites covered by average cabinet area placed at exchange location. Can be ignored for FTTH and assumed captured by first cabinet in e-side cable chain.
Average density of cabinets	Calculated	LE area / number of cabinets
Average Cabinets per Exchange	Calculated	Total sites in LE / (Cabinet size *fill factor). Cabinet size defined by geotype.
Ave. E-side radial distance to cabinets (km)	Calculated	Mean radial distance from LE to cabinets assuming uniform special distribution. = $\frac{2\sqrt{(N_{Cabs}/\pi)}}{3D_{Cab}}$
Average radius of E-side (km)	Calculated	Radius assuming circular LE area, $R_{LE}=\sqrt{(A_{LE}/\pi)}$

Parameter	Type	Notes
Average cabinets spacing (km)	Calculated	This is derived from a simple triangular packing model for cabinets uniformly distributed within the Exchange area. $L_{Cab} = R_{Cab}\sqrt{3}$
E-side cable route degree	Defined	From Geotype
Average Cabinets per Exchange Degree	Calculated	Ave cabinets per LE / E-side degree
Ave. Exch to LE area edge route length/Cab separation		This is calculated from the number of Cabinets in the exchange divided by the number of Cabinets per cable chain.
Average Cabs per E-side cable chain		The method for calculating the average number of Cabinets in an E-side cable chain used for the model is to select the lower of three bounds: • The first bound is calculated for the specific LE by taking the number of Cabinets in the LE area and dividing by the Cabinet degree. This ensures that at least one cable chain per degree (cable route out of the Exchange) will exist but will also produce the longest chains as the Cabinets are shared across the minimum number of cable chains for the LE area. In general these chains would be much longer than derived by the other methods and will be rarely selected within the model • A second bound uses the average cable route length from LE to Cabinets at the edge of the LE area divided by the average separation of Cabinets; this gives a more pragmatic value for the number of Cabinets expected along a typical cable chain route. • A third bound is an arbitrary maximum bound of 8 Cabinets per cable chain a limit not expected to be seen but a bound to keep a limit that aids implementation of the model. The minimum of these bounds is selected for the model value. In practice the calculations in the model nearly always selects the value from the average cable route length divided by the Cabinet separation.
Ave No. E-side cables per exchange area		This is calculated from the number of Cabinets in the exchange divided by the number of Cabinets per cable chain.
Average cables per Exchange degree		Calculated by dividing the number of E-side cables by the Exchange E-side cable degree. It could be considered that if cable branching points exist then **the** cables per degree could be consolidated into larger fibre count cables to reduce cost. The first branching point could be assumed to be the first Cabinet in the chain and a second at the second

Parameter	Type	Notes
		Cabinet in the chains. After the second Cabinet all cable chains could be considered as separate. For simplicity separate cables per chain are assumed in the model
% e-side cable OH		From Geotype
% e-side cable UG		From Geotype
%e-side cable DB		From Geotype
Target Cabinet size (sites)		This target cabinet size matches the binary split size that would give good average PON utilization if it were achievable. Note that this figure includes the growth provided by the fill factor parameter. The value is fetched from the Geotype table.
Average density of DPs (DPs/km^2)		This is calculated by dividing the LE site density by the average number of sites per DP
Average DPs per PCP		This is the number of DPs per cabinet to fully utilise the cabinet serving area given the DP splitter size, note the target cabinet size is the size including the growth provided by the fil factor parameter
Average D-side (PCP to DPs) radial distance (km)		This uses the relationship between DP site density assuming uniform spatial distribution of DPs and the mean radial distance to those DPs from the cabinet. That relationship is Rave =(2/3)Sqrt(Ndp/(pi()Ddp)
Mean radius of cabinet area (km)		This uses the mean radius in the previous cell multiplied by 3/2 to derive the equivalent circular cabinet area radius
Average radial spacing between DPs (km)		This average separation assumes optimal triangular packing of DPs across an assumed circular cabinet serving area and is therefore given by sqrt(3) x DP area radius.
Average number of Cabinets in Exchange area		This is the total sites in the exchange area divided by the cabinet size without the spare site capacity, that is cabinet size x fill factor
Ave. D-side cable route degree		Fetched from the geotype table. This is the average physical degree from the PCP sites (the number of separate physical cable routes into and out of the PCP). This is a difficult figure to derive if not given in the exchange data set but in order to have a figure that relates to the exchange type rather than just a national average figure applied to all PCPs an assumption related to geotypes is made: The degree range is 2, 3 or 4 and is related to the geo-types as follows:

Parameter	Type	Notes
		City, Metro, Urban1 are assumed to be degree 4 Urban2, Urban3, Rural1 are assumed to be degree 3 Rural2, Rural3, Sparse are assumed to be degree 2
Average DPs per Cabinet Degree		Simply calculated from the DPs/cabinet divided by the cabinet degree.
Ave. Cab route length to area edge/DP radial separation		This value is the radius of the cabinet area multiplied by the route factor to represent a maximum cabinet cable length. This length is then divided by the DP radial separation to determine a figure for the potential number of DPs in a cable chain from cabinet to the cabinet area edge. This value is probably a lower bound on the number of DPs along a cabinet cable route capturing DPs as these cable routes may be longer than a radial distance multiplied by a route factor. The radial separation of DPs is used as DPs are relatively close together and the cables will generally run a long a relative straight piece of road to reach the next DP however at each DP there could be a junction that deviates the route hence the total route radius multiplied by a routing factor.
Average DPs per D-side cable chain		The method for calculating the number of DPs in a cable chain used for the model is to select the lower of three bounds: • The first bound is calculated for the specific LE by taking the number of DPs per cabinet and dividing by the Cabinet degree. This ensures that at least one cable chain per degree (cable route out of the cabinet) will exist but will also produce the longest chains as the DPs are shared across the minimum number of cable chains for the cabinet serving area. In general these chains would be much longer than derived by the other methods and will be rarely selected within the model • A second bound uses the average cable route length from cabinet to DPs at the edge of the cabinet area divided by the average separation of DPs, this gives a more pragmatic value for the number of DPs expected along a typical cable chain route. • A third bound is an arbitrary maximum bound of 8 DPs per cable chain a limit not expected to be seen but a bound to keep a limit that aids implementation of the model. The minimum of these bounds is selected for the model value. In practice the calculations in the model nearly always selects the value from the average cable route length divided by the DP separation.
Number D-side cables per Exchange area		Derived by dividing the number of DPs in the exchange area by the number of DPs per cable chain.

