

telelektronikk

2.05

Optical Communications

Teletronikk

Volume 101 No. 2 – 2005
ISSN 0085-7130

Editor:

Per Hjalmar Lehne
(+47) 916 94 909
per-hjalmar.lehne@telenor.com

Editorial assistant:

Gunhild Luke
(+47) 415 14 125
gunhild.luke@telenor.com

Editorial office:

Telenor R&D
NO-1331 Fornebu
Norway
(+47) 810 77 000
teletronikk@telenor.com
www.teletronikk.com

Editorial board:

Berit Svendsen, Vice President Telenor Nordic
Ole P. Håkonsen, Professor NTNU
Oddvar Hesjedal, CTO Kyivstar GSM
Bjørn Løken, Director Telenor Nordic

Graphic design:

Design Consult AS (Odd Andersen), Oslo

Layout and illustrations:

Gunhild Luke and Åse Aardal,
Telenor R&D

Prepress and printing:

Rolf Ottesen Grafisk Produksjon, Oslo

Circulation:

3,600

Networks on networks

Connecting entities through networks – in technological, societal and personal terms – enables telecommunication. Networks occur on different levels, form parts of larger networks, and exist in numerous varieties. The artist Odd Andersen visualises the networks on networks by drawing interconnected lines with different widths. Curved connections disturb the order and show that networks are not regular but are adapted to the communication needs.

Per H. Lehne, Editor in Chief

Contents

Optical Communications

Introduction

- 1 Guest editorial – Optical Communications;
Evi Zouganeli
- 3 Optical networks: From point-to-point transmission to full networking capabilities;
Evi Zouganeli
- 20 Optical fiber transmission: From research to commodity;
Aasmund Sudbø

Access and Metro Networks

- 34 A new architecture for optical networks;
David Payne and Russell Davey
- 49 Fibre-optic techniques for broadband access networks;
Ton Koonen
- 66 Photonic integration;
Meint Smit
- 72 Advancements in metro optical network architectures and technology;
Loukas Paraschis, Errol Roberts and Ori Gerstel

Network control and Core Networks

- 81 Wavelength switches and the automated optical network – Status and outlook;
Claus Popp Larsen
- 87 Deutsche Telekom GSN+ – A comprehensive ASON/GMPLS demonstrator;
Hans-Martin Foisel, Andreas Gladisch, Monika Jäger, Sabine Szuppa and Armin Ehrhardt
- 96 Optical transmission systems in the Telenor network – In retrospect;
Steinar M. Svendsen
- 102 Building a fibre-optic network in Norway;
Harald Jansen
- 109 Cost and benefits of survivability in an optical transport network;
Guido Maier, Massimo Tornatore and Achille Pattavina

Optical Packet Switching

- 126 Why bother with optical packets? An evaluation of the viability of optical packet/burst switching;
Evi Zouganeli, Ragnar Ø. Andreassen, Boning Feng, Astrid Solem, Norvald Stol, Heidi Kjensberg, Aasmund Sudbø, Bjarne E. Helvik and Rolf B. Haugen
- 148 Burst, packet and hybrid switching in the optical core network;
Steinar Bjørnstad, Martin Nord, Torodd Olsen, Dag-Roar Hjelme and Norvald Stol

Fibre characterisation

- 162 Revealing the wide spectrum installed fibre cable loss in Telenor's network;
Svend Hopland

178 Terms and acronyms in Optical Communications

Guest editorial – Optical Communications

EVI ZOUGANELI



*Evi Zouganeli
is Senior
Research
Scientist at
Telenor R&D*

Optical fibre was first installed in telecom networks around two decades ago. At that time most of us, engineers included, had never heard of photonics – binoculars, cameras, spectacles, and contact lenses more or less covered all the optics we cared to know about. Since then the field of fibre optics and photonics has grown to become arguably one of the top ten technological achievements of the previous century and has had an enormous socio-economic impact. The invention of the laser and the commercial use of optical fibre in telecom have been the turning points and the key enablers for the transition from the industrial era of the previous centuries to the so-called information era and the knowledge economy we are witnessing today. The field of photonics covers a large scope, from medical applications, sensors, and CD players, to optical communications. This issue of *Teletronikk* is devoted to optical communications, in particular optical networks.

Optical transmission has led to an increase of the single cable transmission capacity by a factor of 3000 and has the potential to increase it many times over. Fibre has enabled transatlantic transmission without repeaters and many tens of parallel high capacity channels in one fibre cable. This massive capacity increase has meant a complete paradigm shift for the telecom network and the telecom sector as a whole. Cost per bit is dramatically reduced and no longer dictated by transmission distance. Bandwidth is in practice no longer a luxury resource. These have been central factors in order to enable the convergence between telecom and datacom – converging services, terminals, and networks – and the creation of the high capacity networks we are so accustomed to today. Telecommunications has been transformed from being the business of providing voice connections to becoming a very advanced sector that embraces a whole range of services, rather central to our every day activities and transactions, at work and leisure in and out of the home. The way we do business, retrieve and store information, learn, entertain ourselves, shop – communicate at all – has been and is continuously being dramatically changed by the new global digital space.

Imagine a scenario without optical fibre, with copper wires of limited capacity between network nodes. Network ownership and maintenance would still have required a fortune, cost per bit would still have been huge, and abolishing the telecom monopoly would have been obsolete. The Internet Protocol (IP) would have remained a university exercise, satellite tele-

vision would probably have prevailed over cable; we would still have mobile phones for voice and text, however, at considerably higher cost per bit, and the cost of advanced data-based services over fixed or mobile terminals would have been prohibitive for wide public use.

As revolutionizing as it has been, the real potential of optical technologies does not lie with high capacity transmission from A to B. The real potential of photonics lies with its ability to switch large quantities of data efficiently quickly and transparently (i.e. independent of format), and the flexibility it can provide in a network context by interplaying with other technologies or other network layers. Photonics possesses unique characteristics that complement these of IP – it can increase throughput dramatically, rationalize and automate network processes, and provide robustness and quality of service differentiation. The evident prevailing of IP gives a strong incentive for the introduction of more photonics inside network nodes, i.e. besides its use in transmission equipment. It is indeed symptomatic and not accidental that large IP routers have exclusively optical interfaces and that it costs less than one tenth as much to switch interfaces optically than it costs to pass them through an IP router. Also resource management and protection take place increasingly at optical channel rather than IP packet level. Optical technologies and functions are becoming a larger and larger part of the IP router – IP is going optical! Photonics will gradually penetrate the very core of IP routers and replace parts of the switching matrices themselves. Reflections on these, together with an overview of optical networks, are further discussed in the introductory article in this issue. An overview of the historical development and current state-of-the-art of optical transmission is given by Aasmund Sudbø. A number of articles are devoted to core network architecture, the introduction of optical cross connects, reliability, and network control. Emerging and future solutions are in particular discussed. Also, two articles are devoted to optical packet switching, a future technology that can increase flexibility. It is a pleasure to have three articles on the real issues involved in building planning and testing a network, based on the experience from Norway.

The above relate to the evolution in the inner part of the network. However, fibre optics progresses towards the end user to provide broadband connections to the home, and large-scale Fibre-to-the-Home deployments have been taking place internationally in the past five

years or so, especially in Asia. The development is just catching up in the US and is markedly slower in Europe, so it will still be a while before we observe high penetration rates here. Yet the fact remains indisputable – Fibre-to-the-Home is slowly but steadily becoming a reality at long last. Developments in the access network are addressed by two articles, one on existing and emerging technological solutions and one on the underlying market forces and the economics of fibre to the home. Ultimately, in order to press costs really down, efficient photonic integrated circuits will need to be developed, as discussed in a brief review article on photonic integration.

But the world is evidently going wireless, so what is the future of optical fibre, a wire line technology? Mobility poses by no means a threat to optical fibre technologies. The need for mobility combined with the growing requirement for increasingly broadband connections speak for more fibre, closer and closer to the end-user, because wireless technologies have inherent bandwidth limitations. Hence fibre and wireless/mobility are complementary rather than competing technologies, though they may be competing in some cases in the short term. Every mobile network uses a fixed network backbone. Fixed and mobile services and terminals converge and so do the underlying network infrastructures as well as all network mechanisms that need to go seamless between domains. Operators and providers of mobile services will gradually need to address converged broadband services and hence relate to both mobile terminals and homes. The role of fibre and of broadband optical networks will therefore become increasingly central and visible also for operators and providers that currently focus on the mobile telecom market. Metro optical networks are in particular interesting in this context.

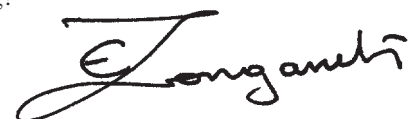
The metro network has become especially important in the last five years or so. A number of players – such as municipalities, internet service providers, network operators, etc. – are present in the metro area. New technological solutions can be easier adopted in the metro area because of its limited extent. Proximity to the end user makes it a determining part of the network with respect to performance, quality of service, and cost. Flexibility in resource allocation and the possibility to reuse the available resources are extremely

important characteristics here, as discussed in a separate article on metro network evolution.

It is an honour and a challenge to present the current topic, optical communications, to an industry that has been reborn because of it and at the same time has developed some form of technophobia in the course of its own transformation and expansion. The sector has undergone radical restructuring in the past ten years as well as experienced a particularly steep bubble period. After 150 years of technology focus in a monopolistic environment, incumbent operators are now on denial when it comes to technology and the importance of having a good grasp on it. Incumbents have had to embrace new skills and adapt to the market challenges of a highly competitive dynamically changing sector. Telecom equipment vendors, the original developers of optical technology, have had to downsize this part of their business at the same time that they embrace the idea of a more central role in designing and running networks and explore the realms of IP. IP router manufacturers on the other hand, have evidently recognized the significance of photonic technology and its ability to disrupt their business and are both investing in it and gradually incorporating it in their products. Computer manufacturers have realized the possibilities opened by optical networks and the extension of their business to enterprise networks and distributed computing. Finally, many key components and modules are developed and produced by specialized small-medium enterprises.

The players that timely recognize and exploit the further potential of photonic technology in the data and telecom sector will have a considerable head start in the race ahead. Photonics can create a second revolution, this time by revolutionizing network architecture, increasing throughput quality and flexibility, simplifying service creation and delivery, and enabling completely new value-added services. So, I dare say – watch this space!

I hope you will find much of interest in this issue – enjoy reading!



Evi Zouganeli holds a PhD in Optoelectronics (1992), an MSc in Telecom (1988), both from University College London, a BSc in Applied Physics from the University of Patras, Greece (1985), and a Master of Management (2001) from the Norwegian School of Management. After postdoctoral work at the Swiss Federal Institute of Technology (ETH, Zurich), she joined Telenor R&D in 1994 where she has focused on high capacity optical networks, broadband access, network migration and upgrading strategies, and technology evaluations – on contract from Telenor Business Units and in a number of European collaboration projects. She is a member of a number of international technical management and evaluation committees and is currently a senior research scientist at Telenor R&D.

email: evi.zouganeli@telenor.com

Optical networks: From point-to-point transmission to full networking capabilities

EVI ZOUGANELI



Evi Zouganeli is Senior Research Scientist at Telenor R&D

Optical technology has led to an increase of the single cable transmission bandwidth by three orders of magnitude in the past twenty years enabling high capacity IP networks and the start of the so-called information era. Yet whereas optical transmission is a mature business, the potential of optical networking is far from being exploited. Optical network functionality appears to be the answer to efficient and reliable data-centric networks, a technology that complements IP. This article aims at giving an overview of the driving forces behind optical networking, the potential offered by it, the challenges encountered, and the current state-of-the-art.

Introduction

Telecommunications has been dramatically transformed in the past couple of decades with a huge impact on the everyday life of the average consumer. Mobile phones are an indispensable item in everyone's pocket or purse, surfing the internet is not just for nerds, hot spots are to be found in hotels and petrol stations, wireless local area networks are broadly used also in our home. One in two has a laptop in western Europe, and downloading files and music from the internet is something many do. The world is going digital, global and on the net and it's become difficult to recall that only ten years ago most of us had neither laptop nor mobile phone ...

The internet and the Internet Protocol (IP) have caught the headlines quite often in the past decade and the importance of mobility and new mobile services is broadly understood. However, the prime enabler behind a lot of this development remains by and large invisible. The advent of optical fiber transmission in the networks has been the key enabler for this evolution. The convergence between telecom and datacom and the prevalence of IP in the network would simply not have been possible were it not for the abundant transmission capacity the optical fiber provides. The introduction of optical fiber in the mid 80s increased the transmission capacity per link from a max of 565 Mbit/s in copper wires to around 1.5 Tbit/s in an optical fiber today, increasing the total bandwidth per cable by a factor of 3000 within twenty years. At the same time the maximum transmission distance has been increased to several hundreds of kilometers from a couple of km in copper wires. It is fair to say that the optical fiber has revolutionized telecom networks and transformed telecommunications to a sustainable business and a driver for the information era.

The idea of transmitting high-capacity signals in the form of light pulses down hair-thin strands of glass may have sounded somewhat spaced out once. It was

early realized that optical fiber has a tremendous bandwidth and that the main issue would be to find technically feasible and economical ways to capitalize on this potential. Also the low electromagnetic interference of the medium makes it ideal for secure communication – something that appealed in particular to the defense industry that provided the vital early sponsoring of the field.

The impact optical fiber transmission has had on the network is by and large well appreciated by telecom engineers, however, the technology is understandably enough not quite known to the average user of telecom services. This is bound to change dramatically in the coming years when broadband connections using fiber to the building or home become more common, a trend that is clearly already on its way.

Competition and changing telecom landscape

The convergence of voice services with data and multimedia services, mobility requirements, and the presence of several competing players in the market has changed the telecom landscape for good. This is a far cry from the monopoly years. We have entered an era where new paradigms are dominating corporate, social as well as private life and where datacom is becoming an important part of our life. Despite some ups and downs in the forecasts, the fact remains that network traffic has continued to increase with a staggering 115 % per year on a global basis in the past years. Voice traffic and the corresponding revenues fall gradually but steadily and more and more services are delivered over IP. Data traffic surpassed voice traffic in volume a while ago and data traffic does pose different requirements on network performance from what voice does. In contrast to traffic generated by telephony, data traffic patterns are highly unpredictable, asymmetric in terms of load distribution, and bursty in nature. Convergence in the applications front, the potential for more sophisti-

cated services and the requirement for tailored billing mechanisms in view of the expanding competition in the liberalized telecom market, create a positive feedback loop that further strengthens the requirement for more dynamic and flexible networks. In addition, because of the dramatic increase in the capacity carried by each cable (of the order of Tbit/s), it is mandatory to have reliable and fast ways to restore the network in case of fiber cut or other failure and to be able to prioritize traffic depending on the carried service. These call for new and improved ways to build our networks.

In the years of the telecom bubble an amazing amount of innovation was generated in the optical space with something like 600 start-up companies worldwide. The demand for high performance high capacity networks appeared to be pressing at that time and a lot of development work needed to take place before we could enjoy the super-duper networks that could support all these brand new services everyone was expecting would emerge. Many of the issues that needed to be addressed in order to create robust long-distance transparent and high capacity networks were addressed and to an extent solved in the years of the bubble. Because of the large amount of money burned on optical network technologies, the bubble period stigmatized the optical field for a while and discouraged investments further. However, on the technical front, although clearly all problems are far from solved, this innovation generated a great technical confidence in the field that came with better knowledge of the potential as well as the limitations of the technology.

In the meantime, in the quiet years that have followed, advancements in other areas and technologies have clarified a number of issues that were unsettled some years ago and that are determining factors for the overall landscape with regard to services, market, competition and technological alternatives. Let us have a brief overview of these in the following.

Broadband penetration

The main bottleneck in the network is in the access part. Ever higher capacity has been introduced in the transport network with the use of wavelength multiplexed systems combined with higher bit-rates per channel. In the access part of the network on the other hand, relatively low bandwidth was delivered to private customers until only a few years ago. The reason for this is that whereas transport network investments address a large number of subscribers, the investment cost per subscriber associated with an upgrading of the access network is tremendous. When it comes to access bandwidth, the details have varied somewhat from continent to continent but the underlying situa-

tion has been somewhat similar everywhere: private consumers have had telephone lines at 64 kbit/s, internet via the telephone line or via ISDN (128 kbit/s), and television via coax cable. This has changed dramatically in the past five years or so, primarily with the introduction of ADSL (xDSL) access but also because of the deployment of a good deal of Fiber to the Home / Building (FtH/FtB) especially in Asia Pacific but also in the US and sporadically also in Europe. The number of FtH subscribers in Japan has surpassed 4 millions (Q3 2005). NTT has announced investments of USD 47 billion in FtH in total, aiming at 30 million Japanese homes passed by 2010 [1].

Broadband adoption is surging throughout Europe as well, however, there is still scope for growth in coming years. The total number of broadband access is expected to quadruple [2] over the next five years. Heavy Reading projects that broadband subscribers in this region will grow from 38 million at the end of 2004 to 128 million at the end of 2009, representing a penetration rate of 69 %. FtH is forecasted to grow to 9 million subscribers worldwide in the next two years [3]. The main driver behind this evolution is triple-play services, i.e. integrated voice, data and video. This service requires an estimated 20 Mbit/s downstream, whereas the upstream bandwidth is rather lower, at least at first, due to the prevailing role of xDSL in the broadband penetration globally. xDSL is an inherently asymmetric technology and its current dominance is bound to have a determining effect on service characteristics in the first years.

The high penetration of broadband access will have a dramatic effect on the total traffic volume in our network as well as on the type of traffic generated, since higher bandwidth will not be consumed by more voice but rather by a gradually and steadily decreasing percentage of voice traffic in the network.

Rationalized network architectures: IP/MPLS prevalence and network convergence

The way networks are built today is by no means efficient. In most current networks the SONET/SDH network infrastructure has provided a guaranteed level of performance and reliability for voice calls and leased lines. Existing networks have been designed for telephony and are thus adequate for handling static traffic patterns but are rather inefficient in handling the new traffic patterns that are dominated by data services. For historical reasons, most networks today comprise a number of layers as services have been based on ATM, SDH, or Ethernet (see Figure 1). In addition, again due to legacy, many operators operate more than one network, e.g. a separate circuit

switched network for voice and leased lines and one for IP/Ethernet networks to deliver data services. At this stage it has become clear that it is uneconomical to have several infrastructures, among other things since it does not allow for multiplexing gains, it requires duplication of operation and maintenance etc. Also, Ethernet has undercut other technologies and there is a massive migration towards Ethernet based services. Overall it has become rather clear that network infrastructures will converge to a single IP/MPLS core to deliver all services over a common platform and capitalize on synergy effects through common standards for protocols and interfaces. An all-IP network can provide a sustainable platform that can grow to support existing and emerging value-added services at optimized cost. Legacy platforms based on ATM, SDH etc. will be phased out and indeed a number of operators have already been out with the modernization of their network (e.g. Telenor).

Service and terminal convergence

Service convergence refers to the trend that it shall be possible to reach the same content from different technical platforms. There is a clear trend towards service convergence across fixed and mobile networks, voice and data platforms, and rich IP applications. The total deregulation of telecom markets and open competition among operators and other players, combined with technological advancements and facilitated by emerging multimedia services and general mobility, have created an environment where it makes sense to provide personalized multimedia services with consistent seamless mobility across different access networks and domains as well as unique user identification and authentication irrespective of terminal, location, and access medium. This notion has been supported and promoted by the ITU-T through its Next Generation Network (NGN) initiative. One of the important aspects of NGN is that services and service related functions are separated from the underlying transport technologies so that the two can evolve independently.

The effect of service convergence is enhanced by terminal convergence that is taking place in parallel. Terminal convergence is manifested e.g. in the coming together of consumer devices such as fixed and mobile phones, TV and PC, Phone and TV etc. The PC terminal is developing into a real competitor of the TV set, and the two are substitutes for a range of services.

Market and industry changes

The whole value chain in the telecom sector has been redefined in the last decade as a result of a series of factors, such as market liberalization and increased




Figure 1 A schematic of the typical network architecture today and the envisaged leaner architecture of the future where optical networking can facilitate higher functionality and overall better performance in IP networks

competition, new regulation, new technological possibilities that enable terminal and network convergence, service convergence and others. The former boundaries between IT, telecom, broadcasting, and other mass media are being gradually wiped out and the sectors have by and large merged. Because of service convergence, content providers may provide content to more than one sector and new opportunities are created across earlier role boundaries. There is also a convergence in equipment production as a result of the need for network and terminal convergence. Traditional datacom equipment providers enter the telecom space, and vice versa. At the same time there has been a clear trend towards outsourcing a number of operations that have traditionally been part of the core business of telecom operators, and there is a trend towards a separation between the role of the service provider and that of the network operator.

On the whole, the market-, business- and technology related landscape of the telecom sector has been considerably reshuffled recently, creating both threats and new opportunities for existing players and entry opportunities for newcomers. The role of network infrastructure in this context has probably been both overestimated and underestimated at times. There is no doubt that for a telecom operator to sustain competitive edge and be able to deliver the new converged consumer services as well as wholesale products in a competitive manner, it will be necessary to have an advanced and cost-efficient network. Among other characteristics the network should be able to adapt to changing traffic requirements and make a good use of the network resources at all times, with fast automatic reconfiguration, efficient traffic engi-

neering, and service differentiation. These will allow a rationalized – i.e. economical – use of the network, respond to the network service and terminal convergence that are taking place as well as the changing business conditions, and enable the introduction of new value-added services.

First generation optical networks

How does optical technology contribute to the vision of advanced next generation networks? Optical transmission is widely used in most parts of the network worldwide, namely from transatlantic and pan-European connections, to national connections between cities as well as within cities. Fiber is also gradually moving closer to the end-user, in Fiber-to-the-Base-ment or Fiber-to-the-Home solutions. Also, in order to provide higher capacity via copper wire it is required to upgrade to VDSL or ADSL2+ broadband technology, and new fiber nodes need to be built close to the subscriber since higher bandwidth limits reach for xDSL based solutions. It is fair to say that telecom

networks are optical fiber networks with wireless, mobile, copper and recently also fiber last mile drops.

The role optical technology has played so far has been that of a “dumb fat pipe”, i.e. providing high capacity point-to-point connections. Layer 1 functionality including protection/restoration has been covered by SONET/SDH and routing/forwarding by IP/MPLS while a range of Layer 2 technologies is implemented to realize different services. The network functionality potential of optical technology is still largely unexploited.

Transmission aspects are discussed in detail in the following article by Aasmund Sudbø. Here let us mention that most commercial transmission systems use Dense Wavelength Division Multiplexed (DWDM) systems with wavelengths in the 1530 – 1565 nm range, the so-called C-band that corresponds to the low-loss window of the optical fiber. This is also the range where the Erbium doped optical amplifier is operative [4]. A grid of allowable wavelengths/frequencies of operation is specified by ITU-T, centered at 193.1 THz or a wavelength of 1553.3 nm and all frequencies spaced at multiples of 50 GHz (= 0.4 nm) around that center frequency. The advent of Raman amplifiers [5] in the past years has enabled operation in the 1565 – 1625 nm range, the so-called L-band. This has enabled a doubling of the total capacity that can be carried by conventional systems to just over 160 channels at 10 Gbit/s per channel. The tight spacing between these channels requires temperature control of the lasers and increases cost.

A lower cost system is obtained with Coarse Wavelength Division Multiplexing (CWDM). Here the channel spacing is increased to 200 GHz, wavelength precision requirements are relaxed such that both the fabrication and the temperature control requirements of these lasers are lower, and the overall cost dramatically reduced. CWDM systems are broadly used in metropolitan area networks typically combined with relatively low-reach (< 80 km) non-amplified systems that further reduce cost per bit.

Optical transmission has been around for almost twenty years and has evolved with steady incremental steps forward, so that it can be considered a mature business. A lot of fiber was installed in the bubble period and the general understanding has been that by increasing the number of WDM channels and/or the capacity per channel following incremental technology improvements there would be plenty of capacity many decades ahead. As it comes out, this is far from true. Installed fiber increases at a rate of 20 % annually in the USA and 10 % on a world basis [6]. Only

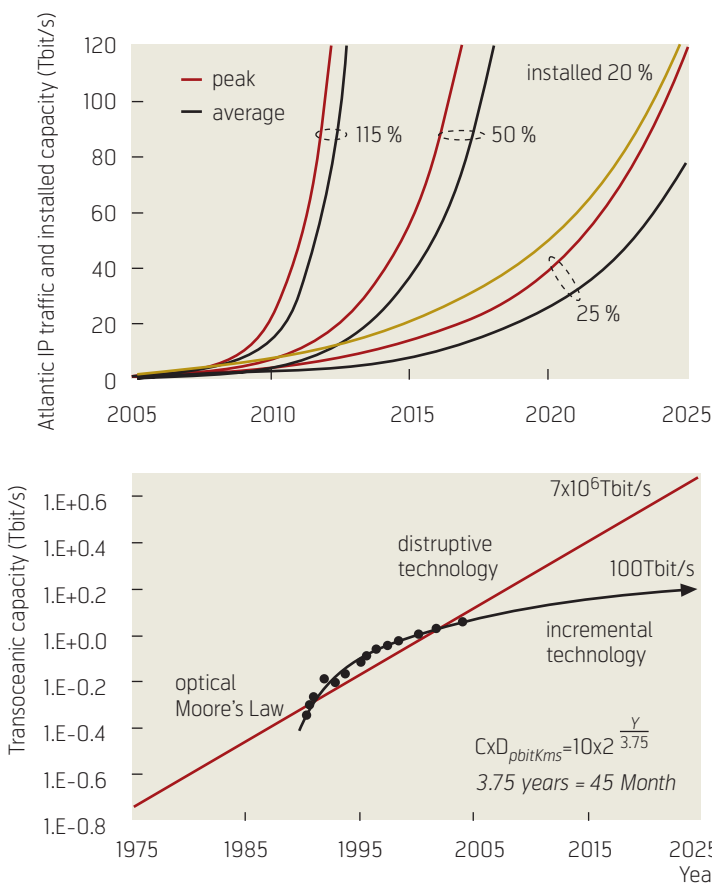


Figure 2 a) Atlantic IP traffic projection for a number of scenarios and installed capacity projection using today's installation rate for the US, b) Moore's Law for optical transmission is not met by incremental technology – a breakthrough in optical transmission technologies is required – after [6]

if capacity demand increases at less than 25 % annually will the total installed fiber manage to deliver the required capacity by year 2025. Incremental transmission technology has already started leveling off so that there is a need for a breakthrough in optical transmission technology! These are illustrated in Figure 2 [6].

Typically close to 70 % of the traffic that arrives at a node is through-passing traffic, i.e. traffic with destination another node – depending on the network area. Instead of processing and switching/routing all traffic at SDH and/or IP/MPLS level at each node, an optical add-drop multiplexer (OADM) in a ring network can allow at least some of the through-passing traffic to be forwarded optically, directly through the node. An OADM uses the wavelength space to distinguish traffic of different destinations and forward it accordingly either through the node or to the node. An OADM has two line ports where the composite WDM signals are present, and a number of local ports where individual optical channels are locally dropped or added. So an OADM takes in signals at multiple wavelengths and selectively drops some of these locally while letting others pass through. The wavelengths that have been dropped may be added back in the link. An example implementation is shown schematically in Figure 3, where signals are aggregated in channels according to destination node, and the appropriate channel is dropped at the destination node using low cost fixed wavelength optical filters. This architecture does not allow re-configurability, and in its simplest form wavelengths are pre-allocated to nodes. Similarly in meshes, OADMs can be used to create a number of selected direct connections between some node pairs with high traffic interest. These systems are static, i.e. non re-configurable, a fact that minimizes the potential of these in a network context exactly because it makes them inaccessible to the management/control system. Optical channels are thus still providing point-to-point connections in these configurations.

Optical network functionality and optical network elements

The potential of optical technology is far from realized when such static networks are used. In WDM networks, the wavelength domain can be used to realize advanced and very efficient network functionality rather than only increase the capacity carried in the fiber, as it is used today. Optical technology can be employed to realize very efficient switching and forwarding systems. These can act in a separate layer and take over a lot of the functionality that is traditionally carried by higher layers. Another way to describe this is that optical technology is currently




Figure 3 A schematic of a ring network employing static OADMs. A set of fixed wavelengths is assigned to, and dropped at each node




Figure 4 With the introduction of optical cross-connects (OXCs) and the optical layer, through-passing traffic can bypass intermediate IP routers that are then significantly reduced in size. Add/drop channels from the node are directed to the IP router whereas the rest of the traffic is handled at optical channel/band level by the OXC

used in the physical layer whereas it can be employed to realize layer 1, 2 and “2.5” functionality and complement or substitute other technologies in their tasks. An alternative approach is that optical technology and functionality is controlled directly by IP routers and becomes an extension of these. Finally, future options include an introduction of optical technology within

the router in order to make larger and/or more efficient and/or more robust IP routers, or in order to realize optical routers switching optical packets. These are briefly discussed at the end of this article. We give a more detailed account of the key optical network elements and their functionality in the following.

Opaque versus transparent

One aspect of optical networks that has been broadly discussed and debated is transparency. Transparency refers to the fact that the optical channels can carry data at a variety of data rates, protocols, and formats and can support a variety of higher layers concurrently. It is a quality of optical signals that potentially gives optical networking a huge competitive advantage over other technologies – electronics – exactly because it makes it inherently scalable in terms of speed, as well as bit-rate-, format- and protocol-blind. These components are called all-optical and are often referred to as transparent. The alternative is to handle signals at optical channel level (rather than IP packet or VC4 level etc.) and to convert them opto-electronically and actually do all the switching or other processing using an electrical core. This type of solution is called opaque and is feasible and mature today. Such opto-electronic solutions do not scale well and will ultimately be displaced by all-optical alternatives. When this will happen may vary from component to component, however, it can be expected that bit-rate will play an important role in the timing.

Indeed, as bit rates increase, electrical solutions need to scale according to the total capacity handled, whereas optical solutions scale with the number of channels and will at some point become more cost-efficient than their electrical counterparts. This is accelerated by the evolution of photonics in general for applications other than optical networking. On the other hand, transparent optical networking requires advanced all-optical components, i.e. where functions take place in the optical domain. This is not always easy to achieve and a lot of development has been going on in the past 15 years reporting both success and fiasco stories. At the same time, all-optical solutions mean in general that the signals are not re-generated and hence pose stringent requirements on the transmission systems. A broadband all-optical regenerator is a very critical component in this respect [7].

The key NEs in order to realize re-configurable optical networks are reconfigurable OADMs (R-OADM) and optical cross connects (OXC). Both these can be realized either as opaque solutions using optoelectronic conversion and an electrical core, or as all-optical solutions. We give a brief account of these two key NEs in the following with some emphasis on all-optical solutions.

Re-configurable add-drop multiplexers

Static OADMs were introduced in the previous section. Some re-configurability of these may be achieved at relatively low cost if simple (slow and cheap) 2x2 switches are inserted in the add- and drop-

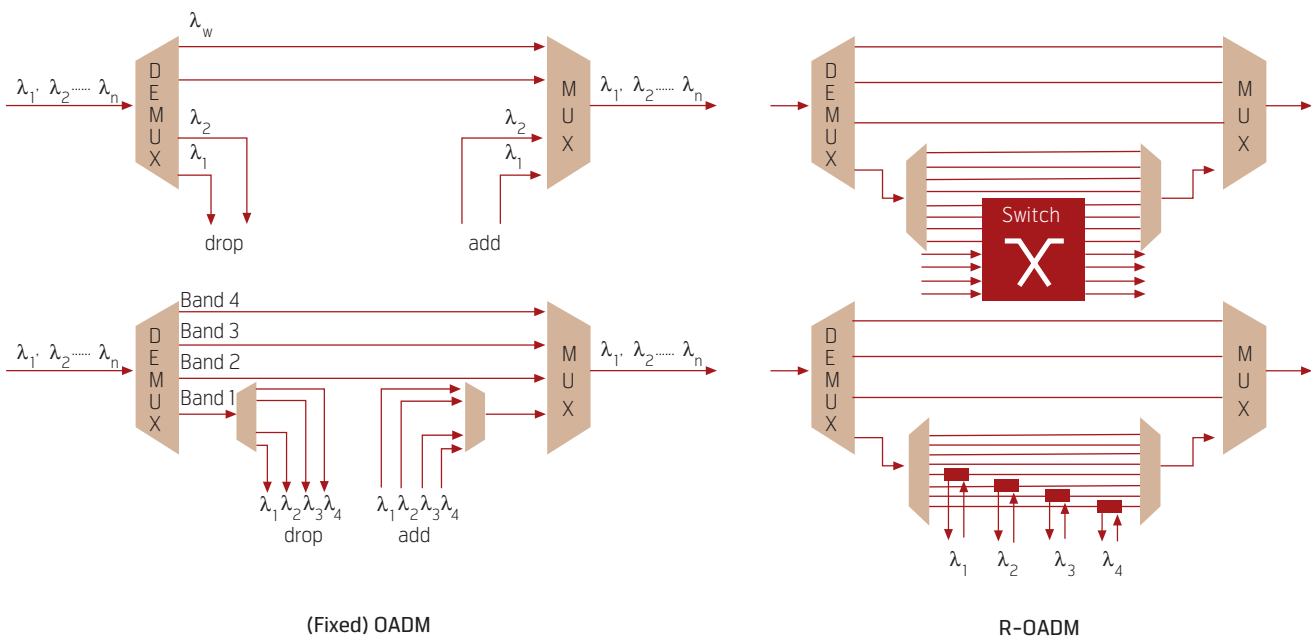


Figure 5 Generic OADM and R-OADM architectures

path of the system in Figure 3 so that the wavelength may or may not be dropped/added at the node under question. This makes the OADM controllable by the network management and/or control system and hence a valuable network element (NE).

Some simple OADM architectures are shown in Figure 5. Several technologies have been explored and employed to realize R-OADMs. These include liquid crystal based switches combined with multiplexers and de-multiplexers, micro-electro-machining systems (MEMS), and a range of tunable filter technologies where Bragg gratings may be the most popular and mature solution among these.

One example of a MEMS-based R-OADM can be made with the optical signal de-multiplexed and then passing through an array of individually controllable reflectors with one wavelength per reflector [8]. The wavelengths that are to be dropped at the node are reflected to the drop output of the R-OADM whereas the rest are reflected back to the optical ring (Figure 6). Another example, by Pirelli [9], is shown in Figure 6 where Bragg gratings are employed to realize tunable filters that are tuned accordingly to tap out the selected wavelengths.

R-OADMs are currently becoming commercially available and simple versions are deployed primarily in metropolitan area networks.

Optical cross connects

Ring networks are simple and relatively low-cost solutions. However, they do not use the bandwidth available in the fiber efficiently and neither do they use the available fiber infrastructure efficiently. Optical networks will consist of a multiple of sub-networks as dictated by administrative geographical and technological factors. These will need to be interconnected by optical links in an arbitrary topology, i.e. in a mesh topology since the physical network is in the general case a mesh network. A mesh is a scalable topology that makes the best use of both the available bandwidth and the available infrastructure, facilitates load balancing in the network, as well as provides multiple and short restoration paths based on 1:N protection schemes. The latter means that protection paths are shared by more than one (namely N) working paths through the network – a fact that limits over-dimensioning of the network and leads to considerable cost savings.

In order to be able to realize mesh networks, optical cross connects are required. An OXC performs in essence the same function as an R-OADM but for a larger number of line ports and directions, in practice also a larger number of local ports. A schematic of an OXC is shown in Figure 7. OXCs have a large number of input and output ports and are able to switch wavelength from one input to another. They can thus be used in meshed networks to direct signals through the network.

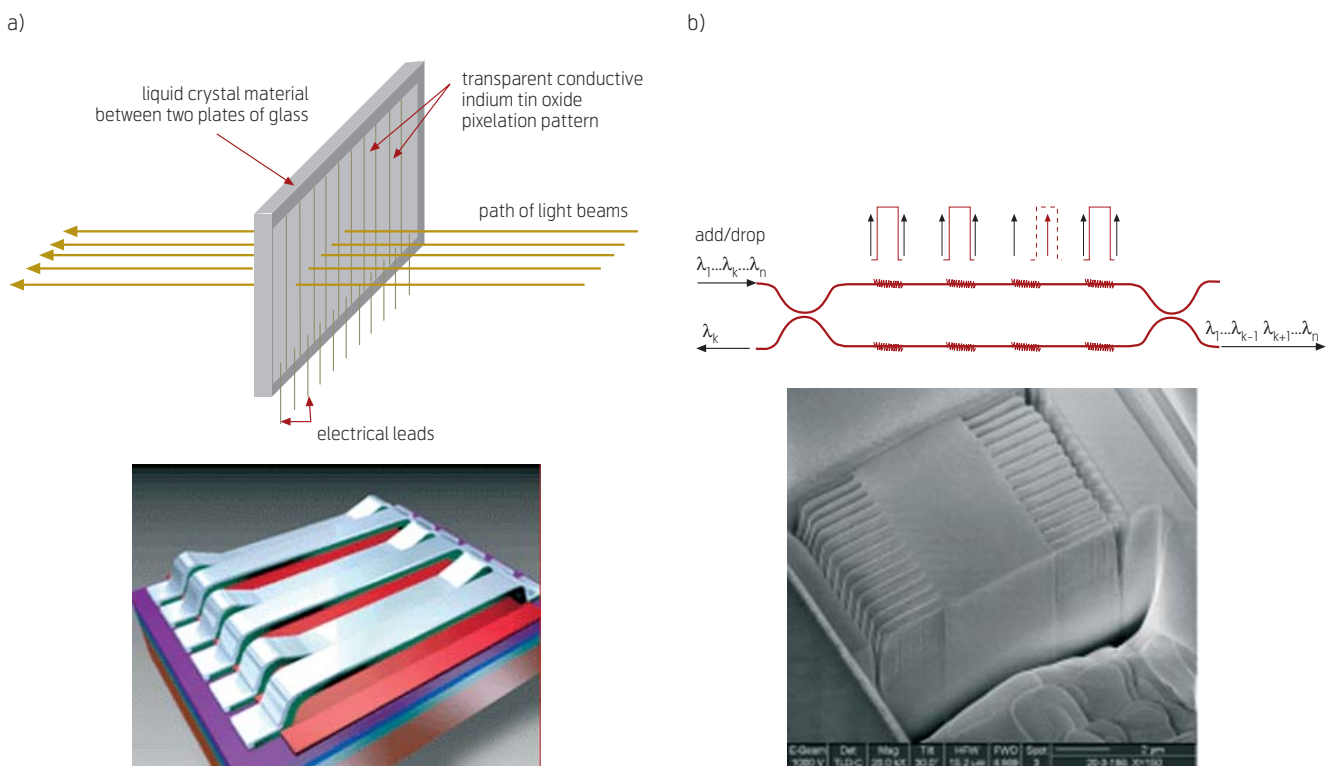


Figure 6 Examples of R-OADM solutions; a) a MEMS based R-OADM [8]; b) a Bragg filter based R-OADM [9]

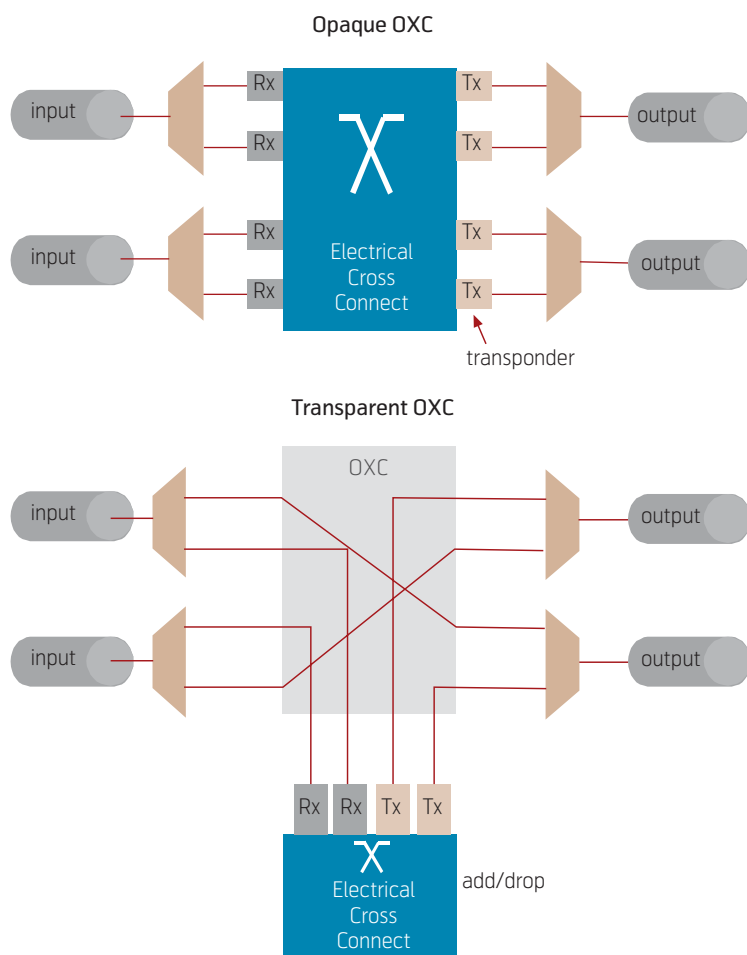


Figure 7 Schematic of an optical cross-connect; optoelectronic (top), all-optical (bottom)

An OXC provides several key functions in a large network [10]:

- **Service provisioning**

An OXC can be used to provide end-to-end lightpaths in a large network in an automated manner. Remotely configurable OXCs allow re-configurability of the network to respond to traffic changes. “Point-and-click” provisioning of optical channels (OCh’s) can be achieved using OXCs. Bandwidth can be allocated on-demand or “created” at the parts of the network where it is required. Note that an OCh is not necessarily all-optical along the end-to-end link; neither is it necessarily one single wavelength along the whole link.

- **Protection and restoration**

Protecting the network against fiber cuts or equipment failure at the node. The OXC combined with monitoring equipment can provide swift restoration of huge amounts of traffic by redirecting wavelengths from the failed paths to alternative paths. This function may be complemented with wavelength conversion.

- **Wavelength conversion**

Lightpaths need not use one single wavelength through the whole network since this complicates wavelength management in the network. The possibility to implement a moderate degree of wavelength conversion gives flexibility especially in dynamic networks. Wavelength conversion may be performed all-optically, e.g. using semiconductor amplifiers in a Mach-Zehnder interferometer [11], a system that results in 2R regeneration of the signal and is a popular one. Another alternative is to use optoelectronic conversion and retransmit at a new wavelength, which has the advantage of fully regenerating the signal but is rather expensive and currently has to take place per channel. The OXC is a natural location to perform wavelength conversion.

- **Multiplexing and grooming**

In its purest form the OXC is all-optical. However, the OXC comprises also an add-drop part where multiplexing and grooming of ingress and egress signals can take place.

OXCs are just emerging and currently relatively expensive network elements. The main systems used today are opaque and employ an electrical core. A large number of alternative all-optical solutions have been explored, however, some time is still required before the winning technologies are identified and OXCs become mature, broadly-used systems. A very straightforward way to realize an all-optical OXC is to use the well-known generic switch architectures, such as Clos or perfect Shuffle, and realize these using optical switches and fiber connections. Similarly, Broadcast-and-Select type switching matrices can be realized optically using semiconductor optical amplifiers (SOAs) that block the signal when they are off and let it pass through when they are on.

One of the most prominent OXC solutions is based on a two-dimensional or three-dimensional array of miniature mirrors (MEMS based) that measure only a few tens of μm width each. One such mirror and an OXC based on two arrays of MEMS mirrors are shown schematically in Figure 8 [12].

In recent years the advent of widely tunable lasers has promoted a new OXC architecture that uses tunable lasers and Array Waveguide Gratings (AWG) as shown in Figure 9. The AWG is a mature component that is used to make optical multiplexers and demultiplexers. Its function is to provide “hard wired” wavelength dependent connections following a table. The port a signal will exit the OXC from, is explicitly determined by its wavelength and the AWG input it arrives at. Hence the OXC is re-configured by tuning the wavelength of the signal in a wavelength con-

verter using tunable lasers [13]. Tunable lasers are becoming the commonly used transmitter and will entirely displace traditional fixed-wavelength lasers. This makes the AWG-based OXC a very viable solution.

Multi-protocol label switching (MPLS) and Generalised MPLS

MPLS is a network concept that places all complex functionality to the edge of the network and is inherently multi-protocol as the name states. Its success has been due to its properties complementing IP. It is widely developed for IP/MPLS networks and is currently being deployed in networks. MPLS is connection oriented and performs forwarding of packets (Layer 2) according to their belonging in Label Switched Paths (LSPs). Packets are assigned an MPLS label (stack) and all processes are carried out based on this label stack. The use of MPLS increases in practice the granularity of IP networks since routing decisions are taken for all packets that belong to an LSP and carried out as forwarding tasks at each router rather than packet-per-packet routing. The main benefits of MPLS lie with the forwarding processes, traffic engineering and restoration mechanisms. Fast reroute in the case of fiber cut or node failure is shown in Figure 10.




