

DISCUS: An end-to-end Solution for Ubiquitous Broadband Optical Access

Marco Ruffini, Lena Wosinska, Mohand Achouche, Jiajia Chen, Nick Doran, Farsheed Farjady, Julio Montalvo, Peter Ossieur, Barry O'Sullivan, Nick Parsons, Thomas Pfeiffer, Xing-Zhi Qiu, Christian Raack, Harald Rohde, Marco Schiano, Paul Townsend, Roland Wessaly, Xin Yin, David B. Payne.

Abstract

Fibre-to-the-premises (FTTP) has promised to increase the capacity in telecommunications access networks for well over thirty years. While it is widely recognized that optical fibre based access networks will be a necessity in the short to medium term future, its large upfront cost and regulatory issues are pushing many operators to further postpone its deployment, while installing intermediate, unambitious solutions such as Fibre-to-the-cabinet (FTTC). Such high investment cost of both network access and core capacity upgrade often derives from poor planning strategies that do not consider the necessity to adequately modify the network architecture to fully exploit the cost benefit that a fibre-centric solution can bring.

DISCUS is a European Framework 7 Integrated Project that, building on optical-centric solutions such as Long-reach passive optical access and flat optical core, aims to deliver a cost-effective architecture for ubiquitous broadband services. DISCUS analyses, designs and demonstrates end-to-end architectures and technologies capable of saving cost and energy by reducing the number of electronic terminations in the network, and sharing the deployment costs among a larger number of users, compared to current fibre access systems.

This paper describes the network architecture and the supporting technologies behind DISCUS, giving an overview of the concepts and methodologies that will be used to deliver our end-to-end network solution.

1 Introduction

Efficient high-speed broadband communications networks are essential for thriving modern communities. In many countries however, persistent exploitation of old copper transmission technology creates a digital divide between users with good copper access, close to operators' electronic infrastructure, and users with poor or no service, in more remote areas. This issue has already seen the intervention of the European Commission that has set a digital agenda to push national operators to improve the speed and reach of their infrastructure by 2020 [1].

The future-proof technology that can offer ubiquitous high speed broadband is fibre-to-the-premises (FTTP), able to deliver capacity and services that are independent of the distance between users and service nodes. However the deployment of such technology around the world is progressing more slowly than expected, mostly due to uncertainties on its financial viability and the high costs that operators need to face. Indeed, deployment of FTTP requires high upfront investment, whose uncertain return in the short term can make it a high-risk investment. Therefore a significant number of operators are opting for an intermediate solution, i.e., Fibre-to-the-cabinet, which reduces deployment costs by partly reusing the existing copper pair based infrastructure. However this technology, besides having significant limitations in maximum peak rate, which is well below what cable operators can already deliver today, makes the digital divide issue even worse, as it accentuates the distance dependency of the broadband services.

In addition, while FTTP could solve the access bottleneck for the foreseeable future, if deployed with the current network architecture, it would only move the bottleneck from the access to the metro and core networks. This would happen because FTTP networks have enormous capacity and can grow user bandwidth to orders of magnitude larger than copper and wireless technologies. Since the traffic entering the metro and core networks will grow proportionally, metro and core capacity will need to scale accordingly. However, growing current electronic-centric architectures does not scale, as their cost and power consumption grows linearly (or worse) with the required capacity.

The network cost issue is only exacerbated by current practices of building proprietary and independent communication infrastructure for different network access services, instead of trying to

share the abundant capacity that one single fibre-access system can provide. Mobile network operators are a typical example, as they tend to build their own optical backhaul to interconnect base stations. As the demand for mobile access capacity grows steadily, operators are forced to deploy a higher and higher number of base stations, making the cost of the optical backhaul a relevant portion of the entire network cost.

One of the most effective ways to reduce costs in telecommunications networks has been, historically, sharing network capacity among multiple users. This has been a natural choice in metro and core networks, where traffic is progressively aggregated into a smaller number of higher capacity links. In the access however, because of the low capacity of copper transmission technology (i.e., twisted copper pair), networks have been deployed as point-to-point links connecting users directly to central offices. Recently, the deployment of fibre technology in the access has removed such capacity constraints, making sharing of access link capacity feasible and cost-effective. Indeed this idea has been exploited in the current Passive Optical Network (PON) deployments.