Parameter	Type	Notes
Ave. number of D-side cables per Cabinet		Derived by dividing the number of cable chains per LE area by the number of cabinets in the LE area
Ave. number of D-side cables per Cab. degree		Derived by dividing the number of cable chains per cabinet by the cabinet degree
% d-side cable OH		From Geotype
% d-side cable UG		From Geotype
% d-side cable DB		From Geotype
Target mean DP splitter size		From Geotype
Target Average sites/DP		This is simply the target splitter size multiplied by the general fill factor that allows for spare fibre ends for future growth of the customer base. At this stage it is left as a fractional value to minimise rounding errors later in the model calculations.
Average sites per DP used in model		
Number of DPs in Exchange area		This is simply the total number of sites in the LE area divided by the average number of sites per DP
Average radius of DP area (m)		This is the radius of a circular area defined by the number of customer per DP and the customer site density. Basically assume a radial capture area around the DP.
Mean drop length for Geotype (radial OH) (m)		For a uniform distribution of sites inside a radius R the mean distance to those sites is 2R/3. however this assumes the customers are spatially spread in a single horizontal plane and does not take into account multi floor dwellings and offices. a crude accommodation for this multi floor buildings is provided by ensuring a minimum drop length defined in cell AX12 (the cell above this cell)
Mean drop length for Geotype (UG (m)		To accommodate UG drops which are not radial a routing factor similar to cable routing factor to the radial OH distance for simplicity this routing factor is simply kept at SQRT(2).
Mean drop length (UG+OH)		This computes the average drop length across UG and OH drops for the exchange area using the OH to UG distribution defined in the geotype table. It is only used for estimating average VDSL and ADSL bandwidths for the exchange area. So that copper technologies performance limits can be taken into account when comparing with FTTH technologies.
% drop cable OH		From Geotype

Parameter	Type	Notes
% drop cable UG		From Geotype
% sites taking bespoke fibre ccts		Because the DISCUS architecture has huge potential for providing private ccts etc. (up to 100Gb/s over the LR-PON aces-metro infrastructure) the demand for bespoke fibre networks is assumed to be confined to the large and very large business premises and these are further assumed to be confined to the metro and city geotypes. However in order to have some distribution of bespoke network capability in all geotypes a basic minimum of 0.1% of customer sites in all other geotypes is also provided for. If for small exchanges this yields less than 2 fibres per E-side cable chain then the minimum value of 2 fibre per E-side chain is set.
Metro-node Chain ID		The Id number of the backhaul cable chain that connects the LE to the primary and secondary metro-code nodes
First Metro-node		Either the Primary or Secondary MC node at the end of a backhaul cable chain arbitrarily called the first.
Second Metro-node		Either the Primary or Secondary MC node at the other end of the backhaul cable chain arbitrarily called the second.
Chain length km		
Chain cable size		
Average distance Cabinet to customer		Used for ADSL and FTTCab average bandwidth performance estimates
"Max" distance cabinet to customer		Used for ADSL and FTTCab minimum bandwidth performance estimates
Average VDSL2 BW (Mb/s)		Average FTTCab for Exchange area using average Cabinet to customer distance.
Min VDSL2 BW		This is not the absolute minimum bandwidth in the exchange areas due to line length but merely an indicator of the range of bandwidth based when assuming a circular exchange and cabinet areas.
Average ADSLK2+ BW		
Min ADSL2+ BW		

8.4 Appendix 4 Infrastructure models

Using cable chains assumes that cables will run along streets and drop off fibre to service DPs and Cabinets as required. The number of cable chains and the number of DPs or cabinets in a cable chain are derived within the Exchange data base used within the cash flow model. These values are based on geometrical assumptions and uniform distributions of customers within the serving areas of street furniture and building nodes. For simplicity the first version of the model will use separate cables for each chain. This could be modified later to assume branching points along physical duct routes that share cable chain routes. In these situations there may be advantages of combining smaller cables into larger cables which could reduce per fibre costs further if the cables have not already exceeded the maximum cable size used in the model (276 fibre cables) however this saving will be offset by additional splice housings and splicing at the cable branching points and more complex planning rules.

The basic data assumed to be publically available is exchange locations and the number of sites being served by the exchanges. From the location data size of the exchange areas are computed using Voronoi polygons (this also needs coastal and country boundary data for exchanges adjacent to the edges of the country or region). Once the area data is computed the exchanges are classified into geotypes based on site density. Assumptions are then made about distributions of basic parameters across these geotypes that match any known national statistics. Once these assignments are completed all other parameters are computed using geometrical and average models, again checking against known national statistics, when these are available, for normalisation of values and distributions.

With the exchange data base populated the cable sizes and lengths for all the cable sections in the exchange area need to be computed. The drop section is computed first then the d-side cable chains of DPs then the E-side chains of cabinets and finally the backhaul chains of local exchanges that span a pair of metro-nodes. The cable model needs to be generic enough to accommodate LR-PON, GPON and point to point fibre, For the FTTCab model a point to point fibre e-side model is used although high capacity PONs such as that proposed by DISCUS could also service Cabinet locations, if time allows a LR-PON FTTCab solution will also be added for comparison. FTTCurb or G.Fast solution will be added at a later stage but this is probably beyond the resources available in DISCUS, however the modular structure of the model should readily enable additional technology modules to be added after the end of the DISCUS project.

The main complexity is deciding on the splitter housing structure and the spare fibre handling which includes bespoke fibre and fibre for future PON growth. The fundamental assumption here is that there will be an allowance for growth in the day one build and dimensioning. This is provided for by including a fill factor across the model. Typically this is set at 80% allowing 20% spare capacity for growth. However in PON solutions the number of spare fibre needs to reduce as the cables run from the customer premise towards the LE site through the D-side, then the E-side and through the backhaul network to the core. This is because most growth will be randomly distributed and not all growth provided near the customer needs to be translated up through the network. Also for PON solutions additional splitter nodes would be added which reduces upper network fibre counts. Making this sparing scalable and internally self-consistent adds complexity to the model and it is necessary to get the sparing strategy and methodology clearly defined upfront.

To make sparing effect spare fibre needs to be available at the access edge i.e. at the DP position. However as previously mentioned it is neither necessary nor economic to carry all the spare fibre at the edge deeper than necessary into the network. We therefore will have spare fibre unterminated and not spliced through at intermediate street furniture nodes in the D-side and E-side cable networks.

The basic design rules used within the model are therefore as follows:

1. The number of spare fibres entering the LE site should be ~working fibres *(1 - fill factor).
2. There will be a minimum of two fibres per DP for bespoke fibre networks (generally will only be required by larger business customers)
3. The D-side cable chain carries four categories of fibre:
   - Fibres feeding the PON splitters (includes fibre for additional network side splitter ports if NxM splitters are deployed).
   - Fibres for the bespoke networks
   - Spare fibre for growth (additional PONS)
   - Spare fibre arising from fitting actual cable sizes to the required fibre count (the fitted cable size will invariably be of greater size than the minimum required fibre count which includes fibres for growth).