Figure 8 a) A picture of a miniature MEMS mirror, b) schematic of an OXC using miniature MEMS mirrors [12]

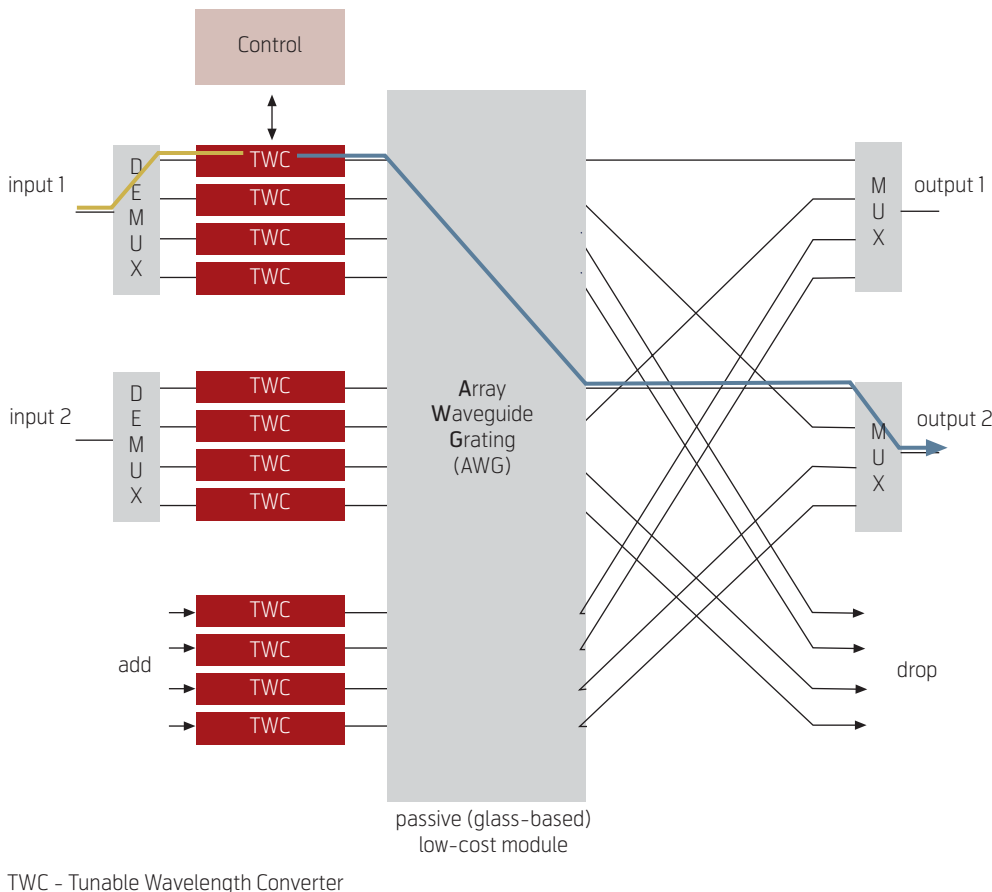


Figure 9 A schematic of an OXC based on an Arrayed Waveguide Grating and tunable wavelength converters




Figure 10 Fast reroute in MPLS requires rerouting of thousands of LSPs in a typical network

An extension of MPLS – Generalized MPLS (GMPLS) – has been proposed by IETF to include the time and optical wavelength domains. The aim is to achieve a more flexible labelling and forwarding mechanism that uses a generalized label, which is applicable to a variety of technologies and networks. For IP routers the labels designate principally input and output ports. For an OXC they designate input and output ports as well as wavelength, or band of wavelengths. The hierarchy of different labels in GMPLS is schematically shown in Figure 11.

The main aims of GMPLS are to [14]:

- provide a framework for real time provisioning of optical channels;
- adopt optical technology and encompass the development and deployment of a new class of programmable OXCs;
- allow the use of uniform semantics for network control in hybrid networks that consist of both OXCs and label switching routers.

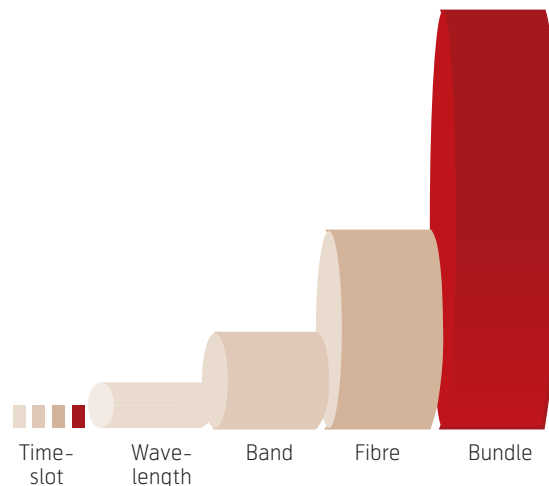


Figure 11 The hierarchy of labels that can be used in GMPLS

The impact of the introduction of GMPLS in the integration of the IP and the optical layer is potentially huge. GMPLS makes it possible for routers to “see” optical wavelengths and wavebands as manageable resources. One single forwarding decision may be taken for a whole wavelength or a whole set of wavelengths that correspond to a large volume of packets. The whole forwarding process can be carried out much more efficiently and the throughput of the IP network dramatically improved. It also makes sense to perform protection and restoration processes directly at optical level, either in a separate control plane or directly by IP/GMPLS routers in a peer-to-peer architecture. These are discussed in a later section.

Why optical functionality

In order to understand why it is important and advantageous with optical functionality, consider the following. IP routers have now reached a processing capacity of the order of 1 Tbit/s with several optical interfaces at 2.5 Gbit/s or 10 Gbit/s. IP packets are typically between 300 and 1500 bytes. Hence a 1 Tbit/s IP router processes 100 – 500 million IP packets per second where packets undergo an opto-electronic conversion at each node, their MPLS label is read and rewritten, and they are retransmitted (Figure 12). The deployment of MPLS has provided IP/MPLS networks with two important features that pure IP networks lack: forwarding capabilities at LSP level, and more efficient restoration using tunneling techniques in fast reroute. Fast reroute can attain a restored path within 50 ms. However, it would be interesting to see how these things scale and how “fast” fast reroute actually is when network size and traffic interest increase. After all even with its fast tunneling techniques, MPLS needs to process several 1000s of LSPs in order to reroute the traffic in one fiber in a typical network.

These underline an emerging clear bottleneck in IP/MPLS networks. IP/MPLS routers have the wrong granularity to handle multi-Terabit networks and are not scalable in a sustainable way. A router with 320 Gbit switching capacity requires half an ETSI rack today. Larger matrices are realized by interconnecting several racks to an electronic switching matrix using optical cables. Power dissipation per cm³ makes further miniaturization of electronics increasingly challenging so that an infinite reduction of size to match increasing processing capacity will not be possible. Optical technologies can eliminate these bottlenecks. For example, multiplex section protection in the optical layer, shown in Figure 13, is carried out by rerouting a drastically reduced number of LSPs as compared with Figure 10.

A quick look at the cost of interfaces today underlines this trend rather clearly. The cheapest 10 Gbit/s interface is an Ethernet interface and it costs around USD 50 – 70,000. A Packet over SONET 10 Gbit/s interface costs around four times that much. By contrast, the current price of optoelectronic OXCs is quoted per interface and lies at USD 6,000 per 10 Gbit/s port, which is at worst one tenth of the cost of the corresponding IP interface. This can be interpreted to mean that the cross-over point for the introduction of optical cross connects today is when the traffic interest between node pairs in the network is around 1/10 of the capacity of the optical channel, i.e. around 1 Gbit/s. Upgrading costs and protection/restoration improvements in an optically cross connected network favors optical networks and push the cross-over point even lower.

Network architectures

With regard to architecture alternatives for an IP over optical network, the determining aspect is whether and to what degree the control plane of the optical network will be integrated with IP or independent from it. The IP and optical control planes can in other words be loosely or tightly coupled in terms of, firstly, the details of the optical network topology, resources, and routing information that is revealed to the IP layer, and secondly, the degree of control IP routers have on optical network elements and thus the degree to which they can determine the exact paths through this optical network. Three architecture options can be identified from this point of view:

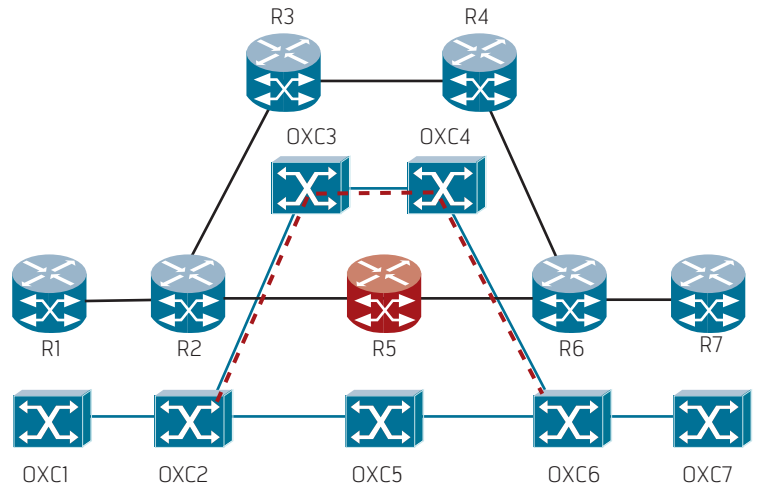


Figure 13 Multiplex section protection at the optical layer

The overlay model

In this architecture option the optical network has full control over its network resources by means of a fully independent optical control plane. The Automatically Switched Optical Network (ASON) has been proposed by ITU-T in accordance with the overlay model. A schematic of ASON is shown in Figure 14. All communication with client networks in an ASON is done via the User Network Interface (UNI) that carries all signaling information exchanged between ASON and its clients as well as the actual signals transported by the network (via the physical layer part of the interface). The client networks request a connection between two edge nodes, requesting also a set of quality related characteristics for this connec-




Figure 12 A schematic of the forwarding process in an IP/MPLS router

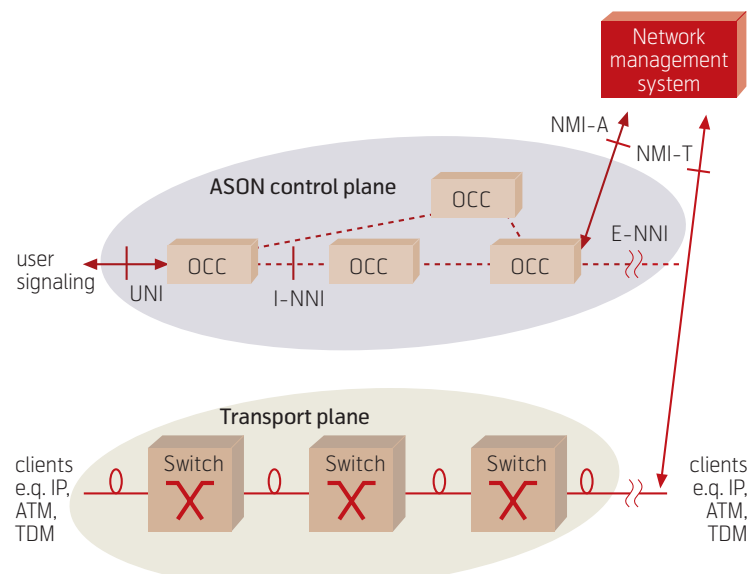


Figure 14 A schematic of the Automatically Switched Optical Network, according to the overlay model

tion. These characteristics do not only regard bandwidth but also e.g. delay, jitter, degree of protection etc. The client networks have otherwise no control over the exact routing and priority received within the optical network. GMPLS is used in the separate control plane.

ASON provides end-to-end OCh connections to its clients with a certain QoS, as agreed via service level agreements (SLA) with the client. These connections can be static, established via the management system, or dynamic. Three types of OCh services can be provided: Permanent OCh connection, soft-permanent OCh connection, automatically switched OCh connection. The latter provides an end-to-end optical channel connection activated by direct signaling from the client network. Establishment and tear down of this service is handled automatically by the ASON control plane and the client is notified accordingly. An extensive field trial of user controlled automatically switched connections has taken place and is up and running in Japan [15]. It is nonetheless unclear whether this type of solution will prevail.

In the overlay model an “intelligent optical network” carries out part of the network functionality. One advantage of this model is that it is a multi-client solution, which can accommodate technologies other than IP that many operators will need to relate to, at least for a while. Separating the two control planes implies also that the two parts may evolve, be adapted, and be optimised independently, which is a good future-proof policy. The optical network provides here a universal platform that is not tied to one specific protocol but is open to any future newcomers. The disadvantage of the overlay model, on the other hand, is that it requires the creation of a new control plane that to an extent may duplicate functionality and may introduce delays – repeating what is the old problem with layered networks. It can be expected that as IP gradually displaces alternative technologies, the overlay architecture will at some point become an anachronism.

The peer model

In this architecture the control planes of the optical network and IP are fully integrated such that IP routers and optical switching nodes (OXC) are




Figure 15 Schematic of a router with widely tunable wavelength interfaces and integrated OXC functionality using an Array Waveguide Grating and the arrangement in Figure 9

peers. The two networks are merged into a new integrated network that is managed in a unified way so that the optical network topology is fully visible to routers. A single protocol is run through all domains and establishes paths through all network elements in a seamless manner.

The view is to keep the optical part unintelligent and rely on IP intelligence to run the network. The advantage of this architecture stems exactly from the fact that it is IP-centric. The architecture is scalable, functionality is not duplicated and conflicts between several control planes do not arise. On the other hand, this architecture demands that information regarding the optical network elements is advertised to routers, resulting in excessive information flows within the network. Finally, this architecture is not inherently multi-client, an aspect that may be important for some operators (e.g. incumbents). Despite its drawbacks, the peer-to-peer model may be the architecture that will be adopted in the longer term since IP indeed is dominating the scene.

The advent of tunable lasers may promote the peer model. As tunable lasers become more mature and cheaper, they are gradually and steadily replacing the most common fixed wavelength lasers – DFBS. IP routers use today 10 Gbit/s optical interfaces that are bound to be tunable lasers in the near future. As soon as tunable laser interfaces are introduced, it will be possible to realize OXC functionality by adding an AWG at the output of the router as shown in Figure 15 according to the arrangement in Figure 9. Wavelengths are steered by the router using GMPLS and hence it may be most natural to employ a peer model.

QoS in optical networks

An optical network can provide service differentiation based upon a set of parameters – such as priority, delay, jitter, etc., as well as dependability aspects – following IP resource handling models [14]. The different classes or bundles of classes are carried by the same channel or set of channels. This way, differentiation can take place and be handled optically, without intervention or knowledge at the IP layer, making some processes – such as restoration – much simpler, swifter and more efficient.

Reference network

The network can be seen as comprising three parts: The core and long-haul part, the metropolitan area network, and the access part. Traditionally the core part is a national network connecting major cities. The changes of the telecom landscape and increased competition have given rise to both larger and smaller

core networks, e.g. extending a whole continent or providing a backbone to a local mobile operator. The metro network has emerged as an important part of the network in the past years. Typically, a metro network extends over a large city or a province, comprises less aggregated traffic than the core as it is closer to customers and end users, and involves lower capacities than the core. Finally the access part is the part that provides a connection between an end user (e.g. a home) or a business to the metro and core networks.

Optical metro networks

As capacity increases, the role of the metro network becomes increasingly important. Metro networks are also small arenas where newcomers can easier establish themselves and where new technology can easier be employed. This increases the dynamics in this part of the network and allows swifter changes. CWDM is broadly used in new metro deployments today, OADMs are quite common and R-OADMs start being deployed. Legacy metro networks have been layer upon layer of SDH rings. Some of these are being upgraded using NG-SDH, however, the trend towards Ethernet based services promotes packet-based networks that are significantly more cost-efficient. The introduction of optical functionality in the metro area is rather evident.

Access bandwidth increases with the introduction of broadband to the home, storage applications, the use of IT by government, hospitals, universities, banks and businesses. In addition, there is an increasing number of public wireless access hot-spots. A metro network shall serve all these accesses and ideally provide bandwidth on demand when required. It is expected that automatic provisioning will be introduced in metro networks, at first controlled by the management system rather than via a UNI. A future optical metro network can provide flexible bandwidth allocation to a plethora of services and access technologies and facilitate seamless mobility independent of access technology. Optical metro networks can e.g. provide feeder functionality to relatively simple wireless antennas by centralizing processes, moving all intelligence to the central office and using RF over Fiber techniques [16]. An example of a flexible broadband optical metro network is schematically shown in Figure 16 where different segments or services may be carried by different sets of wavelengths and where QoS may make use of the optical wavelength domain. Optical metro networks may initially be ring based, however, mesh networks are to gradually replace these, provided cost-efficient OXCs are available.

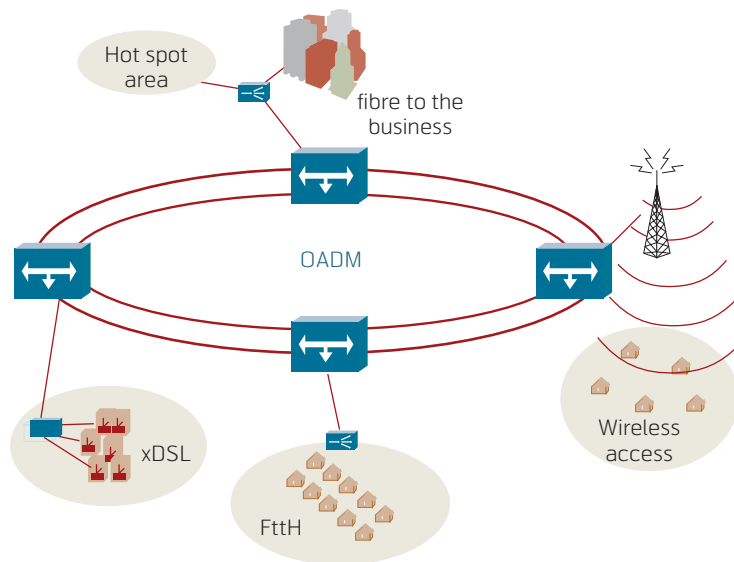


Figure 16 A flexible multi-service broadband optical metro network that provides seamless mobility between a number of access technologies and bandwidth on demand

Fiber in the access

Fiber has just entered the private access space with Fiber to the Home solutions in the past years, and is progressing at high speeds. The drivers behind this evolution are the success of triple play services combined with increased competition, an increasingly maturing optical technology due to its broad application in other parts of the network, and the superiority and futureproofness of Fiber-to-the-Home. The main competing technology in Europe is xDSL (VDSL, ADSL+2). However, the maximum reach of xDSL technology is an inversely quadratic function of the transmitted bandwidth. Triple play services require bandwidths of around 20 Mbit/s and above – a bandwidth that requires a fiber node closer to the subscriber (within 1 km for highly asymmetric VDSL at 26 Mbit/s downstream and dramatically reduced for the somewhat more symmetric varieties). At the same time, fiber to the home terminals and solutions are available at a cost that is comparable to high capacity xDSL solutions (e.g. VDSL), and the cost of installing fiber in green-field scenarios is equal to that of installing VDSL. The main cost with FttH deployments is associated with digging costs in order to put the fiber in the ground. All in all fiber is still a costly alternative for incumbent operators that need to tap out the value of their existing copper infrastructure. New installation techniques in existing pipes (e.g. blowing fiber cables) make fiber a much more realistic solution than it has been earlier. Also fiber can reduce operational costs (OpEx) considerably since it can allow passive networks as well as a reduction of the total number of central offices due to its long reach – passive optical networks (PON) with a reach

of 60 km are available (see articles on the access network in this issue). When OpEx savings are taken into account the picture becomes brighter for FttH.

Strategic decisions regarding the technology shift to optical fiber in the access need to be based on market and competition dictated conditions. Given that triple-play services are taking off and given the bandwidth-distance product of xDSL, it becomes rather clear that xDSL is not very far from the end of its lifecycle. A plan for a gradual introduction of fiber should be at its place soon, also for incumbent operators. Indeed, the higher capacity and symmetric capabilities of FttH grant it a significant competitive advantage. Upgrades will cost very little in the case of fiber while they will become more and more of a challenge and very costly for xDSL. These are illustrated in Figure 17. This means that incumbents ought to consider carefully the right timing for starting to introduce fiber in the access, to ensure that it does not become a disruptive technology in this part of the network. Due to the somewhat proprietary solutions in the actual implementation of FttH solutions, customer churn between FttH providers is very low. Also existing access fiber investments create an entry barrier for newcomers since they reduce the potential for high take-rates as well as limit margins and pricing options. Regulation regarding unbundling of the local loop is a main show-stopper for incumbents in Europe where FttH deployments happen more or less exclusively by newcomers.

But isn't the future wireless?

Mobility in datacom is like color TV in broadcasting – once wireless mobility is experienced one can never go back to wire-line connections to phones and PCs. Fiber is a broadband access wire-line technology, so what is the future for fiber in the access? Mobile and wireless technologies suffer more or less inherent bandwidth limitations. Higher bandwidth per subscriber can be accomplished when wireless technologies are employed in point-to-point directional connections – that do not offer mobility. High frequency radio solutions will have higher bandwidth, however, wireless technologies provide relatively low bandwidth that shall be shared by many subscribers and are not the solution for the provision of triple play services, neither will they be able to provide the high bandwidths required by businesses and the public sector. The way to obtain both wireless access and high bandwidth is to bring the fiber closer to the location and then use wireless locally. Hence, although in some cases wireless and fiber solutions may be competing technologies in the access, especially in the short term, in general the two technologies are com-

plementary and will in fact promote and facilitate the deployment of one another.

Future IP solutions

As discussed earlier (Why optical functionality) IP routers do not appear to be able to scale gracefully to meet the needs of multi-Terabit networks. Larger IP switching matrices are created by interconnecting several routers with a switching matrix. Alternatively, better scalable large switching matrices may be created using both electrical and optical switches. Electrical switches allow buffering and packet-per-packet processing whereas optical switches handle efficiently switching of large quantities. One such architecture is shown in Figure 18 [17]. A stage before this step, optical switches may be introduced to improve the reliability of router matrices [18].

A next step is optical packet or burst switching (OPS/OBS) – or so-called optical routers. These are also discussed in separate articles in this issue. The main principle behind OBS is that an optical burst comprising a series of IP packets is forwarded all-optically through the network using (G)MPLS principles. These optical bursts – or “mega-packets” – are formed at the edge of the network and carry aggregated traffic, e.g. according to destination and class of service, if applicable. An optical label is attached to this mega-packet. This label is detected at each node and processed electronically whereas the signal is switched all-optically. In many cases the payload waits in an optical delay line while the label is being processed. In other cases the label precedes the payload. OBS increases the throughput of IP networks as it combines the high efficiency of optical technologies with the granularity and flexibility of packet technologies.

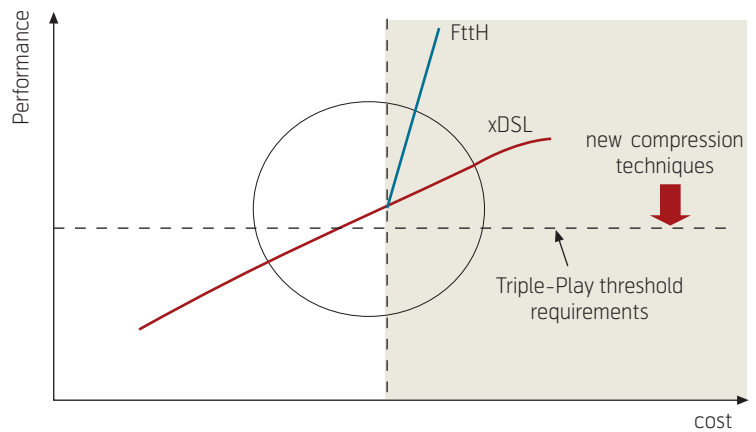


Figure 17 The performance/cost ratio of fiber is significantly superior to that of xDSL. Given that triple play services pose challenges to xDSL technologies already, xDSL is evidently not far from the end of its life cycle and upgrades will be increasingly costly. A gradual technology shift to fiber should be seriously considered – also by incumbent operators

Summary

Although optical transmission is a mature business that has revolutionized telecom and datacom, the potential of optical technology is still far from being exploited. Optical network functionality in wavelength-routed networks is still in its infancy. Optical networks may well be the key to high capacity intelligent networks that utilise their resources in an efficient way and can provide a range of differentiated services. Optical switching provides an economical way to handle large amounts of traffic and to build reliable networks. It leads to a dramatic reduction of the required processing capacity and to rationalised network architectures without duplication of functionality and expensive superfluous interfaces.

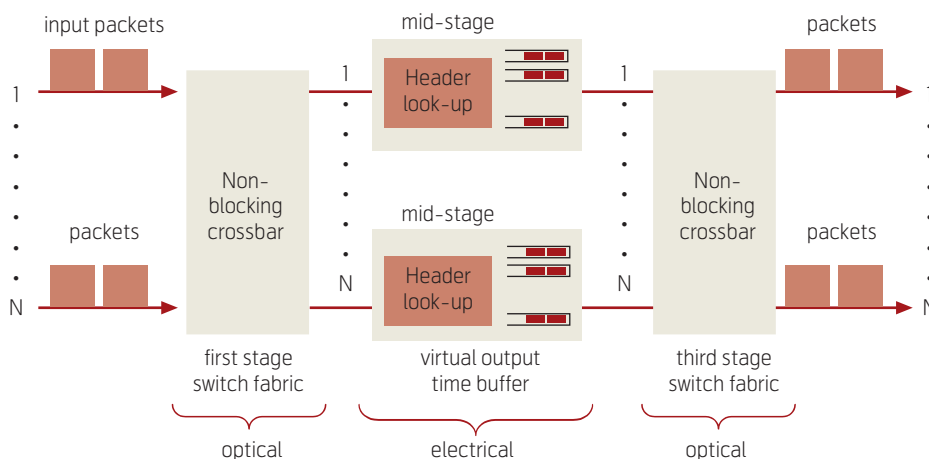


Figure 18 A schematic of a hybrid optical-electrical router, using both optical switching matrices and electrical switches – after [16]




Figure 19 A schematic of an OBS router where optical packets are forwarded all-optically whereas the optical label is processed electronically at each node

Optical functionality can be introduced at a separate layer and complement higher network layers, and/or it can be directly integrated and controlled by IP routers. Optical technology can also be employed to realize larger/higher throughput and more robust IP routers. It is not easy to foresee at this point which way things will evolve. The determining factors lie firstly with what is technologically a feasible and efficient solution but are also dictated by competition and industry consolidation conditions.

The winning technologies, architectures, protocols, and overall solutions require some work before they are more clearly identified. However, one reality is rather clear: Optical network technologies and network functionality are complementary to IP/MPLS and bring unique features to IP networks that no other technological alternatives can provide. Optical network elements and functionality will be introduced in IP networks in the not so far future.

References

- 1 Hagimoto, K. *Terabit network revolution in Japan*. Invited presentation, European Conference on Optical Communication, ECOC 2005, Glasgow, Sept 2005.
- 2 Heavy Reading. *Next Generation Broadband in Europe: the need for Speed*, 3 (5), March 2005.
- 3 Cahners in-Stat, RHK, Corning. *Fiber to the Home Forecast*. 2004. www.instat.com
- 4 Desurvire, E. *Erbium Doped Fiber Amplifiers, Principles and Application*. New York, John Wiley, 1994.
- 5 Bobbs, B, Goldstone, J A, Johnson M M. Phase-matched broadband Raman amplifier. *IEEE J Quantum Electronics*, 30 (8), 1896–1906, 1994.
- 6 Desurvire, E. Optical Communications in 2025. *European Conference on Optical Communication, ECOC 2005*, Glasgow, Sept 2005.
- 7 Morthier, G, Mingshan, Z, Vanderhaegen, B, Baets, R. Experimental demonstration of an all-optical 2R regenerator with adjustable decision threshold and “true” regeneration characteristics. *IEEE Photon. Tech. Lett.*, 12 (11), 1516–1518, 2000.
- 8 *Xtellus Dynamic Optics*. December 21, 2005 [online] – URL: <http://www.xtellus.com>
- 9 Tormen, M, Ghidini, S, Crespi, P, Nottola, D C A, Sardo, S, Crippa, D, Di Muri, M, Zuliani, G, Giacometti, F. Randomly Sampled Apodization in Bragg Gratings. Submitted for publication in *IEEE J Lightwave Tech.*, Oct 2005.
- 10 Ramaswani, R, Sivarajan, K N. *Optical Networks – a practical perspective*. Second edition. Morgan Kaufman publishers, 2002.
- 11 Durhuus, T et al. All-optical wavelength conversion by semiconductor optical amplifiers. *IEEE/OSA JLT/SAC*, 14 (6), 942–954, 1996. (Special issue on multi-wavelength optical technology and networks)
- 12 *Lucent-Optical*. December 21, 2005 [online] – URL: www.lucent-optical.com
- 13 Broberg, B, Rigole, P J, Nilsson, S, Andersson, L, Renlund, M. Widely tunable semiconductor lasers. *Optical Fiber Communication Conference 1999, OFC 99*, 21–26 Feb 1999.
- 14 Vigoureux, M, Berde, B, Andersson, L, Cinkler, T, Levrau, L, Fdez-Palacios, J, Jaeger, M. Multi-layer traffic engineering for GMPLS-enabled networks. *IEEE Communications Magazine*, July 2005, 44–55.
- 15 Otani, T. Nationwide field trial of GMPLS optical networking. *European Conference on Optical Communication, ECOC 2005*, Glasgow, Sept 2005.
- 16 Koonen, T et al. Novel Cost-Efficient Techniques for Microwave Signal Delivery in Fibre-Wireless Networks. *Proc. ECOC 2004*, Stockholm, September 5–9, 2004.

- 17 Neilson, D T et al. Fully provisioned 112x112 micro-mechanical optical cross-connect with 35.8 Tb/s demonstrated capacity. *Proceedings Optical Fiber Communication Conference, OFC 2000*, post deadline paper, 204–206.
- 18 McKeown, N. Scaling routers using optics. *European Conference on Optical Communication, ECOC 2003*, Rimini, Italy, Sept 2003.

For a presentation of Evi Zouganeli, please turn to page 2.

Optical fiber transmission: From research to commodity

AASMUND SUDBØ



Aasmund Sudbø is Professor at the University of Oslo and Unik

Optical communication technology has made the cost of information transmission around the world negligible. This article reviews the advances in technology that have made this possible. Fundamental physical limitations to further advances are also discussed.

Introduction

The Internet is a world where geographic distance has no meaning. We read emails from our next-door neighbor with the same (dis)interest as those coming from overseas. The spammers get to us, regardless of distance. When we visit a Web page, we do not think about where in the world it resides. We are annoyed that the computer screen rarely fills in immediate response to our mouse-click, but these tests of our patience are seldom related to the distance to the source. Most network operators now offer services where the charges are independent of the transmission distance, e.g. allowing us to freely click on a web-link irrespective of where in the world the link points to, and letting us make phone calls to any country in the world via the Internet, without being

charged for the time we spend talking. For the network operator, the cost of tracking the origin and destination of individual messages in order to charge customers fairly can easily exceed the cost of transmitting the messages.

This 'Distance for Free' feature is a quite recent development in telecommunications. Not many years ago, the income from long-distance and overseas telephone calls was the most important source of revenue for the telephone companies. By contrast, in Europe today, the phone companies usually charge more for a call to a mobile phone within walking distance than for a call to a regular phone on the other side of the Atlantic Ocean, and many consider mobile phone users to be the most important customers.



Photo: Steinar Karlsten

Figure 1 In January 2004, a submarine optical cable between Svalbard and mainland Norway was opened. The photograph shows the cable ship as it starts its cable laying operation from Svalbard in July 2003. Submarine cables are major engineering projects, where the entire cable and accompanying signal regenerators are made in a factory, before they are taken on board a cable ship

In our everyday life, we depend on numerous flows of bits, ones and zeros, to enable mobile phone conversations, to play music on MP3 players, to download text and images from the Internet, or to enjoy digital satellite TV. Whenever these bits need to be transmitted more than a few hundred meters along a cable, optical fibers are now used, unless copper wires happen to be in place. The reason is that a fiber can handle a much higher bit rate, i.e. the number of bits transmitted per second. Different services need different bit rates, and the bit rate needed for TV is more than a factor 100 greater than that for speech.

Optical fiber transmission has enabled a dramatic decrease in the price per transmitted bit in telecom networks, and is the main enabler of 'Distance for Free'. With this technology, optical fibers guide the light that carries the data from transmitter to receiver. Initially, fibers were used where the bit rate was highest, in the core of the telephone network, inside and between big cities. Over the years, optical fiber deployment has been extended towards the customers at the edges of the network. For a long time, an optical fiber communication terminal with its light source and receiver was an expensive element of the network, and a major issue for a network operator was how many customers should share a fiber terminal. Now, the technology has reached such a stage of maturity that it is both technically and economically feasible for each customer to have his own fiber terminal at home. The dominant cost of a network connection is the cost of getting the cable in place, especially if it has to go into the ground. Therefore, most of us will probably stay connected to the Internet via copper wires for a while, because we can use old telephone wires that are already in place. Nevertheless, the network operators in many countries, in particular USA, South Korea, and Japan, are now aggressively installing 'Fiber to the Home' (FTTH) to provide high-speed Internet connections, as shown in Figure 2 for Japan [1].

Basic principles of optical communication

A basic optical fiber transmission system consists of a transmitter, an optical fiber, and a receiver. The transmitter converts electrical current pulses representing transmitted bits (either the ones or the zeros) into light pulses, tiny flashes of light. The transmitter is mounted at the end of an optical fiber and sends the light along the center, the core, of the fiber. The material of choice for telecom fiber is silica glass. No solid material has been found that has a lower attenuation per unit length for the transmission of light. To get a flexible glass 'wire', the 'wire' needs to be a hair-thin fiber, and this fragile fiber is always

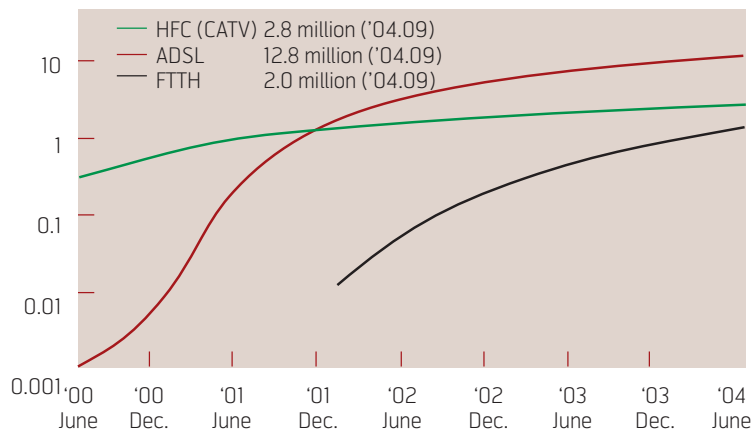


Figure 2 (From [1]) Number of Internet connections (in millions) as a function of time in Japan, for three different technologies, ADSL using telephone wires, HFC using television cable, and FTTH using optical fiber

wrapped in several layers of plastic and other polymer materials, for environmental and mechanical protection. The resulting structure is called a fiber cable. At the other end of the fiber the light emerges, and there is a receiver that converts the light pulses back to electrical current pulses.

As it travels along the fiber, the light is attenuated; a fraction of the optical power disappears for every kilometer traveled. If the distance between transmitter and receiver is so large that the receiver cannot distinguish the light pulses from noise, amplifiers have to be inserted along the line. The loss in optical fibers is exceptionally low, and systems exist where the light pulses can be recovered several hundred kilometers away from the transmitter. The allowed amplifier-free distance between the transmitter and the receiver depends on how sophisticated these units are.

The history of optical transmission is to a great extent the history of the optical transmitter, how it has been optimized through many steps of development to match the properties of silica fiber. In the historical review below we shall go through the most important steps. Key transmitter parameters are optical output power, signal bandwidth, optical frequency, and optical bandwidth. The most advanced transmitters today put out several milliwatts of power, transmit bit streams at 40 Gbit/s, have an optical bandwidth equal to the signal bandwidth, and the optical frequency can be adjusted and controlled with 5 digits accuracy.

Several users can share a fiber by what is called time division multiplexing (TDM), whereby access to the fiber is switched between the users. The more users, the less time per user, the shorter the light pulses for each user, and the higher the bit rate. The number of telephone calls that a single optical fiber can transmit is thus proportional to the bit rate. The bit rate that




Figure 3 Principles of an optical fiber transmission system based on wavelength division multiplexing (WDM) and erbium doped fiber amplifiers (EDFAs)

can be achieved is limited by what electronics can practically handle, typically 10 Gbit/s today.

One very important contribution of optics to communication technology is a way to have different signals carried by different optical frequencies on a single fiber. This way of combining optical signals is called wavelength division multiplexing (WDM)¹⁾. WDM technology allows tens of bit streams, each at 10 Gbit/s, to be carried by a single fiber. Each signal is assigned its own frequency interval, what is called a WDM channel. The International Telecommunication Union (ITU-T) has standardized channel center frequencies that are integer multiples of 50 GHz (0.4 nm), and systems with 100 GHz (0.8 nm) channel separation are common today. The center or carrier frequency of a transmitter depends on temperature, and 50 GHz channel separation is so small that transmitters with temperature control are required. We therefore speak of DWDM, dense WDM, as opposed to CWDM, coarse WDM, where the channel separation is so large that very inexpensive transmitters without temperature control can be used. For CWDM, a channel separation of about 2.5 THz (20 nm) between channels is standard.

The success of DWDM depends on the existence of an optical amplifier, where the incoming light is amplified directly and not converted into an electrical current, and where all the DWDM channels in a fiber can be amplified in parallel in a single amplifier.

Optical amplifiers today have a bandwidth of 4 THz or more, permitting simultaneous amplification of 80 DWDM channels or more. Figure 3 shows a transmission system incorporating all the elements discussed above.

Before optical amplifiers were available, signals were actually carried in parallel by separate fibers in a fiber cable, and amplified by electronic amplifiers in parallel, as shown in Figure 4. Each amplifier was actually a regenerator, containing a full receiver and a transmitter. The maximum bit rate permitted per fiber was then limited by what electronic amplifiers could handle. The advent of optical amplifiers and WDM increased the permissible bit rate per fiber a hundred-fold or more.

The receiver usually has electrical circuitry that performs the so-called 3R function: reamplification, reshaping, and resynchronization, so that the current pulses emerging from the receiver have the same magnitude as the ones sent into the transmitter, that they have the same shape, and that they emerge at regular intervals, equal to the intervals between the input pulses. In addition, the receiver may incorporate an optical preamplifier that amplifies the light pulses before they are converted to electrical pulses.

In what follows, I shall attempt to give a brief review of some of the most important developments in optical communication, and some fundamental limita-




Figure 4 Optical fiber transmission systems with several line amplifiers, before optical amplifiers and wavelength division multiplexing became available. Each fiber in the optical cable carried only one bit stream, and the bit rate was limited by what the electronic line amplifiers could handle

¹⁾ Conceptually, WDM is the same as what electrical engineers call FDM, frequency division multiplexing. Optical engineers have a long history of referring to the vacuum wavelength of a wave in cases where an electrical engineer would refer to the frequency.

tions. This is a tutorial review, not a literature review. The interested reader is referred to the many excellent books on the subject, a few of which are listed in the bibliography. A lot of material can also be found in the proceedings from a variety of scientific conferences. For more than 30 years, many of the most important developments in the field have been reported at the annual European Conference on Optical Communication (ECOC) and the Optical Fiber Communication Conference (OFC) in the USA.

Historical development

This historical review emphasizes practical problems that telecom engineers struggled with at the time, and how optical fiber communication (fibercom) technology helped solve many of these problems. Fibercom technology is the result of the concerted effort of a large international community of scientists and engineers, initially working in the research laboratories of the large telecom network operators around the world, and later also in the laboratories of the telecom equipment manufacturers. Funding for the development of this technology was made possible by large revenues for the network operators from long-distance phone calls. This source of revenue is rapidly disappearing for the network operators, ironically, because fibercom technology has made dramatic price reductions for long-distance calls possible.

The idea of optical fiber communication was conceived in 1966, with the suggestion that glass fibers could be used for signal transmission in the telephone network [2],[3]. At that time, the world already had a single large telephone network where it was possible to place a call between any two phones on the globe, and traffic was increasing steadily. Both frequency domain and time domain multiplexing were in use as means of having many telephone calls sharing a single line.

Used for signal transmission, electrical lines have the problem that as frequency increases, the signal loss per length of line increases, and the distance between amplifiers along the line has to be reduced. Larger-diameter lines can have lower loss, but bulky lines are no less expensive than many amplifiers. To the telephone network planners in 1966 it looked like electrical lines carrying telephone traffic between cities would require major investments.

Attenuation in optical fiber

The original idea in 1966 was to combine glass fibers for signal transmission with light emitting diodes as light sources, and silicon photodiodes for receiving the light [2],[3]. Light sources and receivers were available at the time, but the available glass material

was not suitable, having too much loss. Loss is quantified by the attenuation coefficient, the relative reduction in optical power per length traveled. High-quality glass at the time had an attenuation coefficient of several thousand dB/km, whereas telecom engineers required an attenuation coefficient of less than 20 dB/km, before they would be interested in glass fibers for signal transmission.

Glass scientists knew that the problem was in the impurities, and that a sufficiently pure glass would have the desired properties. After 4 years, scientists working for Corning Glass in the USA in 1970 reported fabrication of the desired optical fibers with an attenuation coefficient less than 20 dB/km. The fiber material was synthetic silica (chemical name: silicon dioxide), a chemical compound occurring in natural rock as the mineral quartz.

The fiber designers knew that to achieve such a low loss, the light in the fiber must be kept away from the fiber surface. Because light tends to stay where the refractive index is the highest, this was achieved through control of the refractive index or dielectric constant in the fiber. At the center of the fiber, in what is called the core of the fiber, the refractive index is higher than in the surrounding cladding. Most of the light is confined in the core, with the light intensity in the cladding decreasing exponentially with the distance from the center. The refractive index depends on the glass composition, and there are many materials that can be added to silica to change the refractive index without affecting the attenuation coefficient significantly.



Figure 5 A telecom fiber is a thread of glass, slightly thicker than a human hair. The beam of invisible infrared light guided along the center of the thread has a diameter of 0.01 mm



Figure 6 (From [6]) Two separate optical cables connect Svalbard to mainland Norway




Figure 7 (From [6]) Outline of the fiber cable installation that connects Svalbard to the mainland

The attenuation coefficient in pure silica decreases steadily as the wavelength increases from the visible range into the infrared. The first sources for optical communication had a wavelength of about 900 nm, corresponding to a frequency of about 330 THz. At this wavelength, the attenuation coefficient is about 2 dB/km, whereas it is less than 0.5 dB/km at 1300 nm, and has a minimum less than 0.2 dB/km near 1550 nm. The effect of impurities on fiber attenuation is greater for the longer wavelengths, however, and in the first fibers the attenuation increased with increasing wavelength. Improvements in fiber fabrication methods gradually eliminated all contributions to fiber attenuation from undesired impurities in the silica glass material, and optical fibers with an attenuation coefficient less than 0.2 dB/km at 1550 nm are now standard. The lowest reported coefficient is 0.16 dB/km [5], in a fiber with a pure-silica core.




Figure 8 (From [6]) Structure of the submarine cable connecting Svalbard to the mainland. In contrast, indoor optical cable typically contains one single optical fiber inside a soft plastic tubing with an outer diameter of about 2 mm

A modern telecom fiber usually has an outer diameter of 125 μm , and a lightguiding core with a diameter of 10 μm (Figure 5). Industrial standards for optical fibers and cables have been worked out, most of them by The International Telecommunication Union (ITU-T). The standard for the most common telecom fiber type is [4].

There is a very important commercial application for low-loss fibers, namely transoceanic cables. Such cables typically consist of more than a hundred identical cable sections, each terminated in a signal regenerator station designed to rest with the cable on the ocean floor. The cost of the regenerators represents a significant contribution to the cable system cost, and the number of regenerators needed is inversely proportional to the attenuation coefficient, because there is a maximum total attenuation permitted between regenerators, determining the maximum permitted distance between them. Transoceanic cables are major engineering projects, where the entire cable and accompanying regenerators are made in a factory, before it is taken on board a cable-laying ship like the one shown in Figure 1. The first transatlantic optical cable came in 1988²⁾. A state-of-the-art submarine optical cable, installed between Norway and

²⁾ The first transatlantic cable was a telegraph cable that went into service already in 1868 [7].




Figure 9 (From [6]) Submarine optical repeater used for the Svalbard cable, weighing 750 kg. By comparison, an optical transceiver designed for indoor use and the same bit rate (10 Gbit/s) is small and lightweight, much smaller than a matchbox

the arctic island of Spitsbergen in the fall of 2003, is described in [6]. The installation consists of two cables running separate routes along the ocean floor,

as shown in Figures 6 and 7. The structure used for most of the cable is outlined in Figure 8, and a submarine amplifier housing is shown in Figure 9.