The DISCUS project **Error! Reference source not found.** extends and broadens this idea, taking advantage of the large capacity of optical fibre to propose important modifications to the current access architecture that allows sharing resources across as many customers as possible so that the cost per customer is minimised. In addition it exploits the typical low transmission loss of fibre to bypass (and thus eliminate) as much electronic equipment from the network as possible, a solution that also addresses the problem of growing power consumption.

The DISCUS architecture builds on the concept of Long-Reach Passive Optical Networks (LR-PON) [3] in the access, and a flat optical core arranged into islands of transparency. The proposed architecture is shown in Figure 1. LR-PON is an evolution of well-known PON technologies (e.g., EPON and GPON), developed to minimise equipment and infrastructure by multiplexing traffic from customers onto a single fibre as close to the customer premises as possible. By introducing a longer fibre reach from the user to the network node (which we will refer to as DISCUS node), and a larger split in the distribution network, a much larger number of customers can be served by the same infrastructure. The longer optical reach (over 100 km has been experimentally demonstrated [4]), allows concentrating fibre terminations from many PONs in a smaller number of “DISCUS” nodes, each covering large areas. Thus, a substantial number of nodes deployed in current telecoms networks, which carry out electronic packet switching and processing, can be consolidated as long reach access systems can by-pass local exchanges and terminate directly on a core node.

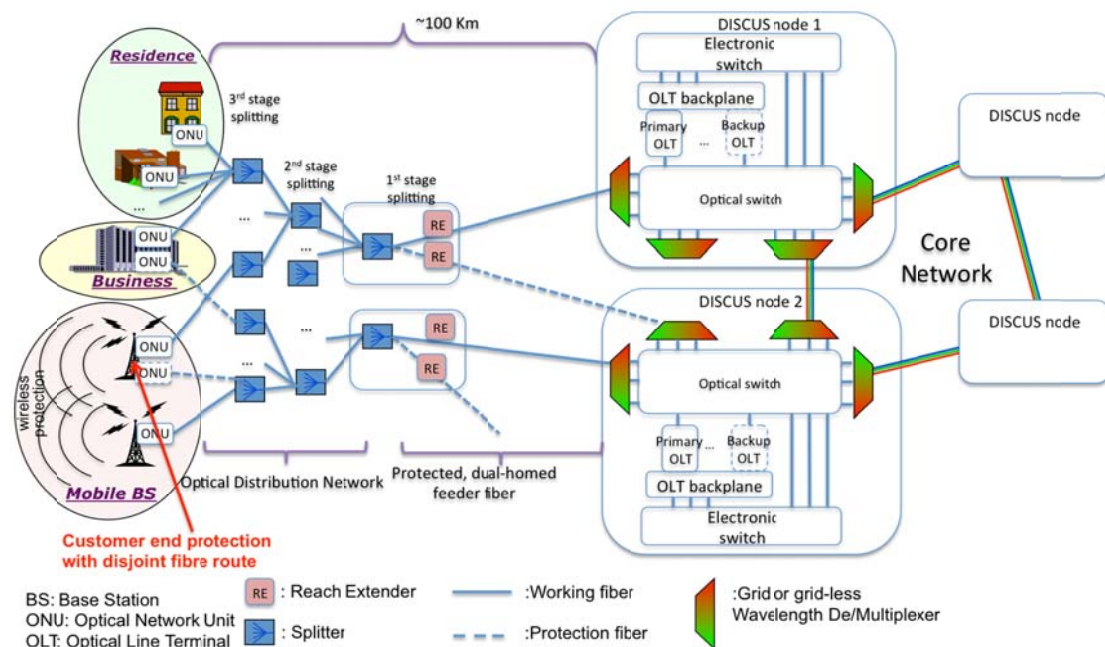


Figure 1 DISCUS initial architecture, based on Long-Reach Passive Optical Network access and flat optical core.

The DISCUS node is a core edge node providing the only electronic packet processing interface between PONs and core transmission network. It is placed in a similar architectural position in the network as what are often called metro-core (MC) nodes in today’s architectures. Each node has an access side, facing the LR-PON and a core side, facing the wavelength-switched optical core. The core is a flat optical network interconnecting DISCUS nodes through a full mesh of wavelength channels.

These wavelengths traverse intermediate core nodes without entering the electronic sub-layers. All nodes belonging to the same flat core form a transparent island. Interconnection among transparent islands can be achieved through signal regeneration, or traffic grooming and processing where required.