Not all these spare fibres need or can to be spliced through at intermediate DP locations in the DP cable chain, for two reasons: Statistical considerations mean that not all the spare capacity needs to be passed up the chain but also the cables on the input and output sides of a DP position will generally not have equal numbers of spare fibres and only the smaller of the input or output spare fibre can e spliced through. However although splicing through spare fibre adds cost at day one installation, to minimise future disturbance, all spare fibre in a D-side cable chain that can be spliced through will be spliced through to the cabinet location, regardless of statistical demand. This is to minimise the number of future network visits required to implement growth and to minimise intervention faults and maximise a "hands off" operational network strategy.

8.4.1 D-side Cable Chain model

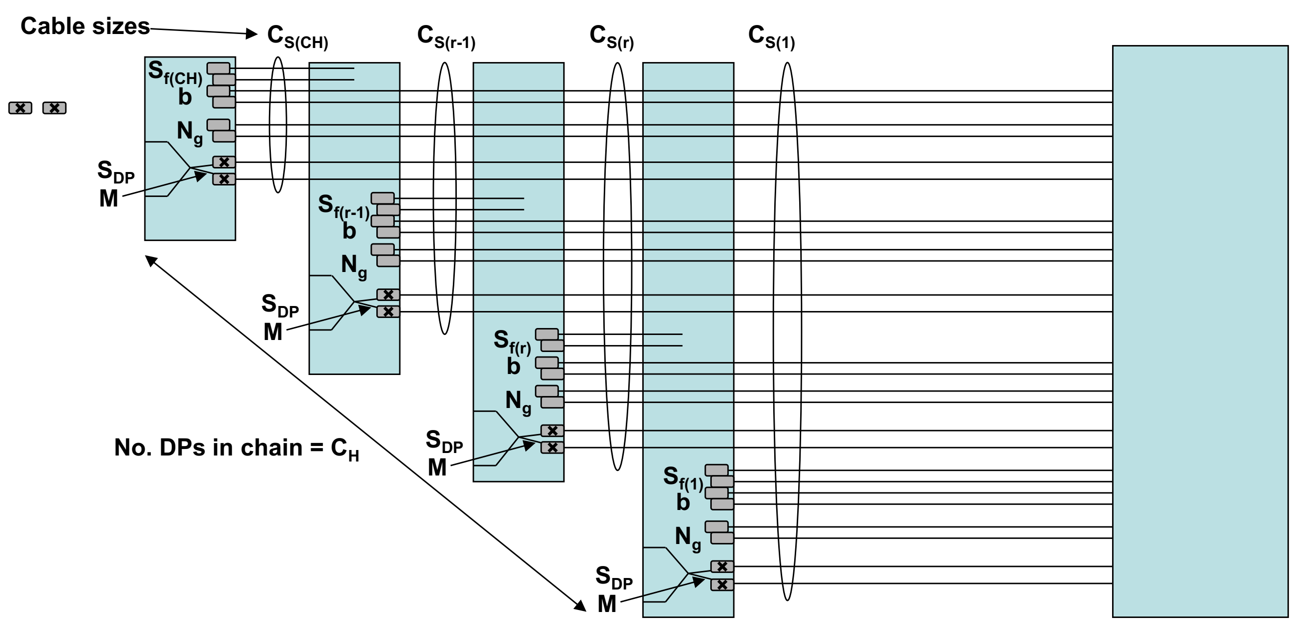

Fibre counts at Cabinet:

Bespoke fibre = $C_H b$ where b= number of bespoke fibres/DP

Spare fibres for growth = $C_H N_g(1+f)$; (1+f) if for two or single fibre working

Fibre for upstream splitter ports = $C_H M_{DP}(1+f)$

(f is a flag, f=1 for two fibre working ODN)

Minimum fibres in cable $C_{S(1)}$= $MinC_{S(1)} = C_H b+(1+f)C_H(N_g+MDP) = C_H(b+(1+f)(N_g+M_{DP}))$

(assumes all spare splitter ports, fibres for growth and DP bespoke fibre are spliced through to the cabinet)

Spare fibres in cable $C_{S(1)} = C_{S(1)} - MinC_{S(1)} = C_{S(1)}-C_{H(}b+(1+f)(N_g+M_{DP}))$

Minimum fibres in cable at r^{th} DP = $MinC_{S(r)} = MinC_{S(1)} - (r-1)C_H(b + (1+f)(H_g+ M_{DP})) = C_H(b+(1+f)(N_g+M_{DP})) - (r-1)C_H(b+(1+f)(N_g+M_{DP}))$

$= (C_H-r+1)(b +(1+f)(Ng+M_{DP}))$

Spare fibres in cable $C_{S(r)} = C_{S(r)} - (C_H - r+1)(b +(1+f)(N_g+M_{DP}))$ and spare fibres in cable $C_{S(r+1)} = C_{S(r+1)} - (C_H - r)(b+(1+r)(b+(1+f)(H_g+M_{DP}))$

The maximum number of spare fibres that span the whole chain due to cable sizes exceeding the required number of fibre either side of a DP is the minimum of the set of absolute vales of the spare fibre count difference about the set of DPs in the chain. These spare fibres should be distributed across the DPs in the chain. However this is a complex planning rule in terms of fibre splicing of these spare fibres at the DPs. Although there is a small cost penalty a simpler planning rule would be to splice all spare fibre through at each DP, this number of splices at the r^{th} DP is:

$ABS(C_{S(r)} - (C_H-r+1)(b +(1+f)(N_g+M_{DP})) - C_{S(r+1)} + (C_H - r)(b +(1+f)(Ng+M))) = ABS(C_{S(r)} - C_{S(r+1)} - ((C_H - r +1) + (C_H - r))(b+(1+f)(H_g+M_{DP}))) =$ $\mathbf{ABS(C_{S(r)} - C_{S(r+1)} -(2C_H - 2r + 1)(b+(1+f)(N_g+M_{DP})))}$**.**

This maximises the number of spare fibres spliced through to the cabinet but they do not all need to be spliced through to the local exchange over the E-side cables.

The DP splitter will be installed with splice trays for the incoming fibre cable/unit and the drop fibre units. Drop fibre splices however are not included in the splitter node cost as they are included in the customer drop costs as a "Just in Time" cost (JIT).

Note: Often the drop fibre unit will be two or four fibres. For this model it is assumed that only one fibre is spliced for single fibre working (typical) or two fibres for two fibre working. Any spare fibres in the drop fibre unit are assumed to be stored on the splice tray for that drop but without splices, two fibres per tray. For the urban areas typical splitter sizes at the DP location will be 32 way with some 16 way splitters, and in the sparser rural areas 8 way. In the current version of the cash flow model the DP splitter size determines the cabinet splitter size and assumes all DP splitters in the exchange area will be the same size. In sparse rural areas the DP splitters could be split into additional tiers with smaller end point splitters as small as 4-way and occasionally 2-way. Optimisation techniques to fit cables and splitters to these areas are being developed, they are not included in the cash flow model at present and sparse rural is accommodated by assuming longer drop lengths dependent on mean customer density.