Figure 10 a) Cross sectional structure of a typical semiconductor diode laser, seen from the light-emitting end. b) Mounting of the laser with electrical leads for driving current, and with a lens for coupling light from the laser into the 10- μm -diameter core of a single-mode fiber. The diode laser itself is a small grain of semiconductor, less than 1 mm in size, and the light is generated in a microscopic rod inside the grain with typical dimensions 0.1 μm by 1 μm by 300 μm




Figure 11 Dispersion makes the different frequency components of a 'one' bit travel with different speeds in the fiber. Each bicycle rider in the figure represents a frequency component, and different colors represent different frequencies, with correspondingly different speeds (from [10])

Sources

The first light sources used for optical fiber communication were light emitting diodes (LEDs) made of the material gallium arsenide, similar to the LEDs commonly used today in remote controls for TV sets and the like. The first diode lasers suitable as sources for optical communication were demonstrated in 1970³⁾ and exhibited a number of very desirable properties:

- a efficient conversion of electrical power to optical power;

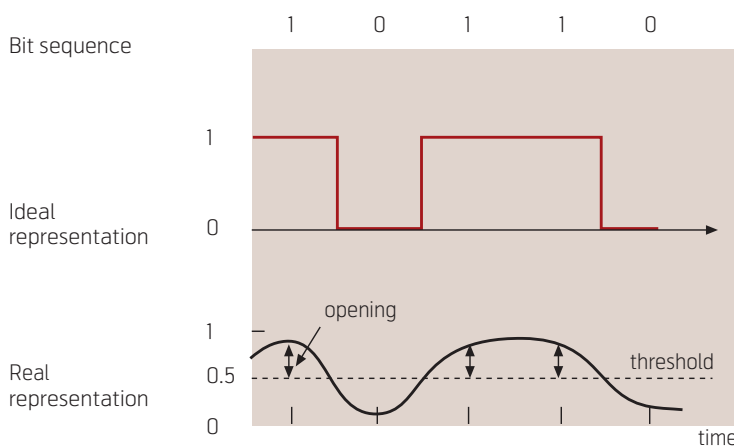


Figure 12 Top: Squarewave pulses representing the bit sequence '10110'. Bottom: Corresponding pulse shapes obtained after some pulse broadening in an optical fiber. The bit sequence is recovered by measurements at regular intervals of the pulse height relative to a threshold

- b being a point-like source, allowing almost all emitted power to be coupled into the optical fiber;
- c providing an optical wave with a well-defined amplitude, frequency, and phase (in other words, what is called coherent source);
- d easy control of the wave amplitude, frequency, and phase via the drive current;
- e rapid response of the amplitude, frequency and phase to variations in the drive current.

These are all properties that distinguish diode lasers from LEDs (see Figure 10).

The first diode lasers were made of the material gallium arsenide. Lasers made of this material cannot emit the wavelengths around 1550 nm needed to exploit the low attenuation of silica fiber. To develop sources and detectors with the required wavelengths, an entirely new semiconductor material was needed. The first material that met the requirements consists of the elements indium, gallium, arsenic, and phosphorous, and is called InGaAsP. The development of this material is one of the major milestones in the development of optical communication.

Pulse broadening and dispersion

In general, the light pulses that travel along optical fibers must be considered as having several components traveling at different speeds in the fiber, so that the pulses broaden as they travel. When the bit rate is high, the light pulses that represent bits are short, and little pulse broadening is allowed before succeeding pulses start to overlap (Figure 11 and Figure 12).

One pulse broadening mechanism is called multimode distortion, whereby rays traveling in different directions relative to the fiber axis have different speeds along the fiber. We distinguish between singlemode and multimode fibers. In singlemode fibers, there is only a single allowed direction for light rays in the core, parallel to the fiber axis. A multimode fiber, by contrast, is a fiber where several ray directions are allowed. Multimode fibers have larger core diameters than the 10 μm typical of a singlemode fiber. There is a significant cost associated with mechanical assembly of parts with the positional accuracy required by singlemode fiber. Initially, therefore, multimode fibers with larger cores were developed and standardized for optical communication. Today, much higher bit rates

³⁾ In the year 2000, the Nobel Prize in physics was awarded to Zhores Alferov, Herbert Kroemer, and Jack Kilby. Alferov and Kroemer were credited for their role in developing diode lasers for optical fiber communications [8][9]. Kilby was credited with giving us the integrated electronic circuit (IC), a very far-reaching contribution. ICs are everywhere, in every electronic device we use in our daily life, including the transmitters and receivers for optical fiber transmission.

are common, and the pulse broadening in multimode fiber is so severe that this type of fiber is now only used for very short distances, mostly indoor.

The most important pulse broadening mechanism in a singlemode fiber is dispersion, whereby light waves with different frequencies have different speeds in the fiber (see Figure 10). A diode laser has inherently a spread of frequencies (optical bandwidth) that is proportional to the bit rate of the emitted signal⁴⁾, so that the pulse broadening is proportional to the bit rate.

In the silica material that telecom fiber is made of, the dispersion vanishes at a frequency of about 230 THz, corresponding to a vacuum wavelength of 1300 nm. A source with a wavelength near 1300 nm is highly desirable, to eliminate pulse broadening from dispersion in singlemode fibers. The gallium arsenide material that was first used yields sources that emit wavelengths around 900 nm, and cannot emit near 1300 nm.

Historically, the development of the new material InGaAsP for the light source was spurred more by a desire to reduce the problems of dispersion than by a desire to have lower fiber attenuation. 1300 nm is now a commonly used wavelength in optical transmission systems over distances of up to a few tens of kilometers. The step from 900 nm to 1300 nm wavelength was a difficult one to complete, requiring development of a new material for sources and detectors. The step from 1300 to 1550 nm wavelength came later, and represented less of a challenge. The same standard fiber [4] is used at both wavelengths, and the InGaAsP material can provide sources and receivers both at 1300 nm and 1550 nm wavelength.

The introduction of 1550 nm as an operating wavelength reintroduced dispersion as a problem. It is possible, however, to move the zero-dispersion wavelength in single-mode silica fiber from 1300 nm to 1550 nm, by increasing the difference in index of refraction between core and cladding while decreasing the fiber core diameter. For a few years, this dispersion-shifted type of fiber was the most common type installed in the network, especially in submarine cables. When WDM technology was introduced, however, fibers with vanishing dispersion turned from an asset to a liability for the network operators (see below).

The pulse broadening induced by dispersion in a fiber can be reversed if the pulse is sent through a second fiber with dispersion of the opposite sign, and a suitable length. Today, such dispersion compensators are

standard, made from specially designed dispersion shifted fiber.

Polarization

Light has a polarization, meaning that there is a field vector direction associated with an optical wave, perpendicular to the direction of propagation. In an optical fiber in a cable, the polarization changes randomly along the fiber. Furthermore, normal cable movements induced by changing weather conditions and temperature are sufficient to change the output polarization completely, even if the input polarization is stable. So in general, the polarization of light emerging from an optical fiber is unpredictable. Therefore, components for optical fiber communication must be polarization-insensitive.

For a given input polarization, the output polarization from an optical fiber will in general depend on optical frequency. If the frequency dependence of the output polarization is so strong that the output polarization is significantly different for the different frequency components of the signal, we have to consider a pulse broadening mechanism called polarization mode dispersion (PMD). The PMD in installed optical fibers has always been so low that it did not represent a problem until the bit rate 10 Gbit/s appeared in the network. Cable manufacturers now routinely deliver cable where the PMD is so low that operation at 40 Gbit/s is guaranteed, but much of the older cable cannot be used with 40 Gbit/s, and some cables also have trouble with 10 Gbit/s [11].

Optical amplifiers

In an optical amplifier, the incoming light is amplified directly. 1550 nm, the wavelength of minimum attenuation in optical fibers, happens to be a wavelength for which an excellent optical amplifier can be built. Of particular importance is the fact that the amplifier may be realized as an optical fiber with a small concentration of erbium atoms in the core. The erbium doped fiber amplifier (EDFA), first demonstrated in 1987 [12], [13], was one of the most significant developments in the history of optical fiber communications. This optical amplifier exhibits a number of very attractive properties for optical fiber communication, such as

- a. polarization independence;
- b. linear response;
- c. an amplifier bandwidth of 4 THz;
- d. ideal noise performance;
- e. being a silica fiber, easy to connect to transmission fibers;

⁴⁾ The constant of proportionality between bandwidth and bit rate is of the order of one, and it can be reduced with careful engineering.




Figure 13 Optical power as a function of distance along an optical fiber transmission line, for a distributed fiber Raman amplifier (FRA) and for an erbium doped fiber amplifier (EDFA). The two cases shown have the same minimum optical power and hence approximately the same signal-to-noise ratio (SNR) after the amplifiers. Evidently, the FRA allows the greater spacing between amplifiers. (See the section about quantum noise for an explanation of SNR)

f. being driven by a diode laser very similar to the sources used for optical transmission.

Diode lasers similar to the ones used as optical sources can also serve as optical amplifiers, but they are not suitable as amplifiers along an optical fiber transmission line, because they are polarization-dependent and nonlinear.

An alternative to the EDFA is the fiber Raman amplifier (FRA). Both amplifier types use diode lasers as pump lasers, i.e. input power to the amplifiers⁵⁾. The FRA does not need a special fiber like the EDFA, the transmission fiber can be used for amplification. The FRA pump laser is typically mounted in front of the receiver or regenerator, sending pump light along the transmission fiber towards the source, resulting in amplification of the incoming light before it reaches the receiver, and permitting the distance between source and receiver to be increased, compared to a

system with an EDFA. The FRA is a truly distributed amplifier, having the gain distributed along several kilometers of transmission fiber. An EDFA, by contrast, is a few meters of fiber just in front of the receiver, by comparison a lumped amplifier. Figure 13 shows the distribution of optical power between transmitter and receiver for the two types of amplifiers.

Whereas electronics today can provide an amplifier bandwidth of a few tens of GHz, the EDFA provides 4 THz of bandwidth in a single amplifier, allowing many wavelength division multiplexed (WDM) channels to be amplified with a single amplifier. Optical filters are needed in WDM to assemble optical waves with different frequencies in one fiber and separate them again at the end of the fiber. An optical filter industry existed long before optical communication, and was ready to provide filters for WDM as soon as the EDFA came on stage.

Sources for wavelength division multiplexing (WDM)

Before WDM was introduced, the key output parameter of the optical source was optical power. With WDM, the spectral properties are equally important. The center frequency has to be accurately locked to one of the allowed output frequencies. The center frequency of a diode laser depends on temperature, so temperature stability is needed in WDM sources. The output power has to be kept within a well-defined spectral band around the center frequency, and the output in other parts of the spectrum has to be suppressed.

The first WDM sources had fixed center frequencies, so that for example a 40-channel WDM system had 40 transmitters that were distinct and could not be interchanged. Some WDM source suppliers have now started to offer general purpose tunable sources, having an adjustable center frequency. These sources incorporate special diode lasers⁶⁾ with several control currents in addition to the one controlling the output power [14]–[16]. An example is shown in Figure 14. The first application of these tunable sources is as all-purpose spares for fixed-frequency sources.

An important application of tunable lasers is for switching, where several output fibers exist for each input signal, and where the output fiber is selected by the carrier frequency of the signal. The simplest opti-

⁵⁾ Before the advent of the EDFA, the FRA was considered the most promising optical amplifier for fiber communication applications. An FRA needs at least ten times more pump power than an EDFA, however, and the high pump power required by an FRA is a major disadvantage.

⁶⁾ Even though the center frequency does depend on temperature, the temperature sensitivity is not sufficient to permit the frequency to be controlled by temperature, considering that the entire 4 THz bandwidth of an EDFA must be covered.

cal switch of this kind is a single tunable laser together with a standard demultiplexer (see Figure 3) for WDM. Tunable lasers offer two important features: They are fast, and a single laser can switch between many outputs. There is at the moment no other optical switching technology that features this combination. The application of tunable diode lasers to switching has been an active research area in recent years, also in Norway [17], [18].

Fiber nonlinearities

With the optical power in the fiber highly concentrated in a core with 10 μm diameter, a few milliwatts of power is sufficient for optical pulses to be distorted by nonlinear processes in the fiber, i.e. processes whereby the light emerging from the fiber has frequency components different from those that went into the fiber. Such processes are in general detrimental to signal transmission in the fiber⁷⁾. The nonlinearly generated optical power is proportional to the square of the transmitted power (as long as the generated power is much smaller than the transmitted one), so the ratio between the generated and the transmitted powers increases with transmitted power.

The nonlinear process most easily observed in optical fibers is stimulated Brillouin scattering. It has the striking effect that almost all of the light sent into the fiber is reflected back to the sender if the optical power exceeds a threshold. The threshold is lowest for narrowband light, having a bandwidth of less than about 100 MHz. The threshold is inversely proportional to the optical bandwidth, so stimulated Brillouin scattering can be suppressed by having a sufficient WDM channel bandwidth well above 100 MHz.

Several nonlinear processes lead to cross-talk between WDM channels. Given that the signals in the different WDM channels are uncorrelated, the cross-talk from the other WDM channels appears as random noise in each channel. This noise grows so rapidly with increasing channel power that it puts a limit on the channel power, where a further increase in the power results in a reduced signal-to-noise ratio (SNR) in the receiver.

The dispersion-shifted fiber that was deployed extensively for single-channel systems at 1550 nm, suffers so severely from cross-talk that this type of fiber can in general not transmit more than one WDM channel. Fiber dispersion reduces the nonlinear cross-talk between WDM channels, however, and the older fiber [4] originally designed for use at 1300 nm has




Figure 14 Top view of a rapidly tuneable laser with four independent drive currents: Output power (Gain), length (Phase) adjustment, left and right reflectors. MMI: Y coupler based on multimode interference. Each reflector is a modulated Bragg grating, and a change in either reflector current results in a change in the output frequency. Adapted from [15]

turned out to be quite well suited for upgrades involving a transition from a single channel to several WDM channels around 1550 nm. In this wavelength region the fiber has the dispersion needed to suppress the nonlinear WDM channel cross-talk. Single mode fibers that were installed in the telecom network more than 15 years ago can still be used with new fiber terminals, even though the bit rate, the operating wavelength, and the number of WDM channels all have changed over the years.

In general, nonlinear signal distortions tend to accumulate through optical amplifiers. This statement holds for self-induced signal distortions in each WDM channel as well as for WDM channel cross-talk. Hence, for an optical transmission line with several amplifiers, nonlinear optical processes put a limit on the amplifier output power that is inversely proportional to the number of optical amplifiers along the line.

The accumulation of noise and nonlinear distortions through optical amplifiers limits the distance that a submarine cable can reach if it incorporates only optical amplifiers and no regenerators. The maximum distance that systems designers find when they pull tape measures around the globe in search of the longest submarine cable routes between the cities of the world is 10,000 km. Considerable effort has been made by the submarine cable industry to design regenerator-free WDM transmission systems capable of such a reach. Present-day conference contributions report on how to achieve this target by combining all the engineering tricks discussed in this article, and more.

⁷⁾ Nonlinear optical processes in the fiber can also be exploited together with dispersion to improve signal transmission, in particular in so-called soliton transmission systems.

Fundamental limitations

In the development of optical communication, engineers have encountered a number of fundamental physical limits to technology development. The first limit, reached about 20 years ago, was the silica material attenuation coefficient. The attenuation in commercially available optical fiber at 1550 nm is now as low as the silica material permits. Over the years, a significant research effort has gone into a search for a fiber material with a lower attenuation coefficient than the 0.16 dB/km of silica, without any success. It seems that we will have to settle for silica fiber for optical communication in the telecom network for the foreseeable future.

The most promising alternative to solid silica so far is hollow silica, where the light travels in holes running along the fiber instead of in the silica glass material. Hollow-core fibers with an attenuation coefficient below 1 dB/km have recently been fabricated [19]. Further work on this fiber type will certainly decrease the attenuation, but it is still an open question if the attenuation coefficient can be pushed below the 0.16 dB/km of solid silica.

Quantum noise

Optical power does not flow continuously, but as particle-like photons. The energy of each photon is inversely proportional to the wavelength, and equal to $1.28 \cdot 10^{-19} \text{ J} = 0.128 \text{ aJ}$ for the 1550 nm wavelength. The flow of optical power in an optical fiber is a random flow of photons. The attenuation in the fiber yields a flow that exhibits shot noise and is very accurately modeled by Poisson statistics, and the signal-to-noise ratio (SNR) is then equal to the number of photons per bit. Several photons are required to represent the presence of a pulse of light with a reasonable certainty, and the standard definition of 'reasonable certainty' in optical communication is with an error probability of less than 10^{-12} , requiring on average 27 photons per light pulse. So the quantum nature of light implies that an optical receiver needs a minimum number of photons per bit to distinguish between zeros and ones. In other words, the power that the receiver needs (the receiver sensitivity) is proportional to the bit rate.

The signal-to-noise ratio (SNR) from any optical fiber transmission line can never be greater than that given by the shot noise at the point along the line where the transmitted power is at a minimum. Typically, there will be one single minimum, usually just before an EDFA, that dominates the noise. If there is a distributed Raman amplifier along the line, the noise-determining minimum is a few kilometers from the pump laser, as shown in Figure 13. Transoceanic submarine transmission lines usually consist of many

identical amplifier units connected by identical cable sections, and then there is more than one power minimum that contributes to the noise at the end of the line.

To detect low-power optical signals, an optical preamplifier (usually an EDFA) is incorporated into the receiver. The laws of physics require an EDFA to add noise, an amount at least equal to the shot noise. Raman amplifiers have the same limitation. This is a serious problem for transoceanic optical transmission, where the signal typically has to go through more than 100 submarine amplifiers to reach the shore, and where each stage adds noise. The signal-to-noise ratio at the end of an optical transmission line is inversely proportional to the number of optical amplifiers along the line, and this limits how far an optical signal can be sent with the help of EDFAs. Nevertheless, there are now several transatlantic optical cables in operation with WDM transmission and EDFAs.

It is possible to build an optical amplifier with less noise than an EDFA, but only if the amplifier is made sensitive to the phase of the optical wave. Such amplifiers have been realized, but they are finicky laboratory animals not suitable for ocean floor installations [20].

Information theory

With clever engineering, the number of photons per bit required by the receiver can be reduced from the estimate given above. According to the information theory of Shannon [21], [22], [23], the maximum bit rate that a receiver can detect is proportional to the signal bandwidth, and the constant of proportionality is given by the signal-to-noise ratio. For a shot-noise-limited optical signal the signal-to-noise ratio is equal to the number of photons per bit. Information theory allows optical receivers to operate with as little as 1 photon per bit, for a receiver bandwidth equal to the bit rate [21], [22]. To approach this limit, one has to code the bits using phase shifts of the optical wave [22], [24] and use error-correcting coding (ECC) [25], [23].

Figure 15 illustrates how bit errors can occur in a digital receiver when the signal-to-noise ratio is too small. Bit errors can be corrected with error-correcting coding (ECC), a highly developed engineering tool with a wide range of applications in computers and communication. The basic idea in error correcting coding is that a longer bit sequence is used to represent a shorter bit sequence, so that many different longer bit patterns are permitted representations of the same shorter bit pattern. Then, the shorter bit pattern can be recovered from the longer one, even if it contains a few corrupted bits. Sensitivities within a fraction of a dB from the information-theoretical limit




Figure 15 Phase coded digital signal with noise. The coding shown is called quadrature phase shift keying (QPSK). Each dot in the two diagrams represents a measurement of the amplitude and phase of the wave at the receiver, plotted in polar coordinates. Each dot represents two sent bits, with the dot color representing the sent bit values: red = 11, black = 10, green = 00, and blue = 01. The difference between the left and right diagrams is a factor of ten in signal-to-noise ratio. A measurement result outside the quadrant where it belongs represents a bit error in the receiver. Whereas the measurements cluster nicely together and no bit errors can be seen when the signal-to-noise ratio is large (left), bit errors occur when the signal-to-noise ratio is small (right). (The figure shows results of computer simulations, not actual measurements)

have been realized experimentally for wireless RF communication systems incorporating ECC, and optical communication systems are not far behind [26].

Nonlinear processes

The signal-to-noise ratio can be increased by an increase in the transmitted optical power. We are not free to increase this power, however, because of the nonlinear effects in optical fibers. With clever engineering we can reduce the nonlinear WDM channel cross-talk, but we are still left with the nonlinear in-channel pulse distortion that puts a limit on the channel power that can be launched into an optical fiber.

Furthermore, there is stimulated Raman scattering in the fiber, a nonlinear process whereby optical power is transferred from short-wavelength channels to long-wavelength channels. As already discussed, this process can be exploited in fiber Raman amplifiers. On the other hand, the process puts a limit on the total power in all the WDM channels in the fiber, a limit that is inversely proportional to the optical bandwidth occupied by the channels. This limit is more important than the limit on single-channel power, for systems with tens of WDM channels spread over the gain bandwidth of the EDFA.

It is the nonlinear properties of silica that matters. Silica is the material of choice for optical telecom fibers, because it has the lowest attenuation coefficient of all known materials. It is actually a very linear optical

material, and there are very few engineering tricks that can be used to reduce the nonlinear signal distortion in optical fibers. One thing that can be done is to reduce the light intensity, by increasing the area of the fiber core where the light propagates. Increasing the core area without increasing the fiber attenuation coefficient is not straightforward, however [5]. A very promising recent development is the low-loss fiber reported in [19], a hollow-silica fiber type having an optical nonlinearity orders of magnitude smaller than standard silica fiber.

Concluding remarks

Almost all of the cities of the world are now connected by optical fiber cable, and in Norway we all depend on optical fibers in our everyday life. Optical fiber transmission has dramatically lowered the cost of sending bits around the world, to the extent that the cost of tracking the bit streams and billing individual senders can easily exceed the cost of actually transmitting the bits, regardless of their destination.

Telecom fibers have turned out to have a service life and upgrade potentials far beyond the wildest hopes of the original fiber designers, and optical fiber cables have proven to be a very valuable long-term investment for most network operators. On the one hand, optical communication is a mature field in telecom engineering. On the other hand, the exploitation of installed optical fibers for steadily growing bit rates

and increasing number of WDM channels remains a serious challenge for telecom engineers and scientists. The telecom business boom did give us an abundance of dark (unused) fiber in the network, but demand for bandwidth in the fibers is growing rapidly, and the network operators have to resume their tradition of innovative exploitation of the fiber infrastructure they own.

Acknowledgements

The author would like to thank Evi Zouganeli and Martin Nord at Telenor for their careful reading and comments that led to a significant improvement of the manuscript.

References

- Shinohara, H. Broadband Expansion in Japan. Plenary talk at the *First Joint Optical Fiber Communication Conference and National Fiber Optic Engineers Conference (OFC/NFOEC 2005)*, Anaheim, California, USA, March 6-11, 2005. (ISBN 1-55752-784-9)
- Kao, K C, Hockham, G A. Dielectric-fiber surface waveguides for optical frequencies. *Proc. IEE*, 133 (3), 1151–1158, 1966.
- Werts, A. Propagation de la lumiere coherente dans les fibres optiques. *L'Onde Electrique*, 46, 967–980, 1966.
- ITU-T. *Characteristics of a single-mode optical fibre and cable*. Geneva, International Telecommunication Union. ITU-T Recommendation No. G.652.
- Tsukitani, M, Matsui, M, Nagayama, K, Sasaoka, E. Ultra low nonlinearity pure-silica-core fiber with an effective area of 211 μm^2 and transmission loss of 0.159 dB/km. *European Conference on Optical Communication (ECOC 2002)*, Copenhagen, Denmark, Sep 8–12, 2002, Paper 3.2.2.
- Gjesteland, E. Technical solutions and implementation of the Svalbard fibre cable. *Teletronikk*, 100 (3), 140–152, 2004.
- Lehne, P H. The transatlantic telegraph cable of 1858 and other aspects of early telegraphy. *Teletronikk*, 100 (3), 203–208, 2004.
- Alferov, Z I. The History and Future of Semiconductor Heterostructures from the Point of View of a Russian Scientist. *Physica Scripta*, T68, 32–45, 1996.
- Kroemer, H. Band Offsets and Chemical Bonding: The Basis for Heterostructure Applications. *Physica Scripta*, T68, 10–16, 1996.
- Bigo, S. Multi-terabit/s optical transmission: where do we stand? *The Norwegian Electrooptics Meeting*, Tønsberg, 2–4 May 2004.
- Tessmann, H J, Breuer, D, Foisel, H M, Reiner, H, Cremer, H. Investigation of the relation between PMD coefficient and cable installation year in Deutsche Telekom fiber plant. *Proc. NIST Symp. Optical Fiber Measur.*, Boulder, Colorado, Sep 24–26, 2002, 187–190.
- Mears, R J, Reekie, L, Jauncey, I M, Payne, D N. Low-noise erbium-doped fibre amplifier operating at 1.54 μm . *Electronics Lett.*, 23 (19), 1026–8, 1987.
- Desurvire, E, Simpson, J R, Becker, P C. High-gain erbium-doped traveling-wave fiber amplifier. *Opt. Lett.*, 12 (11), 888–90, 1987.
- Jayaraman, V, Chuang, Z M, Coldren, L A. Theory, design and performance of extended tuning range semiconductor lasers with sampled gratings. *IEEE J. Quantum Electron.*, 29, 1824–34, 1993.
- Öberg, M, Nilsson, S, Streubel, K, Bäckbom, L, Klinga, T. 74 nm tuning range of InGaAs/InP vertical grating assisted rear sampled grating reflector. *IEEE Photonics Technology Letters*, 5, 735–8, 1993.
- Wesström, J O, Sarlet, G, Hammerfeldt, S, Lundqvist, L, Szabo, P, Rigole, P J. State-of-the-art performance of widely tunable modulated grating Y-branch lasers. *Optical Fiber Communication Conference (OFC 2004)*, Los Angeles, CA, Feb 2004, paper TuE2.
- Buus, J, Morthier, G, Sarlet, G, Baets, R, Zouganeli, E, Mlonyeni, A, Sudbø, A, Broberg, B, Rigole, P J, Ishii, H, Yoshikuni, Y, Borley, S, Rawsthorne, J, O'Dowd, R, Mulvihill, G. Application and control of widely tunable lasers in HD WDM networks. SPIE Annual Meeting 2000, July 30 – August 4, *Proceedings of SPIE*, 4111.
- Feng, B, Stol, N, Zouganeli, E, Andreassen, R O, Solem, A, Haugen, R B, Sudbø, A, Helvik, B E. Direct Comparison between Optical Circuit and Optical Packet/Burst Switching Using Simula-

- tions. *Opto-electronics and Communications Conference and International Conference on Optical Internet (OECC/COIN 2004)*, paper 14A2-2, Yokohama, Japan, July 12–16, 2004.
- 19 Tsujikawa, K, Kurokawa, K, Tajima, K, Nakajima, K, Matsui, T, Sankawa, I. Penalty-free 10 Gb/s Transmission in 1.0 μm Band over 24 km Low Loss PCF. *European Conference on Optical Communication (ECOC 2005)*, Glasgow, UK, Sep 25–29, 2005, Paper Tu4.4.2.
- 20 Marhic, M E, Wong, K K-Y, Kalogerakis, G, Kazovsky, L G. Toward Practical Fiber Optical Parametric Amplifiers and Oscillators. *Optics and Photonics News*, 15 (9), 21–25, 2004.
- 21 Shannon, C E. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27, July and Oct 1948.
- 22 Kahn, J M, Ho, K-P. Spectral efficiency limits and modulation/detection techniques for DWDM systems. *IEEE J. Sel. Top. Quantum Electronics*, 10 (2), 259–272, 2004.
- 23 Øien, G E (ed.). Information Theory and its Applications. *Teletronikk*, 98 (1), 2002.
- 24 Olsen, T, Tveito, O. High performance 140 Mbit/s heterodyne PSK receiver employing a PLL synchronous demodulator. *European Optical Communications and Networks Conference, EFOC/LAN 1992*, 382–4, IGI Europe, Boston, MA, USA.
- 25 Sab, O A. Forward error correction techniques. *Optical Fiber Communications Conference (OFC 2003)*, *Trends in Optics and Photonics Series*, 86, IEEE Cat. No. 03CH37403, 391.
- 26 Mizuocho, T, Miyata, Y, Kobayashi, T, Ouchi, K, Kuno, K, Kubo, K, Shimizu, K, Tagami, H, Yoshida, H, Fujita, H, Akita, M, Motoshima, K. Forward error correction based on block turbo code with 3-bit soft decision for 10-Gb/s optical communication systems. *J. Sel. Top. Quantum Electron.*, 10 (2), 376–386, 2004.

Bibliography

Agrawal, G P. *Lightwave Technology: Telecommunication Systems*. John Wiley & Sons, 2005. (ISBN 0471215724)

Agrawal, G P. *Lightwave Technology: Components and Devices*. John Wiley, 2004. (ISBN 0471215732)

Agrawal, G P. *Nonlinear Fiber Optics*, 3rd edition. Academic Press, 2001. (ISBN 0120451433)

Buus, J, Amann, M-C, Blumenthal, D J. *Tunable Laser Diodes and Related Optical Sources*, 2nd edition. Wiley, 2005.

Green, P E Jr. *Fiber Optic Networks*. Prentice Hall, 1993. (ISBN 0133194922)

Hecht, J. *Understanding Fiber Optics*, 5th edition. Prentice Hall, 2005. (ISBN 0131174290)

Kaminow, I P, Li, T (eds.). *Optical Fiber Telecommunications IV B*. San Diego, Academic, 2002.

Ramaswami, R, Sivarajan, K. *Optical Networks: A Practical Perspective*, 2nd edition. Morgan Kaufmann Publishers, 2001. (ISBN 1558606556)

Aasmund S. Sudbø received an M.E. in physics from Norges Tekniske Høgskole, Trondheim, Norway, in 1975, and a PhD degree in chemical physics from the University of California, Berkeley, in 1979. For two years, he continued this work as a Postdoctoral Researcher with IBM Research Laboratories, New York. From 1981 through 1994, he worked on optical fiber communication at Telenor R&D, Kjeller, Norway, one of the years at Bell Laboratories, Murray Hill, NJ. He also was Adjunct Professor in the Department of Physics, University of Oslo. Since 1995 he has been Professor of Optoelectronics at the University of Oslo and at UniK, University Graduate Center at Kjeller, Norway. From 2002 to 2003 he was a Visiting Scholar at Stanford University, Stanford, CA, where he worked on optical MEMS. He has been a member of the technical program committees of many conferences on optical communication, in particular the annual European Conference on Optical Communication (ECOC), where he was the program chair in 1996.

email: aas@unik.no

A new architecture for optical networks

DAVID PAYNE AND RUSSELL DAVEY



David Payne is Head of the Future Networks team at BT, Martlesham Heath, UK



Russell Davey is Head of BT's R&D activities in optical systems

This paper briefly reviews the history of optical access, in particular the Passive Optical Network (PON) architecture (still the most promising candidate for the mass-market solution), when and how it might evolve further into the future and maybe finally become the access network technology of choice for all service providers by tackling the problem of bandwidth growth costs outstripping revenue growth.

Introduction

During the past few years there has been massive growth in the take-up of broadband with little sign of any slow-down in demand. There is also a clear trend towards higher bandwidths with both Japan and Korea adopting fibre-to-the-home (FTTH) strategies to deliver 100 Mbit/s capability, while in the US some major FTTH deployments are planned over the next few years. Although Europe has not yet embraced fibre as the access technology of choice, serious consideration is now being given to fibre access solutions for future deployment with small-scale deployments and trials taking place.

The relentless progress in consumer technology has also continued unabated with processor speeds continuing to increase roughly in line with Moore's Law. At the same time very high capacity low cost storage has become readily available to match this increase in processing power. Video and imaging services via mobile phones have also seen remarkable growth showing there is real latent demand for new innovative services, and some of these really will require higher bandwidth connectivity as expectations rise and frustrations with delay and low quality limit usage. Image and video file transfer, richer content on the web, increasing popularity of online gaming (not just in Korea and Japan) and the need to back up large amounts of personal data will all drive the demand for increasing access speeds and network capacity.

As operators and service providers move forward trying to meet these ever-increasing demands for high capacity services, a major obstacle to sustainable profitability arises. The issue is that the cost of installing the additional network capacity to meet the predicted growth in demand for bandwidth can easily exceed the subsequent growth in revenues.

Background

Bandwidth growth exceeding revenue growth is not a new phenomenon but has previously been balanced by the normal price decline of equipment that occurs in any industry as product volumes increase. The problem for the telecommunications industry with a

broadband future, particularly when moving beyond today's DSL speeds and contention ratios, is that many of the projected possible service and growth scenarios produce such large bandwidth growths that traditional price declines will not be sufficient to keep cost growth in line with growth in revenues.

The basic problem is that today's networks are built upon a large base of electronic equipment which only price declines in line with the traditional price declines seen within the electronics industry. Historically, electronic goods and products have typically declined with an 80 % learning curve; that is, for every doubling of product volume the unit price drops to 80 % of the previous price. Over the past decade bandwidth has followed the same relationship. As network capacity increases the cost of providing that capacity falls in line with the price decline of the supporting electronic technology and there is no evidence that the overall situation is changing.

If the cost of building greater capacity in traditional networks cannot be solved by the usual expected equipment price declines then either future broadband growth is going to be curtailed or new network architectures that are less dependent on the basic electronics are going to be required. These new architectures will need to eliminate significant tranches of equipment by removal of nodes and minimisation of network port cards and internal interfaces – like BT's 2CN vision – but at the same time enabling the required network capacity growth.

The interesting question for optical networking and optical technology is, can it produce faster unit bandwidth price declines than traditional network developments? There is no reason to suppose that the price of optical technology itself will decline any faster than electronic technology. Many of the manufacturing techniques used within the optical technology industry have evolved from or are developments of processes used within the electronics industry.

The opportunity offered by optical technology and systems is the prospect of new radical approaches to network architecture and design, enabled by exploit-

ing the capacity, reach and flexibility of optical networking to minimise the numbers of tiers, nodes and interface ports required in large scale networks. One approach to this, discussed in this paper, is the design of a network that combines the access network and outer-core/backhaul/metro network into one common high capacity network that can bypass and eventually eliminate the current access nodes. At the same time it reduces the number of customer ports and network interfaces required. This combination of node elimination and port elimination can lead to significant capital expenditure (CapEx.) savings and because of the overall simplification of the network and reduction in network intervention points, there are also major ongoing operational expenditure (OpEx) savings.

Optical access and all optical networks have been topics in the optical communications world for over two decades, yet it has still not happened on any significant scale, compared to the size of legacy networks. For those working in the field it has often seemed that the fibre-to-the-home/office revolution was “just around the corner” albeit the corner being two to five years away and never appearing to come closer.

There have been a number of good reasons for this, discussed briefly later, but primarily FTTH is an almost revolutionary step for incumbent operators to take. The new competitive operators did not really tackle the access space as there was thought to be much richer and easier pickings in the core/backbone and outer-core/metro spaces. Although many of these networks were dominated by optical technology and systems they were largely simple point-to-point transmission systems with the networking technology still being mainly in electronic switching and routing nodes with overall topologies following conventional architectures. The only other operators that built broadband access were the cable operators and at the time of their major build the access fibre solutions were very immature and not readily available. They installed coaxial cable systems followed later by hybrid fibre coax (HFC). The coax remained as the customer premises terminating technology, fibre only being used in the longer distance feeder part of the network. Now like the incumbent telecom operators who have a valuable legacy of copper, the cable operators have a valuable legacy of coax. So what is the opportunity for access fibre?

This paper will briefly review the history of optical access, in particular the Passive Optical Network (PON) architecture (still the most promising candidate for the mass-market solution), when and how it might evolve further into the future and maybe finally become the access network technology of choice for all service providers by tackling the problem of bandwidth growth costs outstripping revenue growth.

Brief history of fibre in access

Fibre-to-the-home and -business was a consideration from the earliest days of optical fibre technology development. In the late 1970s point-to-point replacement of copper by fibre was being considered as a way of delivering broadband (mainly video) services to customers. These early systems were predicated on multimode fibre technology, the only viable solution at that time. More sophisticated versions of these early systems were also studied including remote electronic multiplexers sited at the street cabinet location with fibre feeders from the exchange and point-to-point optical links from the cabinets to the customers. This was not so different from today's fibre-to-the-cabinet proposals using VDSL technology which uses advanced modulation and coding techniques to exploit the legacy copper from the cabinet to the customer rather than installing new fibre all the way from the remote nodes to the customer site.

The first considerations of a passive optical network approach for the access network was around 1982 when single mode fibre technology was being seen as a possible new way forward for optical communications. Single mode fibre offered many advantages compared to multimode fibre, much greater bandwidth being one of the more obvious. However, probably of equal importance for the future evolution of optical networking was the ability to make high performance optical components.

One component that became available fairly quickly with the arrival of single mode technology was the fused fibre directional coupler. These splitters or couplers can be cascaded and any size of splitter or star coupler can be made. These optical splitters were the key component behind the Passive Optical Network concept and can also be manufactured using planar technologies.

In the first half of the 1980s the passive optical network concept was centred on wavelength switched networks. These used star couplers to interconnect network terminations and wavelength selection to route paths [1] across the network. At the same time ideas of using the couplers as simple passive splitters for broadcasting television signals were also being considered [2].

All this early thinking was technology led, it was thought that these networks would be used to deliver broadband services, with video being the obvious application. Little thought was given to revenue streams or the business environment. It was thought that if the bandwidth capability was delivered services and revenues would follow.

In the mid 1980s BT became interested in the possibilities offered by optical access and this led to a refocusing of the passive optical network approach. The operational units brought a much needed business focus to the research and challenged the research teams to develop a system that could be economical for telephony. This was a service with a known revenue structure as opposed to the unknown revenues from future broadband services. This approach became known as a “telephony entry strategy” and led to the invention and development of the “TPON” (Telephony over Passive Optical Network) system [3][4]. Several manufacturers became heavily involved in the development including Fujitsu, Alcatel and Nortel. In the end BT deployed a Fujitsu system to several tens of thousands of customers.

TPON was TDM based and the early system had a limited bandwidth of 20 Mbit/s, adequate for telephony and ISDN but not for broadband. Broadband would be added later, as an upgrade, by the addition of extra wavelengths. To facilitate this a blocking filter was added to the TPON ONUs, which only passed the original TPON wavelength and blocked all others, enabling additional wavelengths to be added to the PON at a later stage without disturbing the original telephony only customers. However, because the system was never rolled out on any significant scale the upgrade system (called BPON at the time) was never developed into a commercial product. At the end of the eighties BT was developing an ATM version of passive optical networks called APON. Also around this time optical amplifiers were emerging as a viable network component. At the end of the eighties and in the early nineties several experiments were performed at BTLabs that demonstrated the real potential of the passive optical networking approach and culminated in the publication of initially a 32 million way split network delivering 12 wavelengths at 2.5 Gbit/s each [5]. This was further enhanced a few months later with the publication of a 44 million way split (~the whole population of the UK) with 16 wavelengths at 2.5 Gbit/s and 500 km of optical fibre [6]. Although no one was suggesting a real implementation of such a network it did illustrate that the TPON and APON architectures and indeed the current FSAN specifications for PON systems were only scratching the surface of the potential of this technology.

During the early nineties BT continued with the design and development of a more practical amplified PON architecture, it became dubbed as SuperPON. This examined the design and implementation options of a passive optical network that could service a split of up to 3000 and have a geographical range up to 100 km. The capacity was 1.2 Gbit/s downstream (from network towards the customer) and 300 Mbit/s

upstream. At the time these bit rates were very ambitious for optical access and were considered to be the limits for low cost consumer equipment.

Interestingly this solution was the lowest cost solution found for fibre-to-the-home/ business and even today there are as yet no FTTH/business solutions being offered with lower projected costs.

The European ACTS PLANET project continued with the SuperPON concept up until 1999 [7]; beyond this point there have been no further developments occurring in the supply industry.

No decisions were made to deploy FTTH systems to the mass market, and by the mid nineties it was realised that without international co-operation and standards there was not going to be a sufficiently large common market to drive down the costs of the base components necessary to enable such systems to become economically viable.

In the mid nineties BT, Deutsche Telekom and NTT decided, with other operators, to set up a consortium to develop and standardise PON requirements and systems, this forum became FSAN. For a fuller account of FSAN, see [8].

In recent years PON system have continued to be developed, largely along FSAN guidelines and mainly in the small/start-up company arena. More recently Japan, Korea and the US have revitalised interest in the supply industry and PON access solutions are once more becoming the access solution of choice as FTTH deployment progresses into the 21st century.

Why so little fibre in the access network?

Given the promise and all the development work into fibre-to-the-home systems why hasn't large-scale deployment already happened?

Probably the most significant reason is simply the scale of what is being tackled. The studies in the early nineties showed that it would have required an investment of the order of £15 billion to fibre the UK. With hindsight this may not have been a bad investment for “UK plc” but for commercially focussed operators having to survive in an increasingly competitive environment the lack of clear returns on investment made such a proposition too risky.

It was unclear what the demand for broadband would be. The public Internet and the World Wide Web had yet to emerge and the only known commercial market was broadcast or stagger-cast entertainment video

and video-on-demand (VoD). These services already had established delivery solutions (although true VoD would be more problematic) that would provide fierce competition to new fixed network providers entering the market.

At the same time compression techniques such as JPEG and MPEG were progressing rapidly and DSL technologies were increasing the transmission capacity of the embedded copper base. The advances in DSL technologies and the reductions in bandwidth required for high quality video transmission were being seen as converging and many believed that most of the envisaged broadband services could be delivered on an enhanced copper network without the need for significant fibre deployment into the access arena. This effectively led to decisions on installations of fibre-to-the-home and small business being deferred.

Emerging drivers for FTTH

What requirements are there for higher speed broadband services that could drive the need for fibre much closer to the customer, and what has changed in the last few years that is stimulating a resurgence of interest in fibre access?

The Internet

It is generally agreed that the most significant phenomenon to impact the broadband debate is the emergence and popularity of the Internet and in particular the World Wide Web. In the early days of service modelling long lists of proposed "services" would be generated with titles like "home banking", "home shopping", "games", "e-mail", "VoD" etc. Many of these are now just applications and services available over the Internet.

Clever encoding and website design enable many of these services and applications to be available even with a 56 kbit/s modem connection while the "broadband" connection $> \sim 256$ kbit/s markedly increases the responsiveness and attractiveness of the experience. Apart from high quality video, most of these services can be provided with an access speed of only a few hundred kilobits per second. However, as the content becomes image and video rich and as the quality of that content increases in terms of resolution, frame rate or reduced latency (important for gaming and conversational services), then the speed of the connection, for a satisfactory customer experience, will need to increase.

Digital imaging and video

Another area to consider is the rapid developments in consumer equipment, one example of which is digital video. Most digital video cameras use the DV format

and now have a "fire-wire" (IEEE1394) or USB2 interface that will deliver a high quality digital video signal at about 30 Mbit/s. High Definition Video (HDV) is also emerging and has been designed to use similar data rates by using real-time MPEG2 encoding rather than the cosine transform coding used by the DV standard. Currently the only place that signal can be transmitted is to other CPE a few metres away. Multiple USB and IEEE1394 interfaces are now standard features on PCs and are capable of speeds up to 400 Mbit/s (although admittedly the internal architecture and limitations of the PC and its operating system will often limit this speed, however file transfers at > 100 Mbit/s are readily achieved).

Digital photography is now rapidly displacing film photography and the ability to distribute these images via e-mail and file transfer is one of the major advantages of this media format. Modern cameras, even at the low end of the market, generate multi-mega pixel images and single image raw file sizes can easily be several Mbytes. The mobile phone camera market is also now taking off and people want to transmit the images captured to colleagues, friends and family. These camera phones can actually capture high quality, high-resolution images but the transmitted quality is very poor due to the very heavy compression applied to the images prior to transmission over current narrow band mobile networks. Linking these mobile phones to other network devices, eg. the PC via Bluetooth, could enable much higher quality and even the large raw, uncompressed image files to be transmitted over the broadband fixed network at full resolution.

Developments in digital CPE

The PC processing speed has continued to follow Moore's Law and 3 GHz plus machines are now commonplace, in the relatively near future this is likely to increase to ~ 10 GHz. The increasing power of these machines with increasing speed of the interfaces will enable large files to be moved between different items of CPE at very high speeds, much higher than the few hundred kbit/s capability of DSL technology to the external network. This mismatch between internal and external networking capability will continue to increase until high capacity optical fibre communications is brought much closer to the customer.

With increasing PC power, the popularity and rapid development of very high resolution imaging devices (including still image cameras with Mega-pixel images), DV cameras becoming commonplace and HDV emerging, users in the consumer marketplace as well as business will be generating huge amounts of data. Storage of these files within the PC is already driving hard drive capacities to the hundreds of Gbytes sizes and will soon be in the Terabytes regime.