Although DISCUS is a project funded in 2012, the concepts of LR-PON and node consolidations are not new. Starting in the early 90s at BT laboratories, the concept has been developing over time in projects such as ACTS-PLANET, PIEMAN, ACCORDANCE and OASE. So why after so many years and different projects we have not seen yet any real implementation of this concept? The main reason is that large-scale deployment of any FTTH technology has not occurred in most countries of the world; the only countries with significant deployment being Japan and Korea. Europe has lagged behind, with nation-wide deployment plans relying instead on ADSL or FTTC technology for broadband service delivery. When FTTP is finally rolled out on a large scale in Europe the solution must be as economically viable for the widest range of different geo-types as possible and we believe the DISCUS LR-PON with flat optical core is the most economic, flexible and evolvable solution available.

The PLANET project established the basic principles of a LR-PON access solution, albeit at lower speeds than the DISCUS proposal. The PIEMAN project pursued technology solutions for the 10Gb/s symmetrical bandwidth solution (which is also adopted for the base system in DISCUS architecture) and established that the basic technology could be realised. However it did not investigate technologies such as burst mode EDC and FEC that are now being developed within DISCUS. The other difference in the DISCUS architecture is that whereas PIEMAN implemented a static WDM system approach, DISCUS has flexible wavelength assignment as inherent design feature for fully flexible bandwidth assignment.

ACCORDANCE and OASE have extended the technologies for LR-PON to coherent fine wavelength granularity and also addressed resilience issues for LR-PON as an access technology. DISCUS extends this work by considering LR-PON as part of a complete end-to-end solution with investment resource redistribution as part of the design criteria. Cost redistribution is going to be key part of a viable FTTP solution for Europe and is achieved by maximising cost reductions in the core and metro networks. This allows for a redistribution of investment towards the access edge, while still exploiting the enhanced sharing capability of LR-PON to minimise access network infrastructure investment, which we believe will also span the digital divide and enable sparse rural communities to receive FTTP services.

Indeed a primary objective of the EU Project DISCUS is to consolidate these approaches into a complete end-to-end network design to enable provisioning of ubiquitous high speed broadband for all customers and all services envisaged in the future, thus highly reducing the digital divide, while remaining economically viable and avoiding the power consumption issues that will arise by simply scaling today's networks. DISCUS aims to produce the evidence of the economical advantage of the LR-PON and flat optical core architecture, by providing end-to-end cost and business models and power consumption models. It also develops the enabling technology filling the gaps of the previous projects, e.g. low-cost tuneable transceivers (with tuneable optical filter), burst-mode EDC and FEC, and higher rate multi-point and point-to-point wavelength services over PON infrastructure, which will make a cost-effective transition to this new fibre-lean architecture possible.

One of the constraints of the DISCUS architecture is that it needs to evolve from today's network solutions, and be capable of adopting new technologies as they become available. While the starting point is a generic reference architecture based on LR-PON and a flat optical core, its design will progress and evolve according to the developments gained throughout the duration of the project, such as: technology developments and results from cost and business models, power consumption models, service and traffic models, regulatory and policy studies.

In the remainder of the paper we describe the DISCUS preliminary architecture, going into the details of the access and core networks as well as the proposed DISCUS node architecture. We then discuss our end-to-end architecture modelling and optimization framework. Finally, before concluding the article, we highlight some potential regulatory issues.

2 Long-reach access

The rationale behind a Long-Reach access for the DISCUS architecture is that it maximises infrastructure sharing, thus potentially minimises infrastructure cost per customer. It technically enables bypassing of the majority of local exchanges/central offices, which in the long term can be closed down, massively reducing both cost and power consumption. The long reach access network


also bypasses the metro transmission network, making that hierarchical layer redundant, and further reducing cost and power consumption.

2.1 Long-Reach PON (LR-PON)

The access network proposed in the DISCUS architecture is a LR-PON, which uses power splitters to share one OLT and a feeder fibre among a large number of users (up to 1,024). The main rationale is that passive splitters, unlike for example wavelength demultiplexers adopted in other PON solutions, avoid the use of wavelength selective devices in the network, which reduce flexibility in provisioning of future services.

Considering Figure 1, the traffic from Optical Network Units (ONUs) at the customer premises is terminated at the OLT in the DISCUS nodes, without any intermediate electronic processing. Thanks to the use of optical amplifiers at the 1st splitting stage, the distance between ONU and OLT can be over 100km and a LR-PON could support up to 1024 customers, providing a very high degree of sharing. LR-PON has already shown ability to offer relatively low cost and low energy consumption with capability for enormous capacity expansion. Recent studies have shown that the proposed DISCUS architecture would reduce the number of nodes carrying out electronic packet processing by almost two orders of magnitude [5]. In a test scenario carried out for the UK network, the number of nodes with electronic processing equipment can technically be reduced from about 5600, down to around 75 (Figure 2) [6]. The architecture that we propose for DISCUS also shows major advantages when compared to typical GPON installations, which, for the UK, would require over 2000 electronic nodes.

