

DISCUS: End-to-end network design for ubiquitous high speed broadband services

Marco Ruffini¹, Nick Doran², Mohand Achouche³, Nick Parsons⁴, Thomas Pfeiffer⁵, Xin Yin⁶, Harald Rohde⁷, Marco Schiano⁸, Peter Ossieur⁹, Barry O'Sullivan¹⁰, Roland Wessäly¹¹, Lena Wosinska¹², Julio Montalvo¹³ and David B. Payne¹

¹ CTVR, Trinity College Dublin, Ireland, Email: marco.ruffini@tcd.ie; ² Aston University, United Kingdom; ³ III-V Lab, France; ⁴ Polatis, United Kingdom; ⁵ Alcatel-Lucent, Germany; ⁶ IMEC, Belgium; ⁷ Nokia Siemens Networks, Germany; ⁸ Telecom Italia Italy; ⁹ Tyndall institute, Ireland; ¹⁰ University College Cork, Ireland; ¹¹ atesio, Germany; ¹² Royal Institute of Technology KTH, Sweden; ¹³ Telefónica I+D, Spain.

ABSTRACT

Fibre-to-the-premises (FTTP) has been long sought as the ultimate solution to satisfy the demand for broadband access in the foreseeable future, and offer distance-independent data rate within access network reach. However, currently deployed FTTP networks have in most cases only replaced the transmission medium, without improving the overall architecture, resulting in deployments that are only cost efficient in densely populated areas (effectively increasing the digital divide). In addition, the large potential increase in access capacity cannot be matched by a similar increase in core capacity at competitive cost, effectively moving the bottleneck from access to core.

DISCUS is a European 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. One of the key features of the project is the end-to-end approach, which promises to deliver a complete network design and a conclusive analysis of its economic viability.

Keywords: Long-Reach Passive Optical Networks, end-to-end, flat optical core, FTTP, ubiquitous broadband.

1. INTRODUCTION

Efficient high-speed broadband communications networks are essential for thriving communities in the modern world. Ideally all users and all communities should have equal access to high-speed broadband so that there is no division of service availability due to customer location. However many countries still exploit old copper transmission technology and are now in danger of creating a digital divide, due to the distance dependent quality of connection of copper transmission systems. Customers close to the operator's electronic infrastructure will indeed experience a much larger capacity than those located in more remote areas, connected via much longer copper lines.

Ultimately the only technology that can deliver ubiquitous high speed broadband access is Fibre-to-the-premises (FTTP). The main problem with FTTP however is financial viability. Firstly, there is the obvious challenge related to cost of network infrastructure and its deployment. Secondly, there is a financial problem arising from the sheer capability of FTTP networks to grow user bandwidth compared to copper and wireless technologies. If FTTP is used to provide the future high speed broadband services that customers require then the bandwidth entering the metro and core networks will have to grow in proportion and the network capacity will need to scale accordingly. Unfortunately, revenue will not scale proportionally (mainly due to the wide acceptance of flat-rate business models) and will not cover the cost of the bandwidth provision required [1]. Furthermore, a reduction in the price of electronics will not be sufficient to reduce telecommunications equipment prices fast enough to ensure economic sustainability [1].

The overall objective of the European Integrated Project DISCUS [2] is to design a network architecture that can deliver ubiquitous high-speed broadband access to all users independently of their geographical location. Furthermore, the topology targets economic viability as user bandwidths grow by orders of magnitude while improved energy efficiency will ensure significant power savings over current network configurations. In addition, DISCUS aims to deliver core network bandwidth capability to the access edge so that direct connection to the core network is available from any access terminal fitted with the appropriate technology.