Average number of DPs in Chain = C_H

The value of C_H needs to be determined if not given in the base exchange data which will usually be the case. For a given exchange area we will have the total sites and a figure for area covered either given or calculated (this will be done using Voronoi polygons). The average number of DPs can be assumed to be the number of customer terminations in the exchange area divided by the target number for the average customers per DP. The target number is defined by the target splitter size multiplied by the design fill factor to allow for growth. The splitter size is determined by geotype. It is generally 32 way split but drops to 16 way split for Rural 2 and then to 8 for the Rural 3 and Sparse geotypes to reduce drop lengths

There are $N_{DP(r)}$ fibre drops at r^{th} DP (actual number fitted depends on take rate but 100% assumed to be passed), $N_{DP(r)}$ = Number customers passed at the r^{th} DP divided by the fill factor f_f to allow spare fibre for growth.

Splitter size at DP = $S_{DP} = 2^n$; n = 3, 4, 5 ways. The splitters are N x M_{DP} where $M_{DP} = 2^m$; m = 0, 1, 2

At each DP the minimum number of D-side cable fibres = M_{DP} + b + Ng; where M_{DP} = splitter D-side ports, b = No. bespoke fibres/DP, Ng = fibres for growth/DP = roundup(M_{DP}*(1-f_f),,0).

To minimise intervention at DPs except for adding new customers all these fibres are assumed to be passed up the cable chain to the cabinet location.

For PONs we assume the DP splitters are all equal and are placed in a position to maximise average utilisation. To accommodate fractional values of the number of DP chains required for a specific exchange area a proportion of the chains will be 1 DP longer.

Total No. fibre drops along chain (100% passed) = $\sum_{i=1}^{i=C_H} N_{DP(i)}$; Number of drop splitter ports = C_H x DP_{sp}

$$\therefore \text{Utilisation} = \frac{1}{C_H DP_{sp}} \sum_{i=1}^{i=C_H} N_{DP(i)}$$

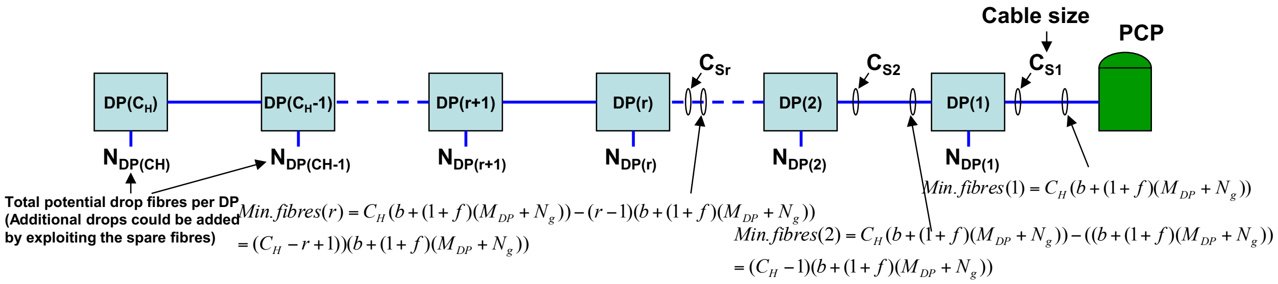

No. Fibres from each cable chain entering cabinet:

Fibres from splitters = $C_H M_{DP}(1+f)$ the factor f = 0 for single fibre working and f=1 for two fibre working

Bespoke fibres = $C_H b$ (Min. b = 2)

Growth fibres = $C_H N_g(1+f)$

Spare fibre from cable = C_{S1} – Min.fibres(1)

At the cabinet C_H(1+f) splitter fibre will be connected to cabinet splitters and the number of E-side fibre will be C_H(1+f)/S_C (S_C = cabinet split size) per D-side cable chain.

An **option** needs to be selected to pass the remaining splitter fibres up to the LE site or just store them in the cabinet for future use as required (this will be set with a flag f_2).

A weighting factor W_b is used to reduce the number of bespoke fibres carried up to the LE site (due to the relatively low probability that any one DP will use bespoke fibres). Bespoke fibres will not be spliced through at the cabinet until required but are stored on splice trays (two fibres per tray).

Growth fibre for future additional PONs would be connected to further cabinet splitters in the future and therefore can be divided by S_C for the E-side fibre count.

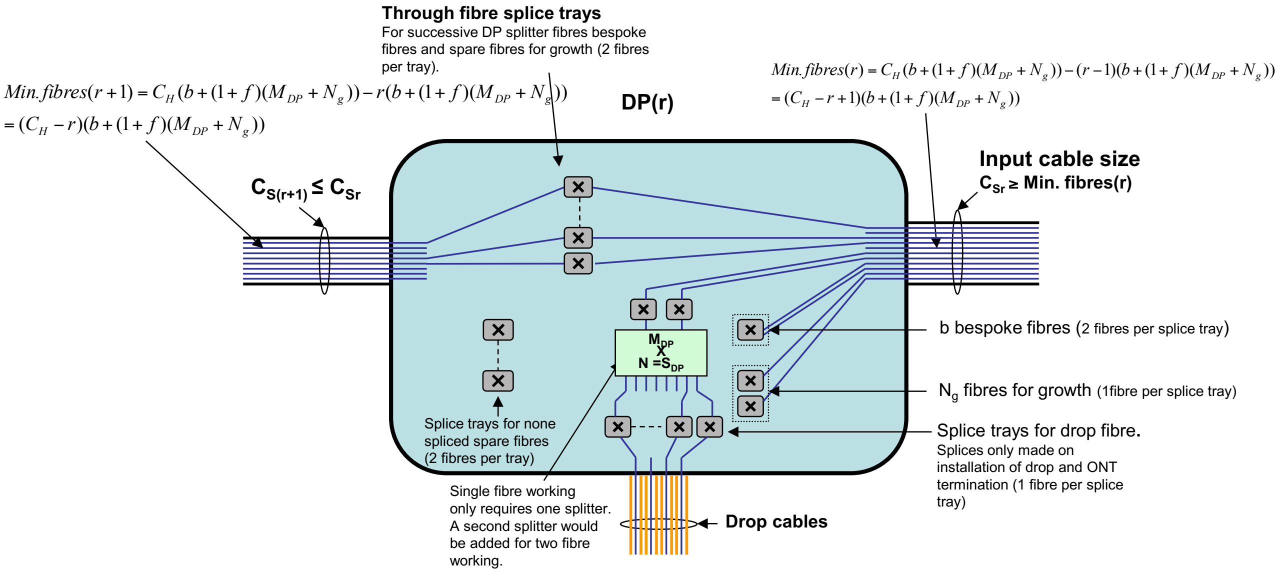

Cost for DP splitter housing =

cost of housing

+ cost of splitters

+ cost of splice trays

+ cost of splicing

+ cost of installation and site visit

Cost of splitters $C_{SP} = (1+f)(M_{DP}+S_{DP})*C_{SPP}$

Where f = 1 or 0 for two or single fibre working.