Much of the earlier work on LR-PON has concentrated on designs based on the ‘lollipop’ model [4], where long fibre connections are in the feeder network, while the distribution network is relatively short. While this approach might fit dense urban areas, DISCUS is considering designs with greater layout flexibility that may be required for more sparse and rural areas. An additional feature addressed by DISCUS is the implementation of the “principle of equality” concept, i.e. that all optical terminations anywhere in the network could have equivalent connection capability, so that customers might be able to request any network service, independently of their geographical location. This is achieved by allowing flexible use of the wavelength domain, with wavelength channels capable of carrying a wide range of services and capacities. So for example the same PON access infrastructure could deliver in parallel 10Gb/s peak broadband services, as well as logical point-to-point 40Gb/s or



the national UK network, from 5,600 local (right-hand figure).

LR-PON can serve big areas with very large for a flexible level of resilience able to satisfy the backhaul systems, business customers and critical services requiring non-interrupted network access [7]. In this context, there are two main challenges associated with providing protection in LR-PON, namely deployment cost and flexibility of upgrading the infrastructure. There is an obvious trade-off between the provided level of resiliency and the deployment cost. From the network design point of view, protection techniques such as OLT dual homing with geographically disjoint feeder fibre routes are foreseen as a part of the standard LR-PON

deployment in order to avoid that all customers in the area will be disconnected as a result of the feeder fibre cut. However feeder backup does not protect the access network below the 1st stage split. For the majority of subscribers this level of protection and, consequently, connection availability can be sufficient. It may not be adequate however for customers such as large businesses and organisations, where uninterrupted access to the network is critical for their operations [8]. In order to satisfy the customers requiring higher levels of resilience, full protection including geographically diverse protection routes installed from customer premises can also be achieved on an individual basis, as shown for the mobile base station and business user in Figure 1.

2.2 Technology evolution

In order to support the proposed network architecture and in particular the LR-PON access, new physical layer technologies will be required both in the OLT and ONU, for example to offer dynamic wavelength allocation capability in a cost-effective manner, while improving energy efficiency. In this regard the technology evolution focuses on three main areas: tuneable transmitters and receivers, advanced burst-mode receiver design, and energy-efficient 40Gb/s downstream links.

Flexible dynamic wavelength allocation requires both tuneable sources and receivers. Thus the most critical components in the wavelength agile ONU are low cost tuneable laser and filter. DISCUS has selected a tuneable laser design based on a multi-section Fabry-Perot (FP) with integrated slots [9]. The FP slotted laser design only requires a single epitaxial growth stage and uses feature sizes that can be reached with standard deep-UV optical lithography. The simple fabrication process results in high yield and is potentially low-cost. Regarding the tuneable filters, the main issue with commercial devices is that they are fabricated on bulk optics (e.g., Fabry-Perot etalon, thin film etc.), and thus result in large footprint that is undesirable in the ONU and OLT. Therefore, DISCUS is proposing low-cost, polarization diversity, tuneable optical filters with low intrinsic insertion loss ($< 3\text{dB}$) and compact size ($< 500 \times 500 \mu\text{m}^2$) based on Silicon Photonics technology.

The upstream direction of a LR-PON is TDMA-based. One of DISCUS' targets is to support access topologies that move away from the previously introduced "lollipop model". This gives the flexibility to place power splitters anywhere in the feeder fibre, generating topologies that can be more cost-effective in aggregating users into PONs, especially in sparsely populated areas. As a consequence, the PON will need to handle a large differential reach, i.e. the difference between the maximum and minimum distance between the OLT and ONUs, beyond the conventional 10 to 20 km. This requires the ability to deal with input at the OLT of rapid (i.e., separated by a few tens of nanoseconds) sequence of bursts whose amplitude could vary by over 25dB, whose phase at the receiver might be unknown, and whose chromatic dispersion and nonlinearities may vary significantly. DISCUS candidate solution is a 'burst-mode', receiver-side electronic dispersion compensation (EDC), whose tap settings can be quickly adjusted from burst to burst under control of the LR-PON protocol. Initial results demonstrate that burst-to-burst adaptation of the tap settings can be achieved using a short sequence of training symbols (only 500 symbols), thus potentially requiring a small overhead [10].