2. PHILOSOPHY OF THE DISCUS ARCHITECTURE

The "philosophy" of the DISCUS architecture, is summarized in **Error! Reference source not found..** This is a diagram showing the logical interconnection between goals (in green), issues (in red) and required actions (in black). Starting from the top of the diagram: if we make the assumption that ubiquitous high speed broadband is a primary requirement in today's networks, current copper-based access technology fails to provide location independency, and bringing fibre to the premises seems the only solution able to guarantee an access speed that is independent of the distance between the user and the network termination. The resulting issue is one of increased cost and power consumption (shown at the bottom of the graph, where the logic flow splits into two). Following the flow in the upward direction, DISCUS proposes two solutions that can reduce both power

consumption and capital and operational costs. One is to decrease the number of nodes (also known as node consolidation), which in the diagram creates a new logical flow linking to the top of the left half circle in the graph. The other is to reduce the number of OEO conversions in the network, which leads to the adoption of a flat optical core, where the amount of electronic packet processing is significantly reduced. Following the flow on the left hand direction from the bottom of the graph, we see that cost can be further reduced by sharing the network resources among a larger number of customers. The issue of decreasing average bandwidth per user can be compensated by using higher rate e.g. 10-40Gb/s access speeds at start, and by increasing the number of wavelength channels in the future. In DISCUS the increase in shared resources is brought about by the Long-Reach PON architecture.

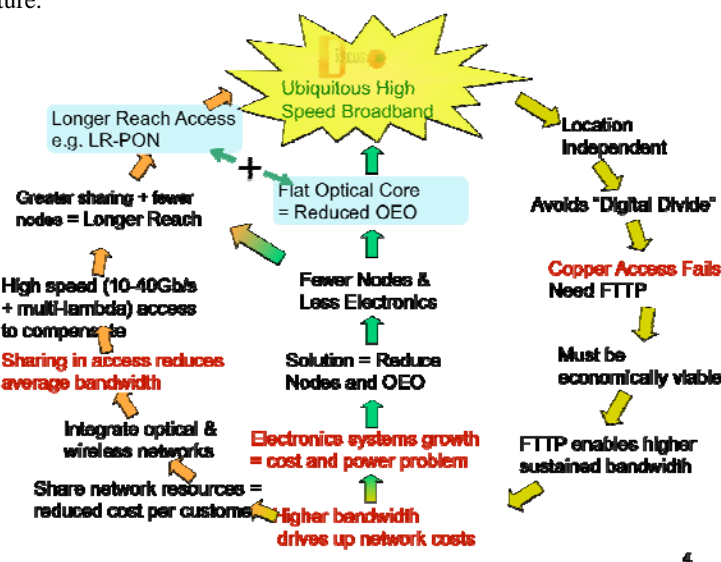


Figure 1 The DISCUS architecture, a logical necessity

The combination of LR-PON plus flat optical core is the basic philosophy of the DISCUS architecture that we believe can economically deliver ubiquitous high speed broadband to all users and will ensure all optical points of presence provide equivalent network and service capability irrespective of geographical location.

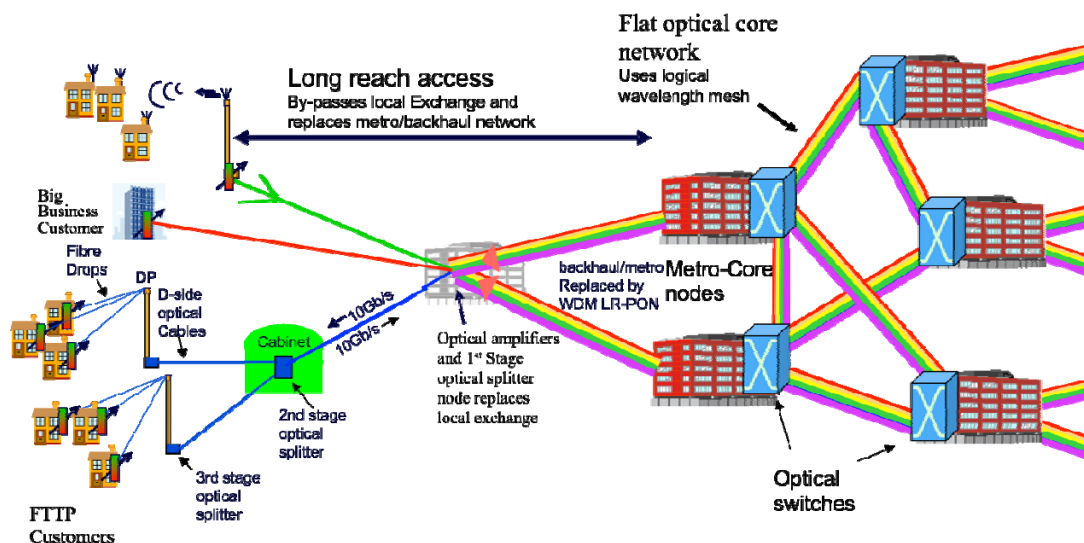


Figure 2 The end-to-end DISCUS architecture with LR-PON and flat optical core network