Figure 1 Time to transmit large files peer to peer

A major problem for owners of such storage is secure back-up of valuable data. Storage area networking is a potentially valuable emerging market for the business sector, currently focussed on large business. But with the huge expansion of low cost mass storage for the mass market, a future service could be protected on-line storage for valuable data or data that customers will want to access wherever they are. Such a service will only be useful and practical if upload and download speeds are sufficiently fast for typical transactions to occur in minutes rather than hours or even days.

Examples of transmission over fibre v xDSL

It can be seen from the above that very large data files are not just the domain of large businesses, consumer equipment can now generate huge data files.

A one hour DV tape can contain over 16 Gbytes of data, single layer DVD can hold ~5 Gbytes, even the humble CD holds ~0.6 Gbytes of data. Transferring such large files electronically over current networks takes an unacceptable length of time, such that the usual method of “large” file transfer is via physical movement of media, that is, the Postal Service. The limitations of access network technologies to move such large data files between distant locations (greater than a few metres) are illustrated in Figure 1 which shows the time taken to transfer large files of various sizes using the various access technology options that are available – note the log scales.

For peer to peer transfer the upstream rate is the limitation for transmit time with ADSL and VDSL tech-

nologies but even for downstream transmission, only fibre-to-the-home provides realistic transfer times for services such as purchase and network downloads, of electronic media such as DVD files or even CD-ROM. The only practical option for say DVD purchase over the Internet or transferring a home movie at DV quality to friends and family is to use the postal service.

In the business community the incidence of teleworking is increasing steadily. Official estimates from a 2001 survey indicate that there were 2.2 million tele-workers in the UK. Tele-workers were defined as people who worked at least one day per week from home [9]. This definition excludes occasional tele-workers (people who have an office location but also work either irregularly or occasionally from home). If these are also included, the number of “tele-workers” could increase significantly.

In many corporate companies employees have access to internal Intranets usually with 10 or 100 Mbit/s Ethernet connections. In order to operate at the same service capability and efficiency as office based workers, tele-workers ideally need a network connection speed comparable to that of the internal office LAN. Current DSL based broadband options go some way towards this but tele-working often needs broadband connections with much greater symmetry of transmission as well as higher data rates.

VDSL technologies or fibre to the home would aid the tele-working environment enormously.

The conclusion from the above brief survey of the drivers for broadband is that DSL technologies can only be an interim solution and will not satisfy the demand generated by emerging consumer technology or future services. In Korea the government is already suggesting that the future broadband convergence network (BcN) should operate at speeds up to 100 Mbit/s to the end customer, ~50 times faster than conventional broadband services.

Financial barriers –

The greatest barrier to greater penetration of fibre in the access network is of course the high investment costs required. As mentioned above, the studies in the early nineties indicated a cost of ~£15 billion to provide fibre to all the homes and businesses in the UK. Although some elements of the costs for fibre-to-the-home/office have reduced over the intervening years, overall the costs, using current technology prices, have not changed significantly and the total cost will still be of the same order.

Early fibre systems targeted specific or niche markets. For example, many of the first generation of commercially available APON and BPON systems were aimed at data services for the medium and small enterprises market and did not support POTS services. The problem with these solutions for mass deployment is that they cause a fragmentation of revenue streams with revenues from different service portfolios being required to support different platforms. For a mass-market solution this is not viable. All possible revenue streams will be needed to support the deployment and operation of a single platform. A fibre-to-the-home/office solution needs to deliver all services: voice (including the regulated PSTN services), video and data to have any chance of being financially viable.

The other issue for broadband access is that historically, overall revenues from the customer base grow relatively slowly. Particular revenues can grow quite fast; however, it is usually via revenue substitution from within the telecommunications business.

On a macro-economic scale this is perfectly reasonable, GDPs of countries, and indeed the world, only grow at a modest rate (typically ~1–3 %) and to have a revenue growth rate much greater than this over an extended period will mean significant underlying economic change; ie. spend in telecomms will mean a down turn in some other sector of the economy. Although such changes do occur they generally occur slowly. In the UK growth in telecomms spend has indeed exceeded growth in GDP for several years implying underlying changes in spending patterns. However, expecting significantly greater and sustainable changes in the spend ratios over extended periods of time is unrealistic.

The conclusion from the relative overall slow growth of revenues for telecommunications services, com-

pared to the potentially very rapid increase in bandwidth demand from future broadband is that the cost per unit of bandwidth must decline very rapidly, much more rapidly than traditional price declines of products in the electronics market.

To illustrate this, consider the three service scenarios outlined below. These are potential growth scenarios, categorised as: Pragmatic Internet, Optimistic + Moderate Video and Very Optimistic & Video Centric. The main service and usage assumptions behind the scenarios are shown in Table 1.


Note these are scenarios not forecasts; the intent is merely to illustrate the impact different service growth and usage demand patterns could have on network capacity requirements and what the possible network architectures and solutions could be to meet the demands economically. The corresponding growth in network traffic for these scenarios is shown in Figure 2.

It can be seen from the results in Figure 2 that there could be big differences in the bandwidth demands that a future network may need to meet. An even bigger issue is that in the high growth scenarios, revenue growth will not match the cost of servicing the bandwidth growth. This will particularly be the case for the optimistic scenarios, unless new ways can be found of significantly increasing the price decline of the cost of bandwidth.


Ideally the industry would like to maintain the Return on Capital Expenditure (ROCE) during a sustained growth period otherwise there is a risk of profitability suffering or even companies going out of business. The interesting question to ask of these scenarios described above is: what price decline or learning curve for bandwidth must be achieved, to enable the growths shown, whilst maintaining ROCE?

	Pragmatic Internet		Optimistic + Moderate Video		Video Centric	
	2009/10	2014/15	2009/10	2014/15	2009/10	2014/15
Total broadband Internet customers (millions) (Includes cable modem)	9.4	12.3	10.7	14.8	10.9	16.2
Number of VDSL/Fibre customers (millions)	0	0	1.65	2.9	2.1	7.7
Video/VoD customers (millions)	0.14	0.54	1.5	3.5	2.0	8.7
Average Internet session time/day (mins)	75	80	96	107	101	126
Average Internet session bandwidth (kbit/s)	77	114	320	1270	488	4780
Average video session time/day (mins)	24	24	58	70	94	128
Average video session bandwidth (kbit/s)	2000	2000	7000	7500	7280	7400


Table 1 Assumptions for example traffic scenarios



This scenario can be considered a pragmatic internet scenario. It only relies on ADSL and Cable Modem technology. Streamed video services are a minority product and most video entertainment is delivered through conventional systems, eg. terrestrial, satellite or cable. Internet growth is assumed to be fairly robust but not particularly ambitious. Customers are assumed to be split between cable modem operators and ADSL operators.



This is a more optimistic scenario but without stretching the bounds of plausibility for the types and usage of the services postulated. The main difference is the use of the network to deliver high quality video entertainment services and fast file transfers for business and e-commerce. Even in this scenario video usage is only equivalent to about 1 feature film per day implying most entertainment video would still be delivered by conventional means. However network delivered video is assumed to be a personalised service rather than broadcast/multicast and is transmitted to individual customers.



This scenario assumes a much greater use of content rich Internet applications, which pushes up average session bit rates. It implies that ADSL, where capable, will enable customers to operate up to ~2 Mbit/s and that some VDSL and/or fibre solutions are also rolled out. The other major change from the "pragmatic" scenario is that a large proportion of customers use their broadband connection for personalised video applications, although it is not the main entertainment video delivery system used by customers.

Figure 2 Projected traffic growth for the service scenarios listed in Table 1

It can be shown that if all growths can be expressed as a compound annual growth rate (CAGR), then there is a simple relationship linking revenue growth, bandwidth growth and the learning curve for the price decline per unit of bandwidth to maintain ROCE, given by:

$$1 + G_R \geq (1 + G_B)^{1+L}$$

Where:

G_R is the CAGR for Revenues

G_B is the CAGR for Bandwidth

and

$$L = \text{Log}(L \% / 100) / \text{Log}(2)$$

Where:

L % is the learning curve (expressed traditionally as a percentage).

A learning curve is defined as the percentage decline in price of a product as the product volume doubles, an 80 % learning curve will mean that the price of a product at a volume V will decline to 80 % of that price at volume $2 \times V$. In this case the product volume is taken as bandwidth.

This is the macro-economic condition that needs to be met for a sustainable business as bandwidth grows in the broadband future.

This relationship is plotted in Figure 3 with contours of constant revenue growth. Typically it is reasonable to expect time averaged revenue growths in the range 4 % to 7 % CAGR. Electronic systems have learning curves of ~80 % (historically transmission systems have been following such a learning curve for the past decade at least) and from Figure 3 this price decline and revenue growth would suggest that overall network bandwidth growths of ~10 % can be sustained. Before the advent of broadband, growths of this order or less were typical and the whole economic system was internally consistent between price declines, revenue growth and bandwidth growth. Note this is a macro-economic argument applied to the total system size; it does not apply to individual products and services or sub-networks within the main network. These could grow at much faster rates by substitution of other products or markets, or they could be small niches and simply cause too small a perturbation to be noticeable in the bigger picture. In the scenarios illustrated above we are looking at large-scale changes to the capacity of the network and though very simple, this model gives a useful guide to the bounds on growths that can be expected to be viable, given a set of revenue constraints.

The envelopes of bandwidth growth for the scenarios illustrated in Figure 2 do not exhibit constant CAGRs due to the many individual parameters and growth and service usage functions that influence the growth curves. However, if an approximation is fitted to the regions, which start where bandwidth from broadband begins to become significant, then approximate CAGRs can be obtained. Using these fitted bandwidth growth rates, assuming a nominal 5 % per annum revenue growth and the relationships shown in Figure 3, then the following estimates of the required learning curves for bandwidth price decline can be estimated:

The “Pragmatic Internet” scenario is on the bounds of economic viability with electronic centric systems and network architectures not very different from

today’s. So the current strategies followed by operators deploying ADSL or Cable Modem for predominantly internet surfing, e-mail etc. with little video services, should be viable and will probably maintain ROCE, particularly if some additional revenue growth can be obtained from the richer service sets that can be offered. In addition operators can control bandwidth growth to some extent by use of contention for backhaul and core bandwidth. From a user perspective this will appear as increased delay for upload and download as the network gets congested in busy periods, effectively reducing the average session bandwidth per user.

However, the case is very different as we go beyond this scenario to the higher bandwidth scenarios that assume greater than a few hundred kilobits per second session rates and greater use of personalised video services. It can be seen that the price declines required to meet the bandwidth growths in these scenarios are much faster than the traditional 80 % learning curves typically arising in the electronics industry. In the case of the Optimistic and Video Centric scenarios the required price declines are so fast that it can be safely claimed they will not be met. From this perspective, future network architectures that continue to use traditional electronic solutions will not be able to price decline sufficiently fast to be able to maintain operating margins and profitability.

Lower cost architectures

The implications of this simple economic relationship are quite profound for the telecommunications industry. It implies that either growth of future bandwidth will have to slow to match that dictated by the price decline achievable in conventional electronic centric




Figure 3 Relationship between revenue and bandwidth growths and the price learning curve for unit bandwidth

Scenario	Bandwidth CAGR fit	Bandwidth learning curve
Pragmatic Internet	10 %	~72 %
Optimistic & Moderate Video	50 %	~54 %
Very Optimistic & Video Centric	80 %	~52 %

Table 2 Learning curves to meet given network bandwidth growths

networks, or we need to develop and implement new alternative architectures that can enable much faster decline in the price of bandwidth. The interesting question addressed in this paper is, can architectures that exploit a much greater use of optical networking technologies produce significantly faster price declines than conventional electronic centric networks?

It is not very likely that optical technology will price-decline any faster than electronic technologies. Many of the manufacturing methodologies are similar and indeed are developed from the electronics world and are therefore likely to follow similar price decline trends. The potential strength of optical technology comes not from lower cost components but its ability to displace electronic sub systems and nodes within the networks, offering much higher capacities with a much smaller equipment inventory, i.e. reducing box and port count while enabling very much higher network bandwidth.

All optical networking in the core, including true optical or photonic switching, has had a long gestation period but practical systems are now emerging. The main outstanding issue to be resolved is the gaining of confidence by operators in the ability to be able to manage and operate systems employing such “transparent” optical core networks. Transparent optical networks that are photonic end-to-end accumulate analogue impairments along the routes. In a dynamically reconfigurable, optically transparent network this is a real problem for an operational network. Network operators require optical networks to be capable of being managed with similar capability and features as today’s networks based on conventional OEO (optical-electronic-optical) architectures.

The promise offered by all optical switching is lower cost transport across the network core, achieved by removing the large number of transponders required to terminate each end of every wavelength channel of the conventional OEO architecture. However this advantage can only be realised if the operational costs do not increase as a result of the added complexity of managing the “analogue” nature of the transmission paths of the all-optical proposition.

By-passing the outer-core/metro network

The approach being evaluated in the following is to extend the “transparent” all optical architecture concept to the access and metro networks with the same aim of removing electronic nodes and port cards. It should be pointed out that this is a radical proposal and is not the mainstream view on how the access network architecture ought to evolve. Traditionally access and outer-core/metro networks are two separate worlds and the usual approach is to design access solutions and metro solutions that are largely independent of each other. This is not too unreasonable because usually there is an electronic access node (the local exchange or remote concentrator unit) placed between the access network and its technologies and the outer-core/metro network. Also historically this “access” node has been placed where it is because of the physical limits of the copper network, indeed this node position has been called the “copper anchor”.

Once optical access is considered this limitation of the first node position is no longer necessary and one interesting architectural option is to use optical access to reach deep into the network and terminate on a core edge node (a Metro Node in the BT 21CN architecture) which becomes the place where traffic grooming, marshalling and concentration etc. are performed. An option proposed for a long reach access network is based on the passive optical network (PON) principles but with optical amplification to boost the power budget. This enables large increases in bandwidth, geographical range and optical split compared to a conventional PON. It can be pedantically argued that this is no longer a passive optical network because of the use of an “active” optical amplifier within the PON. However the optical amplifier is replacing the electronic access node and the original access proportion of the network remains passive and the term PON will continue to be used in this paper (for the purist the P in PON could become Photonic).

Potentially the long reach enabled by the amplifiers would enable bypass of the usual backhaul network which would conventionally be SDH, possibly combined with Metro WDM systems for enhanced capacity. The long reach would also allow the possible bypass and ultimate removal of the local exchange or remote concentrator site and the increased capacity of the system coupled with the increased power budget enables a much greater optical split to be deployed. This allows more customers to be able to share the capacity and costs of the PON and the OLT, located in the 21CN Metro Node.

This combined access and backhaul network terminates on the 21CN Metro Node, which now provide

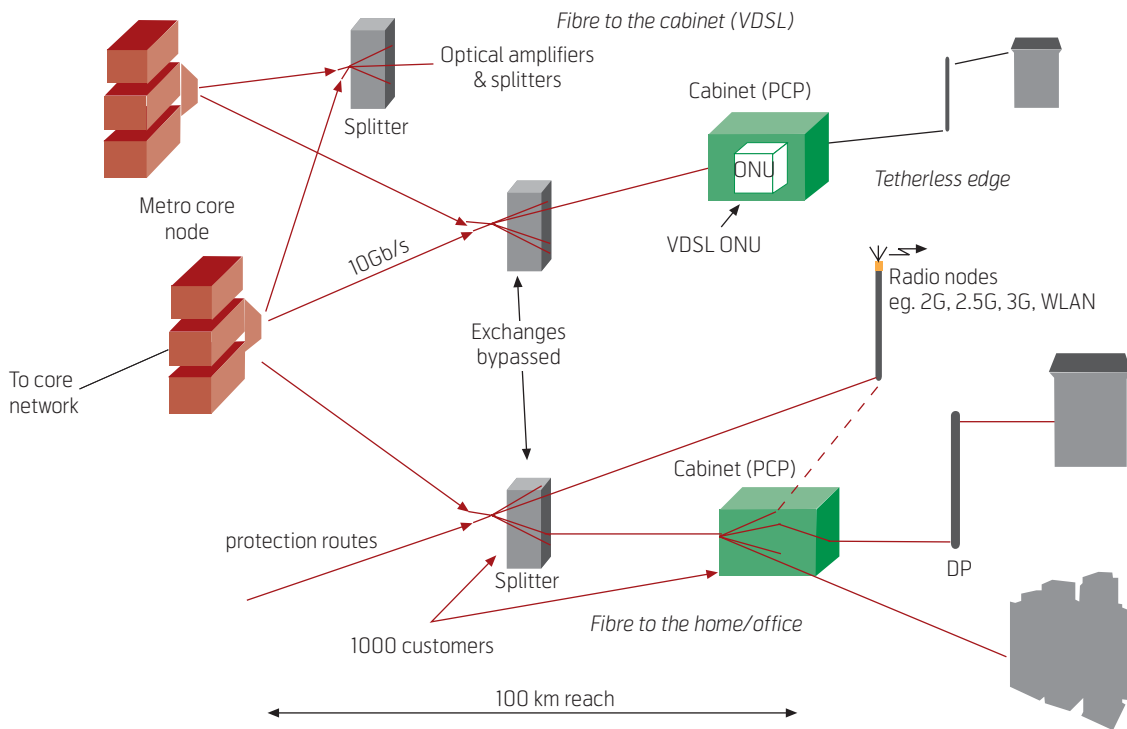


Figure 4 By-passing the outer-core/metro with Deep Reach Access PON

the management and intelligence functions and marshal and groom traffic onto a photonic inner core network. This photonic inner core interconnects these core edge nodes via simple wavelength channels. For the UK it is envisaged that there would need to be about 100 of these 21CN Metro Nodes.

This architecture leads to a highly simplified network that has the potential for significantly reducing unit costs and enables scenarios with very high bandwidth growth. Such a future long reach access network is illustrated in Figure 4, it operates at 2.5 or 10 Gbit/s and could serve ~500 (possibly 1000) customer sites on each amplified PON [10]. The traffic from the long reach PONs would terminate onto the ~100 Metro Nodes which are interconnected by an optical core network.

Although this network can now bypass the local exchange or concentrator, it would be many years before these node locations could be eliminated from the network. Legacy systems, particularly those still utilising the copper network would still need to terminate on them. While these nodes exist they can provide a convenient place for locating and powering the optical amplifiers, however the amplifiers and associated optical splitters have sufficiently low power consumption and are small enough so that street cabinets or even underground manholes/footway boxes could be used to house them. This could lead to the eventual elimination of those buildings.

This architecture would require a PON to be developed with higher split, longer reach and much higher capacity than the current generation of passive optical networks, but if developed it could offer many advantages and options for simplifying and cost reducing the end-to-end network.

Why PON networks?

There is a general assumption that point-to-point fibre solutions are a more “future proof” and therefore preferred access network option, if the capital and OpEx costs compared to PON are not excessive. At first sight this seems intuitively reasonable as there is a “fibre bandwidth” path dedicated to every customer connected.

Historically (and particularly when considering non-optical transmission media) this argument has some validity, however when fibre technology is considered it can be reasoned that the argument for traditional point-to-point access architectures may be logically flawed. This is because for any other transmission technology there has always been a higher capacity, albeit more expensive, transmission technology that the capacity of the access bearer technology can be multiplexed into. When point-to-point fibre access technology is considered and each customer has in principle access to the full fibre bandwidth there is no “higher capacity” technology into which the fibre capacity can be multiplexed. With

the present state of knowledge, fibre is the last of the transmission technologies, there is nothing of higher capacity available and nothing at present envisaged to supersede it.

However if fibre is the “ultimate” transmission system it implies that the “fibre bandwidth” (or even the exploitable bandwidth) capability of point-to-point systems can only be delivered if the network is a mesh, with every node (eg. customer) having a direct fibre connection to every other node (customer) in the network. This requires $N(N-1)/2$ links, which is of course practically and economically impossible.

In practice the fibre bandwidth has to be shared if this N^2 mesh problem is to be avoided and all networks must have concentration, multiplexing and switching/routing nodes contained within them. For traditional point-to-point systems this function is performed by electronics in the exchange nodes, the first of which is the local exchange or Remote Concentrator Unit (RCU). In 21CN this will be the Multi-Service Access Node (MSAN) site and the optical network is only point-to-point up to this location. After this location customer traffic is multiplexed onto shared fibres so even though customers have a dedicated fibre they cannot get a “fibre’s worth” of bandwidth.

For PON systems the customers are connected by point-to-point fibre to the first split point, which is usually closer to the customer than the local exchange site. This reduces the cost of fibre in the access network as well as sharing opto-electronic ports in the exchange. The first multiplexing function is therefore at this first passive optical splitter node. Further splitters introduce more passive multiplexing at the PCP (cabinet) location followed by more conventional electronic multiplexing in the local exchange building. In the long reach access proposal the local exchange multiplexer is also a passive splitter (with some optical gain to overcome system loss limits) and the first electronic multiplexing function is placed at the metro or core edge node.

The question is, given that multiplexing nodes must exist in a practical network and that fibre bandwidth gets shared over multiple customers (arising from this multiplexing and concentration function), where are the best geographical locations for the multiplexing function? Existing RCU sites are located where they are because of the physical and economic constraints of 100-year-old copper pair technology. Those constraints are very different for fibre technology and it would be very surprising if the same node locations were optimum for fibre access networks. Point-to-point fibre systems only terminate on the RCU site locations because of history, not because this is the

optimum location for technical performance and economic reasons.

A further consideration is how efficiently infrastructure capacity is used. Multiplexing increases efficiency by allowing customers to share infrastructure, thereby increasing the average utilisation by statistical multiplexing gains. Point-to-point systems have dedicated capacity to a single customer termination without any ability to share that capacity when it is not required. However PON systems can reallocate unused capacity from one customer termination to another where there is greater demand for bandwidth, thus increasing overall utilisation and efficiency of the network infrastructure. This greater utilisation begins to occur after the first split point and therefore only the infrastructure below that point operates at point-to-point efficiency. All plant above the split point is more efficiently used and each subsequent splitter stage increases utilisation further. This simply qualitative argument suggests that placing the large split points as close to customers as possible will increase overall network utilisation and efficiency. The criteria for the positioning and size of the split point will depend on the customer or access node density being served and the relative costs of fibre cable and splitter infrastructure. However, the advantages of PON over classical point-to-point systems becomes self-evident once it is realised that point-to-point fibre bandwidths cannot be delivered (let alone being required) and multiplexing must be used to share fibre bandwidth between a number of customers in order to implement practical and economically viable networks.

Some of the advantages of a PON architecture over classic point-to-point architectures are:

- Multiplexing is performed as close to customers as possible using low-cost passive splitters.
- Unused bandwidth resources can be dynamically redistributed to match customer requirements as they arise, rather than a uniform distribution of capacity to all the customers regardless of demand.
- The PON architecture has intrinsic multi-cast capability requiring only one copy of multi-destination data to be transmitted (eg. broadcast video, push services etc.)
- Customer line/port cards are shared minimising network termination costs.
- Long reach and lean fibre usage in the feeder/backhaul network enable node by-pass.

- The network is fibre lean, this minimises operational costs associated with installation and maintenance of large fibre count cables, which are required for point-to-point fibre solutions.

The disadvantages of PONs are:

- They require a more complex protocol to operate the network in order to keep customer traffic from colliding within the PON.
- The performance requirements of the optical components are more demanding and therefore initially more expensive.

However these are technical problems that are resolved with the normal march of technological progress. Significantly more complex protocols are now operating in millions of home PCs, and Gbit/s optical systems that a decade ago could only be considered for the inner core of national networks are now used in the low cost LAN environment. In a few years' time 10 Gbit/s electronics will be used in consumer equipment and given volume production. 10 Gbit/s optical systems can also be a low cost, mass-market technology.

Long reach access

Analysis of the UK network suggests that an optical access network with 2.5 Gbit/s to 10 Gbit/s capacity, with up to 512-way (or even 1024-way) split and 100 km reach could be a very attractive option. The 100 km reach would be required because such a network requires protection paths and mechanisms to provide the outer core/metro network connectivity. Protection paths can be up to twice as long as the primary paths. The increased split would be required to minimise the cost per customer of the long reach, protected, portion of the network and the increased capacity would be required to maintain the average bandwidth per customer while serving many more customers. The high bandwidth would enable massively increased burst capability, exploiting the statistical multiplexing gain by providing a "bandwidth reservoir" within the network. This would enhance the customer experience by providing a very responsive and high-speed network.

The deep reach access network would effectively bypass the metro or outer-core collector network and could significantly reduce the cost of backhaul from the access node to the Metro Node. It can also provide a very flexible and highly functional traffic engineering capability. With the right functionality built into the transport protocol, bandwidth could be allocated to customers on demand, broadcast and multi-

cast services could be offered with minimum bandwidth and switching functionality required in the core. It could also act as a distributed concentrator and service switch and simplify the switch and routing functions required in the Metro Node.

It would also be possible to incorporate sophisticated monitoring and diagnostic tools that could give early warnings of problems before they become traffic affecting and also provide terminal performance, identification and location information.

Security is always an issue with shared access media and needs to be addressed carefully to ensure adequate security from eavesdropping and malicious damage to the network integrity. There are many approaches that have been considered during the development of passive optical networks.

Other advantages that the long reach access architecture can offer are:

- *All customer types and services supported:*
This includes residential customers, multi-dwelling units, SMEs and even large business sites. Existing larger business sites already served with point-to-point fibre systems would continue to use existing access fibre but could also be integrated into the long reach architecture by exploiting the wavelength domain. Individual wavelength channels could be delivered to the largest customers and for the large customers not needing the capacity of dedicated wavelengths, sharing a wavelength over a number of business sites could be exploited
- *Symmetrical capacity:*
Although asymmetrical capability could just as easily be offered, large capacity PONs that serve all customer types will probably need to be symmetrical, the proportion of capacity allocated for broadcast services can be a relatively modest fraction of the total capacity and the bulk of the capacity may well need to be symmetrical for future service requirements.
- *Guaranteed QoS:*
The combined access and outer-core network that long reach access architectures span, could provide guaranteed QoS parameters by suitable choice of transport protocol and control system. Real time streamed services and PSTN quality voice (with all the associated delay constraints) together with common packet protocols such as Ethernet and IP could be supported.

Assuming high bandwidth demand and growth the high capacity of the system coupled with traffic

management and dynamic bandwidth assignment produces the lowest cost per unit of bandwidth per customer.

- *Reduced capital expenditure:*
The major saving associated with this architecture is the large reduction in backhaul costs that become possible; the costs for the access portion remain the same if not slightly higher than conventional PONs because of the higher capacity of the optical network units. The saving in backhaul costs arise because the Metro WDM or SDH/SONET equipment is bypassed with a simple node consisting of optical power splitters and a few optical amplifiers. This also means that these systems will also be viable in many more of the rural and low density serving areas.
- *Reduced operational costs:*
The potential operational savings are also significant. The UK network could be reduced to approximately 100 nodes that will require staffing. In theory the long reach access network could bypass the ~5500 local exchange nodes, replacing them with simple amplifier and splitter points that could be located in street cabinets or manholes/footway boxes. This could eliminate the cost of maintaining these local exchange nodes. If fibre was installed to the customer premises all service enhancements/changes could be implemented remotely without the need for visits. This could be the ultimate in “hands off access networks” and could mean that plant fault rates should be determined predominantly by third party dig-ups, producing significant reduction in plant maintenance costs.
- *Enhanced customer experience:*
The customer experience could also be dramatically enhanced by the very high burst speeds that could be offered. Essentially burst rates will be limited by the home/office network interface speeds. If these could be Gbit Ethernet speeds in the future then customers would experience very low delays for most applications (click and it’s there). Even for very large file transfers such as DVD downloads the transfer times could be tens of seconds rather than hours.

An end-to-end vision

Combining long reach access with photonic core networks would lead to highly simplified networks that have potential for significantly reducing unit costs and enabling the very high bandwidth growths predicted by the more optimistic and even video centric broadband scenarios.

Such a future network is illustrated in Figure 5. The deep reach access network serves a customer catchment area of up to 500 or 1000 sites. The traffic is terminated onto a multi-function node with functionality for IP, ATM and TDM traffic and services if required. Note this does not imply that this functionality would be all in one box – although this would be one option – but that the functionality would be co-located at the same geographical node. Traffic from the long reach access networks terminating on the Metro Node would be marshalled into transport containers (eg. this could be concatenated VC4-n containers over SDH or any other adequate transport protocol) and placed onto a wavelength channel destined for another Metro Node. The only circuit manipulation and routing that would take place in the core would be at wavelength granularity. This would avoid optical to electronic conversions followed by electronic processing and then electronic to optical conversion for onward transmission. This function would only be performed in the Metro Nodes. Note that in practice many of the optical nodes in the core would be geographically co-located with a Metro Node containing electronic switching and routing equipment, the optical switches providing through path routing but logically the architecture would be as shown in Figure 5.

This end-to-end architectural proposition could reduce a network the size of the UK to the order of 100 switching/routing nodes with ~30,000 long reach access PONs to serve the customer base. The long reach PONs would have optical amplifying nodes rather than electronic equipment and these could be placed in manholes or street cabinets and it is predicted they would consume only 0.5 % of the electrical power of the equipment and buildings they would replace making this also a very energy efficient and environmentally friendly solution.

Summary – conclusions

The broadband era in telecommunications is going to bring new challenges to the industry. This paper has shown that even fairly conservative scenarios for broadband service growth can generate huge growth in demand for network capacity.

The problem for the industry is that the growth in network capacity and the cost of servicing that growth will not be matched by corresponding growth in new revenues even though the price of network capacity will continue to decline as more capacity is provided. The paper describes a relationship between the bandwidth growth, the corresponding revenue growth and the price decline of bandwidth (defined as a learning curve). This model is then used to derive the “learn-

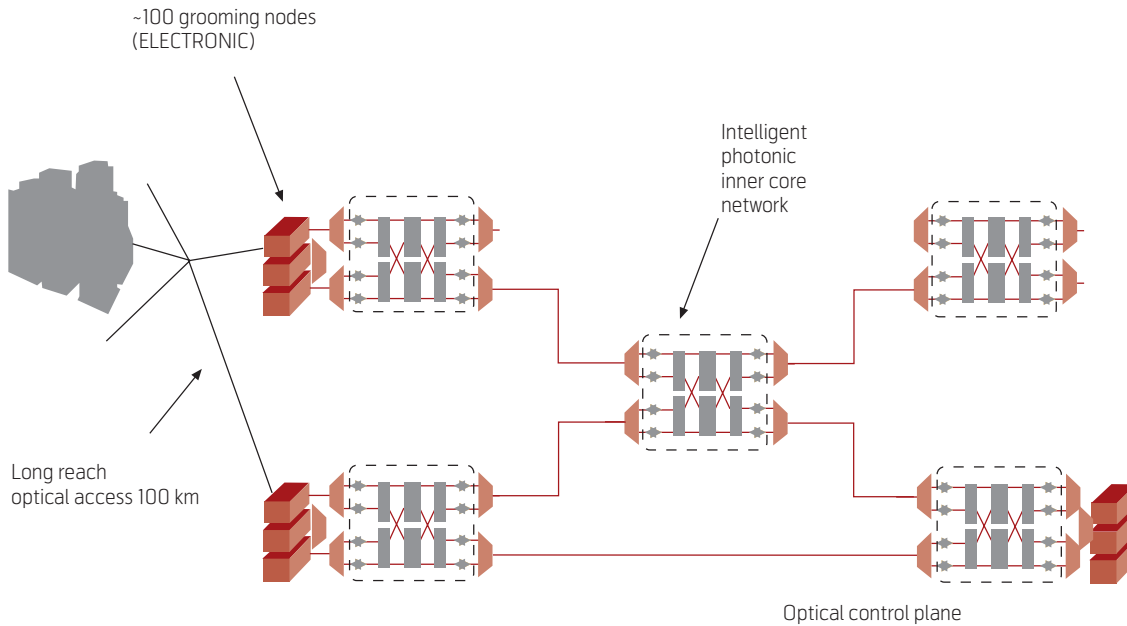


Figure 5 An end-to-end network vision using long reach access with photonic core networks

ing curve” required to sustain business margins for a range of future service scenarios.

The results from this analysis show that for many future scenarios the learning curve required for price decline of bandwidth (network capacity) is not going to be achievable if network operators rely only on historical equipment price declines. This leads to the conclusion that if equipment price declines cannot reduce the cost of bandwidth sufficiently to economically meet the future bandwidth demand then the only option will be to significantly reduce the number of network nodes, port cards and switching/routing equipment required for future network architectures. At the same time the new architectures must enable massive increases in network capacity.

To meet the challenge of these two opposing sides of the business environment – slow revenue growth but huge latent bandwidth demand – it is necessary to find network architectures that can radically reduce unit bandwidth costs. Optical networking may offer a solution, certainly it is clear from the analysis described in this paper that the price decline of network architectures based on today’s electronic intensive solutions show no historical trend to be able to reduce bandwidth costs fast enough.

In the core network, all-optical networking may be the way forward although there are still technical and management issues to be resolved. In the access network, migrating legacy copper networks to DSL architectures is only a step along the road to true broadband and as service bandwidths continue to increase, the push to drive fibre closer to the cus-

tommer will also increase. One way of impacting the cost of providing this greater bandwidth offered by a fibre rich access environment is to integrate the outer-core or metro network with the access network and combining it with an all optical core. The long reach access network using high capacity amplified Passive Optical Networks (PON) may be a solution. It can offer higher capacities, symmetrical access at much lower end-to-end unit bandwidth costs, than any electronic intensive solution.

This network architecture addresses the problem of the cost of bandwidth outstripping revenue growth by massively reducing the unit cost of bandwidth, at the same time by radically simplifying the network it also offers the prospect of significantly reducing network operational costs. Without some such architecture to produce much faster declines in the cost of bandwidth than conventional network solutions can produce, future bandwidth growth will necessarily be constrained if networks are to be operated with reasonable returns on investment.

Restraining bandwidth growth however is unlikely to be an acceptable option; the pace of development of low cost, high bandwidth generating equipment shows no sign of abating. Fast transfer of large files is likely to become a necessity in the broadband future. The paper reviews some of these drivers for increasing broadband service demand and suggests that fibre closer to the customer and eventually all the way is inevitable, the only question is the timescales and mechanisms by which it happens. It also argues that the PON architecture is an effective way of pushing fibre closer to the customer by minimising cost of

fibre access and making the best utilisation of network infrastructure. The long reach access proposal outlined in this paper is a radical extension of the PON architecture that enables bypass of local exchanges and replacement of metro transmission systems. This architecture can radically reduce the cost structure of future broadband network and is being actively pursued as one of the options in the next generation access working group in FSAN. It now remains to be seen if it can be turned into a practical reality.

References

- 1 Payne, D B, Stern, J R. Wavelength Switched Passively Coupled, Single Mode Optical Networks. *ECOC*, Venice, October 1985.
- 2 Stern, J R. Optical wideband subscriber loops and local area networks in the UK. *ICC*, Amsterdam, May 1984.
- 3 Stern, J R et al. Passive optical networks for telephony applications and beyond. *TPON Electronic Letters*, 23 (24), 1987.
- 4 Stern, J R, Hoppitt, C E, Payne, D B, Reeve, M H, Oakley, K. TPON – A Passive Optical Network For Telephony. *Proceedings ECOC 88*, Brighton, UK, September 1988, 203.
- 5 Hill, A M et al. 39.5 million-way WDM broadcast network employing two stages of erbium-doped fibre amplifiers. *Electronic Letts*, 26 (22), 1990.
- 6 Forrester, D W et al. 39.81 Gb/s, 43.8 million-way WDM broadcast network with 527 km range. *Electronic Letts*, 27 (22), 1991.
- 7 *ACTS PLANET project final report*.
<ftp://ftp.cordis.lu/pub/infowin/docs/fr-050.pdf>
- 8 Faulkner, D W et al. The Full Service Access Network Initiative. *IEEE Comms Magazine*, April 1997.
- 9 *Teleworking in the UK*.
http://www.statistics.gov.uk/articles/labour_market_trends/Teleworking_jun2002.pdf
- 10 Nettet, D et al. Demonstration of 100km reach amplified PONs with upstream bitrates of 2.5 Gbit/s and 10 Gbit/s. *ECOC 2004*, Stockholm, September 2004, paper We 2.6.3.

David Payne has worked in optical technology and networks since 1978. He started at BT Labs, UK, on single-mode fibre splicing and connectors and was instrumental in moving single-mode fibre transmission systems to practical reality in the early 1980s. He then moved on to optical networking, with particular focus on access networks, exploiting single-mode fibre technologies. In the mid 80s he was involved in the early design and experiments with wavelength switching and routing networks, leading him and his colleagues to invent the TPON (Telephony over Passive Optical Networks) concept and then further develop the PON architecture during the late 80s and early 90s. This work resulted in an experiment in 1991 involving a 50 million way split, amplified optical network with a reach of over 500 km carrying 16 wavelengths at 2.5 Gb/s each. The use of amplifiers led to the Super-PON idea (Passive Optical Networks employing amplifiers) to further reduce costs for FTTH systems. In the latter part of the 1990s David became much more involved in business and traffic modelling looking at the drivers of bandwidth that could lead to the economic justification for large-scale deployment of optical networks. This included core networks and complete end-to-end solutions. David Payne now runs the "Future Networks" team at BT Aadastral Park, Martlesham Heath, UK.

email: dave.b.payne@bt.com

Russell Davey graduated from Balliol College, Oxford in 1989 with a first class honours degree in Physics. He also obtained a PhD from Strathclyde University in 1992 for his research into mode-locked erbium fibre lasers. After an eighteen month spell working as a software engineer with Logica, he joined BT in 1994. He was heavily involved in the first introduction and applications of WDM to the BT network. In 2001 he obtained an MSc in Telecommunications Engineering from University College, London. Since 2001 he has managed BT's R&D activities in optical systems where his primary interests are in future optical access systems. He is a regular speaker on these topics at major international conferences. He is active in the Full Service Access Network (FSAN) initiative where he is co-chair of the next generation access task group.

Fibre-optic techniques for broadband access networks

TON KOONEN



Ton Koonen is Professor at the COBRA Institute, Eindhoven University of Technology, The Netherlands

The fast growth in capacity demand in access networks, driven by the increasing thirst for service bandwidths, service variety and number of service providers, asks for ever further penetration of fibre towards the end user residences. A number of key technologies are discussed for broadband service delivery through fibre access network infrastructures, encompassing multiple access techniques exploiting the time, the frequency, and the wavelength domain. Particular attention is given to the currently most popular time division multiple access techniques in passive optical networks (APON, EPON, GPON), and to subcarrier multiplexing in hybrid fibre coaxial networks. Also some trends in access network research are highlighted, in particular dynamic network reconfiguration by means of optical routing, and fibre-wireless network techniques for meeting the booming needs of wireless services.

1 Introduction

Residential users have growing needs for all kinds of telecommunication services, such as traditional voice services, but also (increasingly personalized) video services, fast internet, fast peer-to-peer file exchange, high-quality audio, etc. The capacities needed per service vary widely: from 64 kbit/s for traditional voice telephony to beyond 100 Mbit/s for high-speed internet and data. The last link to connect to the end customer, the so-called first mile (or last mile, depending on the point of view), may be bridged with various types of transport media exploited by various network

operators. Coaxial copper cable transports broadcast television and radio services, and increasingly also data services via cable modems. Twisted copper pair cables carry voice telephony, and data services via voice modems or high-speed ADSL and VDSL modems. Wireless systems bring mobile voice telephony via the GSM standard, and also data services via GPRS, UMTS, and Fixed Wireless Access (FWA). Optical fibre to the home/building is entering the market, but still has to surpass some cost barriers. It can offer the full set of integrated broadband services, from broadcast high bandwidth video services to

Medium	Bearer service	Bitrate (down/up)	Reach (km)
Twisted pair	analogue line	rates up to 56 k / 56 kbit/s	
Twisted pair	ISDN	144 k / 144 k data incl. 64 k / 64 k bit/s voice or data circuits	<6
Twisted pair	SDSL	768k / 768 kbit/s	<4
Twisted pair	ADSL	1.5 M to 6 M / 64 k to 640 kbit/s	<4 to 6
Twisted pair	VDSL	26 M to 52 M / 13 M to 26 Mbit/s	<0.3 to 1
Coaxial cable	CDMA/OFDM + QAM/QPSK	<14 M / 14 M (net 8.2 M) bit/s in 6 MHz slot	
Fibre (single mode)	ATM	150 M to 622 M / 150 M bit/s shared up to 1:32 (FSAN ATM-PON); up 1.24 G / 620 Mbit/s (FSAN/ITU B-PON)	<20
Fibre (single mode)	Gbit Ethernet	1 Gbit/s (1.25 Gbit/s 8 B/10 B coded)	<5
Fibre (multi mode)	Gbit Ethernet	1 Gbit/s (1.25 Gbit/s 8 B/10 B coded)	<0.55
Wireless (mobile)	GSM	13 kbit/s (at carrier freq. 900 and 1800 MHz, freq. duplex)	<16
Wireless (mobile)	GPRS	115 kbit/s	
Wireless (mobile)	UMTS	144 k to 2 Mbit/s (at carrier freq. 2110–2200 / 1885–2025 MHz, freq. duplex)	
Wireless (fixed)	MMDS	6 Mbit/s (at carrier freq. >17 GHz)	
Wireless (fixed)	LMDS	45 Mbit/s (at carrier freq. >17 GHz)	

Table 1 First-mile network technologies

Gigabit Ethernet data services. A list of first-mile media with the bearer services, bitrates and reach is given in Table 1. The need for more bandwidth in the access network is growing continuously, due to the increasing amount of bandwidth required by each customer, mainly fueled by video-based services and high-speed internet, the tailoring of services to individual customer needs, and the emergence of more competing operators due to liberalisation. This spurs the introduction of optical fibre (mainly single mode fibre, being a future-proof solution with its virtually infinite bandwidth) into the access network. As the network installation and equipment costs of fibre-to-the-home (FTTH) are still quite high in comparison to the traditional copper wired access lines, hybrid fibre access networks are the first step to introduce fibre. Fibre is used in the upper feeder part of the access network. There it runs from a local exchange (headend station) to a cabinet along the street (fibre-to-the-cabinet, FTTCab) or to the basement of a building (such as an apartment building with many living units; fibre-to-the-building, FTTB). At that point, the optical signals are converted back into electrical ones which are then brought via copper-based first mile links or wirelessly to the end customers.

In the next sections, after discussing some basic fibre-optic access network topologies and basic multiple access mechanisms, time- and frequency-slotted access techniques will be discussed in more depth which have been developed and are currently deployed for point-to-multipoint networks. Subsequently, some techniques which are still in the research labs for next-generation fibre access networks will be discussed: wavelength routing for dynamic capacity allocation, and radio over fibre for offering broadband wireless services. Finally, some concluding remarks are made and speculative prospects for the more distant future are given.

2 Fibre access network architectures

Basically, three architectures may be deployed for the fibre access network:

- 1 *Point-to-point* topology, where individual fibres run from the local exchange to each cabinet, home or building. Many fibres are needed, which entails high first installation costs, but also provides the ultimate capacity.
- 2 *Active star* topology, where a single fibre carries all traffic to an active node close to the end users, from where individual fibres run to each cabinet/home/building. Only a single feeder fibre is needed, and a number of short branching fibres to the end users,

which reduces costs; but the active node needs powering and maintenance.

- 3 *Passive star* topology, in which the active node of the active star topology is replaced by a passive optical power splitter/combiner that feeds the individual short branching fibres to the end users. In addition to the reduced installation costs of a single fibre feeder link, the completely passive outside plant avoids the costs of powering and maintaining active equipment in the field. This topology has therefore become quite popular for introduction of optical fibre into access networks, and is widely known as the Passive Optical Network (PON).

In the point-to-point topology and the active star topology, each fibre link is carrying a data stream between two electro-optic converters only and the traffic streams of the users are multiplexed at these terminals, so there is no risk of collision of optical data streams. In the point-to-multipoint passive star PON topology, however, the traffic multiplexing is done optically by merging the data streams at the power combiner. Collision of the individual data streams needs to be avoided by well-designed multiple access techniques.

3 Multiple access techniques in PONs

The common fibre feeder part of the PON is shared by all the optical network units (ONUs) terminating the branching fibres. The traffic sent downstream from the optical line terminal (OLT) at the local exchange is simply broadcast by means of the optical power splitter to every ONU. Sending traffic from the ONUs upstream to the local exchange, however, requires accurate multiple access techniques in order to multiplex collision-free the traffic streams generated by the ONUs onto the common feeder fibre. Four major categories of multiple access techniques for fibre access networks have been developed:

- Time Division Multiple Access (TDMA)
- SubCarrier Multiple Access (SCMA)
- Wavelength Division Multiple Access (WDMA)
- Optical Code Division Multiple Access (OCDMA)

3.1 TDMA

In a *TDMA system*, as shown in Figure 1, the upstream packets from the ONUs are time-interleaved at the power splitting point, which requires careful synchronisation of the packet transmission instants at the ONUs. This synchronisation is achieved by means of grants sent from the local exchange, which instruct the ONU when to send a packet. At the local exchange in the OLT, a burst

mode receiver is needed which can synchronise quickly to packets coming from different ONUs, and which can also handle the different amplitude levels of the packets due to differences in the path loss experienced.