In order to support increasing capacity in the down-stream direction, a 40Gb/s downstream upgrade from a 10Gb/s LR-PON is being investigated. The state-of-the art in this field is a hybrid 4 stacked PON (4x10G downstream and 4x10G upstream) [11]. DISCUS is evaluating a single carrier 40Gb/s downstream solution based on electrical duo-binary modulation of a high-power, external cavity laser, monolithically integrated with an electro-absorption modulator and booster semiconductor optical amplifier. A combination of optical and electrical dispersion compensation techniques are considered to achieve the required long reach. In order to overcome the problem of high power consumption, DISCUS adopts a new bit-interleaving multiplexing PON concept [12], capable of reducing the energy consumption of downstream electronics by almost an order of magnitude.

2.3 PON backhauling of mobile networks

Although DISCUS is focused on a new architecture for the future fixed network, it recognises that mobile and wireless will be major, potentially dominant, access technologies. DISCUS is thus investigating optical-wireless interworking and convergence as we move to ultra high speed FTTP access networks. The main focus of optical-wireless integration is making the DISCUS access architecture suitable to support services such as backhauling for mobile networks. However, since the LR-PON architecture replaces the access and metro networks, it can potentially generate longer access network round-trip delays compared to today's networks (up to 2ms propagation time for the longer paths plus protocol and data processing delays). These delays may be too long for a number of services requiring frequent and timely interactions between base stations. One example is the scheduling negotiation among base stations in ICIC (Inter-Cell Interference Coordination); another example is the

Coordinated Multi-Point (CoMP) transmission, where the base stations regulate their local spectrum usage and transmitter power levels to minimise mutual interference. Options to reduce the delays incurred by long reach of the fibre access network are being investigated, including passing messages across lower physical levels in the LR-PON network, by using additional intra and inter LR-PON physical paths. In addition, delay can be further reduced by integrating bandwidth scheduling functions of ONUs and base stations.

3. Flat optical core architecture

DISCUS recognises that a nationwide deployment of ultra-high capacity access network has the potential to increase the ingress traffic load to the core network by several orders of magnitude. Current evolutionary approaches to network development, based on adding IP routing/switching capacity into the network, are not scalable. Indeed, without substantial advances in technology, the cost and energy consumption will grow linearly (or worse) with the traffic load increase. However, over the past 10 years the relationship between revenues and service capacity growth has been effectively broken with the move away from usage charging to flat rate models. That is revenues will not be able to cover the increasing network cost. In addition the typical reduction of cost over time experienced by electronics over the past 40 years or so will not be sufficient to reduce telecommunications equipment prices fast enough for economic viability [13]. Energy consumption is also becoming a growing concern since, according to the estimation in [14][15], power consumption related to the ICT sector already causes about 2% of global CO₂ emissions (which is equivalent to the aviation industry) and ICT equipment in the UK accounts for about 10% of the total electricity consumption in the country.

As traffic continues to increase, driving up the amount of switching/routing equipment in the network nodes, eventually capital expenditure, energy consumption and operational expenditure will limit continued traffic growth. The only way forward out of this dilemma is a complete bottom-up redesign of network architectures aimed specifically at addressing these problems. The DISCUS approach to core network design is based on two pillars: (i) a flat optical core partitioned into “optical islands”, which minimises IP routers and switches, (ii) and adopting flexible grid optical transmission and switching technology.

A flat, wavelength switched, optical core interconnects all DISCUS nodes through a full mesh of wavelength channels within an optical island. One of the first problems that DISCUS is targeting is the issue of channel exploitation, which is typical for flat optical networks. In order to understand the scale of this problem we are developing a service and application-based traffic model which will help clarify how current and future sources of traffic will shape traffic patterns among the network nodes. We believe this is a better approach compared to current practices of scaling existing traffic matrices based on extrapolation of past traffic growth, which can be very unreliable. Our traffic and service analysis result will suggest whether a full-mesh is justifiable within an optical island or else different hierarchical architectures are required to operate electronic traffic grooming.

The maximum size of an island is generally limited by its number of nodes (as the full-mesh of wavelength channels grows quadratically with the number of nodes), rather than optical reach. Different optical islands need thus to be interconnected, which can be achieved by electronic routing, wavelength regeneration and conversion, or simply by transparent switching if the quality of the optical signal is sufficient. Two architectures showing intra and inter-island interconnections, one based on multi-gateway and the other on multi-hub hierarchical architecture, are shown in Figure 3.