S_{DP} = splitter size, M_{DP} = Number D-side splitter ports

and C_{SPP} = cost per splitter port.

Number of splice trays: single fibre circuit management are assumed for the Drop fibre splices i.e. I splice per tray when single fibre working is used and two splices per tray for two fibre working

No. Drop trays D_T = N+roundup(b/2,0)+N_g ; Two fibre working is assumed for bespoke fibres

No. through fibre trays T_T = roundup((Min.fibres(r+1) + Min(S_{in}, S_{out})/2,0)

Where S_{in} = spare fibres in input cable = $C_{S\text{®}}$ – Min.fibres(r)

S_{out} = spare fibres in output cable = $C_{S(r+1)}$ – Min.fibres(r+1)

S_p = trays for spare fibres not spliced through = ABS((S_{in} – S_{out})/2); Stores two fibre per tray

Total splice trays $T_{TOT} = D_T + T_T + S_P$

$\therefore T_{TOT}$ = N+roundup(b/2,0)+N_g + roundup((Min.fibres(r+1) + Min(S_{in}, S_{out})/2,0) + ABS((S_{in} – S_{out})/2).

Total splices at installation = T_{SP} = Min.fibres(r+1) + Min(S_{in}, S_{out}) + M_{DP}

Within the model it is assumed that the DP cable chain is installed as one job and site visit travelling time and installation time are averaged over all DPs and therefore a single value for the installation time per DP is used. Splicing labour time is included within cost per splice figure.

Variables:

C_H = No. DPs in D-side cable chain

S_{DP} = splitter size

M_{DP} = No. network side splitter ports

b = No. bespoke fibres/DP

N_g = No. fibres for growth/DP

r = position in chain

f = a flag for single or two fibre working f = 1 for two fibres and 0 for single fibre working

8.4.2 Point to Point fibre DP cable chain model

Total DPs in Chain = C_H

$N_{DP(r)}$ fibre drops at r^{th} DP (actual number fitted depends on take rate but 100% assumed to be passed)

$N_{DP(r)}$ = Number customers passed at the r^{th} DP divided by the fill factor f_f to allow spare fibre for growth. Assume bespoke fibres are now included within spare fibre count.

Total No. fibre drops along chain (100% passed) = $\sum_{i=1}^{i=C_H} N_{DP(i)}$

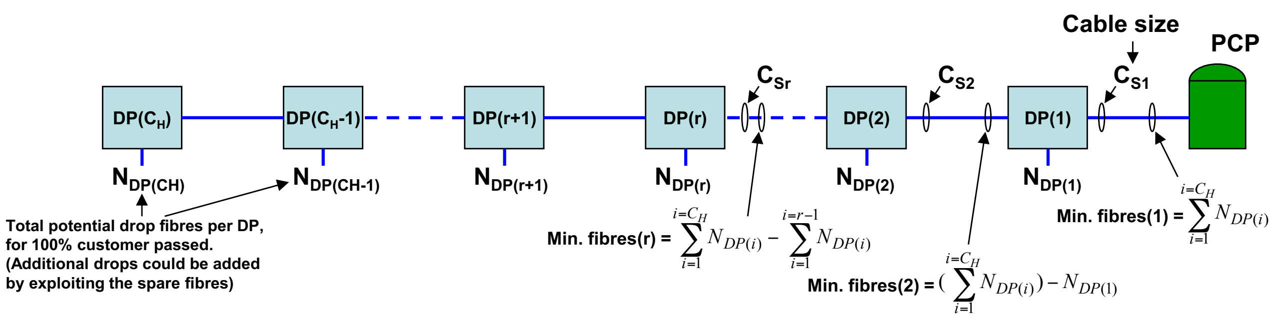

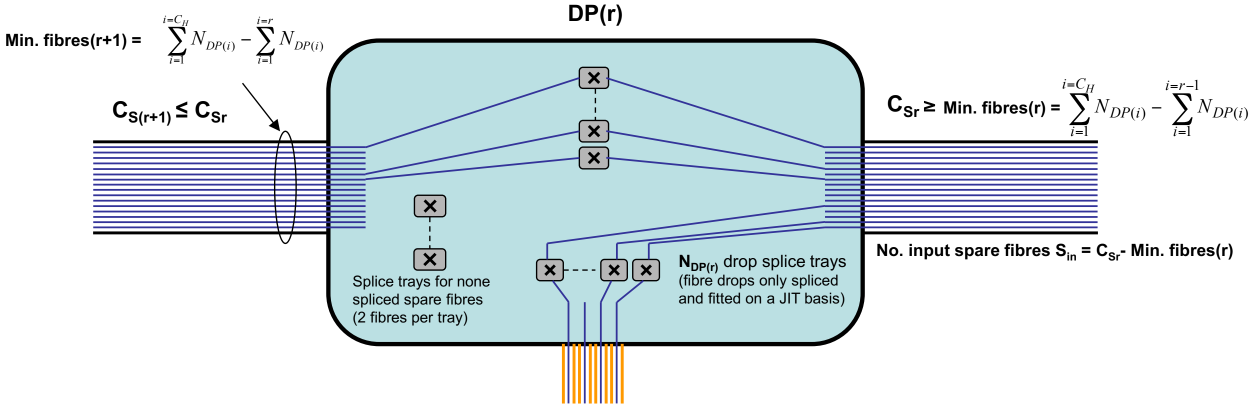

No. output spare fibres $S_{out} = C_{S(r+1)}$- Min. fibres(r+1)

Cost for DP splitter housing =

cost of housing + cost of splice trays

+ cost of splicing + cost of installation and site visit

Number splice trays for non-splice spare fibres S_p = roundup(ABS(($S_{out} - S_{in}$)/2),0)

Number through fibres for potential drops = Min. fibres(r+1)

Number spare fibres spliced through = MIN(S_{in}, S_{out})

Total splice trays for spliced through fibres = roundup((Min.fibres(r+1) + MIN(S_{in}, S_{out}))/2,0)

Total splice tray for DP(r) = roundup((Min.fibres(r+1) + MIN(S_{in}, S_{out}))/2,0) + roundup(ABS(($S_{out} - S_{in}$)/2),0) + $N_{DP(r)}$

Total splices = Min.fibres(r+1) + Min(S_{in}, S_{out}); Drop fibre splices are included within the drop installation cost

Note: to relate the equations to the PON splitter case assume an average number of customers per DP so that $N_{DP(i)} = N_{DPave} = S_{DP}f_f$

and Min.fibres(1) = $C_H S_{DP}$

Min.fibres(r) = $C_H S_{DP}$ - (r-1)S_{DP} = $S_{DP}(C_H$ - r+1)