TDMA techniques are deployed in ATM-PON, Ethernet PON and Gigabit PON architectures.

3.2 SCMA

In an SCMA system, illustrated in Figure 2, each ONU modulates its packet stream on a different electrical carrier frequency, which subsequently modulates the light intensity of the ONU's laser diode. The packet streams are thus put into different frequency bands, which are demultiplexed again at the local exchange. Each frequency band constitutes an independent communication channel from an ONU to the OLT in the local exchange and may thus carry a signal in a format different from that in another channel (e.g. one channel may carry a high-speed digital data signal, and another one an analogue video signal). No time synchronisation among the channels is needed. The laser diodes at the ONUs may have nominally the same wavelength. When the wavelengths of the

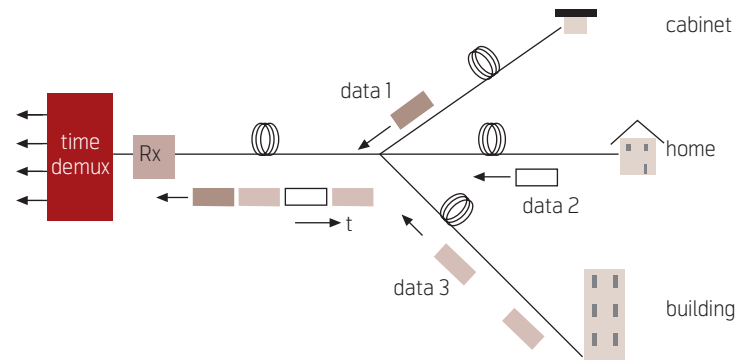


Figure 1 TDMA passive optical network

lasers are very close to each other, the frequency difference between them may result in beat noise products due to optical beating at the photodetector in the receiver. These noise products may interfere with the packet data spectrum. The wavelengths of the laser diodes have to be adjusted slightly differently (e.g. by thermal tuning) in order to avoid this optical beat noise interference.

Subcarrier multiplexing techniques are being deployed in hybrid fibre-coax CATV networks.

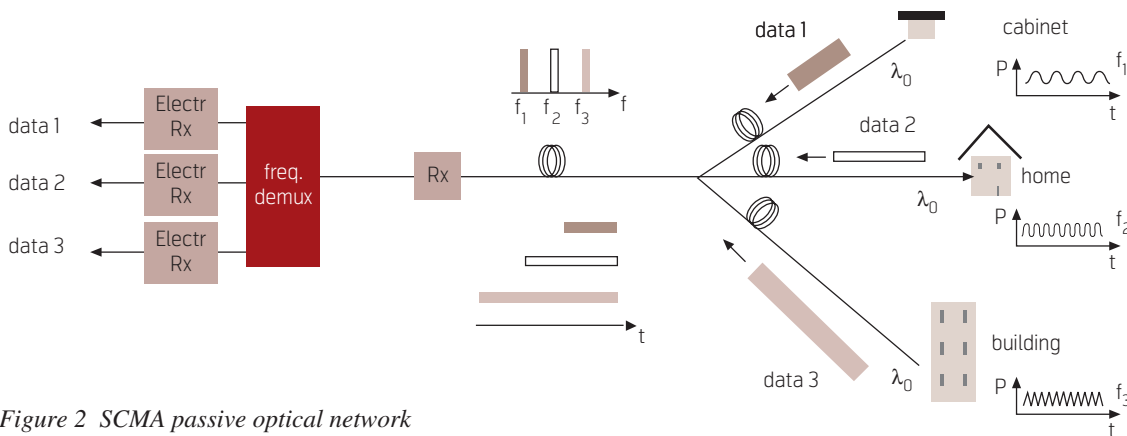


Figure 2 SCMA passive optical network

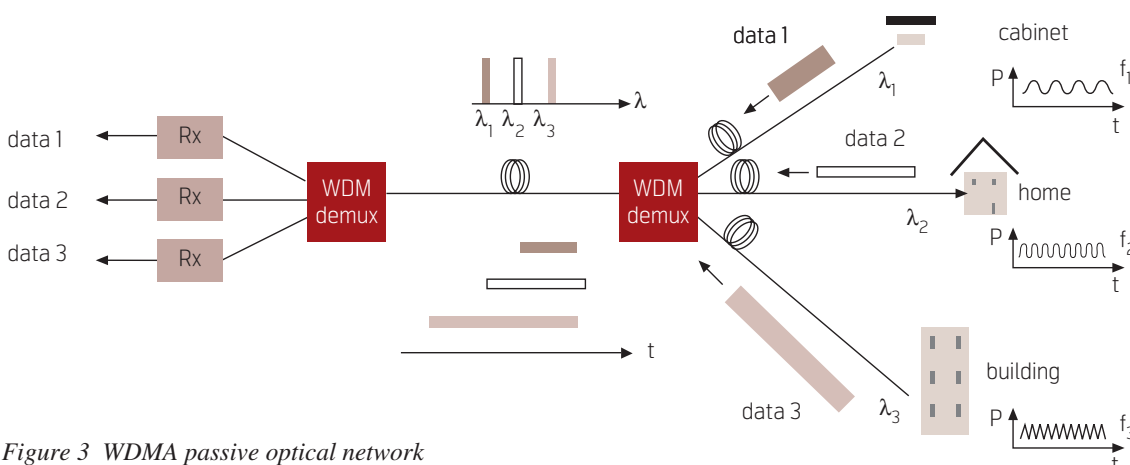


Figure 3 WDMA passive optical network

3.3 WDMA

In a *WDMA system* (see Figure 3), each ONU uses a different wavelength channel to send its packets to the OLT in the local exchange. These wavelength channels constitute independent communication channels and may thus carry different signal formats; also no time synchronisation is needed. The same wavelength channel may be used for upstream communication as for downstream. The isolation requirements of the wavelength demultiplexer should be high to sufficiently suppress crosstalk, e.g. when high-speed digital data and analogue video are carried on two different wavelength channels. The channel routing by the wavelength multiplexer at the network splitting point prohibits broadcasting of channels to all ONUs, as needed for instance for CATV signal distribution; a router bypass needs to be implemented for that. Every ONU needs a wavelength-specific laser diode, which increases costs and complicates maintenance and stock inventory issues. An alternative is to use a light source with a broad spectrum at the ONU (e.g. a superluminescent LED), of which the in-field multiplexer cuts out the appropriate part of the spectrum. This "spectral slicing" approach reduces the inventory problems, but also yields a reduction of the effective optical power available from the ONU and increased intensity noise due to spectral instabilities of the source, and thus limits the reach of the system. Another alternative is to use a reflective modulator at the ONU, which modulates the upstream data on a continuous light channel emitted at the appropriate wavelength by the OLT and returns it to the OLT [1]. Thus no light source is needed at the ONU, which eases maintenance; but again the power budget is limited.

3.4 OCDMA

In an *OCDMA system*, each ONU uses a different signature sequence of optical pulses, and this sequence is on-off modulated with the data to be transmitted. The duration of the sequence needs to be equal to that of a data bit, and thus a very high-speed signature sequence is needed to transmit moderate-speed data. This limits the reach of the system due to the increased impact of dispersion and the decreasing power budget at high line rates. In the OLT at the local exchange, the received signals are correlated with the known signature sequences, in order to demultiplex the data coming from the different ONUs. As the signature codes may not be perfectly orthogonal, some crosstalk may occur.

3.5 Comparison of multiple access techniques

TDMA systems have received the most attention for broadband access networks, as they are most suited for high-speed data transmission at relatively moder-

ate complexity, and the required digital signal processing can be readily accommodated in electronic integrated circuits. Three types of TDMA passive optical networks have been addressed extensively in standardisation bodies: the ATM PON (APON) carrying native ATM cells following the G.983 standard series of ITU-T SG15, the Ethernet PON (EPON) carrying Gigabit Ethernet packets in IEEE 802.3, and recently the Gigabit PON (GPON) able to carry ATM as well as Ethernet packets with high line rates and high efficiency (up to 2.4 Gbit/s up- and downstream) in the G.984 standard series.

Subcarrier multiplexing is particularly attractive for downstream broadband broadcasting, such as in hybrid fibre-coax CATV networks. In upstream direction, broadband communication in individual separate frequency bands requires a quite extended frequency range and high linearity of the user equipment, plus additional precautions need to be taken to avoid beat noise interference. Hence, SCMA is not commonly deployed in interactive fibre access networks.

WDMA offers the most powerful solution for multiple access, as it creates a virtual point-to-point topology. It is, however, also the most costly due to the additionally required wavelength selective functions. Some cost reductions may be achieved by using 'colourless' ONU techniques, such as a reflective modulator (e.g. reflective semiconductor amplifier) cooperating with a remote source, or spectral slicing of a broadband source. Using wavelength-based dynamic optical routing, very flexible future-proof access networks can be implemented, which can readily accommodate service upgrades, reallocation of traffic capacity, etc. Dynamic WDMA is an attractive solution for next-generation access networks, being addressed in research (see section 6).

OCDMA puts high speed requirements to the electro-optical terminals, due to the line rate being a multiple of the data rate. This leads to costly terminal equipment, and hence has not become popular for fibre-optic access networks.

4 TDMA PON systems

4.1 ATM PON

The Full Service Access Network (FSAN) group, a committee of presently 21 major telecommunication operators around the world, has since 1995 been promoting the ATM PON (also termed APON or BPON) for broadband access networks.

4.1.1 ATM PON system architecture

As laid down in the G.983.1 Recommendation of ITU-T [2], an ATM PON may have a downstream bitrate of 155 or 622 Mbit/s and an upstream one of 155 Mbit/s. The maximum optical splitting ratio is 32 (may grow to 64), and the maximum fibre length between the OLT in the local exchange and an ONU is 20 km. The range in which this length is allowed to vary is from 0 to 20 km. Standard single mode fibre (G.652) is foreseen. Coarse wavelength multiplexing is used for separating the bi-directional traffic: the downstream traffic is positioned in the 1.5 μm wavelength band, and the upstream traffic in the 1.3 μm band (using cheap Fabry Perot laser diodes in the ONUs).

In the downstream direction of a 155 Mbit/s down / 155 Mbit/s up system, 54 ATM cells of 53 bytes each are fitted together with 2 PLOAM cells (Physical Layer Operation, Administration, and Maintenance) of 53 bytes in a frame [2]. The PLOAM cells contain 53 upstream grants each. A grant permits an ONU to send an ATM cell. By sending these grants, the OLT controls at each ONU the transmission of the upstream packets, and can therefore assign dynamically a portion of the upstream bandwidth to each ONU. In a 622 Mbit/s down / 155 Mbit/s up system, a frame contains four times as many cells (i.e. 216 ATM cells and 8 PLOAM cells). The downstream frame is broadcast to all ONUs. An ONU only extracts those cells that are addressed to it.

In the upstream frame, both for the 155 Mbit/s down / 155 Mbit/s up system and for the 622 Mbit/s down / 155 Mbit/s up system, 53 ATM cells are fitted of 53 bytes each plus an overhead of 3 bytes per cell. This overhead is used as guard time, as a delimiter and as preamble for supporting the burst mode receiver process in the local exchange.

The power budgets needed to bridge the fibre losses and the splitter losses are denoted by three classes of optical path losses: class A 5-20 dB, class B 10-25 dB, and class C 15-30 dB. At the ONU, a launched optical power of -4 to +2 dBm is specified for class B, and -2 to +4 dBm for class C [3]. The ONU receiver sensitivity at 155 Mbit/s should be better than -30 dBm for class B, and -33 dBm for class C.

The ONUs are usually positioned at different distances from the local exchange. Therefore the upstream transmission of the packets from each ONU should be carefully timed, in such a way that the packets do not collide at the network splitter [2] [5]. The OLT has to measure the distance to each ONU for this, and then instructs the ONU to insert an equalising transmission delay such that all distances

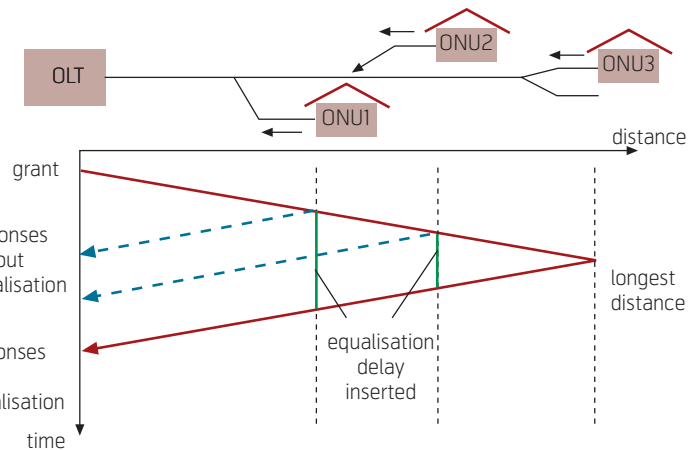


Figure 4 Time ranging in a TDMA PON

from the ONUs to the OLT are virtually equal to the longest allowable distance (i.e. 20 km); see Figure 4. To measure the distance to each ONU, the OLT emits a ranging grant to each ONU, and on receipt the ONU returns a ranging cell to the OLT. In this distance ranging process, the OLT can deduce the distance to each ONU from the round trip delay.

Each ONU sends an upstream cell upon receipt of a grant. Because the path losses from each ONU to the OLT may be different, the power of the cells received by the OLT may vary considerably from cell to cell. The burst mode receiver at the OLT should therefore have a wide dynamic range, and should be able to set its decision threshold quickly to the appropriate level to discriminate the logical ones from the zeros. Also the power of the ONU transmitter can be varied over a certain range to limit the requirements on the receiver's dynamic range. In this amplitude ranging process, the overhead to each ATM cell is used for supporting the fast decision threshold setting at the OLT burst mode receiver and the power adaptation at the ONU burst mode transmitter.

4.1.2 Network protection

Four types of network protection have been described in Recommendation G.983.1 [5], as shown in Figure 5. Type A protection involves protection of the feeder fibre only by a spare fibre over which the traffic can be rerouted by means of optical switches. After detection of a failure in the primary fibre and switch-over to the spare fibre, also re-ranging has to be done by the PON transmission convergence (TC) layer. Thus only limited protection of the system is realised. Mechanical optical switches are used up till now; when optical switching becomes cheaper, this protection scheme may become more attractive. Type B protection features duplication of both the feeder fibre and the OLT. The secondary OLT is on cold standby, and is activated when the primary one fails. Due to the high sharing factor of the duplicated resources by the

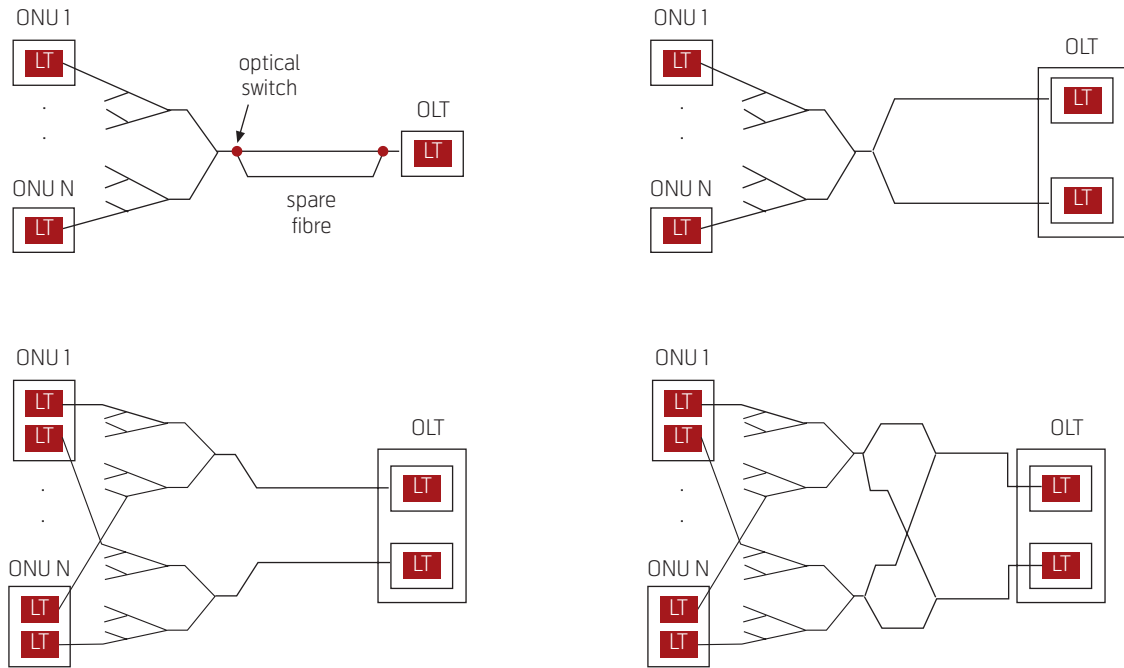


Figure 5 PON protection schemes

Type A: Feeder fibre protection; Type B: OLT and feeder fibre protection; Type C: Full PON duplication
 Type D: Independent duplication of feeder and branch fibres

ONUs, this approach offers an economical yet limited protection. Type C protection implies full duplication of the PON, and all equipment is normally working, which allows fast switch-over (within 50 ms) from the primary equipment to the secondary one. The branch fibres as well as the ONUs are protected; also a mix of protected and unprotected ONUs can be handled. Type D protection features independent duplication of the feeder fibres and the branch fibres. It cannot offer fast restoration. It is less attractive than C, as it requires more components but not a better functionality. In summary, types B and C are the most attractive schemes in G983.5.

4.1.3 Extensions of ATM PON

To further increase the speeds laid down in Recommendation G.983.1, research has been done into 622, 1244 and 2488 Mbit/s line rates, both for upstream and downstream. A key technical issue is the development of faster burst-mode circuitry to adequately retrieve the timing and set the decision threshold level, which becomes increasingly more difficult at higher line rates. Operation of 622 Mbit/s burst-mode circuitry has been achieved recently [3]. In January 2003, ITU set standards for Gigabit-capable PONs (G-PONs). Further details are given in section 4.3.

The G.983.1 ATM PON was initially mainly designed for high-speed data communication. However, in the residential access networks there is also a clear demand for economical delivery of CATV services, for which subcarrier multiplexing techniques are quite appropriate. In the enhanced Recommendation G.983.3 [3], room has been allocated in the optical spectrum to host video services or additional digital services next to the ATM PON services. As shown in Figure 6, the APON upstream services remain in the 1260 to 1360 nm band (as in G.983.1), but the band for downstream services is narrowed to 1480–1500 nm (1480–1580 nm in G.983.1). Next to those, an enhancement band for densely wavelength multiplexed bi-directional digital services (such as private wavelength services) is foreseen, or an enhancement band for an overlay of video delivery services. The latter is used in downstream direction only, and coincides with the C-band as economical erbium-doped

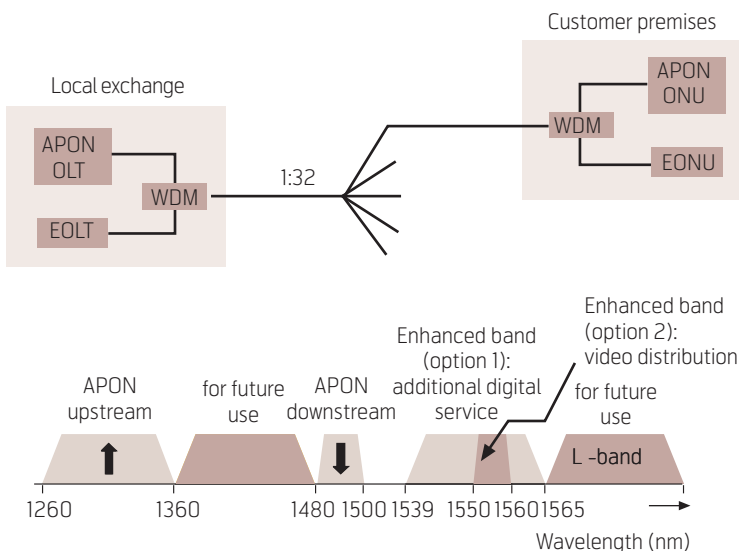


Figure 6 WDM enhancement G.983.3

fibre amplifiers can thus be deployed for the power boosting required. When positioning an overlay of CATV distributive services in the C-band, stringent crosstalk requirements have to be put on the wavelength multiplexers and demultiplexers, to prevent noticeable interference of the CATV signals into the digital ATM signals, and vice versa [4].

In order to further improve the economics of ATM PON systems, an extended PON system with an increase of the network splitting factor to 128 and even 256 has been developed, while still maintaining a passive outside plant and compatibility with G.983.1 compliant ONUs [5]. This extended split is achieved by creating a larger optical power budget. In the downstream direction, at the OLT a high power laser diode or an erbium-doped fibre amplifier (EDFA) is used to boost the power. In the upstream direction, the sensitivity of the burst-mode receiver is improved by applying an avalanche photo diode (APD). Also eight single-mode feeder fibres (each feeding a 1:16 or 1:32 power splitter in the field) are at the OLT coupled to a multimode fibre yielding a low-loss coupling to the receiver.

Even further extensions of the split factor and of the reach of an ATM PON have been realised in the SuperPON system [6]. An extension to a splitting factor of 1:2048 has been achieved; this however needs active equipment in the field. In the downstream direction exploiting the 1530–1560 nm wavelength window, EDFAs are used for overcoming the large path losses. In the upstream direction, gated semiconductor optical amplifiers (SOAs) are deployed. Each SOA gate is opened when upstream packets arrive, and is shut otherwise in order to avoid funneling of the amplified spontaneous emission noise towards the OLT. This SuperPON approach is not compliant with present standards, and may be economically feasible only in the long term [5].

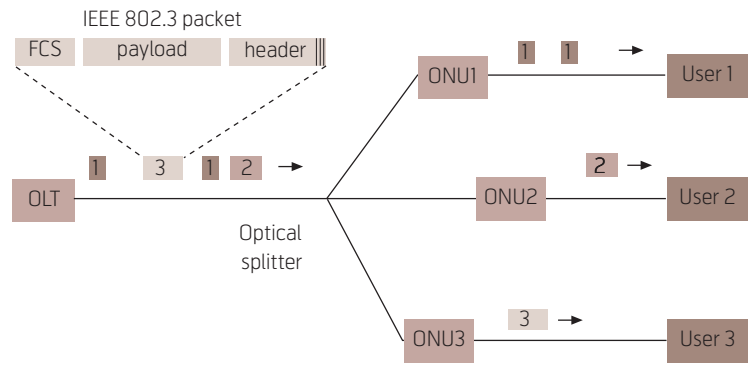


Figure 7 Downstream traffic in an EPON

4.2 Ethernet PON

With the rapid penetration of Ethernet-based services, Ethernet PON (EPON) techniques are receiving increasing attention and are promoted by the IEEE 802.3 Ethernet in the First Mile (EFM) group. The major difference from ATM PONs is that an EPON carries variable-length packets up to 1518 bytes in length, whereas an ATM PON carries fixed-length 53 bytes cells. This ability yields a higher efficiency for handling IP traffic. The packets are transported at the Gigabit Ethernet 1.25 Gbit/s speed using the IEEE 802.3 Ethernet protocol.

The EPON features full-duplex transmission similarly to the ATM PON, with downstream traffic between 1490 and 1510 nm and upstream traffic at around 1310 nm. As shown in Figure 7, standard IEEE 802.3 Ethernet packets are broadcast downstream by the OLT to all the ONUs. Each ONU inspects the headers, and extracts the packets that are addressed to it. Several variable-length packets are put into a fixed-length frame of 2 ms duration, and each frame begins with a one-byte synchronisation marker. In the upstream direction, also 2 ms frames are used. A frame contains time slots that each are assigned to




Figure 8 Upstream traffic in an EPON

one of the ONUs (see Figure 8). Each ONU puts one or more of its upstream variable-length IEEE 802.3 packets into a time slot; if it has no packets to send, the time slot may be filled with an idle signal. No packet fragmentation takes place. The time slot overhead consists of a guard band and indicators for timing and signal power. The OLT thus allows only one ONU to send at a time, and no collisions occur. The time slot size is 125 or 250 μ s.

4.3 Gigabit PON

In order to extend the capacity of PONs into the Gbit/s arena, the ITU has set standards for the Gigabit PON (GPON) in the G.984.x series. The GPON architecture, set in Recommendation G.984.1, is much alike the ATM PON one: the maximum optical splitting ratio is 128, and the maximum fibre reach from OLT to ONU is 20 km whereas its minimum is 0. Protection schemes have also been foreseen, similar to those shown in Figure 5.

The GPON Physical Media Dependent layer has been set in G.984.2; it includes downstream line rates of 1244.16 or 2488.32 Mbit/s, in the wavelength range

1480–1500 nm. In upstream direction, line rates foreseen are 155.52, 622.08, 1244.16, or 2488.32 Mbit/s, in the wavelength range 1260–1360 nm.

In GPON Transmission Convergence Recommendation G.984.3, a framing format of 125 μ s length is used which can host a lot of different packetised traffic formats. This GPON Encapsulation Method (GEM) may host Ethernet packets, and/or native ATM packets, and/or native TDM, as illustrated in Figure 9 [7]. Thus, a GPON system may operate in an Ethernet-packet-only mode, or in an ATM-only mode, or in a mixed mode. Ethernet frames may be fragmented among a number of GEM cells, which is not possible in the native IEEE 802.3 technology. Hence, GPON using GEM can obtain a high efficiency for transport of IP data payload by utilizing up to 95 % of the available bandwidth in the transmission channel.

GPON also supports Quality of Service, as it enables Service Level Agreement (SLA) negotiations between the OLT and the ONU through the ONU Management and Configuration Interface set out in G.984.4.

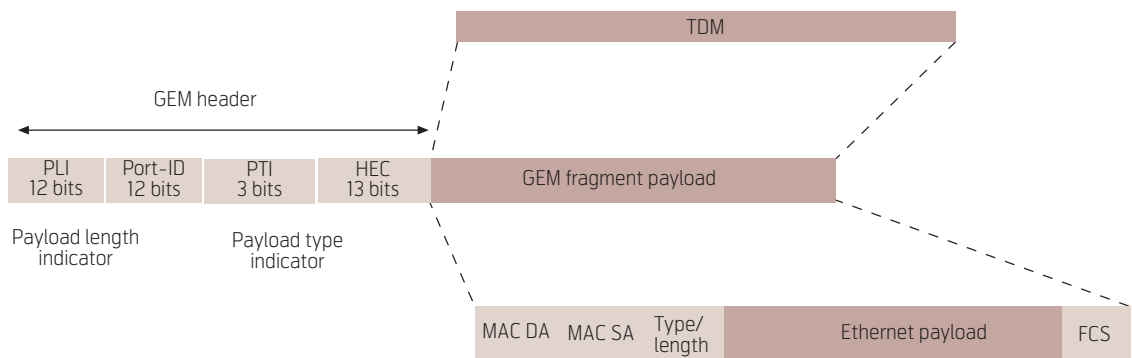


Figure 9 GPON encapsulation method according to ITU-T Rec. G.984.3




Figure 10 Hybrid fibre coax network

4.4 Comparison of TDMA systems

By using ATM techniques, ATM PON offers built-in Quality of Service for all traffic classes, whereas EPON through using native Ethernet does not. EPON thus cannot support voice services with quality of service as provided in the traditional public switched telephone network (PSTN), and also the support of real-time services still has issues due to latency and packet jitter. On the other hand, ATM suffers from the cell tax (5 bytes header per 53 bytes cell), and thus EPON is more efficient and simple for transporting variable length IP packets. The recently introduced GPON can carry ATM as well as Ethernet traffic in any mixed mode, with high efficiency, and hence may nicely combine the QoS advantages of ATM with the efficiency of Ethernet.

5 Hybrid fibre coax networks

CATV networks are usually laid out over large geographical areas and are mainly designed for downstream broadcasting of analogue TV channels (or digital TV channels, multiples of which fit into one analogue TV channel frequency slot) that are frequency-division multiplexed in a carrier frequency grid extending up to 1 GHz. As shown in Figure 10, in a hybrid fibre coax (HFC) system a CATV headend station collects the CATV signals, remodulates them into a specific frequency grid, and sends them via single-mode fibres to fibre nodes. Each fibre node converts the composite optical signal into an electrical one, which is carried via a coaxial cable network including several RF amplifiers to the residential homes. A single headend may thus serve hundreds

of thousands of customers, and a fibre node some thousands of customers. In particular during transmission in the coaxial cable network, the signal quality deteriorates due to the addition of noise from the electrical amplifiers and intermodulation products from non-linearities in the system. On the fibre part of the network, the signals are carried with subcarrier multiplexing; see Figure 11. Each TV channel is amplitude-modulated on a separate frequency, and after summing all these modulated signals, a highly linear high power laser diode (or laser diode followed by a linearised external modulator) generates an optical signal which is intensity-modulated with the composite CATV signal. At the receiver site, the optical signal is converted into the electrical CATV signal by means of a highly linear PIN photodiode, and subsequently the signal can be passed to the coaxial cable network or to a selective receiver. When using a laser diode with low relative intensity noise and high linearity (or followed by a carefully linearised external modulator), the CATV signal can be transported with very little loss of quality. If a 1.5 μm wavelength laser diode is used, erbium-doped fibre amplifiers can be used to boost the power at the headend and to compensate for the splitting losses; thus very extensive networks feeding thousands of ONUs can be realised. In this wavelength region, however, with direct laser modulation second order intermodulation products may arise due to laser chirp in combination with fibre chromatic dispersion; with an external modulator, however, the chirp is small enough to avoid these intermodulation products.

The CATV signal quality that can be maintained in HFC networks is very high due to the fibre's low losses and high bandwidth in comparison with coaxial cable. Therefore in HFC networks fibre is gradually brought deeper into the network, and fibre nodes have to serve fewer customers through a coaxial cable network of limited size (i.e. mini fibre nodes, each serving in the order of around 40 customers).

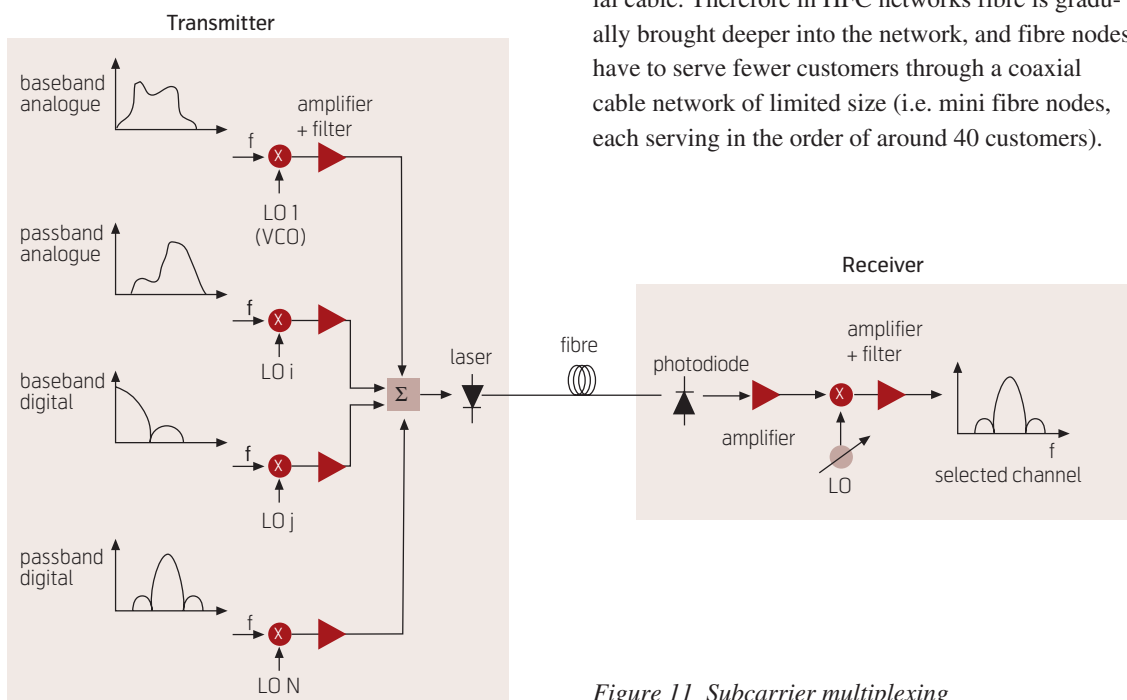


Figure 11 Subcarrier multiplexing

HFC networks are nowadays not only carrying CATV and FM radio broadcast services, but cable operators are also exploiting them for voice telephony and data transport by using cable modems. For the upstream traffic which is involved with these interactive services and which is carried upstream in SCMA mode, parts of the spectrum unused for CATV and FM radio broadcast can be used. In Europe, typically the 5 to 65 MHz band is used for this; in the US, the 5–42 MHz range. For downstream data, e.g. the 300 to 450 MHz range is used, taking into account that Internet traffic is usually highly asymmetric (much more downloading traffic than uploading). Downstream per 8 MHz CATV channel, 30 to 50 Mbit/s data can be accommodated deploying 64 or even 256 Quadrature Amplitude Modulation (QAM). Upstream due to ingress noise less complicated modulation schemes are to be used; DQPSK offers about 3 Mbit/s per channel.

6 Dense wavelength multiplexing in access networks

In general, access networks have to meet a fast growth in capacity demand for several reasons. Customers are asking for second and more telephone




Figure 12 Dynamic wavelength routing in hybrid access networks




Figure 13 Dynamically allocating wavelength channels to ONUs

lines; internet data traffic is booming with higher data rates, more users and longer sessions on-line (even always on); intense peer-to-peer file transfer traffic; multiplayer on-line gaming; an increasing amount of video-based services; fast growth in number of mobile phone users and session frequency; new operators entering asking to rent capacity on existing access networks; etc. This hunger for more capacity and the strive for convergence of services on a single network can most adequately be met by bringing fibre ever closer to the end users, from where only a short copper cable based (or wireless) link has to be bridged to reach the customer. Ultimately, when installation and equipment costs have come down sufficiently, the most powerful network is achieved when fibre is running all the way to the customer's home (fibre to the home, FTTH).

Upgrading installed fibre plant to higher capacities while protecting the investments already made can efficiently be done by introducing wavelength multiplexing techniques [8]. Wavelength channels may be allocated to specific sets of services (for service unbundling), and/or to separately host (new) service operators (leasing of network capacity).

6.1 Dynamic capacity allocation by flexible wavelength assignment


To cope with variation in service demand by the users and the sometimes quickly changing operator conditions, it is efficient to flexibly allocate the augmented available network capacity across the access network. Dynamic wavelength routing techniques can be used for this, thus making more efficient use of the network's resources and generating more revenues. Figure 12 illustrates the principle: from the OLT in the headend station of the network, multiple wavelength channels are fed to the ONUs via a tree-and-branch PON. By wavelength-selective routing in the PON, or wavelength selection at the ONU, wavelength channels can be assigned to a number of specific ONUs. Thus capacity can be specifically shared between these ONUs. Each ONU subsequently transfers its capacity share to its first-mile electrical network con-

necting the end users. The mapping of the network capacity resources to the first-mile networks can thus be changed by changing the wavelength channel assignment among the ONUs. Basically two approaches can be followed for this, as illustrated in Figure 13: a wavelength router in the field, or wavelength selection at the ONUs. As shown in Figure 13.a, a tunable wavelength router directs the wavelength channels to specific output ports, and this routing can be dynamically adjusted by external control signals from the headend. In order to also support the delivery of broadcast services to all ONUs, extra provisions have to be made for enabling broadcast wavelength channel(s) to bypass the router. As the wavelength channels are routed to only those ONU whose customers require the associated services, no optical power is wasted. As shown in Figure 13.b, another approach is to broadcast all wavelength channels to every ONU, and subsequently tune each ONU to the wavelength channel wanted. Clearly the power of the non-wanted wavelength channels is wasted by the ONU, and losses at the broadcasting power splitter are significant. An optical amplifier is usually needed

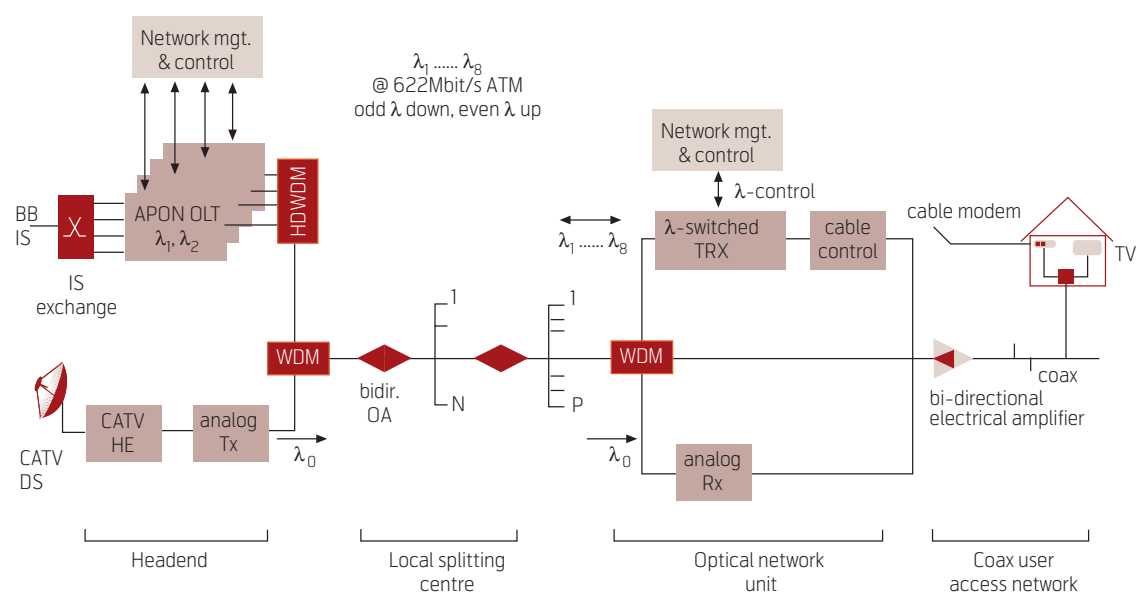
to make up for these losses; the amplifier needs to operate bi-directionally to handle downstream as well as upstream traffic. This approach inherently supports broadcast services.

6.2 Wavelength broadcast-and-select access network

Figure 14 presents a multi-wavelength overlay of a number of ATM PON networks on a HFC network, following the *wavelength channel selection* approach [9]. Figure 14.a shows a fibre-coax network for distribution of CATV services, operating at a wavelength λ_0 in the 1550–1560 nm window where erbium-doped fibre amplifiers (EDFAs) offer their best output power performance. Thus, using several EDFAs in cascade, an extensive optical network splitting factor can be realised and a large number of customers can be served. For example, with two optical amplifier stages and typical splitting factors of $N = 4$ and $P = 16$, and a mini-fibre node serving 40 users via its coaxial network, a total of 2560 users is served from a single headend fibre. For interactive services, the upstream frequency band in a standard HFC network



a) fibre-coax network for distribution of CATV services



b) upgrading of the fibre-coax network with multi-wavelength APON system for delivery of broadband interactive services

Figure 14 Flexible capacity assignment in a multi-wavelength fibre-coax network by wavelength selection at the ONUs

(with a width of some 40 to 60 MHz) has to be shared among these users, thus allowing only limited bitrates per user for narrowband services such as voice telephony.

An upgrade of the system in order to provide broadband interactive services can be realised by overlaying the HFC network with a number of wavelength-multiplexed APON systems, such as have been developed in the ACTS TOBASCO project [9] and are shown in Figure 14.b. Four APON OLTs at the head-end site are providing each bi-directional 622 Mbit/s ATM signals on a specific downstream and upstream wavelength pair. These eight wavelengths are positioned in the 1535– 1541 nm window, where the up- and downstream wavelength channels are interleaved with 100 GHz spacing. The APON wavelengths are combined by a high-density wavelength division multiplexer (HDWDM), and subsequently multiplexed with the CATV signal by means of a simple coarse wavelength multiplexer (thanks to the wide spacing between the band of APON wavelengths and the CATV wavelength band). The system upgrade implies also replacement of the uni-directional optical erbium-doped fibre amplifiers by bidirectional ones which feature low noise high-power operation for the downstream CATV signal, and for the bidirectional ATM signals a wavelength-flattened gain curve plus a nonsaturated behaviour (to suppress crosstalk in burst-mode). At the ONU site, the CATV signal is first separated from the APON signals by means of a coarse wavelength multiplexer and subsequently converted to an electrical CATV signal by a highly linear receiver and distributed to the users via the coaxial network. The APON signals are fed to a wavelength-switched transceiver, of which the receiver can be switched to any of the four downstream wavelength channels and the transmitter to any of the four upstream ones. The wavelength-switched transceiver may be implemented by an array of wavelength-specific transmitters and receivers, which can be individ-

ually switched on and off; this configuration allows to set up a new wavelength channel before breaking down the old one (“make-before-break”). Alternatively, it may use wavelength-tunable transmitters and receivers, which in principle can address any wavelength in a certain range; this eases further upgrading of the system by introducing more wavelength channels, but also implies a “break-before-make” channel switching. The network management and control system commands to which downstream and upstream wavelength channel each ONU transceiver is switched. By issuing these commands from the headend station, the network operator actually controls the virtual topology of the network, and is thus able to allocate the network’s capacity resources in response to the traffic demands at the various ONU sites. The network management command signals are transported via an out-of-band wavelength channel in the 1.3 μm wavelength window. The APON signals channel selected by the ONU is converted by the transceiver into a bidirectional electrical broadband data signal, which is by a cable modem controller put in an appropriate frequency band for multiplexing with the electrical CATV signal. The upstream data signal is usually put below the lowest frequency CATV signal (so below 40 to 50 MHz), and the downstream signal in empty frequency bands between the CATV broadcast channels. The signals are carried by the coaxial network (in which only the electrical amplifiers need to be adapted to handle the broadband data signals) to the customer homes, where the CATV signal is separated from the bi-directional data signals. The latter signals are processed by a cable modem that interacts with the cable modem controller at the ONU site.

By remotely changing the wavelength selection at the ONUs, the network operator can adjust the system’s capacity allocation in order to meet the local traffic demands at the ONU sites. As illustrated in Figure 15, the ONUs are allocated to the four upstream (and downstream) wavelength channels, each having a maximum capacity of 622 Mbit/s for ATM data. As soon as the traffic to be sent upstream by an ONU grows and does not fit anymore within its wavelength channel, the network management system can command the ONU to be allocated to another wavelength channel, in which sufficient free capacity is still available. Obviously, this dynamic wavelength re-allocation process reduces the system’s blocking probability, i.e. it allows the system to handle more traffic without blocking and thus can increase the revenues of the operator.