A simple schematic of the DISCUS architecture is presented in **Error! Reference source not found..** It shows a dual homed LR-PON bypassing existing local exchanges and terminating on the metro-core nodes. The metro-core nodes are interconnected with an optical circuit switched wavelength layer. The overall architecture consists of three major parts: the **LR-PON access network**, the **metro-core node (or DISCUS node)** and the **flat optical core-network**. DISCUS main aim of promoting seamless end-to-end design of the network architecture drives the design of these three parts, described in the following sections, into an integrated and coherent system.

3. LONG REACH ACCESS NETWORK

The strategy of increasing infrastructure sharing and reducing amount of network nodes inevitably leads to the Long Reach Passive Optical Network (LR-PON) architecture for the access and metro networks. Using a LR-PON enables sharing fibre infrastructure as close to the customer as possible using a simple passive optical splitting element. This passive splitter requires no power and is a highly reliable component. The LR-PON uses optical amplification to support greater total split (in the order of 500-1000), longer reach (over 100 km) and higher bit rates (10Gb/s were demonstrated in [3]) than today's PON solutions.

The higher split means that it is possible to have multiple split stages to further increase infrastructure sharing and minimise cost per customer. Suitable locations for splitting points are the distribution points (DP, referred to as 3rd stage optical splitter in **Error! Reference source not found.**) close to the customer premises, primary cross connect (PCP, referred to as 2nd stage optical splitter) or cabinet locations, typically less than 1km from customers, and the local exchange (LE, the 1st stage splitter in the figure) or central office site, where the optical amplifiers will also be located, as this node has electrical power available. The average distance of customers from the local exchange will be of the order of 2 to 3km but the tail of this length distribution can go out as far as 10km so the optical distribution network (ODN) from the old Local Exchange (LE) site is designed to support at least 10km, as shown in Figure 3(a). The figure also shows a basic feeder fibre protection mechanism based on the LE.

In sparse rural areas the ODN reach may need to be extended and alternative configurations are considered. A possible solution is shown in Figure 3(b), where a longer ODN side is traded for a shorter backhaul. The advantage of a longer ODN is that it allows covering a larger number of access points (which are normally sparsely distributed in rural areas), thus allowing sharing network resources among an adequate number of users. The major trade-off is that, since the losses in the ODN are now much larger and diverse than those in the backhaul, the number of users supported is smaller. Figure 4 shows a practical relation (analytically calculated) between reach of the ODN and maximum number of users per PON, assuming a conservative value of 512 users for the default 10 Km case.