Min.fibres(r+1) = $C_H S_{DP}$ - rS_{DP} = $S_{DP}(C_H - r)$

8.4.3 *E-side Cable Chain models*

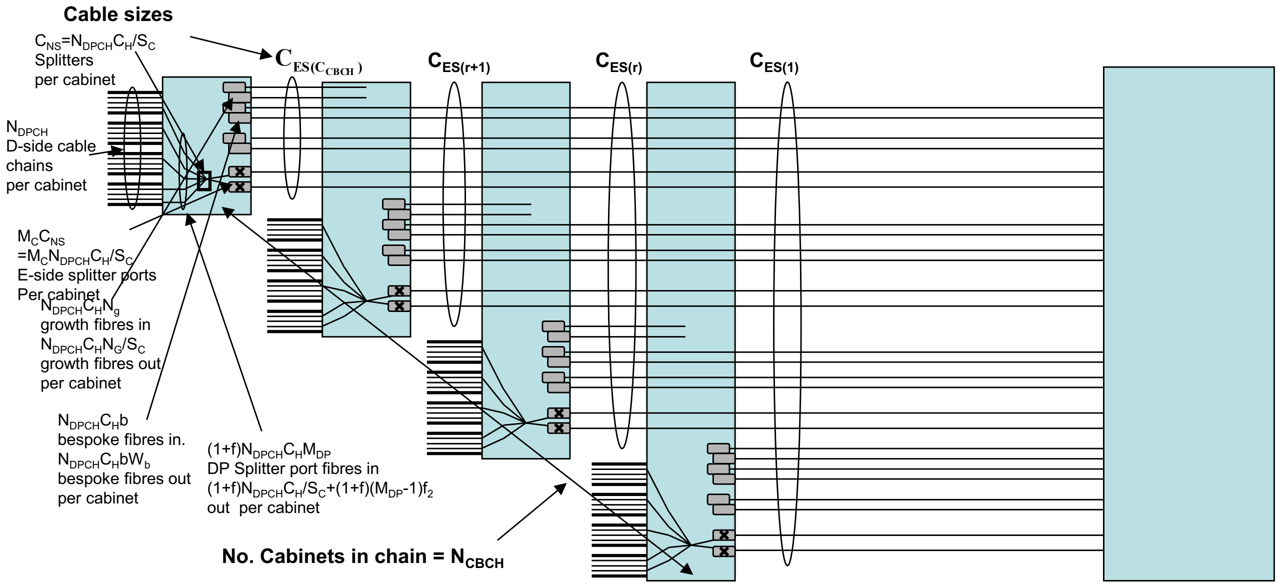

Fibre counts at LE per Cabinet Chain:

Minimum fibres at Cabinet from all DP chains = $N_{DPCH}\ C_H(b+(1+f)(N_g+M_{DP}))$; N_{DPCH} = Ave. No. DP chains/cabinet; C_H = the average number of DPs per DP chain, b is the number of bespoke fibres per DP, f is a flag, f=1 for two fibre working ODN, N_g is the spare fibre for future PONs and M_{DP} is the number of D-side splitter ports from each DP splitter.

The three categories of fibres in this expression do not all need to be carried over the E-side cable:

The bespoke fibres (b) do not all need to be spliced through as only a small proportion (W_b) of all the customers will ever take bespoke fibre services, W_b is calculated by geotype and placed into the exchange data base, the number of E-side bespoke fibres required per DP chain is therefore $C_H bW_b$. and per cabinet = $\mathbf{N_{DPCH}C_H bW_b}$.

The spare fibres are reserved for future PONs and would feed additional DP splitters, they in turn would be connected to cabinet splitters and therefore the number of E-side fibres required will be the D-side growth fibres divided by the cabinet splitter size (S_C). The number of E-side fibres from the cabinet to service the DP spare fibres = $\mathbf{(1+f)N_{DPCH}C_H N_g/S_C}$.

The D-side splitter ports from the DP splitters (M_{DP}) in the DP chains needs to be split into two parts, one of the splitter ports will be spliced to the cabinet splitter, the M_{DP}-1 additional splitter ports (if M_{DP}>1) can either be terminated in the cabinet or connected to E-side fibre for transport to the LE site. To accommodate this a flag f_2 is incorporated (f_2=1 if fibres are provided in the E-side cable and f_2=0 if terminated in the cabinet), Therefore the number of fibres to support the DP splitter ports $=(1+f)N_{DPCH}C_H/S_C+(1+f)N_{DPCH}C_H(M_{DP}-1)f_2$.

here are also the Mc ports from the cabinet splitters, there are $(1+f)N_{DPCH}C_H/S_C$ cabinet splitters each with Mc E-side ports, at least one of these ports/splitter must be spliced through to the LE site. The other Mc-1 ports can be optionally connected through as in the case for the additional DP splitter ports. The same flag f_2 can therefore be used (these flags could be kept independent if required) so the number of cabinet E-side Cabinet splitter port fibres $C_{NS} = (1+f)N_{DPCH}C_H/S_C + (1+f)(Mc-1)f_2N_{DPCH}C_H/S_C$ = $\mathbf{(1+f)N_{DPCH}C_H(1+(Mc-1)f_2)/S_C}$.

The total E-side fibre from the cabinet chain = $N_{DPCH}C_HbW_b$ $+(1+f)N_{DPCH}C_HN_g/S_C$ + $(1+f)N_{DPCH}C_H/S_C+(1+f)N_{DPCH}C_H(M_{DP}-1)f_2+(1+f)N_{DPCH}C_H(M_C-1)f_2/S_C$

$=N_{DPCH}C_H[bW_b+(1+f)(N_g/S_C+1/S_C+(1+(M_{DP}-1)f_2+(M_C-1)f_2/S_C)]$

$$= N_{DPCH}C_H(bW_b + \frac{(1+f)}{S_C}[N_g + 1 + f_2(S_C(M_{DP}-1) + M_C - 1)])$$