6.3 Wavelength routing access network

Figure 16 presents the *dynamic wavelength channel routing* approach in a fibre-wireless network to allo-




Figure 15 Re-allocating ONUs to wavelength channels

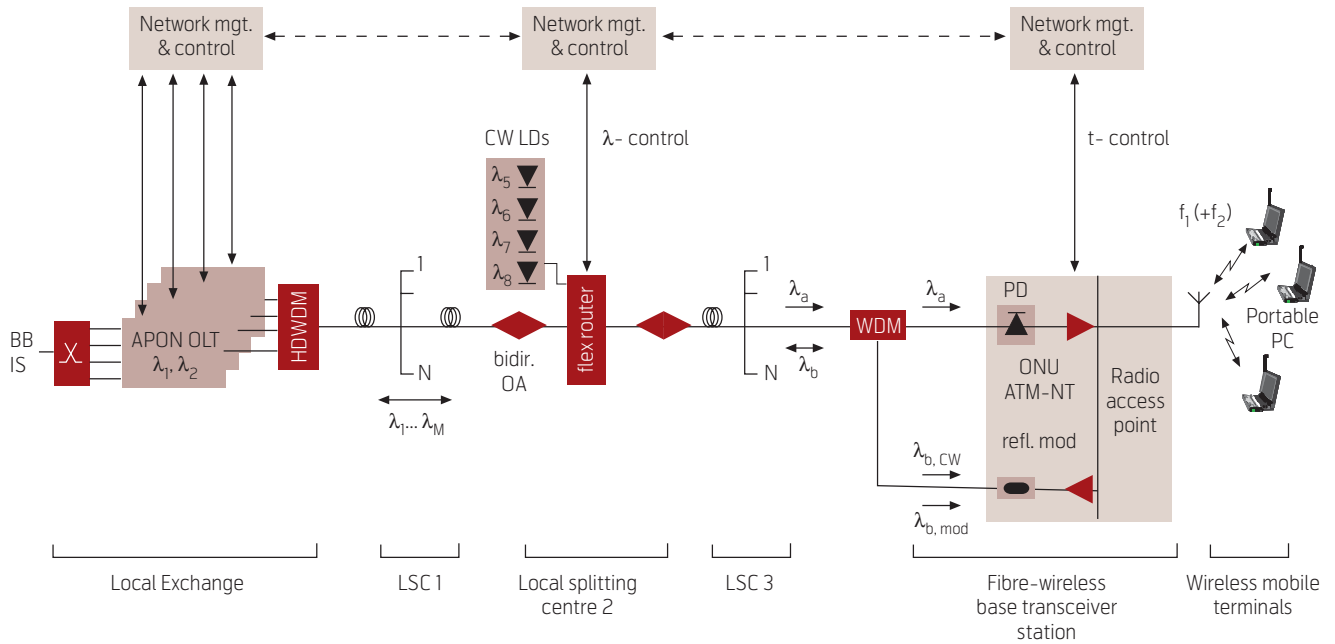


Figure 16 Flexible capacity assignment in a multi-wavelength fibre-wireless network by wavelength routing in the field

cate flexibly the capacity of a number of ATM PON systems among ONUs in a single fibre split network infrastructure [10]. Each ONU feeds a radio access point (RAP) of e.g. a wireless LAN, which wirelessly connects to a variable number of users with mobile terminals. These users move across the geographical area served by the network (e.g. a business park), and they may want to set up a broadband wireless connection to their laptop at any time anywhere in this area. When many users are within a wireless cell served by a certain RAP, this cell may have to handle much more traffic than the other cells; it has become a “hot spot” which has to be equipped with additional capacity. The corresponding RAP may switch on more microwave carriers to provide this additional capacity over the air, and also has to claim more capacity from the ONU. This local extra capacity can be provided by re-allocation of the wavelength channels over the ONUs, which is done by a flexible wavelength router positioned in the field. Similar to the architecture of the wavelength-reconfigurable fibre-coax network in Figure 14.b, the architecture in Figure 16 (which was developed in the ACTS PRISMA project) has four 622 Mbit/s bi-directional APON OLTs with a specific downstream wavelength and an upstream one each. The four downstream wavelengths are located in the 1538–1541 nm range, with 100 GHz spacing, and the four upstream ones in the 1547–1550 nm range with the same spacing. The flexible wavelength router directs each downstream wavelength channel to one or more of its output ports, and thus via a split network to a subset of ONUs. The RAPs could operate with up to 5 micro-

wave carriers in the 5 GHz region, each carrying up to 20 Mbit/s ATM wireless LAN data in OFDM format. At the flexible router (or at the local exchange) a number of continuous-wave emitting laser diodes are located, providing unmodulated light power at the upstream wavelengths. The flexible router can select one of these upstream wavelengths and direct it to the ONUs that can modulate the signal with the upstream data and return it by means of a reflective modulator via the router to the local exchange. Thus no wavelength-specific source is needed at the ONU, the downstream light sources are shared by a number of ONUs, and all ONUs are identical, which reduces the system costs and the inventory issues. The flexible wavelength router can be implemented with a wavelength demultiplexer separating the wavelength channels, followed by power splitters, optical switches and power couplers in order to guide the channels to the selected output port(s). Depending on the granularity of the wavelength allocation process, the flexible router may be positioned at various splitting levels in the network.

Using a similar strategy as depicted in Figure 15 to assign wavelength channels to the ONUs, a statistical performance analysis has been performed of the blocking probability of the system. It was assumed that the total network served 343 cells, of which 49 were “hot spots”, i.e. generated a traffic load two times as large as a regular cell. It was also assumed that the system deployed seven wavelength channels, and that the calls arrived according to a Poisson process where the call duration and length were uni-




Figure 17 Improving the system performance by dynamic wavelength allocation

formly distributed. Figure 17 shows how the system blocking probability depends on the offered load (normalised on the total available capacity, which is 7 times 622 Mbit/s), using various system architecture options. When wavelength re-allocation would not be possible (i.e. static WDM) and all the 49 hot spots were positioned at cells served by ONUs assigned to the same wavelength channel, the blocking probability is obviously the worst. On the other hand, in the static WDM case when the 49 hot spots were evenly spread over the seven wavelength channels, the blocking probability is much lower (i.e. best case). Unfortunately, a network operator cannot know beforehand where the hot spots will be positioned, so in this static WDM situation the system blocking probability will be anywhere between the best case and the worst case, and no guarantee for a certain blocking performance can be given. When dynamic re-allocation of the wavelength channels is possible, however, the system can adapt to the actual hot spot distribution. Figure 17 shows that when the flexible wavelength router is positioned at the second splitting point in the network, the blocking performance is better than the best-case static WDM performance; but more importantly, it is also stable against variations in the hot spot distribution, and thus would allow an operator to guarantee a certain system blocking performance while still optimising the efficiency of his system's capacity resources. The blocking performance may be even better and more stable when positioning the flexible router at the third splitting point; however, this implies that the costs of the router are shared by less ONUs. Locating the router at the second splitting point is a good compromise between adequate improvement of the system blocking performance and system costs per ONU.

7 Microwave signals over fibre

Wireless communication services are steadily increasing their share of the telecommunication market. Next to their prime feature, mobility, they need to offer growing bandwidths to the end users. This entails also an increase of the radio carrier frequencies, which leads to smaller radio cell coverage (due to the increased propagation losses and line-of-sight needs). Wireless LANs in the 2.4 GHz range according to the IEEE 802.11b standard carry up to 11 Mbit/s, evolving up to 54 Mbit/s in the IEEE 802.11g standard. The IEEE 802.11a and the HIPERLAN/2 standard provide up to 54 Mbit/s in the 5.4 GHz range. Research is ongoing in systems that may deliver more than 100 Mbit/s in the radio frequency range well above 10 GHz (e.g. LMDS at 28 GHz, HyperAccess at 17 GHz and 42 GHz, MVDS at 40 GHz, MBS at 60 GHz, Fixed Wireless Access up to 60 GHz, etc.). Due to the shrinkage of radio cells at higher radio frequencies, ever more antenna sites are needed to cover a certain area such as the rooms in an office building, in a hospital, the departure lounges of an airport, etc. Thus more Radio Access Points (RAPs) are needed to serve e.g. all these rooms in an office building, and hence also a more extensive wired network to feed the RAPs. Instead of generating the microwave signals at each RAP individually, feeding the microwave signals from a central headend site to the RAPs enables to simplify the RAPs considerably. The signal processing functions can thus be consolidated at the headend site. Thanks to its superb broadband characteristics, optical fibre is an excellent medium to bring the microwave signals to the RAPs.

7.1 Heterodyning systems

Carrying multi-GHz analogue signals over fibre requires very high frequency optical analogue transmitters and receivers, including careful fibre dispersion compensation techniques. An attractive alternative avoiding the transport of multi-GHz intensity-modulated signals through the fibre is to apply heterodyning of two optical signals of which the difference in optical frequency (wavelength) corresponds to the microwave frequency. When one of these signals is intensity-modulated with the baseband data to be transported, and the other one is unmodulated, by optical heterodyning at the photodiode in the receiver the electrical microwave difference frequency signal is generated, amplitude-modulated with the data signal. This modulated microwave signal can via a simple amplifier be radiated by an antenna; thus a very simple low-cost radio access point can be realised, while the complicated signal processing is consolidated at the headend station.

This approach, however, requires two light sources with narrow spectral linewidth and carefully sta-

bilised difference between the optical emission frequencies. An alternative approach requiring only a single optical source is shown in Figure 18 [11]. The optical intensity-modulated signal from a laser diode is subsequently intensity-modulated by an external Mach Zehnder modulator (MZM) which is biased at the inflexion point of its modulation characteristic and driven by a sinusoidal signal at half the microwave frequency. At the MZM's output port a two-tone optical signal emerges, with a tone spacing equal to the microwave frequency. After heterodyning in a photodiode, the desired amplitude-modulated microwave signal is generated. The transmitter may also use multiple laser diodes, and thus a multi-wavelength radio-over-fibre system can be realised with a (tunable) WDM filter to select the desired wavelength radio channel at the antenna site. The system is tolerant to fibre dispersion, and also the laser linewidth is not critical as laser phase noise is largely eliminated in the two-tone detection process.

7.2 Optical frequency multiplying systems

An alternative approach to generate microwave signals by means of a different kind of remote optical processing, named optical frequency multiplying, is shown in Figure 19 [12]. At the headend station the wavelength λ_0 of a tunable laser diode is swept periodically over a certain range $\Delta\lambda_{sw}$, with a sweep frequency f_{sw} . Alternatively, the wavelength-swept signal can be generated with a continuous-wave operating laser diode followed by an external phase modulator that is driven with the integral of the electrical sweep waveform. The intensity of the wavelength-swept signal is on/off modulated with low frequency chirp by the downstream data in a symmetrically

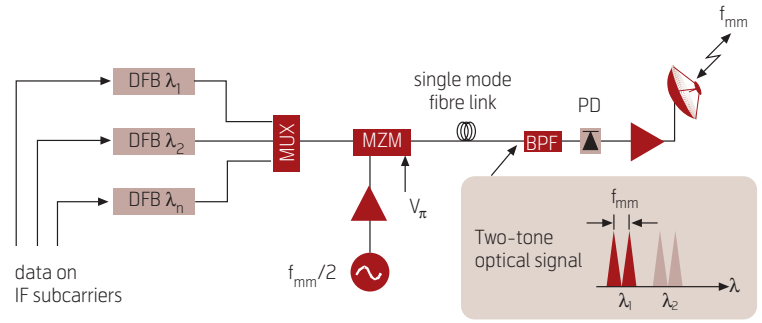


Figure 18 Generating microwave signals by heterodyning

driven Mach Zehnder modulator. After travelling through the fibre network, at the receiver the signal transverse an optical filter with a periodic bandpass characteristic. When the wavelength of the signal is swept back and forth over N filter transmission peaks, the light intensity impinging on the photodiode fluctuates at a frequency $2N \cdot f_{sw}$. Thus the sweep frequency is multiplied, and a microwave signal with carrier frequency $f_{mm} = 2N \cdot f_{sw}$ plus higher harmonics is obtained. The intensity-modulated data is not affected by this multiplication process, and is maintained as the envelope of the microwave signal. The microwave signal is subsequently transverse an electrical bandpass filter (BPF) which rejects the unwanted harmonics. Theoretical analysis and simulations have shown that the microwave signal is very pure; the inherent cancellation of the laser's optical phase noise causes the spectral linewidth of the microwave signal to be much smaller than that of the laser diode. Experiments have shown e.g. the generation of a microwave linewidth less than 50 Hz,

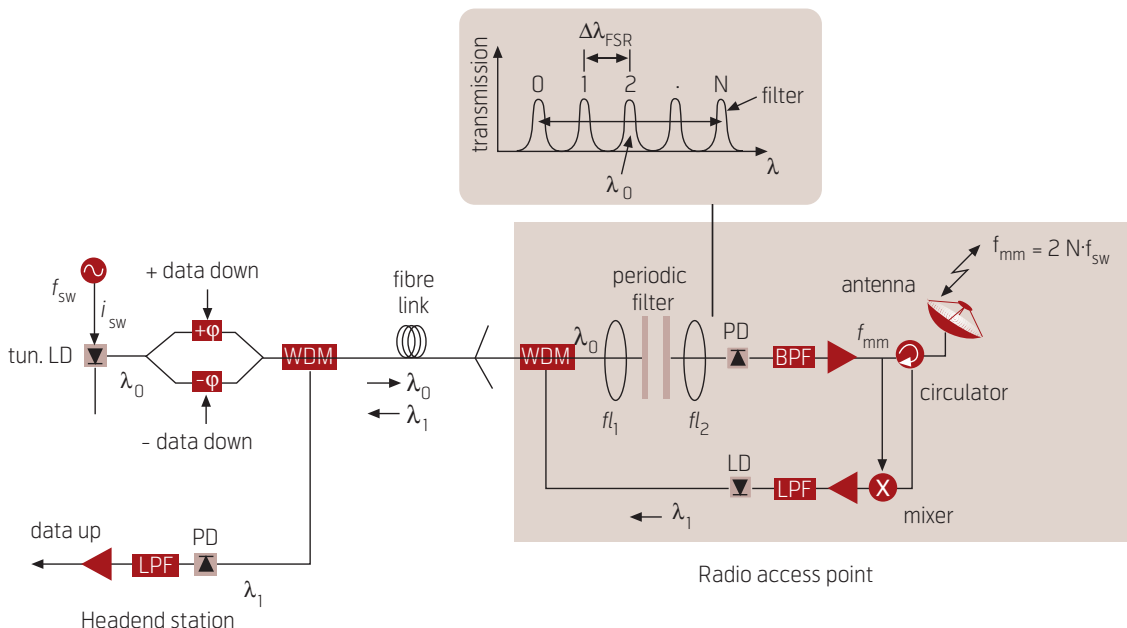


Figure 19 Generating microwave signals by optical frequency multiplying

whereas the laser linewidth exceeded 1 MHz [13]. The periodic optical bandpass filter can be advantageously implemented by a Fabry Perot filter with a free spectral range $\Delta\lambda_{FSR}$ which is N times as small as the wavelength sweep range $\Delta\lambda_{sw}$. The microwave signal can also carry advanced data modulation schemes, due to its spectral purity; e.g. 16-level Quadrature Amplitude Modulated (16-QAM) signals may be modulated on a subcarrier first, and then drive the Mach Zehnder Modulator. The main advantage of this optical frequency multiplying method is that only relatively moderate sweep frequency signals are needed at the headend site (e.g. an f_{sw} up to 1 GHz), which can easily be generated with low phase noise, while at the antenna site low phase noise microwave signals with carrier frequencies in the tens of GHz region are generated. The system does not rely on heterodyning, and thus may also operate on multimode fibre networks (such as polymer optical fibre, which is easy to install inside buildings).

Assuming a linear behaviour of the fibre, it can be shown that the periodic optical filter may also be positioned at the headend site, yielding the same optical frequency multiplication at the receiving end. Thus, the complexity of the antenna site is reduced further, and the characteristics of the filter can be readily tuned to the frequency sweep of the laser diode [14].

The system can also transport upstream data from the antenna station to the headend. By using a simple tunable local oscillator and a mixer, the microwave carrier generated at the antenna station by the optical frequency multiplying process can be down-shifted, and be used for down-converting the upstream microwave signal from the mobile terminal to a tunable intermediate frequency. This downconverted signal can subsequently be transmitted upstream on a separate wavelength by an IF-frequency optical transmitter [14]. Tuning to a different IF per antenna station enables the use of an SCMA upstream transmission protocol.

8 Concluding remarks

Optical fibre is now generally recognised to be the most powerful medium for transporting information, due to its very low losses and extremely wide bandwidth. Next to space and time multiplexing, the wavelength dimension offers unprecedented opportunities to extend not only the data traffic transport capacity, but also the traffic routing possibilities in networks.

In access networks, fibre is penetrating steadily towards the end user. Infrastructure costs are the major nut to crack here. Shared-feeder concepts such as the passive optical network tree-and-branch one

greatly reduce the installation costs of the fibre network, and can support various multiple access techniques (ATM and Ethernet for time multiplexed access, among others). In the past, operators have invested a lot in various last-mile networks (e.g. twisted pair, coaxial cable) to reach their residential customers. In upgrading these networks to higher capacity and larger service variety, fibre can support a wide range of last-mile technologies by hybrid combinations such as fibre-twisted pair, fibre-coax, and fibre-wireless. By means of wavelength multiplexing techniques, the fibre feeder part of such hybrid networks can very flexibly host different operators and service categories. Augmented with wavelength routing, capacity-on-demand can be realised for e.g. handling hot spots, while respecting Quality of Service requirements. By deploying advanced techniques for transporting microwaves over fibre, the antenna sites of broadband wireless systems can be significantly simplified and thus system costs be reduced, whereas also more sophisticated signal processing (for e.g. antenna diversity systems) is facilitated.

9 Future prospects

With the ongoing improvements in fibre characteristics and development of novel optical amplifier structures, the wavelength range available for communication will stretch from below 0.8 μm to beyond 1.6 μm , covering a bandwidth of some 200 THz. Further improvements in signal coding yielding a higher spectral efficiency, in ultra-dense wavelength division multiplexing and in ultra-high speed optical time division multiplexing will enable us to exploit this huge bandwidth, and will push the transport capacity of a single fibre beyond 100 Tbit/s.

The end user is to benefit from all these ultra-broadband communication possibilities. Therefore the optical fibre will not only reach up to his house, but it will penetrate into it as well. It will reach close to his personal area network, which thanks to the ongoing miniaturisation may consist of a myriad of small wireless power-lean terminals, sensors and actuators. These wireless devices will be incorporated not only in his residential living environment, but also in his clothes, his car, etc. Next to the traditional wired terminals such as the TV, these wireless devices will be connected to the fixed in-home and access network by a myriad of small intelligent antennas. Radio-over-fibre techniques, augmented with optical routing to accommodate dynamically the hot spots, will provide the best match of the ultimate capacity of fibre with the user freedom of wireless. Even in the wireless domain, optics may penetrate by means of intelligently-steered free-space light beams providing the ultimate in wireless transport capacity.

Which in the far foreseeable future will make the communication world an end-to-end globally transparent one, with nearly unlimited communication capacity for anybody, anytime, anywhere, for any kind of service ... the ultimate global crystal ball!

References

- 1 Frigo, N J et al. Demonstration of performance-tiered modulators in a WDM PON with a single shared source. *Proc. ECOC'95*, Brussels, Sep. 17–21, 1995, 441–444.
 - 2 Ueda, H et al. Deployment Status and Common Technical Specifications for a B-PON System. *IEEE Communications Magazine*, 39 (12), 134–141, 2001.
 - 3 Effenberger, F J, Ichibangase, H, Yamashita, H. Advances in Broadband Passive Optical Networking Technologies. *IEEE Communications Magazine*, 39 (12), 118–124, 2001.
 - 4 Schoop, R, Fredericx, F, Koonen, T, Hardalov, C. WDM isolation requirements for CATV in BPON. *Proc. ECOC'2002*, Copenhagen, Sep. 8–12, 2002.
 - 5 Vetter, P et al. Study and demonstration of extensions to the standard FSAN BPON. *Proc. ISSLS 2002 (XIVth International Symposium on Services and Local Access)*, Seoul, April 14–18, 2002, paper 4-2, 119–128.
 - 6 van de Voorde, I et al. The Super PON demonstrator: a versatile platform to evaluate possible upgrades of the G.983 APON. *Proc. ISSLS 2000 (XIIIth International Symposium on Services and Local Access)*, Stockholm, June 18–23, 2000, paper 14-3, 10.
 - 7 Van Caenegem, T et al. Prototyping ITU-T GPON, the new efficient and flexible FTTP solution. *Proc. ISSLS 2004*, Edinburgh, March 21–24, 2004.
 - 8 Koonen, A M J. Flexible wavelength multiplexing techniques for broadband fibre access networks. *Teletronikk*, 95 (2/3), 119–128, 1999.
 - 9 Koonen, T et al. TOBASCO: An Innovative Approach for Upgrading CATV Fibre-Coax Networks for Broadband Interactive Services. *IEEE Communications Magazine*, 35 (4), 76–81, 1997.
 - 10 Koonen, T, Steenbergen, K, Janssen, F, Wellen, J. Flexible reconfigurable fiber-wireless network using wavelength routing techniques: the ACTS project AC349 PRISMA. *Photonic Network Communications*, 3 (3), 297–306, 2001.
 - 11 Griffin, R A, Lane, P M, O'Reilly, J J. Radio-over-fibre distribution using an optical millimeter-wave/DWDM overlay. *Proc. OFC'99*, San Diego, February 22–25, 1999, paper WD6.
 - 12 Koonen, T et al. In-house networks using Polymer Optical Fibre for broadband wireless applications. *Proc. ISSLS 2002 (XIVth International Symposium on Services and Local Access)*, Seoul, April 14–18, 2002, paper 9-3, 285–294
- and
- Koonen, T et al. In-house networks using multi-mode polymer optical fiber for broadband wireless services. *Photonic Network Communications*, 5 (2), 177–187, 2003.
 - 13 Ng'oma, A et al. High-frequency carrier delivery to graded index polymer optical fibre fed next generation wireless LAN radio access points. *Proc. ECOC 2003*, Rimini, September 21–25, 2003.
 - 14 Koonen, T et al. Novel Cost-Efficient Techniques for Microwave Signal Delivery in Fibre-Wireless Networks. *Proc. ECOC 2004*, Stockholm, September 5–9, 2004.

Ton Koonen is since 2001 a full professor at the COBRA Institute at Eindhoven University of Technology in the Netherlands. Prior to that, he spent more than 20 years at Bell Labs in Lucent Technologies as a technical manager of applied research. He also is a Bell Labs Fellow. Next to his industrial position, he was a part-time professor at Twente University from 1991 to 2000. His main interests are currently in broadband fibre access networks and optical packet-switched networks. He has initiated and led a number of European and national R&D projects in this area, among others in dynamically reconfigurable hybrid fibre access networks for fibre-to-the-curb and fibre-to-the-building. Presently, he is involved in several projects in the European FP6 IST programme on access networks.

email: A.M.J.Koonen@tue.nl

Photonic integration

MEINT SMIT



Meint Smit is leader of the Opto-Electronic Devices group at the COBRA Research Institute

Since 2000 global network traffic is dominated by internet (IP) data, with growth rates in the order of 100 % per year. This explosive growth would have been impossible without the introduction of optical technology in the network. With the rapidly increasing bit rates required in the access network optical technology penetrates rapidly to the users and an increasing number of FTTH (Fibre-to-the-Home) pilot projects can be observed, with Japan and Korea taking the lead in large scale deployment. Close to the end users of the network cost is a major issue and we will have to provide a steadily increasing amount of bandwidth against an essentially constant price. This is only possible through a steady process of technical innovation. Photonic integration is the key technology for providing this steady innovation over a long term period and it is, therefore, a key enabler for a broadband access network.

Introduction

The field of Information and Communication Technologies is showing a development speed which is unprecedented in history. Over a period of more than thirty years key features like processor speed and memory size are roughly doubling each 18 months, and experts believe that this development will continue in the coming decade. It is known as Moore's law and enabled by the development of micro-electronic integration technology. Due to this development a steadily increasing performance of components and systems can be offered at an essentially constant price, a development that we recognise in the fact that the price of our new PC does not differ much from what we paid three years earlier for the previous one, which had only a fourth of the speed. This dynamics has become a major driver for the ICT market, which would collapse without this steady innovation.

Moore's law in micro-photonics

For Photonics to play a substantial role in the ICT market it will have to follow the same dynamics. The increase in functionality prescribed by Moore's law can only be sustained by applying an integration technology that supports a steady reduction of circuit size and fabrication costs. As such Indium-Phosphide (InP) based integration technology is the most powerful technology because it supports the integration of almost all functions required in ICT applications: light generation and detection in the wavelength window from 1300–1600 nm in which the attenuation and dispersion properties of optical fibres are optimal. And optical amplification, modulation, switching, signal regeneration and wavelength conversion, to mention just a few examples. And, not to forget, high speed electronic functions. Indium Phosphide is a compound semiconductor material with a structure very similar to silicon, but a number of advantages: it has much better properties than silicon for generation and ampli-

fication of light, and it has also better properties for realising high-frequency electronics. The question is whether InP-technology has the same potential as micro-electronic integration technology for reduction of device dimensions and fabrication costs over a longer period. To study that question we have put the development of Photonic ICs in our own lab in a graph with the potential integration density in devices per square centimetre on the vertical axis (Figure 1). The open circles mark the first publication of a device or a circuit. It starts with the invention of the Arrayed Waveguide Grating (AWG) wavelength demultiplexer in 1988 [1]. The next circle is the first InP-based Optical Add-drop Multiplexer (OADM) by Vreeburg in 1997 [2], a device that is capable of adding or dropping one wavelength out of a wavelength multiplex. This device integrated a single AWG with four optical switches on an area of 0.2 cm²; i.e. 25 components per square centimetre. As a next step we developed a technology for reducing the size of our AWGs using deep etching technology. This brought us the world's most compact Optical Cross-Connect (OXC) [3], a device that is capable of switching four wavelengths in a wavelength multiplex applied to the input ports individually to one of the two output ports. The device integrated four AWGs and four Mach-Zehnder switches on an area of 5 mm², i.e. an integration density of more than 150 components per square centimetre. Since then we succeeded in a further reduction of AWG-size close to the limits of conventional deep etched waveguide technology in InP: 250 x 350 μm² [4]. As can be seen the devices fit to a straight line with a slope slightly higher than Moore's law. A similar trend is observed in the development of the maximum demonstrated and installed transmission capacity on optical links. The last circle in the graph is a photonic flip-flop, consisting of two deep-etched micro-ring lasers, that was recently published by Hill et al. [5]. This flip-flop can be switched between two different output states (1 or 0) and without external

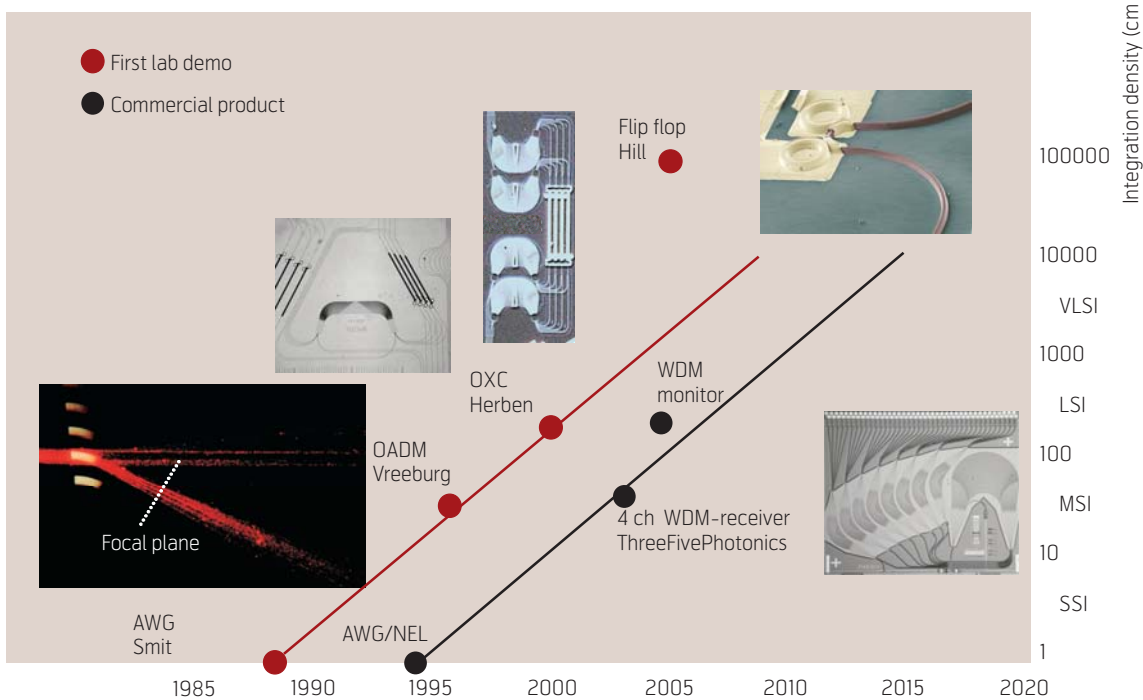


Figure 1 Moore's law in micro-photonics illustrated with devices developed by COBRA – TU Eindhoven. The development started with the realization of the first AWG in 1988, followed by a number of devices in which the AWG was integrated with optical switches in an Optical Add-Drop Multiplexer (1997) and a very compact Optical Cross Connect (2000). Commercial products based on this technology show a time lag around six years with respect to the first lab demonstrators. The most recent COBRA achievement, a fast and low-loss photonic flip-flop (2004), based on micro-ring lasers instead of AWGs and optical switches, does not fit into this scheme, but is ahead of the development by two orders of magnitude, which holds a great promise for the development of ultra fast and low power digital photonics

signals applied it keeps its state. It has a high potential for ultra fast photonic logic circuits operating at very low power levels. The flip-flop, with dimensions of $20 \times 40 \mu\text{m}^2$, allows for integration of more than 100,000 components per square centimetre, in principle. It will be explained in more detail later.

To complete the picture we have also put some (sparse) commercial products in the graph. The first one is the single AWG brought to the market by NEL in 1994. The second and third point are a Wavelength Division Multiplexing receiver (an AWG and four detectors) and a WDM Channel monitor (nine AWGs and 40 detectors) brought to the market by ASIP/ ThreeFivePhotonics in 2003 and 2004, respectively. Again the points fit more or less to a straight line which is delayed by around six years relative to the first demonstration in the lab. Although the sample is small, these examples suggest that Moore's law is also valid in micro-photonics integration technology, with the exception of the flip-flop, which is considerably ahead of the development predicted by Moore's law. It makes clear that the development of Photonics, which has a much larger diversity in devices and technologies than electronics, shows a more complex behaviour

than electronics. The difference in speed is in the positive direction, however: Photonics develops faster than Moore's law. This holds a great promise for the future.

Photonic integration philosophy

The power of micro-electronic integration technology is that a broad class of electronic functionalities can be synthesised from a small set of elementary components, such as transistors, resistors and capacitors. A technology that supports integration of these elementary components can, therefore, be used for a broad class of applications, and investments made in its development are paid back by a large market.

Although photonic integration has much in common with micro-electronic integration, a major difference is the variety of devices and device-principles in photonics. In photonics we have couplers, filters, multiplexers, lasers, detectors, switches, modulators, to mention just a few devices, and for each of these devices a broad variety of different operation principles and materials has been reported. It is impossible to develop a monolithic technology which is capable of realising even a modest subset of all these devices. The key to the success of integration in photonics is,




Figure 2 Three different schemes for integration of active (laser and amplifier) and passive waveguides. In the first one (left) active and transparent waveguides are in the same plane and are coupled via a so-called butt-joint. The structure requires three epitaxial growth steps, of which one is selective (i.e. the wafer is masked, so that the layers are only grown where they are needed). In the second (middle) and the third (right) scheme the whole structure is grown in a single epitaxy step and the active layers are removed afterwards, while covering the active regions with a mask. The single-step epitaxial growth schemes have the advantage of being simpler from a fabrication point of view. Both schemes have been successfully applied for development of commercial products. A disadvantage is the restricted flexibility: because the whole structure is grown in one step it is not possible to use different compositions and doping levels in different regions. This reduces the design freedom and the number of different components that can be integrated

therefore, reduction of the broad variety of optical functionalities to a few elementary components. We made a choice for 3 elementary devices: 1) A passive waveguide structure that allows for low-loss interconnection of devices and for realization of miniaturized components like couplers, filters, multiplexers, polarisation and mode converters. 2) An element for manipulating the phase of optical signals. A choice has been made for a fast electro-refractive modulator (ERM) that can be easily integrated in our passive waveguide structure. Main applications so far are fast optical switches and modulators (for both phase and amplitude). 3) An element for manipulating the amplitude of optical signals: the Semiconductor Optical Amplifier (SOA), a powerful but complicated device (it also modulates the phase) that allows for both linear and non-linear signal processing, e.g. in WDM light sources, femtosecond pulse lasers, wavelength converters, signal regenerators and ultra fast optical switches.

A generic technology for integrating these three elementary components can be used for realising photonic chips for a broad class of functionalities. A

major issue for Photonic Integration is the way in which active (laser and amplifier) sections are integrated with transparent (waveguide or modulator) sections. In principle three different schemes can be distinguished, each with a number of variants. They are illustrated in Figure 2. In the first one active and transparent waveguides are in the same plane and are coupled via a so-called butt-joint. The structure is realized by first growing the active layer stack, selectively removing it while covering the active regions with a mask, selectively regrowing the transparent regions, while covering the active regions with a mask, and covering the whole structure with a cladding layer in a third epitaxial growth step. The whole process requires good equipment and skills, because all layers should get an almost perfect crystalline structure with well controlled composition, morphology and dopant level. In the second and third scheme the whole structure is grown in a single epitaxy step and the active layers are removed afterwards, while covering the active regions with a mask. In the second scheme a thin active layer is applied on top of the transparent waveguide in such a way that the optical fields at both sides of the junctions are not too much different so that the transmission loss through the junction is small. Special provisions are necessary to avoid reflections. In the third scheme the light is coupled from the lower transparent layer to the upper active layer. This scheme needs special provisions for efficient coupling from the lower to the upper layer. The single-step epitaxial growth schemes have the advantage of being simpler from a fabrication point of view. Both schemes have been successfully applied for development of commercial products. A disadvantage is the restricted flexibility: because the whole structure is grown in one step it is not possible to use different compositions and doping levels in different regions. This reduces the design freedom and the number of different components that can be integrated. We have chosen for a three-step epitaxial regrowth process as shown at the left for integration of active (SOA-based) and passive devices because of the flexibility that it offers for the design of the regrown wafer stack in combination with the low loss and the low reflection level at the butt-joints [6]. Electro-refractive modulators are fully compatible with the passive waveguide structures and need no additional regrowth step [7].

Reducing dimensions

A major issue in the development of photonic integration technology is reduction of device dimensions. The key to reduction of device dimensions is the application of high lateral index contrast (i.e. the index contrast in the plane of the waveguide layer) realised by deep etching of the waveguides. Because

of the high index contrast between the semiconductor waveguide (typical refractive index values around 3.3) and air, deep etched waveguides provide strong light confinement. This allows for application of small waveguide widths, small bending radii, compact waveguide couplers and compact AWGs. The price that is paid for the high contrast are increased propagation losses due to scattering at rough waveguide edges, and increased insertion losses at junctions between slabs and waveguide arrays as used in an AWG where the finite resolution of the lithography causes abrupt closure of the tapered interwaveguide gaps. With a low index contrast the optical discontinuities caused by this closure are small; with deep etched waveguides they can introduce a few dB loss per junction. A solution is the use of a technology that supports combination of low loss shallow etched waveguides for interconnect purposes and at waveguide junctions, with locally deep etched regions, where short bends or other high-contrast functions are required. Losses smaller than 0.1 dB per junction have been shown at transitions between deep and shallow etched regions, so that several transitions can be included in a circuit. Using such a technology AWG dimensions have been reduced below 0.1 mm² [4], see Figure 3. This is close to the limits for InP-based ridge-type waveguide technology.

A further reduction of device dimensions may be achieved by moving to photonic wire and photonic crystal technology. These technologies allow us to come close to the fundamental size limits for photon confinement: around half a wavelength (200 nm) for a single defect micro-cavity; i.e. a cavity in which the light is confined in a volume with half wavelength dimensions. For full 3D confinement we should also confine the light in the third dimension, which is extremely difficult from a fabrication point of view. 2.5D solutions, i.e. application of membranes with a high vertical index contrast, may bring us close to the fundamental limits and have the advantage that they relax the requirements on the aspect ratio of the etching technology.

Resonant enhancement of electro-optic and opto-optic interactions

Once we have succeeded in reducing the dimensions of our circuits to the scale of the wavelength the next hurdle is the magnitude of our electro-optic or opto-optic interactions, e.g. in optical amplifiers or electrorefractive modulators. These are such that hundreds to thousands of wavelengths are required for achieving the π phase shift required for optical switching operations. Here a solution can be found in resonant enhancement in cavity-like structures. These are structures in which the light is reflected a few times




Figure 3 Deep etched InP-based AWG wavelength demultiplexer with record small dimensions. The device consists of two so-called diffraction regions (at both sides of the letters 35) and an array of waveguides which behaves as a lens with strongly dispersive properties, like a prism. The device dimensions have been reduced by etching the waveguides very deep, so that the light cannot escape from the waveguide. In this way very short bends can be realized. Without special precautions large reflections and transmission losses occur at the interface between the diffraction regions and the deep etched waveguide array. In the COBRA-device this problem has been avoided by inserting a short shallow edged region between the diffraction regions and the deep etched array

before it leaves the cavity. In this way it crosses the interaction region a number of times, which increases the efficiency, albeit at the price of reduced bandwidth.

Resonant enhancement is a key feature in the operation of the micro-ring laser flip-flop [5] shown in Figure 4. A flip flop is a device that can be switched between two states and that keeps its state without an externally applied signal. A micro-ring laser is a laser consisting of an active waveguide ring with a small radius ($< 100 \mu\text{m}$). By injecting sufficient carriers over the whole length of the waveguide ring the waveguide material will amplify the light travelling through the ring so that it becomes stronger after each round trip until it saturates at a high power level. In the flip-flop both rings can lase in the clockwise (CW) or counter-clockwise (CCW) direction. They




Figure 4 A photonic flip-flop based on coupled micro-ring lasers realised at COBRA. Both ring waveguides are covered with a gold contact layer, which is used to inject carriers into the rings in order to achieve amplification of the light. In the flip-flop both rings can lase in the clockwise (CW) or counter-clockwise (CCW) direction. They push each other in one direction by injection through the coupling waveguide section between them. Using a set pulse they can be changed from CW to CCW and vice versa. If the cavity size is sufficiently small (a few microns) response times smaller than 1 ps are feasible. By etching deep through the active waveguide layer cavity radiation losses can be kept sufficiently small even for very small ring diameter

push each other in one direction by injection through the coupling waveguide section between them. Using a set pulse they can be changed from CW to CCW and vice versa. This switching can be achieved with a signal that is much smaller than the output signals of the ring because of resonant enhancement of the input signal in the amplifying rings. Without this resonant enhancement the input signal should be larger than the output signal which is prohibitive for cascading devices. The device shown in Figure 4 needs only 5 femtoJoule light energy to switch to the other state. The price that is paid for this enhanced sensitivity is a longer response time. The device shown in the figure has a switching time of 15 ps. If the cavity size is sufficiently small (a few microns) response times smaller than 1 ps are feasible. By etching deep through the active waveguide layer cavity radiation losses can be kept sufficiently small even for very small ring (or disk) diameters.

Active or passive

An alternative to the active flip flop shown in Figure 4 is a passive micro-resonator; this is a small resonator in which the light is strongly confined so that the intensity of the light in the resonator can be much

larger than the intensity of the input light. Bistable operation in passive resonators needs high optical fields to build up in the resonator, especially if weak non-linearity like the Kerr effect is used. Input power levels required exceed those in the active flip-flop by more than two orders. The high optical powers also cause a signal induced heating of the cavity. So far, only thermal bistable action has been shown, which is very slow. When using lasers with injection locking, only a small amount of optical power need be injected to switch from one mode to the other. So warming due to changes in input signal will not be so important. Also the CW pumping of the cavity raises the cavity temperature, making cooling more efficient, and so any input light has a harder time raising the cavity temp. Lasers also have the advantage that the gain medium counteracts cavity losses, so that the input light can be used very efficiently. All this provides active devices with an important, if not decisive, advantage over passive resonators for application in large scale integration.

From analog to digital

An important feature of micro-lasers is that they can interact with each other to form low-power digital optical functions. And using the resonant enhancement of the input light (injection locking), they can drive each other. Therefore, coupled micro-cavity lasers are ideal candidates for development of digital photonic circuits. If successful, this will open the way for realizing an increasing class of functions in a digital way, in the same way as it has happened in microelectronics. This will relax the requirements on the integration technology: the number of elementary components may be reduced to flip-flops and interconnecting waveguides. Such a development will stimulate the penetration of photonic integration technology: it reduces the complexity of the integration technology and it enlarges its application range.

To the limits

The flip-flop as shown in Figure 4 has been designed conservatively in order to enhance the chance for successful fabrication. A next step will be the reduction of its size to a few micron diameter (micro disk). This will increase the switching speed into the ps region and reduce the switching energy below a femtoJoule. Ultimate dimensions may be reached by realization of the flip-flop in photonic crystal technology. This will bring us close to the ultimate limits achievable for the photons: confinement of the photons on a scale of half wavelength. Another important issue that has to be solved is the electrical properties of the active material in the micro-cavity, which should be superior in order to keep threshold currents sufficiently

low for allowing dense integration. Here the combination of ultimate confinement of the photons in photonic crystal micro-cavities and the ultimate confinement of the carriers in Quantum Wells and Quantum Dots will bring us close to the fundamental limits and may open new opportunities for photonic signal interactions.

Conclusions

Deep etching technology is the key to reduction of device dimensions. By improving the performance of lithography and etching technology circuit complexities into the VLSI range are feasible and, in the limit of photonic crystal technology, even larger. Deep etching technology is also the key for increasing electro-optic and opto-optic interactions, by resonant enhancement of these fundamentally small effects. In a broader sense control of dimensions on a nanometre scale, both via etching and via selective epitaxial growth, will bring us to the fundamental limits of photonic technology. Moore's law will provide the dynamics for achieving that goal.

References

- 1 Smit, M K. New focusing and dispersive planar component based on an optical phased array. *Electronics Letters*, 24, 385–386, 1988.
- 2 Vreeburg, C G M, Uitterdijk, T, Oei, Y S, Smit, M K, Groen, F H, Metaal, E G, Demeester, P, Frankena, H J. First InP-based reconfigurable integrated add-drop multiplexer. *Photonics Technol. Lett.*, 9 (2), 188–190, 1997.
- 3 Herben, C G P, Leijtens, X J M, Leys, M R, Groen, F H, Smit, M K. Extremely compact WDM cross connection InP. *Proc. 5th Ann. Symp. of the IEEE/LEOS Benelux Chapter*, Delft, The Netherlands, Oct. 30, 2000, 17–20. (ISBN 90-9014260-6)
- 4 Barbarin, Y, Leijtens, X J M, Bente, E A J M, Louzao, C M, Kooiman, J R, Smit, M K. Extremely small AWG demultiplexer fabricated on InP by using a double-etch Process. *IEEE Phot. Techn. Lett.*, 16, Nov. 2004, 2478–2480.
- 5 Hill, M T, Dorren, H J S, de Vries, T J, Leijtens, X J M, den Besten, J H, Smalbrugge, E, Oei, Y S, Khoe, G D, Smit, M K. A fast low-power optical memory based on coupled micro-ring lasers. *Nature*, 432, 11 November 2004, 206–209.
- 6 Binsma, J J M, van Geemert, M, Broeke, R G, Bente, E A J M, Smit, M K. MOVPE waveguide regrowth in InGaAsP/InP with extremely low butt-joint loss. *Proc. 2001 Annual Symposium IEEE/LEOS Benelux Chapter*, Brussels, 3 December, 2001, 245–248.
- 7 den Besten, J H, Caprioli, D, van Dijk, R, van Vliet, F E, Pascher, W, Binsma, J J M, Smalbrugge, E, de Vries, T J, Broeke, R G, Oei, Y S, Bente, E A J M, Leijtens, X J M, Smit, M K. Integration of MZI-modulators and AWG-based multiwavelength lasers in InP. *Proc. 2004 Annual Symposium IEEE/LEOS Benelux Chapter*, Gent, Belgium, 2–3 December 2004, 95–98.
- 8 Gong, Q, Nötzel, R, van Veldhoven, P J, Eijkemans, T J, Wolter, J H. Wavelength tuning of InAs quantum dots grown on InP (100) by chemical-beam epitaxy. *Appl. Phys. Lett.*, 84, 275–277, 2004.

Meint K. Smit was born in Vlissingen, the Netherlands, in 1951. He graduated in Electrical Engineering from the Delft University of Technology in 1974, and received his PhD in 1991, both with honours. In 1974 he started work on radar and radar remote sensing. Joined the Delft University of Technology in 1976. Switched to optical communications in 1981, research field: Integrated Optics. Became leader of the Photonic Integrated Circuits group of the Delft University in 1994. He is the inventor of the Arrayed Waveguide Grating for which he received a LEOS Technical Achievement award in 1997. Appointed professor in 1998. In that year his group, together with the groups of profs. Wolter and Khoe (COBRA, TU Eindhoven) received a research grant of 40 M€ for establishing a National Research Center on Photonics. In 2002 he moved with his group to the Technical University of Eindhoven, where he is presently leader of the Opto-Electronic Devices group at the COBRA Research Institute. In 2002 he was appointed LEOS Fellow for contributions in the field of Opto-Electronic Integration.

email: m.k.smit@tue.nl

Advancements in metro optical network architectures and technology

LOUKAS PARASCHIS, ERROL ROBERTS AND ORI GERSTEL



Loukas Paraschis is in the Optical Technology CTO group at Cisco Systems

This paper summarizes the innovation in network architectures, and optical transport that has enabled metropolitan networks to meet the diverse service needs of enterprise and residential applications, and cost-effectively scale to hundreds of Gb/s of capacity, and hundreds of km of reach. A converged metro network, where IP/Ethernet services and traditional TDM traffic operate over a common intelligent WDM transport layer, has become the most appropriate architecture for significantly reduced network operational cost. At the same time, advanced technology and system-level intelligence have improved the deployment and manageability of WDM transport. The most important application drivers, system advancements, and associated technology innovations in metropolitan optical networks are being reviewed.