Figure 3: Multi-gateway (top) and multi-hub (bottom) hierarchical architectures for inter-island interconnection. The metro-core (MC), referred to as DISCUS nodes, are transparently interconnected within an optical island.

In very high capacity networks, such as the proposed core architecture formed of optical islands, one of the most important challenges to be addressed is network survivability. To avoid service interruption and loss of huge amount of data (and revenues) the core network design and control must include very efficient recovery mechanisms, from cable dig-ups and equipment malfunctioning as well as natural disasters causing large network segment disruptions (such as earthquakes, tsunamis, floods and power outages). Our approach to resiliency modelling considers both interconnections between optical islands and resiliency within optical islands. Service differentiation with respect to reliability requirements is explored in order to minimise the amount of backup resources that are necessary to meet the SLA for each connection request. Furthermore, trade-offs such as survivability versus cost/energy efficiency are considered by analysing which level of resiliency should be offered, and also how to maximise the use of restoration strategies and how to replace 1+1 with 1:N protection mechanisms, while maintaining very high levels of network availability.

As mentioned above, the second pillar of DISCUS core network design is the flexible and efficient use of transport resources, which is achieved through widespread use of flexible grid optical networking technology and the related control plane. As transmission rates increase above 100 Gb/s towards 400 Gb/s, the 50 GHz fixed ITU grid becomes insufficient for spectral widths at standard modulation formats. Either coarser grids with wider bandwidths or preferably flexible grids for greater spectral efficiency must be used. In such a scenario we incorporate modulation formats and spectrum allocation in routing and wavelength assignment (RWA) algorithms.

4. DISCUS node architecture

As mentioned in the introduction, DISCUS nodes are the only point, within an optical island, with electronic packet processing functionality. Flexibility of the architecture is one of the main design goals, in order to assure that different layers and functions can evolve and if necessary displace other layers. The node architecture, shown in Figure 4, consists of two switching layers: an optical and an electronic layer. The electronic switching capabilities can be implemented as separate Ethernet and IP layers or could converge into an OpenFlow/SDN (software defined network) enabled electronic switch, as shown in Figure 4. The OLTs backplane is terminated at the electronic switch on the access side of the node.

The large port-count optical switch allows maximum flexibility of operations, as any incoming fibre or wavelength channel can be terminated, after any required demultiplexing, at any OLT or electronic router/switch port card. Wavelength channels can also be amplified or regenerated to create a link to another access ONU connected to the DISCUS node or sent over the optical core network to another node. The design of the optical switching layer is therefore a major activity in the DISCUS project and it is of paramount importance that the optical switch is highly scalable, while offering a very low optical power loss (ideally < 3dB). In this regard multi-stage non blocking fibre switch structures that can economically scale to over 10000 ports are investigated.

Finally, one of the main targets of the DISCUS node development is the implementation of an OpenFlow-based control protocol. The node controller is responsible for supervising all dynamic wavelength and sub-wavelength provisioning operations within the PONs terminating at that node. Examples are: operating dynamic bandwidth and wavelength allocation to improve efficiency of resource utilization, and operating dynamic wavelength or sub-wavelength services on requests to create flexible virtual networks among ONUs.