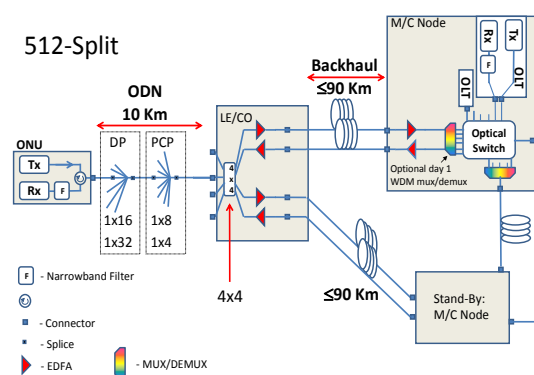


Figure 3(a) Basic LR-PON design




Figure 3(b) LR-PON solution for sparse rural areas

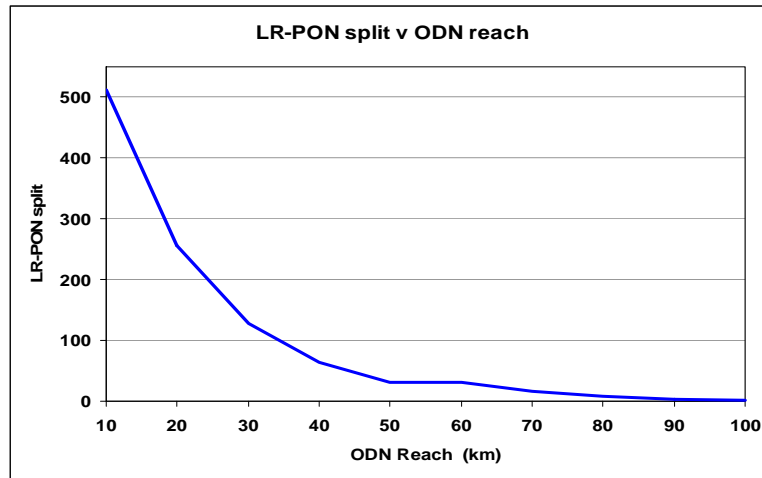


Figure 4 PON split reduction as ODN length increases

4. METRO/CORE or DISCUS NODE

The metro/core or DISCUS node, shown in Figure 5, interconnects the LR-PON at the access side with the flat optical core network. The main principle is to have a transparent optical layer in the form of an optical circuit switch where all other network functions, links and electronic switching/routing layers are connected. The optical layer allows any layer below the optical switch to grow or shrink independently as the network evolves and traffic grows or shrinks. It also does not distinguish between access ports and core connection ports and also enables direct connection to and between optical paths in access and core networks. One of the benefits of the optical switch layer separating the access and core physical layer from the electronic processing layers is that new network protocols and routing and switching systems can be tried in parallel with the conventional systems. Even experimental systems could be given limited access to the network. If these new systems prove successful they could gracefully displace old legacy systems. Such flexibility is a key design feature of the DISCUS architecture and is an important component of the evolution capability of the network architecture.

DISCUS will consider different possibilities for the electronic switching layer. Starting from a legacy architecture that physically differentiates Layer 2 and Layer 3 functionalities (Figure 5(a)), it will consider evolution towards an OpenFlow/Software Defined Network (SDN) based solution, merging L2 and L3 operations into the same hardware (Figure 5(b)).

Hardware integration of L2 and L3 data planes brings a number of advantages, including:

- Better sharing of DWDM interfaces for L2 and L3 services
- Avoiding grey interfaces to connect L2 and L3 hardware
- Reduction of number of optical switch ports required in the node
- Integration of L2 and L3 network control, implemented on a separate computation element.

From an architectural point of view, merging the L2 and L3 data plane can bring issues of scalability and resiliency. As the number of users served by a metro-core node grows in the order of hundred thousand or more, it is not scalable to consider a monolithic switch carrying out all electronic packet processing. Solutions using CLOS or tree architectures, like those in Figure 5(b), can be implemented to scale up the size of the electronic switch inside the DISCUS node.




Figure 5(a) DISCUS node with physical separation of L2 and L3 functionalities

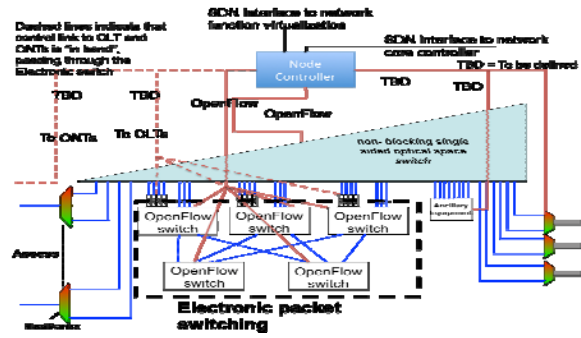


Figure 5(b) DISCUS node architecture with OpenFlow based control plane and integrated L2-L3

5. FLAT OPTICAL CORE NETWORK

In Section 2, we mentioned that current evolutionary approaches to network scaling, mainly based on increasing IP routers capacity in the network, would not be able to support the required growth of core capacity, i.e., 2 to 3 orders of magnitude. The only way forward out of this dilemma is a complete bottom-up redesign of network architectures aimed specifically at addressing these problems.