Errol Roberts is Distinguished Systems Engineer at Cisco

1 Introduction

Service providers are being challenged by the increasingly higher cost of building and maintaining networks that serve a plethora of new applications with diverse quality of service requirements. Moreover, overall traffic continues to grow exponentially [1], but service provider revenue do not increase accordingly, as bandwidth for most of the new (broadband) services has a significantly lower price (\$/b) [2]. Therefore, the main priority for network operators has been to rationalize their cost structure, reducing capital (CapEx) and operational (OpEx) expenditure.



Ori Gerstel is leading the Advanced Optical Technology Groups at Cisco Systems

Network CapEx and OpEx increased as most operators initially developed separate networks for the rapidly growing (packet-based) data services, and the traditional (circuit-switched) leased line, and voice services (Figure 1). Most applications have been increasingly converging to a common IP layer, allowing a unified network architecture based on an IP/MPLS core, with a packet-aware SONET/SDH

and native Ethernet access over a common optical transport infrastructure (Figure 2), to significantly improve the network cost.

Metropolitan area networks (MANs) have been the most appropriate initial convergence point for multi-service architectures (Figure 3). The significant growth of applications with extensive metropolitan networking requirements has further placed increased emphasis on the scalability of multi-service MANs [3]. MAN architectures have thus evolved; inter-networking multiple “access” traffic-collector fiber rings in “logical” star or mesh through a larger “regional” network (Figure 3) [4]. A regional metro ring often extends to hundreds of km, interconnecting many (typically 5–10) access-ring hub nodes [5]. A successful generation of systems addressed the initial metro needs for efficient bandwidth provisioning, leveraging the advancements of “next-generation” SONET/SDH transport [6, 7]. Recently, wavelength-division-multiplexing (WDM) has been additionally utilized in MANs [8]. We presently make the case for a converged metro network architecture comprising of Ethernet/IP and traditional TDM services operating over a common intelligent WDM layer, that can help reduce CapEx, and even more importantly OpEx, by enabling easier deployment and manageability of services [9, 10]. In the following sections, we review the applications and services that motivated the MAN evolution. We then evaluate the innovation in network architectures and optical transport that enables metropolitan networks to meet the diverse service needs of enterprise and residential applications, and cost-effectively scale to hundreds of Gb/s of capacity, and hundreds of km of reach.

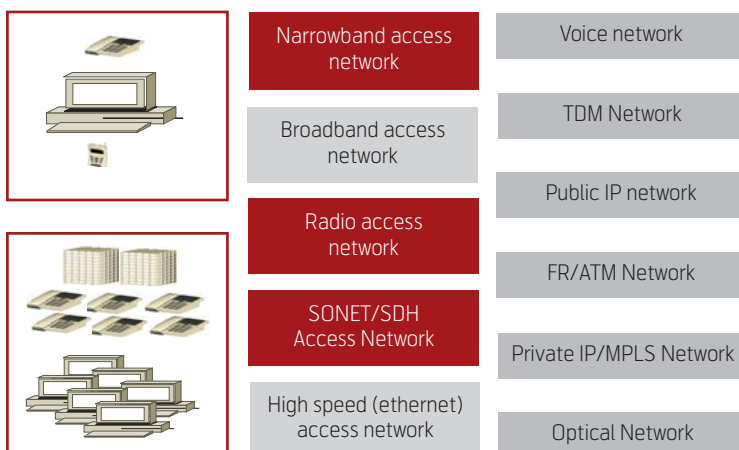


Figure 1 Traditional architecture with multiple service-specific networks, leads to high operational cost (OpEx)

2 Metro network applications and services

Service providers have traditionally relied on different networks (and multiple layers) to address consumer and enterprise market needs. POTS, TDM/PSTN, and to some extent ISDN, have typically served the voice dominated consumer applications, while Frame Relay, ATM and TDM leased-line networks have served the more data intensive enterprise applications. At the same time, video has been mostly distributed on separate, extensively analog, networks. Over the last few years, however, most of these applications have migrated to internet protocol (IP) based services. Voice over IP (VoIP), Ethernet based VPN data services, along with the fast emerging digital video over Ethernet have increasingly been using MPLS/IP L2/L3 backbones and metro Ethernet transport.

Moreover, the explosive growth in e-commerce, data warehousing, and supply chain management applications has fueled significant growth in the enterprise network connectivity and storage needs. The business critical nature of most of these applications also called for uninterrupted and unconstrained connectivity of employees and customers. To best support business continuance and high bandwidth applications, server consolidation in data center facilities, and application hosting by every location (often extending to metropolitan reach), most enterprises have upgraded their networks, replacing ATM, Frame Relay, and TDM private line with more efficient Ethernet (GE and 10GE) transport. In addition, regulatory mandates in the financial and insurance industries significantly increased the bandwidth needed to support data back up and disaster recovery. The large financial firms became the early adopters of enterprise WDM metro networks, driven by the need to

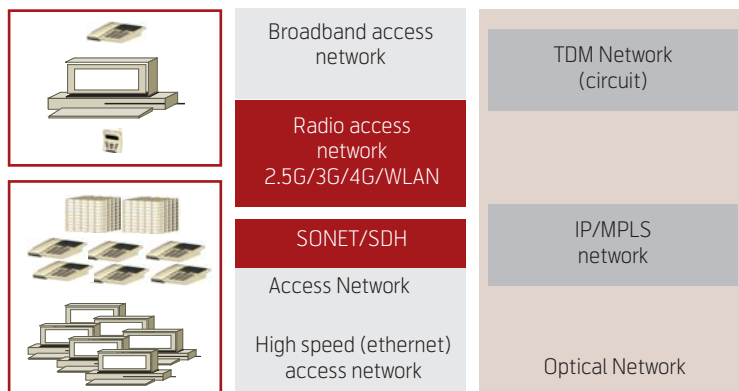


Figure 2 Emerging converged optical multi-service network architecture, significantly improving operations (OpEx)

support very high bandwidth storage applications for disaster recovery and synchronous data replication over metro (often long) distances.

WDM offers the most scalable and transparent transport for the different storage protocols (ESCON, Fibre Channel, Ethernet and FICON), without the additional latency introduced in storage extension over SONET/SDH (with intermediate GFP encapsulation). The introduction of enterprise WDM networks coincided also with the emergence of “dark” fiber providers in the late 1990s, which offered enterprise customers the ability to own or lease fiber assets or purchase managed wavelength services. Enterprises have mostly chosen to outsource optical transport (not being part of their core competence) to managed WDM service providers, introducing additional competition to traditional carriers and their TDM leased-line revenue.

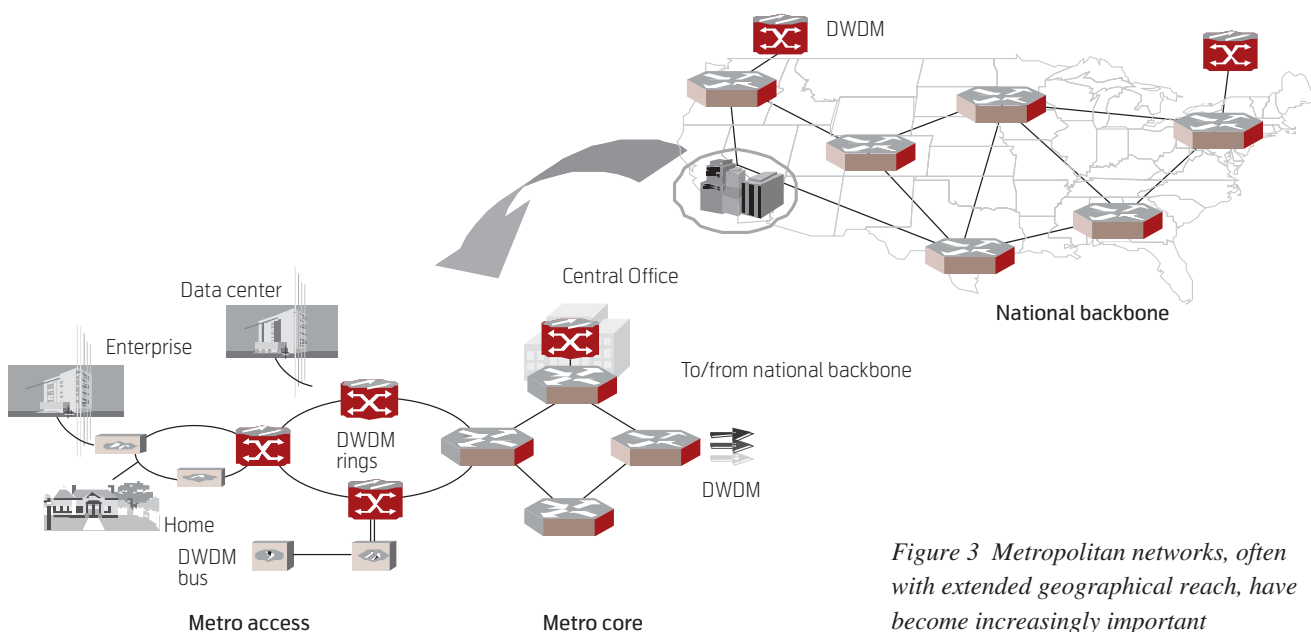


Figure 3 Metropolitan networks, often with extended geographical reach, have become increasingly important

The explosive growth of all these applications has called for the re-engineering of the traditional metro networks which had initially been designed primarily to efficiently transport 64 kb/s voice channels and support a rather slow and predictable traffic growth. The new network aggregates traffic from different “access” points; today typically from IP/GE DSLAMs over twisted pair copper lines in the case of traditional local exchange carriers (ILECs), or from cable modem termination systems (CMTS) over hybrid fiber coax (HFC) lines in the case of cable operators, and in the future from passive optical network (PON) links. Such residential traffic, along with traffic from the enterprise networks, forms the multi-service, predominantly packet-based, metro optical backbones which eventually leverage WDM transport. Metro WDM allows different services including the traditional (slower growth) TDM traffic, to be carried on separate wavelengths as part of the same packet-optimized transport network.

3 Metro optical transport convergence

WDM was early acknowledged as the most promising technology for scalable metro networks [11]. Its initial deployment, however, remained rather limited, addressing mostly fiber exhaustive applications. Metro optical transport initially leveraged technology primarily developed for long-haul (LH). In LH WDM transport, however, the main objective has been to maximize fully deployed network performance, typically measured in “cost per bit/s/km” or even “cost per bit/s per km per Hz”. Also, LH networks typically have well-defined, rather simple point-to-point topologies, allowing system performance to be maximized in advance, and often optimized on a case-by-

case analysis. On the other hand, in metro networks, service flexibility and operational simplicity are very important. Moreover, metro WDM cost has to account not only for a fully deployed network, but also for its ability to scale with the amount of deployed bandwidth; as most metro networks do not employ all (or most) WDM channels at the initial deployment phase, but rather “lit” unused channels when needed to serve future (often unpredictable) network growth.

The first generation of metro WDM systems, mostly based on LH technologies, had high fixed initial capital cost, inflexible (stranded) bandwidth provisioning, and complex operations. Subsequent “point-product” solutions attempted to lower the fixed-capital cost (i.e. CapEx) at the expense of bandwidth scalability, but suffered from even higher operational complexity – hence higher OpEx). More generally, initial metro WDM was limited by lack of “plug-and-play” operational simplicity, requiring service interruption for unplanned network growth and high CapEx due to unnecessary optical-electrical-optical (OEO) conversions and the associated high inventory cost of “fixed” WDM. As a result, network operators delayed WDM deployment in their metro networks until these problems were solved with mature technologies from a stable supply base.

At the same time, the evolution of the SONET/SDH transport standards [6] enabled a successful generation of systems that enabled efficient bandwidth provisioning, addressing most of the initial MAN needs, leveraging the advancements in electronics, and 2.5 Gb/s (STM-16) and 10 Gb/s transport (STM-64) [6, 7]. These “next generation” SONET/SDH systems further allowed significantly improved packet-based

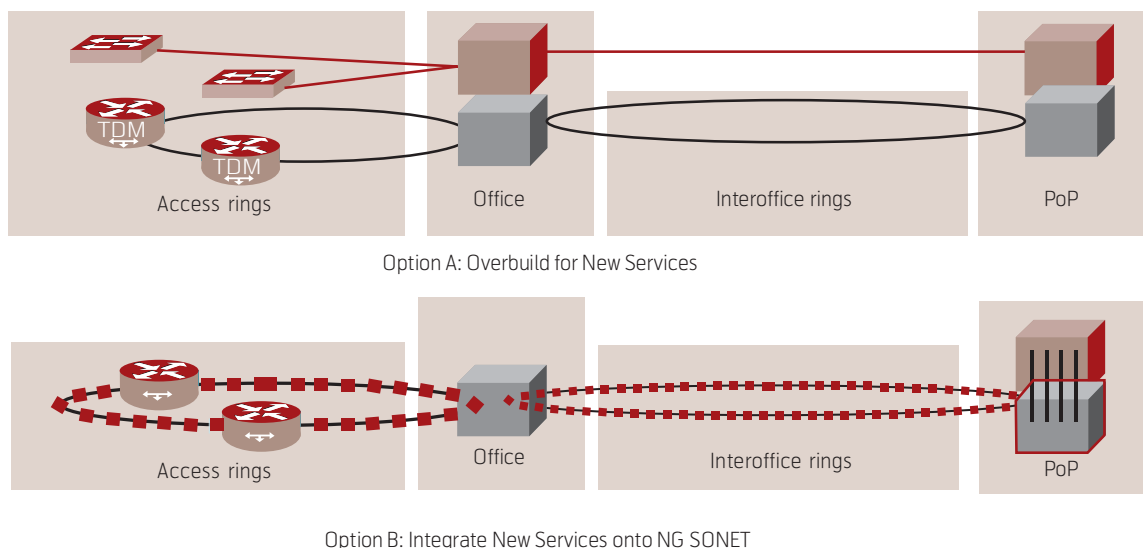


Figure 4 Integration of TDM and Data service in Metro networks

VCAT only	Required VC4:	ML + VCAT
91	Ring 1	30
84	Ring 2	30
49	Ring 3	26
21	Ring 4	8
245	Total VC4	94
(6) 75 % full	STM64 Rings	(4) 40 % full

transport over the existing time-division-multiplexed (TDM) infrastructure (Figure 4B), based on the new data encapsulation protocols (GFP, VCAT, LCAS) [11]. Packet-aware service provisioning enabled Ethernet “virtual” private network (VPN) over a common service provider MAN. The initial rate-limited best-effort Ethernet service architectures evolved to offer quality of service guarantees for Ethernet, as well as IP services (like voice-over-IP) and packet-aware ring architectures, like the Resilient Packet Ring (RPR) IEEE 802.17 standard, provided bandwidth spatial reuse. Such layer 2, and eventually layer 3, intelligent multi-service provisioning has enabled significant statistical multiplexing gains, enhancing network scalability. The table in Figure 4 illustrates an example [2] for a network with VC-4 granularity that serves a VPN with four Gigabit Ethernet (GE) sites, six additional point-to-point GE connections, and a storage area network with two GE and two fiber-channel (FC) services. A “purely” optical solu-

tion would require at least six STM-64 rings that, even after leveraging VCAT, would be at least 75 % full. An advanced multi-layer implementation (ML+VCAT) could be based on just 4 STM-64 rings, each with less than 40 % of capacity utilization, saving more than 50 % in network capacity. This new generation of multi-service platforms allowed for the first time different services to be deployed over a common network infrastructure, instead of separate networks (Figure 4 A), also significantly improving network operations – and reducing OpEx.

As traffic needs grew beyond 10 Gb/s per fiber, however, WDM transport became the best alternative for network scalability. To this end, multi-service systems evolved to “incorporate” WDM interfaces that connect them directly on the metro fibers, thus eliminating any unnecessary optoelectronic (OEO) conversions and the associated OEO cost. The integration of WDM interfaces in the service platforms also changed the traditional “service demarcation point” in the network architecture. This seemingly straightforward convergence of the transport and service layers has introduced additional requirements for improved manageability in the WDM transport. The introduction of many different “wavelength services” amplified the value for “open” WDM architectures that provide robust and flexible transport. In this sense, a converged, flexible WDM metro transport architecture that supports all the different services with the lowest possible OpEx, leveraging elaborate planning and operational tools, and enabling standards-based interoperability, has become increasingly important (Figure 5).

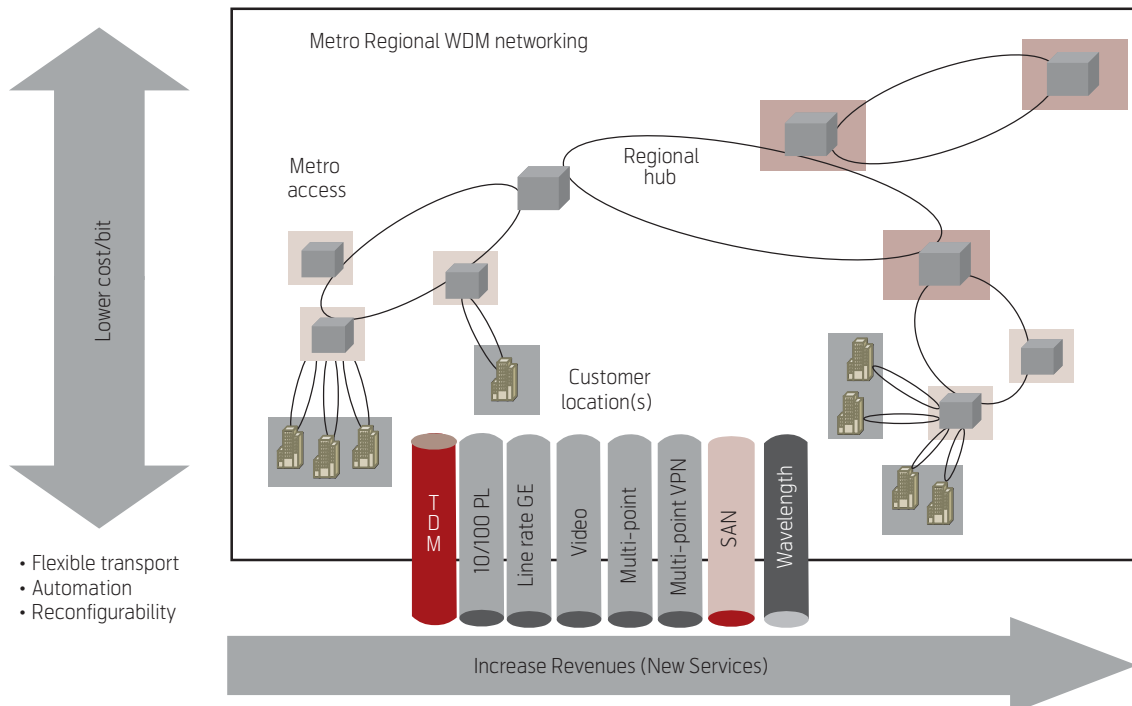


Figure 5 Converged Multi-service Metro Optical Networks

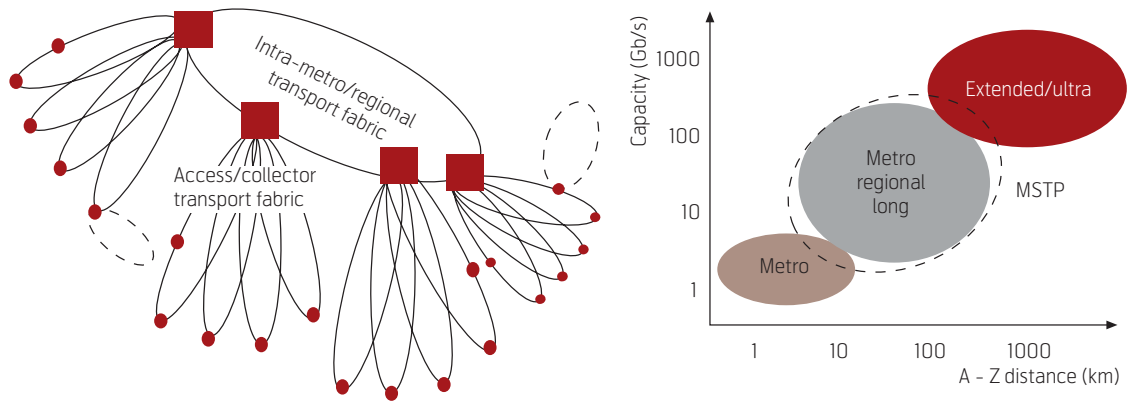


Figure 6 Metro Regional Network benefit from flexible multi-service optical transport that could extend to hundreds of km and Gb/s

4 Flexible WDM transport

A new generation of metro-optimized WDM transport (often referred to as Multi-Service Transport Platforms or MSTPs) has contributed significantly to the recent progress in MAN WDM deployments. WDM transport enables elaborate optical add-drop multiplexing (OADM) architectures that transparently interconnect the different MAN nodes (Figure 6 – left) [11, 4]. MSTP WDM systems have further scaled cost-effectively to hundreds of Gb/s, and to hundreds of km [9]. Even more so, these metro-optimized systems have significantly enhanced ease of deployment and operation, by automated control and integrated management of the optical transport layer.

Although the present analysis is not intended to elaborate the full technical details of MSTP innovation (which are available in [9, 12, 14, 17]), it is important to identify that the evolving nature of metro networks, from simple access to regional reach, places an additional significant challenge on the design of a WDM system that offers span flexibility, operational simplicity, and low cost even at a large number of nodes. The key to a robust and cost-effective implementation of a system carrying hundreds of Gb/s over hundreds of km, is to leverage mature WDM technologies, most notably erbium-doped-fiber amplifiers. In this sense, optical amplification noise (ASE) and chromatic dispersion, at 10 Gb/s or higher line rates, are the primary physical impairments that limit the system performance. More complex transport designs, with multiple OADM nodes, could also be limited by amplifier gain ripple or transients, filter concatenation and loss variation, span loss variation due to ageing or repairing, and in the case of mesh networks, ASE lasing [9]. Therefore, metro-optimized solutions need to address the trade-off between service flexibility and optical signal-to-noise-ratio (OSNR), guaranteeing the bit-error-rate (BER) performance requirements (typically $BER < 10^{-15}$),

while minimizing inflexible channel banding and cost.

Most notably, the LH practice of having tight control of the channel power along the link using per channel gain equalizer would lead to a significant cost penalty. Building on the fact that most physical impairments of interest impact the total optical power, we have designed a cost-effective WDM system that meets the metro regional performance requirements by monitoring and regulating the multiplex power, unless the power of individual channels is already available, at selected critical points along the optical paths [9, 14]. Simple optical tap couplers and photodiodes are used to monitor the optical power, and automatic control software has been developed to regulate optical power within each node by adjusting appropriately the settings of in-line variable optical attenuators (VOA) and the optical amplifier gain. This metro-optimized WDM design and control, along with currently available EDFA subsystems with variable gain and variable-loss mid-stage access, have enabled 32 x 10 Gb/s WDM transport to cost-effectively scale to > 1000 km [14]. Furthermore, the control layer is able to self adjust each node, enabling automatic “plug and play” node setup during the installation phase. The same system design innovation enables automatic adjustment for ageing effects and/or traffic changes.

In [9] we reported in detail results from this innovative and practical WDM system that leverages such an automatic optical power control layer. We further analyzed the performance of a 32 x 10 Gb/s system in an actual 400 km ring (no recirculation loops) with eight optical nodes. We also demonstrated “stressed” system performance for a 10 Gb/s 480 km link in another network with 16 OADMs and highly non-uniform spans (Figure 7). Both results use commercially available optical components, and standard sin-

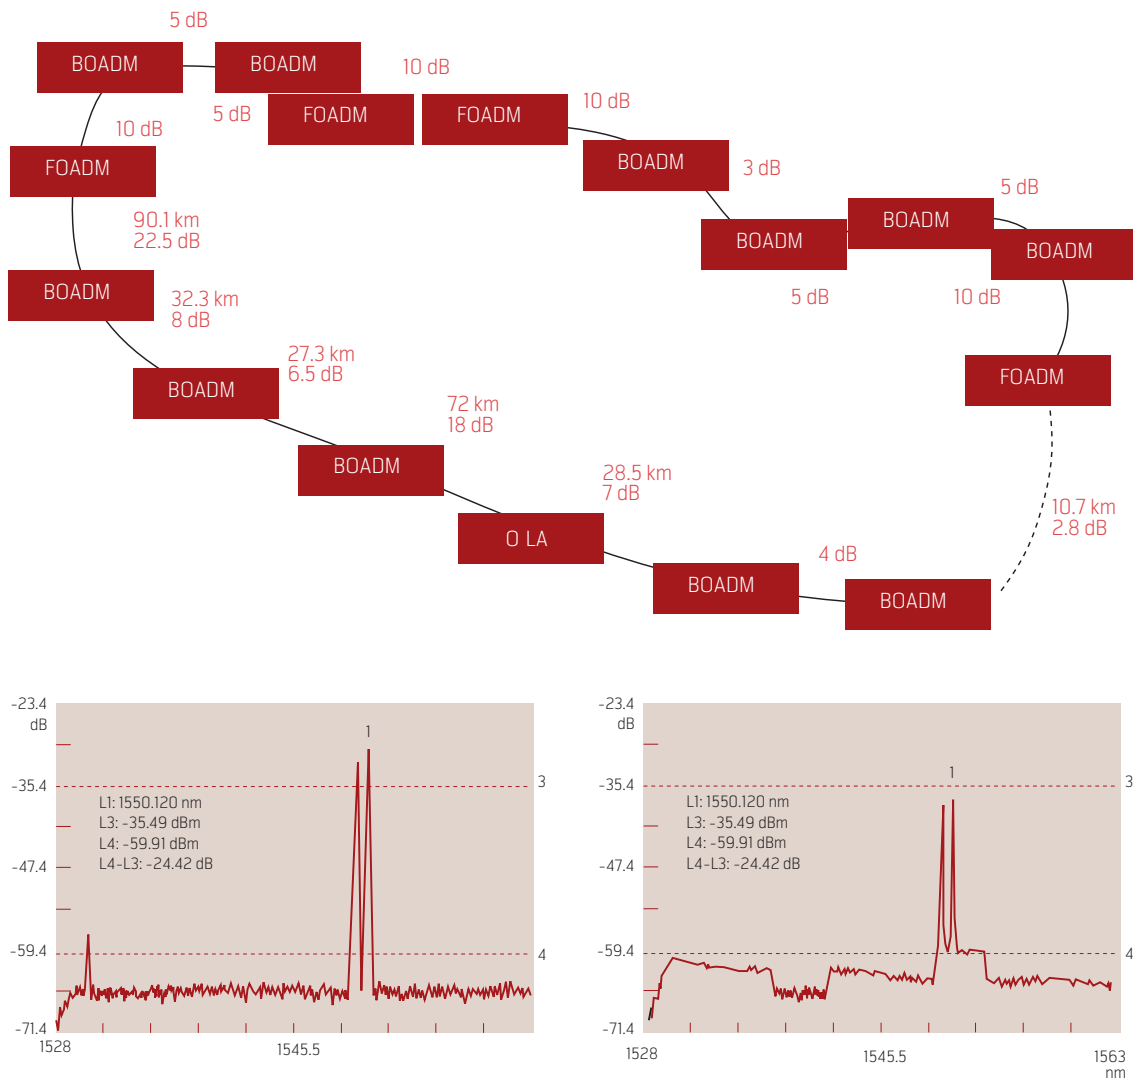


Figure 7 Spectra for two 10G channels in the worse-case path (120 dB loss) in the 16 node MRAN (top)

gle mode (G.652) fiber, and well exceed all previously published related demonstrations. This WDM system offers optical add/drop flexibility, span non-uniformity, and installation/maintenance simplicity, with 32 x 10 Gb/s transport performance that scales to 15 x 20 dB, for a total reach that could exceed 1000 km [14]. We also demonstrated such extended system reach with reconfigurable ROADMs (ROADMs) and analyzed the performance improvements arising from the corresponding better channel optical power control with ROADMs [14].

In addition to automated and scalable link engineering, ROADM flexible network design is also very critical to the success of metro WDM architectures. More specifically, wavelength-level add/drop and pass-through, with automated reconfigurability (ROADM) at each service node, is the only operationally robust solution for WDM deployments that support the uncertain (often unpredictable) future traffic patterns in MANs that scale to hundreds of

Gb/s. ROADM network flexibility also provides the ability to set up a wavelength connection without visiting any intermediate sites, thus also minimizing the risk of erroneous service disruptions during network upgrades. A few different optical switching device technologies have been proposed, including liquid crystals, MEM structures, tunable thin films, or etalons. Each technology offers different performance and functionality characteristics [14, 15]. In the context of the present analysis, it is more important to identify and distinguish between their two main functional characteristics at the network level: 1) ROADM solutions allow for switching of each individual wavelength between the WDM ingress and egress, potentially among more than one fiber facility. 2) More elaborate solutions could also allow extraction or insertion of any client interfaces to any wavelength of any fiber [14]. This latter solution has often been proposed in combination with predeployed tunable transmitters and/or receivers, to realize advanced network automation, via a GMPLS control

plane, enabling dynamic bandwidth provisioning and fast shared optical layer mesh protection. Current network deployments, however, are primarily interested in the ROADM functionality and the most cost-effective related technologies (rather than in the most advanced solution that would meet any conceivable future need, irrespective of price). In this sense, PLC-based wavelength-selective switches, with some switching capability among few fibers, have sufficient functionality to meet most customer needs while being the most mature and thus cost-effective technology [15].

Tunable laser technologies are also being increasingly employed in Metro WDM systems to reduce inventory cost and improve operations [16]. The choice of the appropriate transmitter technology is particularly important, as its cost usually dominates the total cost of a fully-deployed transport system [17]. Pluggable transceiver technology has also been increasingly maturing, offering significant advantages and promising to soon address the needs of even dense (100 GHz) WDM 10 Gb/s MANs [18]. At such high data rates, however, optical performance, predominantly dispersion tolerant (chirp minimized) modulation, becomes increasingly important. Further advancements in transmitter technology and electronics, most notably EFEC and EDE, promise to enable lower cost multi-service transceiver [17]. Such next generation transmitters that would be more easily integrated into the different service platforms are very important, not only for their enhanced performance or lower cost, but also for further simplifying the network architecture and thus reducing the overall network cost.

5 Further integration of IP over WDM

As IP traffic outpaces any other traffic type in networks worldwide, a tighter “integration” of WDM with the routers and switches (Figure 8) allows the most significant network simplification. Unlike TDM, which has unified management with optical transport, and also included WDM interfaces on SDH/SONET multiplexers and cross-connects, the IP layer has yet been mostly designed separately from the transport layer. Traffic between core IP routers, however, has increasingly started to exceed 2.5 or even 10 Gb/s in most metro and regional networks. Therefore, such IP trunks can be economically transported through multiple intermediate sites without requiring any (layer 1-3) electrical processing and additional cost due to electrical grooming between the routers and the WDM layer. Hence, a much lower cost solution is achieved by originating the WDM wavelengths directly from the router and transport




Figure 8 The network convergence at the IP and the WDM layers best minimizes CapEx and OpEx

them through the MSTP network, thus minimizing OEO (costly) conversions between the optical and electronic domains.

Pluggable WDM modules in routers and switches have recently enabled the first such “converged” deployment in the emerging IP-video MAN architectures [19]. The advances in optical switching and transmission technologies discussed in section 4 further allow a flexible optical infrastructure that efficiently transmits and manages the “optical” bandwidth, enabling advanced network architectures to leverage the IP-WDM convergence, and to realize the associated CapEx and OpEx savings. These architectures would ideally be based on open “WDM” solutions (like the ones described presently), so that the same WDM infrastructure can also continue to support traditional TDM traffic, as well as wavelength services or other emerging applications that may not converge over the IP network, e.g. high speed (4 or 10 Gb/s) fiber channels. Such open WDM architectures (Figure 8) further need to offer performance guarantees for the different types of wavelength services, including “alien” wavelengths, support mesh network configurations, and also benefit from coordinated management, and network control. These objectives currently constitute the most important advancements in the next generation of WDM transport [19].

6 Summary

In this paper, we discuss the innovation in network architectures and optical transport that enables metropolitan networks to cost-effectively scale to hundreds of Gb/s of capacity, and to hundreds of km of reach, meeting the diverse service needs of enterprise and residential applications. A converged metro network, where Ethernet/IP services, along with the traditional TDM traffic, operate over an intelligent WDM transport layer is increasingly becoming the most attractive architecture addressing the primary need of network operators for significantly improved capital and operational network cost. At the same time, the optical layer of this converged network has

to introduce intelligence and leverage advanced technology in order to significantly improve the deployment and manageability of WDM transport. We have reviewed the most important operational advancements and the technologies that cost-effectively enhance the network flexibility, and advance the proliferation of WDM transport in multi-service metro networks.

References

- 1 Coffman, K G, Odlyzko, A M. Growth of the Internet. In: Kaminow, I P and Li, T (eds.). *Optical Fiber Telecommunications IV B*. Academic Press, 2002, pp. 17–56. URL: <http://www.jpix.ad.jp/jp/technical/traffic.html>
- 2 Ramaswami, R. Market watch. In: *Optical Fiber Communication Technical Digest Series*, Conference Edition, Optical Society of America, 2004.
- 3 Special Issue on “Metro & Access Networks”. *J. Lightwave Technol.*, 22 (11), Nov. 2004.
- 4 Ghani, N et al. Metropolitan Optical Networks. In: Kaminow, I P and Li T (eds.). *Optical Fiber Telecommunications IV B: Systems and Impairments*. Academic Press, 2002, p. 329–403 (Chapter 14).
- 5 Madamopoulos et al. Design, transport performance study... In: *OFC Technical Digest Series*, Optical Society of America, 2002, Paper ThH6.
- 6 Cavendish, D. Evolution of Optical Transport Technologies ... *IEEE Comm. Magazine*, 38 (6), 164–172, 2002.
- 7 Paraschis, L et al. Challenges in ... In: *Proc. of Contemp. Phot. Techn.*, Japan, 2001, WA1
- 8 Plote, S. Metropolitan service provider business model for profitability. In: *Optical Fiber Communication Technical Digest Series*, Conference Edition, Optical Society of America, 2004. Paper WG3.
- 9 Macchi, M, Piccacia, S, Paraschis, L, Gerstel, O. Design and Demonstration of a Reconfigurable Metro-Regional WDM 32 x 10 Gb/s System Scaling beyond 500 km of G.652 Fiber. In: *Optical Fiber Communication Technical Digest Series*, Conference Edition, Optical Society of America, 2004. Paper OTuP2.
- 10 The authors would like to acknowledge the contributions of many people at Cisco Systems, as well as at service providers and partner companies.
- 11 Saleh, A et al. Architectural Principles of ... *J. Lightwave Technol.*, 17 (10), 2431–2448, 1999.
- 12 ITU-T. *Generic Framing Procedure*. Geneva, 2001. (ITU-T Draft Rec. G.7041)
- 13 Paraschis, L. *Optical Technology Innovation in Metropolitan Networks*. SPIE, 2003, 5282-01.
- 14 Ferrari, C, Pastorelli, R, Piciaccia, S, Paraschis, L, Gerstel, O. Flexible and Reconfigurable Metro-Regional WDM Scaling beyond 32x10 Gb/s and 1000 km of G.652 fiber. In: *LEOS 2005 – The 18th Annual Meeting of the IEEE Lasers and Electro-Optics Society*, November 2005, TuG5.
- 15 Viscardi, V et al. Asymmetric reconfigurable OADMs for Metro-DWDM networks. In: *Optical Fiber Communication Technical Digest Series*, Conference Edition, Optical Society of America, 2004. Paper TuH 3.
- 16 Paraschis, L. *Advancements in Metro Optical Network Architectures*. SPIE, 2004. 5626-45.
- 17 Paraschis, L, Gerstel, O, Ramaswami, R. Tunable Lasers Applications in Metropolitan Networks. In: *Optical Fiber Communication Technical Digest Series*, Conference Edition, Optical Society of America, 2004. Paper MF 105.
- 18 Paraschis, L. Innovations for Cost-effective 10 Gb/S Optical Transport in Metro Networks. Invited Presentation in *LEOS 2004: The 17th Annual Meeting of the IEEE Lasers and Electro-Optics Society*, Puerto Rico, 7–11 November 2004, p 633–634, vol 2. Paper WU1.
- 19 Theodoras, J et al. Emerging Optical Technologies in Multi-service Metro Networks. Invited Presentation in *LEOS 2003: The 16th Annual Meeting of the IEEE Lasers and Electro-Optics Society*, Tucson, Arizona, USA, 29–30 October 2003. Paper WP1.
- 20 Ramaswami, R. IP + Optical in the Next Generation Network. *Intl. Conf. on IP + Optical Network (iPOP2005)*, Tokyo, Japan, Feb 21, 2005.

Loukas (Lucas) Paraschis is in the Optical Technology CTO group at Cisco Systems, responsible for the evaluation of next generation optical multiservice transport technologies, network architectures, and R&D projects. At Cisco, he has primarily worked on metropolitan WDM networks, IP-optical and broadband architectures and techno-economical analysis of converged transport evolution, and multi-service network reliability evaluation. Prior to Cisco, Loukas completed his graduate studies at Stanford University (PhD 1999, MS 1998), and worked as an R&D engineer. He has (co)authored more than 30 peer-reviewed technical publications, multiple invited, keynote, and tutorial presentations, many technical reports, and three patent disclosures. He is the editor for optical networks in JCN, was guest co-editor for JLT (Nov. 2004), panelist in National Science Foundation (NSF), and has been in technical conference organizing committees, and a member of IEEE and SPIE. Loukas was born and raised in Athens, Greece, where he completed his undergraduate studies.

email: loukas@cisco.com

Errol Roberts is a Distinguished Systems Engineer. Errol has been with Cisco for over eight years with technology focus in IP, Optical Transmission and Data Center Technologies.

email: eroberts@cisco.com

Ori Gerstel (SM'01) is leading the Advanced Planning and Technology Team, which is part of Cisco's Optical Networking Group. In this role he is responsible for looking at how new technologies influence optical networking, including SONET/SDH, SAN transport and DWDM. Until 2002 he was a Senior Systems Architect for Nortel Networks' MEMS based photonic crossconnect product. Before joining Xros/Nortel, Ori was the Systems and Software Architect for the Optical Networking Group at Tellabs, where he architected the first commercial mesh DWDM system. Prior to that, he performed early optical networking research at IBM Research. Ori has authored over 50 papers and patents on optical networks and served on the program committee of OFC, INFOCOM and other conferences and journals. He also teaches the Optical Protection short course at OFC. Ori holds a PhD degree from the Technion, Israel.

email: ogerstel@cisco.com

Wavelength switches and the automated optical network – Status and outlook

CLAUS POPP LARSEN



Claus Popp Larsen is Project Manager at Acreo, Stockholm, Sweden

Wavelength switches have faced an extraordinary rise and a grand crash during the last 5–10 years. They never got a commercial breakthrough, but now the circumstances have changed and wavelength switches hold a great promise for the emerging convergence of datacom and telecom because of their ability to carry all common signal formats. Wavelength switches with electrical switch cores can be readily deployed in existing networks (opposed to switches with optical switch cores), and with the ASON/GMPLS control plane they are ideal as the key elements in the automated optical network. Such a network will not appear overnight, but wavelength switches offer a gradual upgrade path towards the vision that is slowly but steadily coming true.

Switches, crossconnects, routers, exchanges ...

Switches, crossconnects, routers, exchanges – from a functional point of view they all serve one main purpose: To connect incoming traffic on a given input port to a given output port.

Deploying switches (or crossconnects/routers/exchanges) in a network leads to more automated handling of the connections. More automated handling means less manual interaction, which in turn leads to time and personnel savings. On top of that a proper management system will offer a better overview of the network and its resources. And all this is eventually translated into a reduction of the operational expenses, OpEx. History has shown that all telecom/datacom network technologies will eventually become fully automated, from telephone exchanges through IP routers to the emerging wavelength switches.

The switch is the key element of any network technology. While the generic functionality is equivalent, the internal switching differs a lot for the different technologies. In IP-routers packets are switched, and in telephone exchanges and wavelength switches whole circuits are switched. Even within wavelength switches – which are the topic of this article – there are different switching technologies and philosophies. A non-exhaustive classification of different switches is shown in Figure 1.

Observe that the terminology around wavelength switches is not clear. Frequently they are referred to as optical crossconnects or optical switches, but even SDH digital crossconnects are occasionally called optical switches.

Wavelength switching in a historical perspective

Fibre optic switching has been the focus for research groups across the world ever since the advent of fibre optic transmission. Nevertheless, wavelength switches have not yet faced their great breakthrough regarding commercial deployment. Tellium, a North American start-up company which no longer exists, was *the* wavelength switch pioneer with the first commercial offerings in the late 90s. Their equipment was based on electrical switch cores. This was followed by a few years of dramatically rising attention to optical switching, and the optical industry was experiencing an unprecedented optimism.

At one stage CNN reported almost daily about the latest rumours, mergers and acquisitions, at NASDAQ optical switching companies were red-hot, and most of the larger system houses acquired optical switch start-ups for huge amounts of money. This development culminated in the spring 2000 where Nortel paid \$3.25 billion in stock for Xros, a 90 staff start-up. Xros had designed an all-optical wavelength

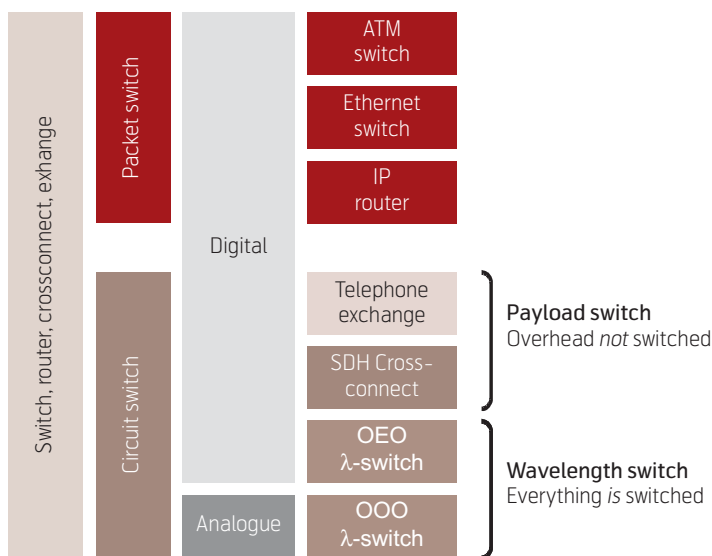


Figure 1 A non-exhaustive classification of different types of switches. The telephone exchange does not really fit in as it is not used in a telecom network, but it is nonetheless included for comparison

switch with allegedly more than 1000 ports. The preferred technology for these monster switches was (with the exception of Corvis) 2D or 3D MEMS (micro electro-mechanical system) switches consisting of tiny (typically a couple of μm^2) movable mirrors deflecting the light in the required direction.

In retrospect we can now clearly see that neither the technology nor the market were ready for such large switches. The combination of inadequate yield, low long-term reliability and high cost issues caused the vast majority of switch fabric vendors as well as system vendors based on these MEMS switches to disappear from the market. Also, there was no need for such port numbers, and the management systems and control planes were not mature enough to support automated operation of the network – and thus the OpEx savings otherwise making the business case for switches were in fact absent.

The situation is different today. The amount of traffic is getting sufficiently high, the management systems and control planes are becoming mature, and reliable and low-cost wavelength switches are commercially available. Operators and network owners today are extremely cost sensitive and they still remember the situation only a few years ago – wavelength switches live with a slightly shaded reputation. This is a shame because a new generation of switch companies has appeared that focus more on the business case than on the technology itself. The commercial applications are considerably less ambitious than what was envisioned during the heydays. Rather than delivering huge core network nodes, wavelength switches today are used for applications such as equipment testing (typically routers or optical instruments), protection of optical connections, high-capacity Internet exchanges and smaller networks handling different types of traffic. More sophisticated applications are found in research networks and test-beds at academic institutions and major operators worldwide.

Switch fabric technologies

There are two main types of wavelength switches. One type contains an *electrical* switch fabric and is called an OEO switch (meaning optical input, electrical switching, optical output), while the other type is based on an *optical* switch fabric and is correspondingly called an OOO switch (the middle “O” denoting optical switching). OOO wavelength switches are also referred to as “photonic switches” or “all-optical crossconnects” or something similar.

The great attraction of OOO wavelength switches is that they do not need to terminate the optical signal; everything passes right through. By saving optical/

electrical conversion (lasers and receivers) and electronics within the switch, the node price can potentially be considerably lower than other types of switches.

Assume that the optical switch fabric itself will once become comparable with the electrical switch fabric concerning yield, reliability and cost – this is not the case today but could well be so in a not too distant future. In that case the OOO wavelength switch will have some significant advantages over the OEO switch with regard to port count and bit rate upgrade. The cost of an OEO wavelength switch scales almost linearly with *port count*. This is because the optics and electronics associated with the ports scale with port size, whereas this is not the case for the switch fabric, the backplane, assembly etc. For an OOO wavelength switch there are no active components at the ports (perhaps apart from monitoring equipment) and the cost will scale with a factor considerably less than the port number. Moreover, when the *bit rate* increases by a factor of four, the OEO switch cost is multiplied by about a factor of four (also here because of the port costs) while the OOO switch cost remains virtually unchanged.