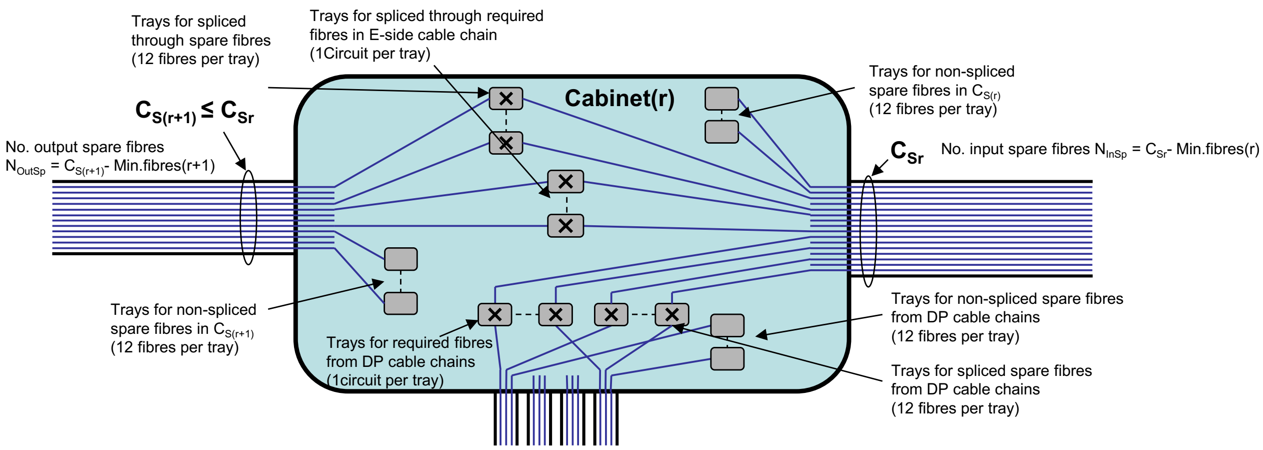

8.5 Appendix 5 Voronoi polygons for exchange area estimates

A simple method for Voronoi polygons given relatively equidistance points i.e. local exchanges have similar separations in a given region is described.

Consider the exchange at xi,yi and a nearest neighbour exchanges at x_j,y_j. J = 1, 2, 3, 4, ….

The point xp,yp is the circumcentre of the triangle (x_i,y_i), (x_j,y_j), (x_{j+1},y_{j+1}).

The circumcentre lies inside the triangle if it is acute and outside if obtuse,

Cartesian coordinates U_x, U_y of the circumcentre for general triangle ABC are:

$U_x=((A_x^2+A_y^2)(B_y-C_y)+(B_x^2+B_y^2)(C_y-A_y)+(C_x^2+C_y^2)(A_y-B_y))/D$

$U_y=((A_x^2+A_y^2)(C_x-B_x)+(B_x^2+B_y^2)(A_x-C_x)+(C_x^2+Cy^2)(B_x-A_x))/D$

where $D = 2(A_x(B_y-C_y)+B_x(C_y-A_y)+C_x(A_y-B_y))$.

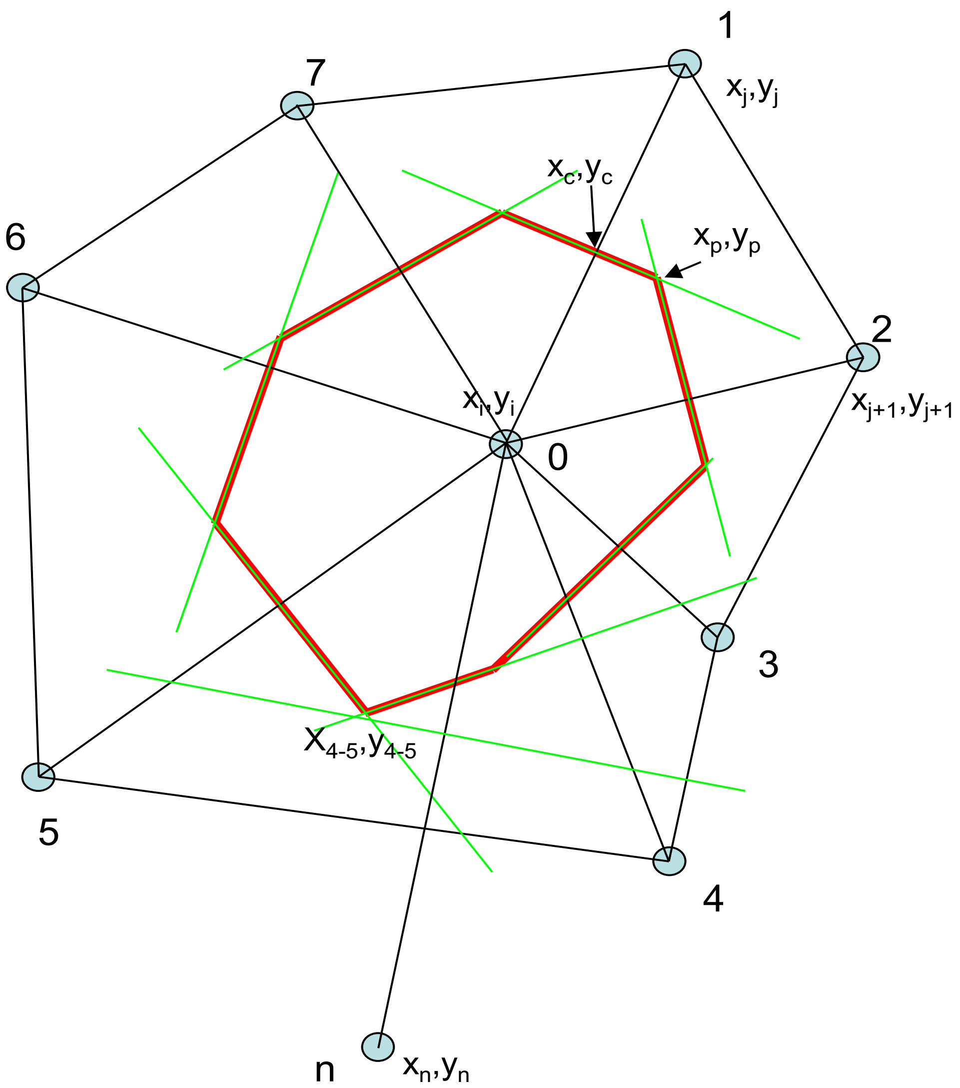

Figure A5-1 Local Exchange at point x_1, y_1 surrounded by set of neighbouring LEs

The sides of the Voronoi polygon is formed by the set of sides that are the bisectors of the radius vectors from x_i,y_i to the nearest neighbour set of surrounding LEs and points x_j,y_j. The question is which nearest neighbours to count into the Voronoi computation and which to leave out without having to check all exchanges. It can be seen from the figure that the point x_n,y_n has a bisector just outside the red Voronoi polygon if it had been closer it would have passed inside the vertex of the Voronoi polygon at $x_{4\text{-}5},y_{4\text{-}5}$ and formed another side.

So if the distance to a point outside the Voronoi polygon is greater than twice the projection of the adjacent vertices onto the radius vector to the point x_n,y_n then the point can be ignored. Also any further additional points that add sides will tend to reduce the Voronoi vertices radius vectors so the point can always be ignored.