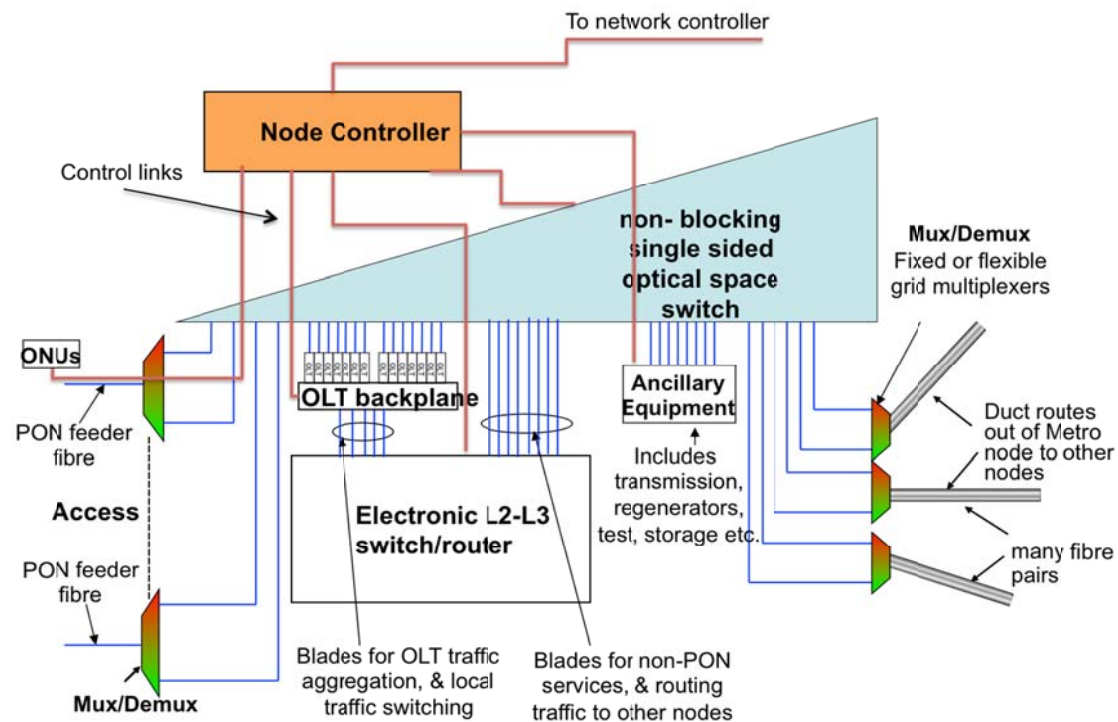


Figure 4 DISCUS node functional scheme with node controller.

5. End-to-end architecture modelling and optimization

From a network planning and deployment perspective, it is important to demonstrate the practical feasibility of the DISCUS architecture with respect to objectives such as investment cost and energy-efficiency. The end-to-end modelling and optimization work aims at finding cost- and energy-efficient solutions, calculated over a variety of scenarios and conditions that might represent the situation of different operators. Major goals are to find the optimal number of nodes where traffic is electronically processed; the optimal size of an optical island and their interconnection; and the required separation between working and protection routes to ensure adequate network availability and disaster recovery capability.

These problems have a high-degree of interdependency and models are coupled so that an end-to-end optimisation can be obtained. In addition, regional and national differences make it difficult to find generalised solutions, because for each network operator they are dependent on some specific circumstances and scenarios (such as: fibre and duct infrastructure; central office, metro and backbone infrastructure; geographic spread of the service area including the number and distribution of end-customers; CAPEX and OPEX trade-offs; traffic and service development; resilience requirements and local policy and regulatory framework etc). However, a common framework for the end-to-end solution, which can be moulded to the specific requirements and characteristics of different regions and countries, has been devised on the DISCUS architecture and reference scenarios.

The DISCUS reference scenarios are composed of:

- Nation-wide fibre network topologies for different countries including today's central office locations, based on data provided by the network operators within DISCUS, and combined with publically available information (e.g., open-street map data);
- Detailed street cabinet, and buildings locations for selected regions;
- Service and traffic models able to capture potential evolution of traffic requirements over time;
- Component and system models for access and core network equipment/systems, including Optical (LR-PON, WDM, Coherent, switching etc.), IP, Ethernet, and other ancillary functions;
- CAPEX and OPEX models, describing the evolution of capital and operational expenditures over time, including price decline of equipment, reinvestment and depreciation time periods as well as the evolution of energy prices.

6 Regulatory approach and standardization

One of the main approaches DISCUS advocates for reducing cost per customer is to dynamically share access infrastructure by using large split LR-PON technology, merging today's access and metro networks. Dynamic, on demand resource sharing has implications for regulation and policy in telecommunications systems.

Current regulatory policy in Europe lags significantly behind other countries [16], as it builds on the past success of local loop unbundling (LLU) of the copper pair, which focused on competition at the physical layer. However such a physical layer regulatory regime has a number of potential difficulties for FTTP deployment:

- As customers migrate from the copper network to FTTP the operational cost of the copper pair network hardly changes, thus cost per line increases as customers leave the copper platform.
- Physical layer competition for FTTP is very inefficient and uneconomic, as resources are not shared.
- LLU locks customers into contracts for typical time periods of 1 year or more. Within this period there is no competition, effectively producing temporal monopolies.

Within DISCUS we are promoting debates **Error! Reference source not found.** about tailoring the regulatory environment to a fully dynamic and flexible bandwidth allocation and service provisioning regime. We aim to maximise the economic benefits of dynamic allocation of network resources while also enabling full competition at the service layers. For example, a customer could instantly and on demand change the provider for a given service (e.g., Internet connection, TV-on demand, etc.) thus obtaining services from multiple providers simultaneously.