The idea behind a flat optical core is that a sufficiently small number of metro-core nodes (up to about 80-100) can be directly connected together via circuit switched optical channels without add-drop functions used at transit nodes. The metro-core nodes are interconnected by a mesh of cable routes which for economic reasons will not be a full mesh. Overlaid on this is a full mesh of optical (wavelength) channels. In Figure 6 the blue lines represent cable routes, and the coloured lines are the logical optical wavelength routes, shown only from metro node 1 to all the other nodes. These optical paths traverse other nodes but do not carry traffic for those intermediate nodes and simply pass through.

Following this approach packet processing is only operated to add or drop packets respectively in and out of the core. Thus, in principle any packet can be delivered in the network passing only two electronic hops, i.e., at the nodes where the sending and receiving users are connected.

The main issue of a flat core with full-mesh wavelength interconnection is that the number of channels grows with the square of the number of nodes. While this might be adequate for a number of countries in Europe (for example we have previously shown that Ireland can be covered with about 20 nodes [4], and the U.K. with about 80 [5]), other countries (especially outside Europe) will require a much larger number of nodes. In such cases a hierarchical core network is required. This leads to the introduction of “Optical Islands” (Figure 6), defined as that set of nodes that are fully interconnected by a set of optical channels (usually wavelength channels) without using a hierarchical structure. Different optical islands can be interconnected, depending on local requirements by electronic routing, wavelength regeneration and conversion, or simply by transparent switching if the quality of optical signal is sufficient.

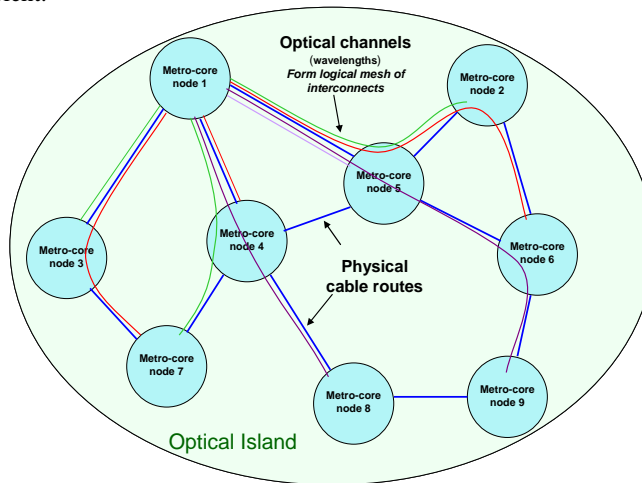


Figure 6 Full mesh interconnection of wavelength channels in an optical island

6. BUSINESS AND REGULATORY MODELS

DISCUS recognizes that it is of paramount importance that regulatory issues and business models are considered at design stage, as failing to do this could totally undermine the viability of the proposed architecture. Regulatory policies have indeed strongly influenced the evolution of fiber access networks in recent years, and the implementation of different policies have produced very different outcomes in e.g., Japan, US, and Europe [6]. DISCUS aims to fill in a gap on the regulations regarding the use of newly installed optical fibre, in particular shared access fibre that does not lend itself to simple physical layer unbundling policies, which have not yet been fully clarified for Europe.

Within DISCUS we are promoting debates 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.

7. CONCLUSIONS

In this paper we presented the main ideas behind the DISCUS end-to-end network design for ubiquitous high speed broadband services. We highlighted the challenges that we are going to address along the project duration. We strongly believe that the outcome of the project will bring a great contribution to the development of future optical networks and will speed up the rollout of the optical access networks.

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”).

REFERENCES

- [1] D. B. Payne and R. P. Davey, The future of fibre access systems?, BT Tech Journal, (20) 4, 104-114, 2002.
- [2] Project web page is available at <http://www.discus-fp7.eu/>
- [3] 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. JLT, vol.29, no.24, pp.3705,3718, Dec.15, 2011.
- [4] M. Ruffini, D.B. Payne, L. Doyle. Protection Strategies for Long-Reach PON. ECOC 2010.
- [5] M. Ruffini, et al. Deployment Strategies for Protected Long-Reach PON. JOCN 4, No. 2, Jan 2012.
- [6] S. Beardsley, et al. Creating a Fibre Future: The Regulatory Challenge. The Global Information Technology Report 2010-2011, World Economic Forum.