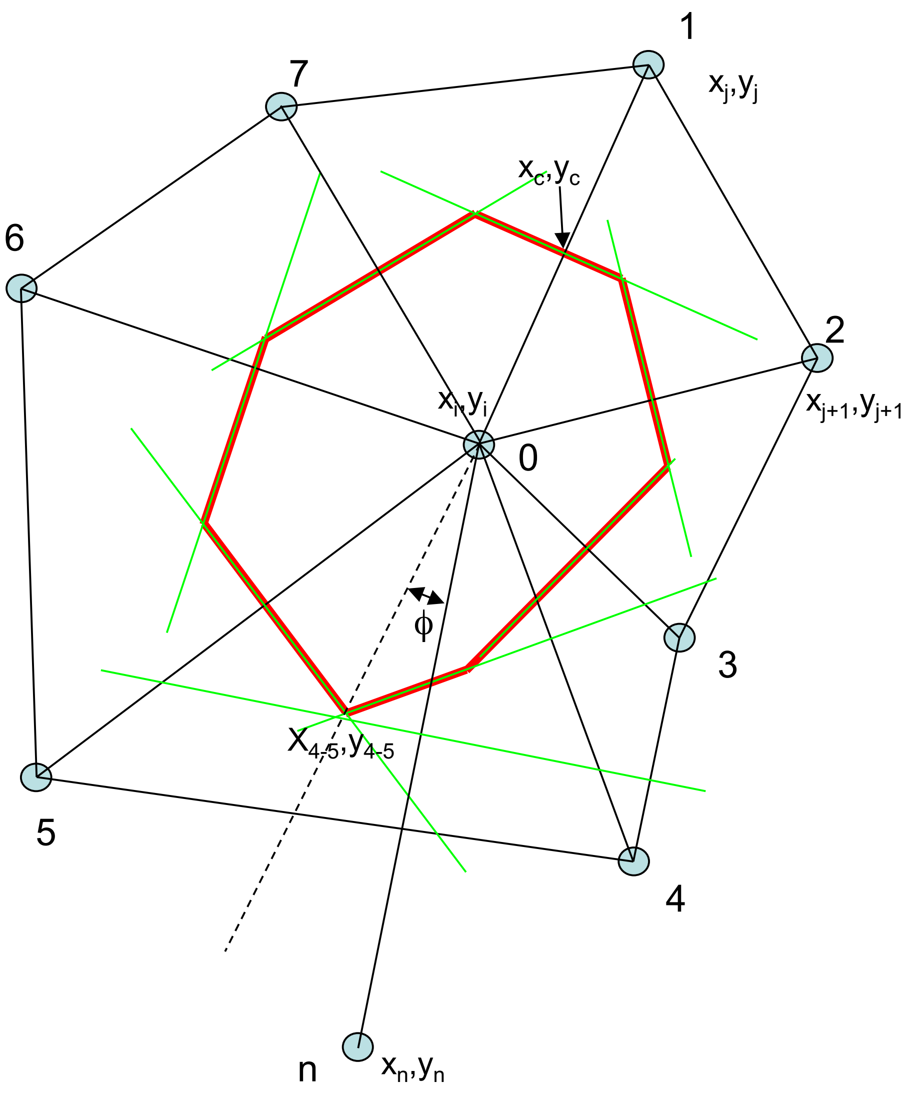

Figure A5-2 The LE at point Xn, Yn lies outside Voronoi polygon

A process for constructing Voronoi polygons for any middle exchange i.e. an exchange surrounded by other exchanges and not adjacent to a boundary line (e.g. coast line or estuary) can be:

1. Find N nearest neighbours within a search range of approximately 2 x mean exchange spacing from target exchange at exchange location x_i,y_i and calculate the radius vectors from the point x_i,y_i to these exchanges i.e. position angle and radius.
2. Find the M (M~6 will be a typical number) nearest and sort in position angle order (rotation of co-ordinate system is arbitrary so let 0 = north).
3. Calculate the circumcentre coordinates for the triangle (using the equation in the previous slide $(x_i,y_i),(x_j,y_j),(x_{j+1},y_{j+1})$ starting at j=1, repeat for all j back to j=1. This should produce a maximum of M vertices of an M sided Voronoi polygon around the point $x_i.y_i$. Store these coordinates in an array.
4. Sort all the remaining N-M nearest exchanges in position angle order and calculate the projection of the Voronoi vertices onto the adjacent radius vectors of the nearest of the N-M exchanges (e.g. onto exchange at x_n,y_n in the figure. This projection for the x_{4-5},y_{4-5} radius vector onto the x_n,y_n radius vector is given by: $r'=((x_{4-5}-x_i)^2+(y_{4-5}-y_i)^2)^{1/2}$ Cos(f);
5. Check the that the distance to the surrounding N-M more distant exchanges (e.g. x_n,y_n in the figure) is greater than twice the adjacent vertices projections.
6. If the above distances to the N-M exchanges are greater than twice the projections of the adjacent Voronoi polygon vertices then those exchanges can be discarded as they do not affect the Voronoi polygon.
7. If the above distances are less than twice the projections then those exchanges need to be include in the set of M closest exchanges that affect the Voronoi polygon and an additional side is added.
8. The next step is to calculate the enclosed area of the Voronoi polygon. This can be done by calculating the area of the triangles formed by the point (x_i,y_i) and two adjacent vertices of the Voronoi polygon. Note the height of these triangles is half the radius vector to the M nearest exchanges. The length of the sides of the Voronoi polygon form the bases of the triangles and can be calculated from the x,y coordinates of the Voronoi polygon vertices calculated previously via the circumcentre calculations. Length of the j^{th} side is therefore: $b_j = ((x_{vj}-x_{vj+1})^2+(y_{vj}-y_{vj+1})^2)^{1/2}$. Where x_{vj},y_{vj} are the coordinates of the j^{th} vertex of the Voronoi polygon.

9 Abbreviations

BAU	Business as usual
CAGR	Compound annual growth rate
CO	Central office
DDCND	Diameter and Degree Constrained Network Design problem
DP	Distribution point
EDFA	Erbium doped fibre amplifier
GPON	Gigabit PON
ICU	Idealist cost unit
ISN	Initial Size Network
LE	Local exchange
LR-PON	Long-Reach PON
JIT	Just in time
MC	Metro core
MSN	Minimum Size Network
MPLS	Multi protocol label switching
MPLS-TP	MPLS Transport profile
NTE	Network termination equipment
OCh	Optical channel
ODN	Optical distribution network
OH	Overhead
ONU	Optical Network Unit
ONT	Optical Network Terminal
OPM	Optical performance monitoring
OXC	Optical Cross Connect
PON	Passive optical network
PCP	Primary Cross Connect
ROADM	Reconfigurable add drop multiplexer
S-BVTs	Sliceable bandwidth variable transponders
UG	Under ground
VLSI	Very large scale integration

WDM	Wavelength division mutliplexing
WSS	Wavelength selective switch

Version	Date submitted	Comments
V1.0	05/02/2016	First version sent to the EU.