Standardization efforts are mainly directed to FSAN and ITU-T. DISCUS follows existing standards while proposing, where possible, changes to address functionality required to support the DISCUS architecture. For example we propose to extend the XG-PON protocol to support longer reach and dynamic wavelength assignment.

7. CONCLUSIONS AND FUTURE WORK

In this paper we have outlined the rationale for the architectural changes that are required if future networks are to remain economically viable as service requirements and network bandwidth grow. This article can be seen as a positioning paper where we present our proposed network and node architectures, which aim to be economically viable, sustainable and evolvable. We strongly believe that these architectures together can offer the features needed to support the emerging services and growing traffic demand in an economic manner. The network architecture is based on LR-PON to bypass and consolidate the majority of central offices, while also eliminating the separate metro transmission network. To minimise cost in the core network a flat wavelength circuit switched optical island structure is proposed, where wavelengths traverse the network from ingress to egress node, bypassing transit electronic routers, thus keeping traffic in the optical domain.

We highlighted a number of interesting research topics and challenges to be addressed in order to support our architectures, also providing some initial solutions.

This paper is initiating a series of articles and will be followed by future submissions reporting the results of the studies highlighted here. In our future work the DISCUS architecture will be validated

through a framework model investigating cost and energy consumption for a number of realistic scenarios. Such feasibility studies will be complemented with regulation and standardization activities, in order to assess the market viability of the proposed architecture, and with development of new technologies enabling high-speed dynamic multi-wavelength services.

ACKNOWLEDGEMENT

The research leading to this article has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. 318137 (Collaborative project “DISCUS”), and Science Foundation Ireland under grants no. 12/IA/1270 and 10/CE/I1853.

REFERENCES

- [1] European Commission, A digital agenda for Europe, 2010.
- [2] M. Ruffini et al., DISCUS: End-to-end network design for ubiquitous high-speed broadband services (**Invited**). ICTON 2013, paper We.B3.1.
- [3] D. B. Payne and R. P. Davey, The future of fibre access systems?, BT Tech Journal, 20(4), 104-114, 2002.
- [4] P. Ossieur et al. Demonstration of a 32 512 Split, 100 km Reach, 2 32 10 Gb/s Hybrid DWDM-TDMA PON Using Tunable External Cavity Lasers in the ONUs. IEEE JLT, Vol 29(24), pp.3705, 3718, Dec.15, 2011.
- [5] M. Ruffini, et al., Protection Strategies for Long-Reach PON. ECOC 2010, paper Tu.5.B.2.
- [6] M. Ruffini, et al. Deployment Strategies for Protected Long-Reach PON. IEEE/OSA JOCN Vol. 4(2), p118 – 129, Jan 2012.
- [7] P. Chancelou et al, Network Operator Requirements for the Next Generation of Optical Access Networks, IEEE Comm. Mag., Vol 6(2), pp. 8-14, 2012.
- [8] OASE Project, Deliverable D3.2: Description and Assessment of the Architecture Options.
- [9] S. O’Brien, et.al. Design, characterization, and applications of index-patterned Fabry-Perot Lasers, IEEE J. Sel. Topics Quantum Electron., Vol. 17(6), pp. 1621-1631, Dec. 2011.
- [10] S. Porto, et al., Requirements for adaptive electronic dispersion compensation in burst-mode systems, OFC 2013, Paper OTh3B.5.
- [11] P. Iannone, et al., Bi-directionally amplified extended reach 40 Gb/s CWDM-TDM PON with burst-mode upstream transmission. OFC 2011, Postdeadline paper PDPD6.
- [12] H. Chow, et al., Demonstration of Low-Power Bit-Interleaving TDM PON. ECOC 2012, paper Mo.2.B.1.
- [13] D.B.Payne, FTTP Deployment Options and Economic Challenges, ECOC 2009, tutorial paper 1.6.1.
- [14] Global Action Plan Report, An inefficient truth, 2007, <http://www.globalactionplan.org.uk>
- [15] SMART 2020 Report, Enabling the low carbon economy in the information age, <http://www.theclimategroup.org>, 2008.
- [16] S. Beardsley, et al. Creating a Fibre Future: The Regulatory Challenge. The Global Information Technology Report 2010-2011, World Economic Forum.
- [17] M. Ruffini et al., DISCUS white papers: “Business and ownership models for future broadband networks” and “Wavelength usage options in access networks”, available at <http://www.discus-fp7.eu/white-papers>.