However, there is a price for the transparency – or rather for not terminating the optical signal: With no access to the electrical copy of the signal, monitoring becomes much more complicated, and regeneration becomes a huge challenge. Indeed, there is currently no commercial all-optical 3R regenerator although a lot of research is taking place in that area. An OEO wavelength switch, on the other hand, offers the possibility for 3R regeneration – amplification, signal reshaping and retiming – just like all other types of switches. The fact that monitoring becomes complicated and that the signal cannot be easily regenerated puts instead more severe requirements on the transmission system and subsequently the network design for OOO switches.

An OEO switch can be relatively easily integrated in an existing network because it can interface with the already existing transmission system, e.g. a DWDM link. This cannot be done with an OOO switch, as the whole idea is to *avoid* the termination of the optical signal. Instead, the transmission system has to reach much further so it can bypass at least one switch or preferably more in a larger network. And if this network deploys any kind of dynamic switching (which is a main driving force behind wavelength switches) there are all kinds of issues with power budget, power balance between channels, non-linear effects, etc., and least but not last: A long-haul transmission link could easily offset the savings obtained by the switches.

All the above mentioned issues can be solved technically, but not at a price that today justifies commercial deployment in large, dynamic networks. This does not mean that OOO switches lack practical applications; it rather means that they do not currently compete in the same segments as OEO switches. It should be mentioned that hybrids between OEO and OOO switches exist, for instance with electrical termination of the optical signal at one of the sides only. Switch designs combining optical and electrical switch fabrics have been devised, but neither of these are used commercially.

OOO wavelength switches *are* indeed used commercially in applications such as simple 1+1 protection switching, high capacity Internet exchanges and the emerging light exchanges, and for automated testing of optical instruments and IP-routers. But again, for large, dynamic networks, OOO switches are today inadequate. OEO switches as key elements in such networks are discussed later in this article.

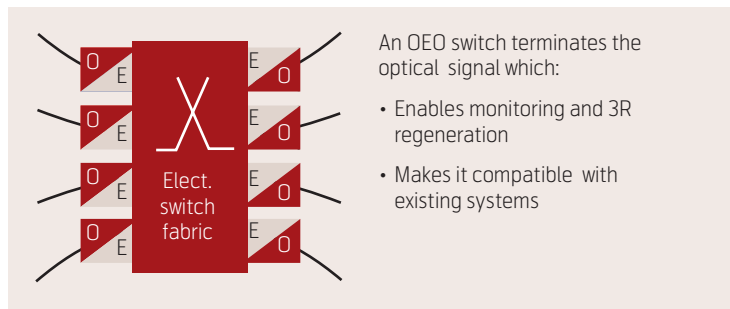
Wavelength switches and competing technologies

Wavelength switches can in principle be deployed anywhere whole wavelengths are handled. The competing technologies in commercial networks today are the Layer 3, 2 and 1 switches. These are typically IP-routers, Ethernet switches and SDH crossconnects, respectively. Also manual patching can be seen as a competing “technology” for some applications but will not be further pursued here.

The greatest advantages with wavelength switches (both the OEO and OOO versions) are most likely the following:

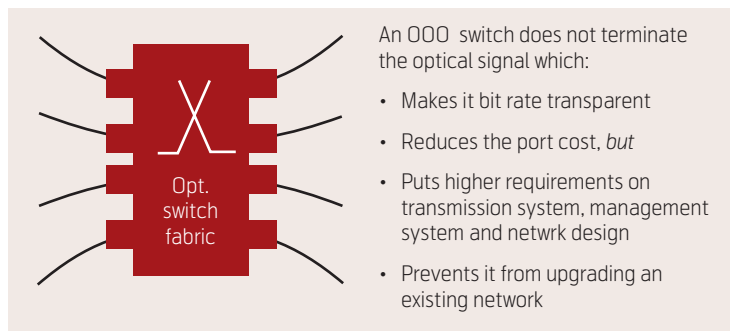
- They provide the simplest and most economical way to bypass a node with a whole wavelength;
- They are protocol agnostic – can be used to converge all types of traffic.

Simple & economical. In Layer 3, 2, and 1 switches the wavelength is divided into smaller units that are individually processed and switched. In Layer 3 and 2 switches the units are packets, and in Layer 1 switches (SDH) the units are virtual containers, VCs. Switching a multitude of units rather than the whole wavelength adds to the complexity of the switch and the amount of electronics involved. This in turn makes the wavelength switch potentially more cost effective than the higher layer counterparts – at least when whole wavelengths are bypassed or otherwise handled with no required access to the individual units.



An OEO switch terminates the optical signal which:

- Enables monitoring and 3R regeneration
- Makes it compatible with existing systems



An OOO switch does not terminate the optical signal which:

- Makes it bit rate transparent
- Reduces the port cost, *but*
- Puts higher requirements on transmission system, management system and network design
- Prevents it from upgrading an existing network

Figure 2 Comparison between OEO and OOO wavelength switches

Protocol agnostic. Operators across the world are in the process of converging their network layers into preferably only one backbone layer that should be able to carry all types of telecom and datacom services. Opposed to all other types of switches a wavelength switch is protocol agnostic in the sense that it can readily carry all common digital formats. While Ethernet cannot carry SDH traffic in a meaningful way, SDH can carry Ethernet traffic. Using generic framing procedure, GFP, an SDH signal can also aggregate lower rate Ethernet signals, which is very useful at the edge of the network, and the equipment is called next generation SDH or multi-service

Network layer	Type of switch
3	IP router
2	Ethernet switch
1	SDH digital crossconnect
“0”	Wavelength switch
	Manual patching

Figure 3 Each network layer has its own switch

The automated optical network should possess

- The flexibility of datacom
- The signal quality and reliability of telecom
- The granularity of a wavelength

switching platform. However, no matter what this extended SDH is called, it is still not nearly as transparent as a wavelength, and above all, it is inherently inefficient and costly when it comes to switching whole wavelengths rather than fractions of wavelengths.

In practice, wavelength switches would be used to *complement* existing equipment rather than replacing it. In the long term by letting the wavelength switch handle the core traffic including protection & restoration and wavelength routing – and leaving the services (in the shape of Layer 2 and 3 traffic) to Ethernet switches and IP routers at the edges of the network. In the short to medium term such an infrastructure can be built up step-wise, gradually deploying more and more wavelength switches in the core network as the network owners feel economically, technologically and strategically confident. This is the topic of the next section.

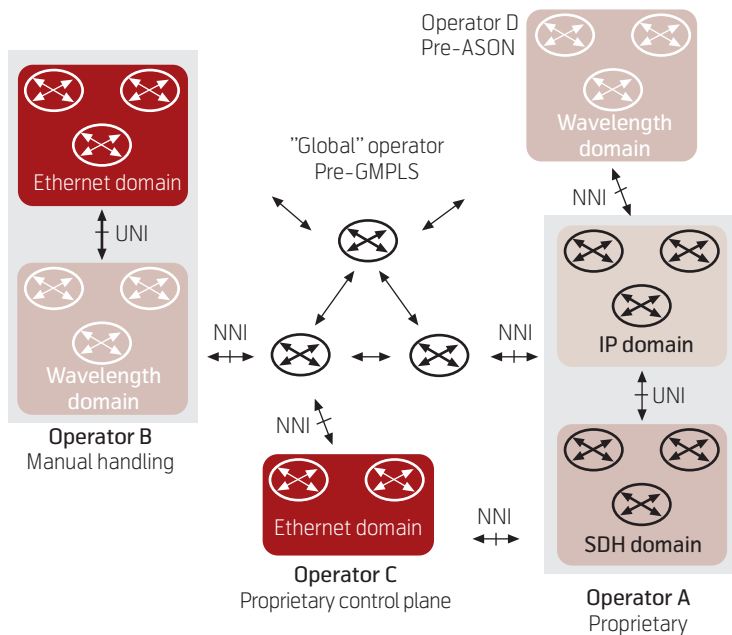


Figure 4 The figure shows a number of autonomous subnetworks forming the global, automated optical backbone network. Each operator uses a different control plane. Containing different technologies, the subnetworks are interconnected by standardised interfaces (NNI) and wavelength switches (the “global” operator). Only the wavelength switch can carry all traffic types in a future-proof and economical way

The vision: An automated optical network

It is commonly accepted that an automated infrastructure based on optical transmission is necessary to support the abundance of bandwidth consuming services. However, there is little consensus on how to achieve this automation – or rather at which network layer to do the switching of the optical signals. The choice stands between Layer 2/3 switching (typically using Ethernet switches and/or IP routers), SDH switching, and wavelength switching. In large backbone networks SDH switching is by far the most popular today whereas the closer to the end-user one gets, the more popular Layer 2/3 switching becomes. New access network installations are completely dominated by Ethernet switches and IP routers.

Optical control plane

An automated infrastructure requires an optical control plane. There are all preferences from no control plane (manual operation) through proprietary control planes to standardised solutions. Despite increasing collaboration between the standardisation bodies no clear consensus has yet materialised, but there are strong preferences for the IETF’s GMPLS (generalised MPLS) in combination with ITU-T’s ASON (automatically switched optical network).

The different views concerning switching and control plane do not in themselves threaten the vision of the automated optical network, but the “religious” discussions may delay the process and make the transition unnecessarily costly. We believe that the quickest, most sustainable and least expensive path to an automated optical backbone is to combine the different network technologies in a framework of switches that use the principle of the lowest common denominator for these different technologies with respect to the following:

- The *flexibility* of handling the network
- The *quality* of the signals
- The *granularity* of the switches

Network and equipment requirements

An automated optical network has to be converged or unified in the sense that it should carry datacom as well as telecom traffic. Further it should support the quickly growing segments of dark/lit fibre and storage.

Datacom possesses historically very high requirements to flexibility, and it is widely accepted that a unified network should also strive towards this degree of flexibility. So concerning flexibility the lowest common denominator is determined by datacom.

Telecom has traditionally had very high requirements to signal quality and should therefore contribute with “carrier class” as its requirement to the automated optical network. This point of view is not shared by all Layer 2/3 equipment vendors (partly because such equipment seldom holds “carrier class quality”) and not even by some operators coming from the datacom side. However, the continuous growth in IP-based services and society’s increasing dependence on secure connections will doubtlessly lead to very high quality requirements to the backbone networks.

What is left to discuss is the switching granularity. The lowest common denominator is the wavelength that is the only format that can carry all other digital formats (SDH, GbE, Fibre Channel, raw wavelength etc). This – combined with the fact that the size of services and aggregated transmission bandwidths continues to grow – also makes the wavelength switch the only really future-proof solution. Bear in mind that a state-of-the-art SDH switch divides a 2.5 Gbit/s signal to 16 units, which are switched and processed individually. A wavelength switch, in comparison, switches and processes only one unit. Therefore – despite its popularity today – SDH is neither a sustainable solution from a cost perspective, nor from a switching granularity perspective regarding the emerging automated optical network.

Subnetwork interoperability with standardised interfaces

We envision the automated optical backbone network combined by a large number of subnetworks. These subnetworks could be between operators, between different vendor equipment or domains within the same operator, and between different network layers. Each sub-network can use its own solution depending on technical requirements, historical preferences, political and strategic choices, etc. A subnetwork can contain Ethernet switches, it can be an SDH ring, it can be a wavelength switched network, etc.

The important requirement is *interoperability* between the subnetworks, and interoperability is best assured through standardised interfaces such as the user-network interface, UNI, and the network-network interface, NNI, at *wavelength* level. ITU-T, OIF and IETF all work with these.

Proprietary vs standardised solutions

The choice of the optical control plane brings us to the dilemma of standardisation: Interoperability and the automated optical network require standardisation, but sometimes only by relieving the compliance to standards is a rapid development possible.

Whereas standardised interfaces are *crucial* in the automated backbone, a common, distributed optical

Important global requirements:

- Autonomous subnetworks with standardised interfaces
- SW upgradeable λ -switches

control plane like for instance GMPLS is just *useful* – at least in the short to medium term. Exactly how the automation *within* a subnetwork is done is less important.

ASON/GMPLS as well as probably the majority of proprietary control planes are based on what has already been developed for Layer 2 and 3 switching and routing – in particular MPLS. Existing protocols need to be extended and new protocols are needed, but the basic idea is the same, just applied on another layer. In a global, automated, interoperable network there should ideally be only one standardised, distributed control plane. Until it becomes clearer which flavour of ASON/GMPLS or a third control plane becomes the winner, the important issue is that whatever switch is used it will need to be software-upgradable in control plane functionality with minimum traffic interruption.

Proprietary control planes will continue to exist even after a single (or maybe more) control planes have been standardised, but this is not a problem. Optical networks are far from automated and interoperable today, and operators will have widely different requirements for the speed and the degree of automation – and there should be room for these differences. The principle with multiple autonomous subnetworks interconnected by standardised interfaces is by far sufficient for many years to come.

An automated optical layer will have only little impact on existing Layer 2/3 network equipment. In its mildest form, automation of the wavelength will just replace manual patching, and in its fullest form wavelength switches can be used to offload traffic from the Layer 2/3 switches and routers and take over functionality as, for instance, switching and protection of wavelengths. Ethernet switches and IP-routers will of course still be needed at the edges of the networks, but the IP and Ethernet based services will be completely indifferent to how the core wavelengths are switched. Wavelength switches will form a server network, eventually serving all other signal formats – including the remaining SDH traffic.

Summary

Wavelength switches never got a commercial breakthrough during the IT-bubble a few years ago, but today the traffic has increased, and the switches and the management systems have become more mature. Meanwhile, competing technologies have got a strong foothold – IP-routers, Ethernet switches and SDH digital crossconnects. However, they all have shortcomings when it comes to supporting the future converged network – technically, practically or economically. Instead the wavelength switch stands out as the

key element gluing together the emerging automated optical network. It offers a pragmatic path towards the automated optical backbone infrastructure and can relatively easily interface with existing network technologies. The idea is that using standardised interfaces the wavelength switches should “globally” interconnect autonomous subnetworks based on, e.g., IP, Ethernet or SDH equipment. The network owner can then in turn gradually upgrade his network towards more sophisticated functionality in his own pace.

Claus Popp Larsen is Project Manager at the Swedish contract research institute Acreo. He works with activities related to the next generation broadband networks. Claus received the MSc degree from studies at the Technical University of Denmark (DTU) and the University of Strathclyde. He holds a PhD in optical switching from DTU in a collaboration with Ericsson.

Claus has worked at different Ericsson sites in Denmark and Sweden; first doing research in components and networks for WDM applications, and later studying reconfigurable optical networks. In 2001 he joined Wavium, a Swedish start-up making wavelength switches, where he had a number of responsibilities ranging from transmission design to sales & marketing. Claus has been with Acreo since 2005.

Claus has furthermore been involved in various European research projects including being a project coordinator, and he has actively participated in standardization work – primarily at the ITU in the area of optical networking.

email: claus.popp.larsen@acreo.se

Deutsche Telekom GSN+ – A comprehensive ASON/GMPLS demonstrator

HANS-MARTIN FOISEL, ANDREAS GLADISCH, MONIKA JÄGER,
SABINE SZUPPA AND ARMIN EHRHARDT



Hans-Martin Foisel is working with T-Systems, Berlin



Andreas Gladisch is head of a department in T-Nova



Monika Jäger is senior scientist at T-Systems, Berlin



Sabine Szuppa is an R&D Engineer at T-Systems, Berlin



Armin Ehrhardt is a senior expert working for T-Systems, Berlin

This paper reports on the continuous efforts of Deutsche Telekom Innovation to set up ASON/ GMPLS network demonstrators to foster control plane enabled network deployment. In 2002, Deutsche Telekom Innovation started building comprehensive ASON/GMPLS network demonstrators and has been pursuing practical evaluations ever since. In a joint effort with system vendors, the so-called Global Seamless Network (GSN) Demonstrator is being set up and continuously extended. The GSN Demonstrator includes promising building blocks of future network architectures and comprises innovative applications in a control plane enabled network. The tests are focused on experiments with prototype implementations related to new standards and specifications from ITU-T, IETF and OIF.

1 Introduction

Since the telecommunication market has become highly competitive, carriers are under continuous pressure to increase their network efficiency and to create new branded services. On the one hand, traffic volumes are still growing, due to the increasing number of broadband Internet accesses, such as DSL users. A variety of new high bandwidth services continuously generates the need for more capacity in the metro and backbone areas of the transport network. On the other hand, the profit that can be realised for certain transmitted traffic volumes is decreasing dramatically. Thus, the limited profit margins force incumbent network operators to implement networking functions and solutions which increase the network efficiency, and which allow exploiting new market shares by offering new services. As data and optical network convergence issues were arising, carriers began to evaluate and introduce intelligent Control Plane (CP) mechanisms that promised operational benefits in a multi-vendor environment. A number of standardisation bodies and forums [1–3], among them ITU-T (Automatic Switched Optical Networks/ASON), IETF (Generalized Multi-Protocol Label Switching/ GMPLS) and the Optical Internet-working Forum/OIF (CP interfaces: User-Network-Interface/UNI and External Network-Network-Interface/E-NNI), are aiming at making this interoperability happen in a standard compliant way. Moreover, new framing and mapping procedures have been developed to enable flexible data mappings into the SDH/SONET infrastructure. Beside specification, the OIF performs interoperability tests of new network functions which are an important step towards implementing new functions.

In 2002, Deutsche Telekom Innovation started the setup of comprehensive ASON/GMPLS network demonstrators and has been pursuing practical evaluations ever since. They are focused on experiments with prototype implementations related to new stan-

dards and specifications from ITU-T, IETF and OIF. In a joint effort with system vendors, the so-called Global Seamless Network (GSN) phase one Demonstrator was built, comprising applications in a CP enabled network. Additionally, a Management System with an integrated view on the ASON/GMPLS and Ethernet-MAN configuration, and end-to-end monitoring capability has been developed. The focus of the GSN was the UNI and the ASON/GMPLS functionalities within one operator domain. In 2004, Deutsche Telekom Innovation evolved towards the current GSN+ phase two Demonstrator configuration, which mainly focuses on NNI functionalities between domains and new services. Besides several ASON/ GMPLS backbone network domains, this field test-bed comprises the mainly envisaged client networks IP, Ethernet (metro and access networks), Storage Area Networks (SAN), and broadband video applications for visualising the manifold network functions. The GSN+ demonstrator activities have been additionally embedded in the OIF World Interoperability Tests and Demonstration in 2004.

GSN+ includes promising building blocks of future network architectures, such as the widely applied Ethernet technology, GFP for flexibly mapping different signal formats, optical transmission and switching for low delay and high bandwidth transmission, and ASON/GMPLS functions for IP centric network control. The GSN and GSN+ are important steps towards the implementation of GMPLS/ASON networks.

2 GSN network configuration and components

The phase one GSN Demonstrator included all components of a modern telecommunications network; i.e. core and metropolitan networks, IP and Ethernet clients, and several kinds of applications. It had been realized on deployed fibre-infrastructure between the

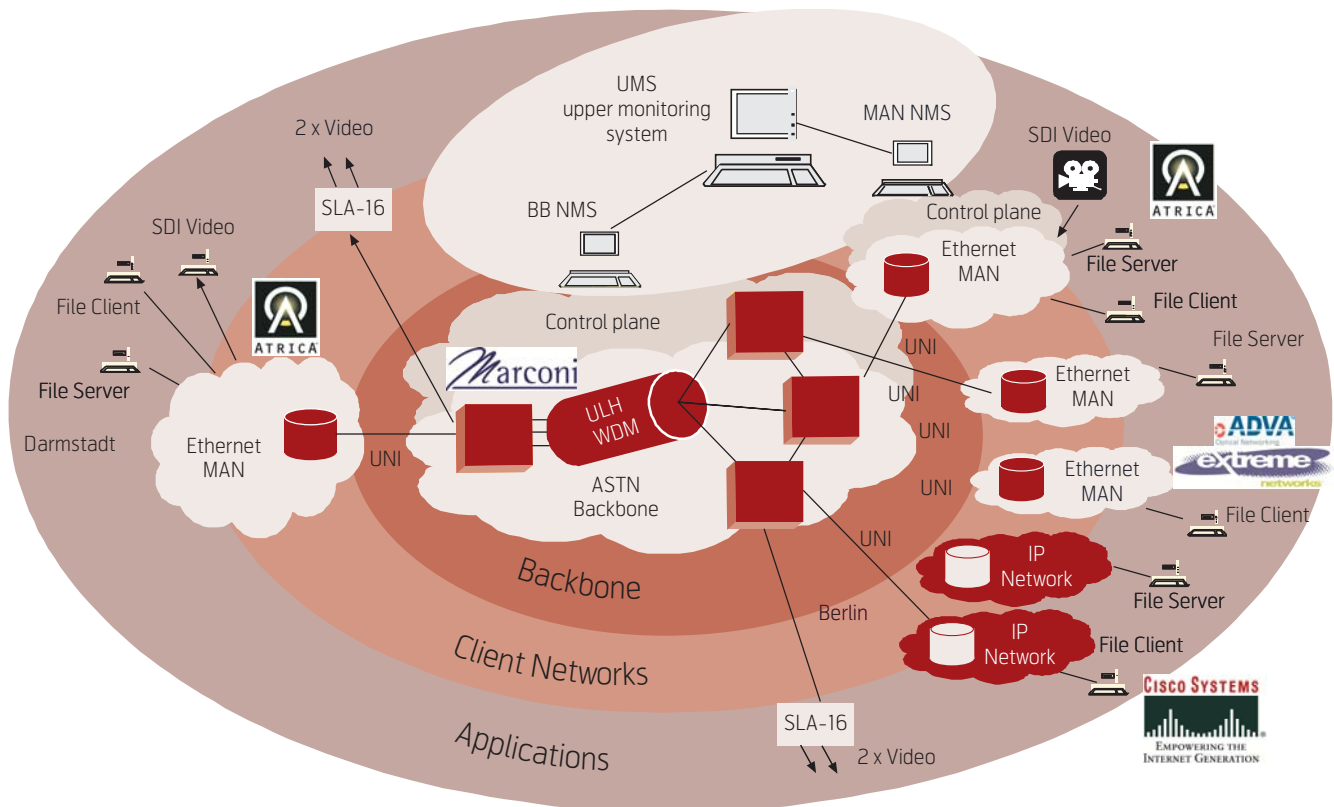


Figure 1 Configuration of the GSN Demonstrator

German cities of Berlin and Darmstadt. The total configuration is depicted in Figure 1.

The backbone consisted of four optical nodes, three of them located in Berlin, one in Darmstadt. On the 750 km distance between Berlin and Darmstadt, an Ultra-Long-Haul (ULH) WDM-transmission system where 3×10 Gbit/s in each direction was installed. Each OCh was treated as a virtual fibre. Ethernet-based metropolitan optical networks (MAN) were set up in the showrooms of Deutsche Telekom in Berlin and Darmstadt, respectively. For visualizing the network functionalities, broadband applications (e.g. video, HDTV) using Ethernet formats demonstrated the end-to-end connectivity over the entire GSN demonstrator. Additionally, E4 applications were connected directly to the ASON/GMPLS backbone network. In this way the impact of network functions such as restoration, could be evaluated and real end-to-end networking functionality could be demonstrated. Finally, a Network Management System (NMS) generated an integrated view on the ASON/GMPLS backbone and the MAN client networks. To set up this comprehensive ASON/GMPLS demonstrator, strong vendor partners had joined the GSN Demonstrator project and had provided equipment to the project, such as the ULH-WDM system, backbone cross connects, backbone IP routers, and Metro Ethernet Switches, supplied with ASON/GMPLS control plane and UNI functionality.

3 ULH transmission system

The efficiency of ULH system deployments depends on several factors, such as the increase of traffic volumes, the network topology and the network architecture; i.e. which technologies and switching granularities are present in the network infrastructure, and which grooming strategies are applied. With the ongoing traffic growth in access, metro and backbone networks, ULH systems are getting more and more efficient in backbone networks, since high bandwidth traffic streams are transmitted in the backbone in coarse wavelength granularity, and cost intensive electrical regenerators can be saved by bypassing the electrical switching nodes. The Deutsche Telekom backbone optical network has a highly meshed topology, due to the high population density in Germany. Therefore, in the Deutsche Telekom transport network data transmissions with transparency lengths of up to 1600 km seem reasonable. In the densely meshed network topology, flexible switching of traffic seems suitable at wavelength granularity in the optical layer.

For testing ULH systems in the GSN Demonstrator, Marconi's ULH WDM-system UPLx160 [4] has been installed on the link between Darmstadt and Berlin. In general, RZ data formats and optimised dispersion management along the link enable a maximum regenerator-free transmission distance of 3000 km at 10 Gbit/s channel data rate. Such kinds of systems are

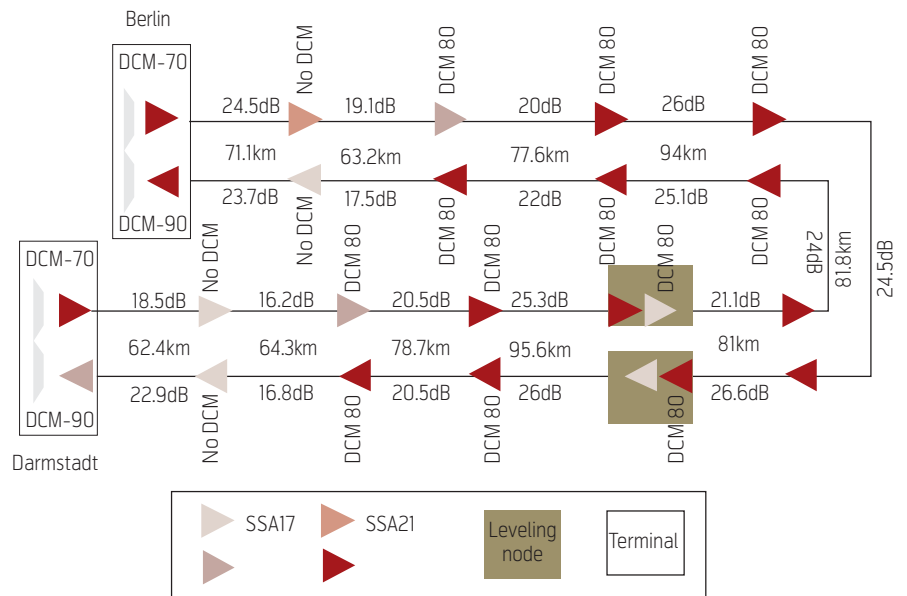


Figure 2 WDM system installed between Berlin and Darmstadt. A transparent loopback switching in Darmstadt is performed for transmission experiments

already in operation in Australia between Adelaide and Perth, covering there 2850 km of un-regenerated transmission distance. Between Berlin and Darmstadt 3x10 Gbit/s for each direction have been realised. The WDM-link setup is depicted in Figure 2.

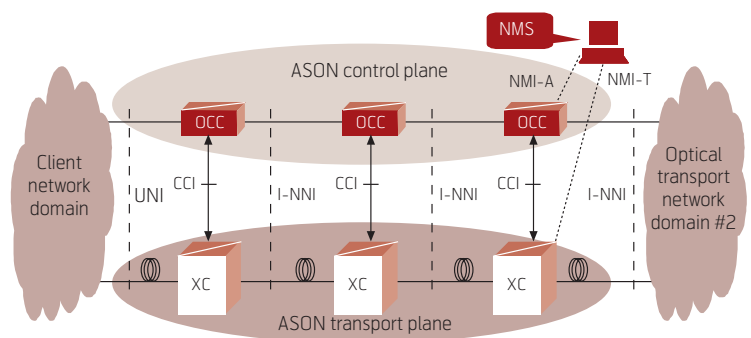
The system is purely Erbium Doped Fibre Amplifier (EDFA) based. Both Dual-Stage Amplifiers (DSA) embedding dispersion compensating fibre modules, as well as Single-Stage Amplifiers (SSA) have been used. In the middle of the link, levelling nodes were implemented to equalise the powers of the propagating optical wavelengths. By transparent loop-back-switching as depicted in Figure 2, transmission experiments over 1500 km have been performed, with a transmission quality of all WDM channels at BER values of 10^{-10} up to 10^{-12} before Forward Error Correction (FEC) could be verified. The total Differential Group Delay (DGD) due to Polarisation Mode Dispersion (PMD) of the fibres of the Berlin-Darmstadt-Berlin loop amounted to 4.3 ps, which was shown not to be critical for a 10 Gbit/s transmission. To evaluate the PMD sensitivity of the system a PMD emulator was installed in Berlin, both before and after the loop.

4 GSN backbone network

As described in the previous section, the backbone of the first phase GSN Demonstrator was composed of the WDM transmission systems and four optical switching nodes. The nodes were SDH VC4-based multi-layer switches. They were supplied with STM-64 network interfaces and STM-16 tributary interfaces. The switching was performed at VC-4 granu-

larity. Every node was equipped with an ASON/GMPLS-like controller, thus realizing a *decentralized control plane* in the backbone network (see Figure 3), featuring signalling and routing protocols for decentralized routing of connections across the network, automatic discovery functions and restoration switching in meshed networks, taking into consideration the actual traffic load of the links.

Between client and core nodes a carrier grade *User Network Interface*, compliant to OIF UNI 1.0 [3] was implemented, including all mandatory and optional functions like in-band signalling and automatic neighbour and service discovery. This interface is based on STM-16 signals. Upon client requests for connections via the UNI interface, the switched connection setup and tear down was performed



OCC: Optical Connection Controller
 CCI: Connection Control Interface
 NMS: Network Management System
 NNI: Network Network Interface
 UNI: User Network Interface
 NMI-A/T: Network Management Interfaces

Figure 3 ASON Control Plane and UNI interfaces of the optical backbone

-VC-12-Xv (n * 2,5 Mbit/s)
 -VC-3-Xv (50 Mbit/s, 100 Mbit/s)
 -VC-4-Xv (n * 150 Mbit/s)




Figure 4 First improvement of SDH: VCAT of VC-12, VC-3 & VC-4

autonomously by the backbone CP. Also soft permanent connections (triggered by the NMS) were realized by the CP. The actual information about the network topology was delivered to the NMS, so that the operator always had an updated view of the network.

The VC-4 switching capability of the backbone nodes enabled multiplexing/grooming of the ASON/GMPLS client signals at VC-4 granularity and Virtual Concatenation (VCAT) (Figure 4). As the client interfaces likewise delivered data in VC-4 granularity, the accumulated ingress data rate of a customer could be adjusted in VC-4 steps via the UNI interface. Therefore the increase of the transport network capacity utilization is considerable, compared to today's used contiguous concatenation formats for data transport (Figure 4).

Generic Framing Procedure (GFP) (Figure 5) is a new advanced encapsulation mechanism with a fixed amount of encapsulation overhead, independently from the binary content of the client. GFP is a robust and scalable frame based encapsulation for PDUs, which provides transparent mapping for the most common 8B/10B encoded signals of data interfaces.




Figure 5 Second improvement of SDH: GFP; Frame based (GFP-F) and transparent (GFP-T)

5 OIF interoperability tests

Based on the results and experiences made with the phase one GSN demonstrator, Deutsche Telekom has initiated and actively participated in the OIF World Interoperability Tests and Demonstrations in the first half of 2004 [5]. First, local CP and Data Plane (DP) tests of UNI, E-NNI and Ethernet over SDH/SONET adaptation were performed at Deutsche Telekom labs. Then, global interoperability tests were carried out in a multi-domain network scenario, covering two main areas:

- CP interfaces: UNI (Client network – transport network domain) and E-NNI (between transport network domains), based on OIF specifications;
- Efficient Ethernet over SDH/SONET adaptation, based on ITU-T Recommendations (GFP-F, VCAT, LCAS).

The goals of these tests were to globally connect and interwork on CP level between network domains. Some of network domains were partly interworking on DP level as well, based on VC-4/STS-3c, connecting local testbeds built as one or more separate vendor domains.

The interoperability tests of the CP interface functions were based on OIF specifications (Implementation Agreements (IA)), IA UNI 1.0 Signalling Specification, Release 2, IA E-NNI-01.0 Signalling, and the draft specification of E-NNI-01.0 Routing [3].

The CP interfaces allowed the configuration of connections from the client side, initiated via the UNI as switched connections or by the EMS/NMS as soft permanent connections, over multiple ASON/GMPLS domains by using the CP functionality (Figure 3). As such, multi-domain connections could be set up without involvement of the EMS/NMS of the intermediate domains. These new functions will significantly ease network operation in future networks, especially connection provisioning in multi-domain environments. Table 1 lists the performed interoperability tests, which were carried out first locally at Deutsche Telekom's labs, second per continent, and finally on a global scale.

Ethernet over SDH/SONET adaptation interoperability tests were based on the ITU-T Recommendations G.7041/GFP-F, G.707/VCAT, and G.7042/LCAS, and were carried out locally at the labs of Deutsche Telekom. The results of these tests show clearly that current standard compliant implementations of efficient Ethernet over SDH/SONET adaptation are interoperable. The local test network at Deutsche Telekom labs was composed of three ASON/GMPLS

network domains from different vendors, interconnected by E-NNIs. An IP client network was attached to these “backbone” networks via UNIs 1.0R2, with VC-4 channelized interfaces. The network elements involved Ethernet adaptation over SDH from different vendors.

6 GSN+ demonstrator

Based on the OIF activity results, the enhancements of the GSN Demonstrator towards the phase two GSN+ Demonstrator was started [6]. The enhancements are focused on most efficient data signal transport from client networks, e.g. IP, Ethernet MAN and Storage Area Networks (SAN) over multiple ASON/GMPLS transport network domains and related customer management solutions (Figure 6).

The GSN+ Demonstrator backbone is composed of three ASON/GMPLS network domains from different vendors, interconnected by an E-NNI, which allows automatic connection configuration and provisioning over multiple network domains. The client networks are interconnected to the backbone domains via OIF UNI 1.0 Release 2 interfaces. The IP client network is equipped with UNI-C 1.0R2 CP and Packet over SDH (POS), channelized at VC-4 granularity data plane interfaces. The most efficient data transport over ASON/GMPLS networks is accomplished by using GFP-F, VCAT and LCAS functions, which had been successfully tested at the OIF interoperability event 2004, and which enable fast provisioning of Ethernet and SAN services. Customer Network Management (CNM) solutions will be implemented for those client networks which have no UNI functionality implemented, such as the SAN clients. In this way, SAN services are enhanced by the ASON/GMPLS function and integrated into the dynamically configurable ASON/GMPLS GSN+ backbone. Hence, new services based on temporary, reconfigurable SAN interconnections are shown. The Demonstrator has been extended to Ethernet based access networks. The interworking of these access solutions with the Ethernet MAN will be practically evaluated. Furthermore, future access network applications will be implemented for visualising the MAN and access network functions.

Different from the phase one GSN Demonstrator, the interconnection between the nodes in Berlin und Darmstadt of the phase two GSN+ demonstrator is based on a transparent DWDM-link with normalized network sections [9]. This configuration allows for an absolutely flexible transmission of data rates ranging from 10 Mbit/s up to 10 Gbit/s on each WDM-channel. The signals on the line side will be transmitted without Digital Wrapper (DW) and Forward Error

Test	Test case	UNI-C	UNI-N	E-NNI
1	Basic routing functionality	-	-	X
2	Routing functionality for virtual links	-	-	X
3	Connection initiated by UNI	X	X	-
4	Dual-domain connection initiated by EMS	-	-	X
5	Dual-domain connection initiated by UNI	X	X	X
6	Multi-domain connection initiated by UNI	X	X	X

Table 1 OIF CP interoperability tests

Correction (FEC). Thus, the signals between the nodes in Berlin and Darmstadt can exhibit any data format and do not necessarily have to be fixed SDH-based data rates.

Further main extensions of the GSN+ Demonstrator will be the interconnections to other project test beds, e.g. the German national project “Vertically Integrated Optical Test Bed for Large Applications in DFN” (VIOLA) [7] and the European IST project “Multi-Partner European Test Beds for Research Networking” (MUPBED) [8].

Finally, the introduction of automatic intelligence of the network nodes by the implemented decentralized control plane enables the realization of various protection and restoration mechanisms. Today’s dedicated protection schemes are fast and robust but very resource intensive. Shared protection mechanisms are able to reduce the backup resources significantly. Depending on the network topology, the traffic demand and the applied network planning and opti-

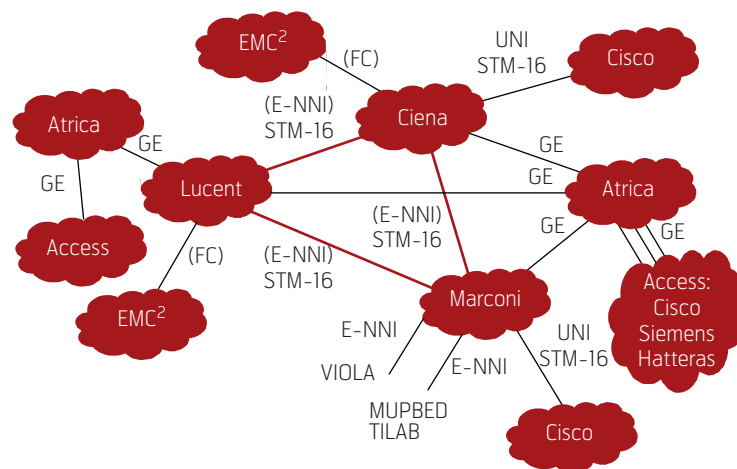


Figure 6 GSN+ demonstrator topology

UNI: User-network-interface

E-NNI: External network-network-interface

FC: Fibre channel

GE: Gigabit ethernet

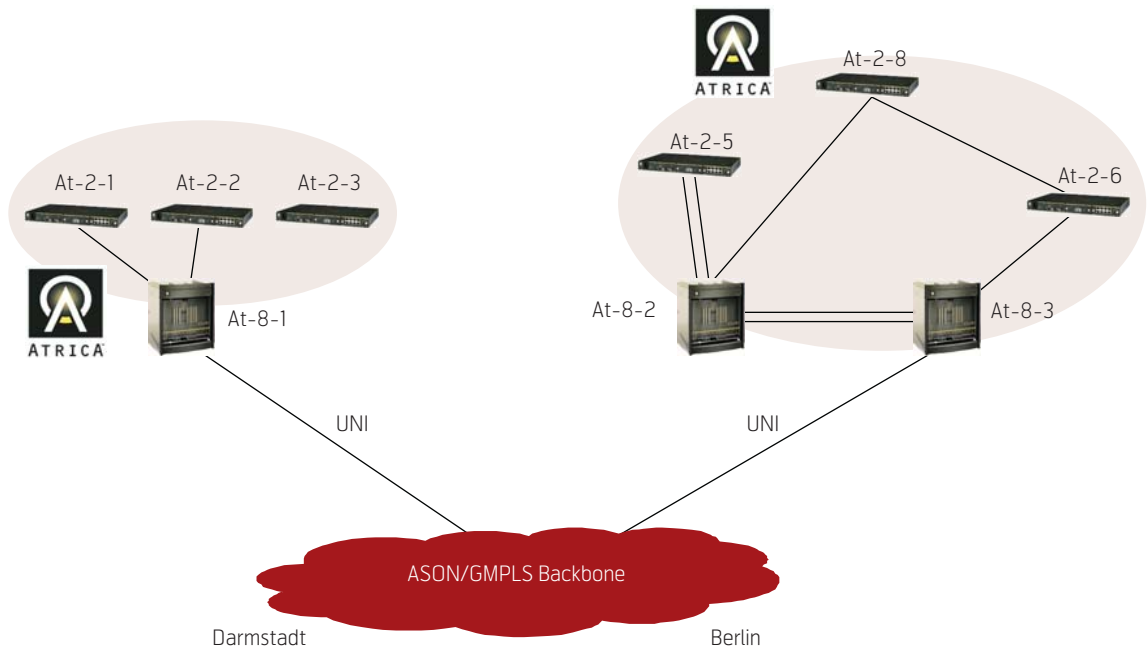


Figure 7 Ethernet Metro Network

Application	Class of service	Protocol/bitrate	ASON/GMPLS solution
Video conference	Real time	E4 140 Mbit/s	VC-4 fixed Restored
DVD Video server	Streaming	TCP/IP 2-10 Mbit/s	
E-Cinema HDTV	Streaming	TCP/IP 2-10 Mbit/s	
File transfer "Print templates"		TCP/IP – 200 Mbit/s	N * VC-4 virtual concatenation Restored
Video on demand	File transfer	TCP/IP – 200 Mbit/s	N * VC-4 virtual concatenation Restored
Video conference (over Ethernet)	Real time	SDI 270 Mbit/s over TCP 300 Mbit/s	N * VC-4 virtual concatenation Restored
Standard video streaming	Streaming	UDP ... 10 Mbit/s	
Storage access	File transfer	Fibre channel N * 200 Mbit/s 1000 Mbit/s	N * VC-4 virtual concatenation Restored

Table 2 Implemented applications

mization strategy, resource reductions of more than 30 % are possible. Restoration mechanisms use dynamic routing and signalling functions available in CP-based networks. Hence, they allow efficiently finding and setting up alternate paths in real-time after failure occurrences, without pre-occupying resources and based on the current network topology and resource status. That is, restoration is very resource efficient and due to its high flexibility provides increased service availability. Up to now, four resilience classes were implemented in the Demonstrator:

- Unprotected
- 1+1 protection
- Pre-planned shared path restoration
- On the fly path restoration

It was shown that in case of system failure or a cable break the backbone autonomously switches to a (here pre-planned) restoration path. The shared restoration which was implemented in the Demonstrator based on pre-planned backup paths, recommends the inverse operation as soon as the working path is accessible again, so the system automatically switched back after an adjustable time.

7 Networks connected to the UNI

Several client networks are attached to the ASON/GMPLS backbone via the UNI (see Figures 1 and 6). They represent the envisaged main future ASON/GMPLS clients, IP and Ethernet. The *IP network* is based on backbone routers, providing IP functionality in the GSN and GSN+ Demonstrator. At the UNI side the routers use POS, VC-4 channelized interfaces, enabling most effective data transport, as the ASON/GMPLS backbone capacity follows the IP payload in VC-4 steps. This is not feasible in case of contiguous concatenation of the SDH payload.

The largest client networks are based on a novel *Ethernet Metro Network* concept (Figures 1 and 7). The Ethernet switching system supports interface data rates from 10 Mbit/s client interfaces up to 10 GE LAN interfaces. At the UNI side the Ethernet networks have POS interfaces supporting the mapping of a GE stream into various kinds of virtual and contiguous concatenation formats. In the Demonstrator, e.g. GE signals are (GFP compliantly) mapped into VC-4-7v. The interconnection of Ethernet MAN and backbone via UNI allows a seamless bandwidth and connection adaptation of the whole network based on the bandwidth requirements at the customer side.

A CNM enables a residential customer to configure his access network (bandwidth and connectivity) in an Ethernet based metro network (MAN) in a pre-defined range given by the Service Level Agreement (SLA). The Ethernet MAN is carrier grade; i.e. connection oriented services with end-to-end SLA can be established. High availability is guaranteed due to sub-50 ms protection switching.

8 Applications

To complete the demonstrator network and to visualize the network functionalities, broadband applications (e.g. video, HDTV) using Ethernet formats demonstrate the end-to-end connectivity via the entire GSN demonstrator. Additionally, E4 applications are connected directly to the ASON/GMPLS backbone network. The whole bundle of applications attached to the GSN Demonstrator are summarized in Table 2.

In this way the impact of network functions, e.g. restoration, has been evaluated and real end-to-end functionality has been demonstrated.

9 Conclusions

The GSN+ Demonstrator represents a further step towards a comprehensive implementation of carrier grade ASON/GMPLS functions in a telecom field environment. It includes not only multiple ASON/

GMPLS backbone domains, but also the mainly envisaged client networks – IP, Ethernet MAN and SAN as well as Ethernet based access networks.

GSN+ comprises many innovative ASON/GMPLS functions. For example, automatic service and neighbour discovery are implemented in the Demonstrator, which simplifies network setup and operation. Furthermore, the Demonstrator tests increased ASON/GMPLS network flexibility on multiple network levels such as SDH and Ethernet, targeting at increased network resource utilization. Reduced provisioning and customer management complexity is demonstrated taking advantage of the enhanced automation in ASON/GMPLS networks. The usability of innovative resilience mechanisms is investigated. Finally, new high broadband flexible services are supported in the Demonstrator. New Ethernet and UNI/control plane based dynamic services are demonstrated.

The GSN+ demonstrator activities have successfully shown the feasibility of introducing CP functions based on ITU-T/OIF specifications into the transport network. However, when introducing the functions into the transport networks, influences on network operator internal processes have to be carefully considered. For example, traditional management systems are based on ITU-T recommendations, which have to be harmonized with management concepts coming from the IETF based GMPLS protocols.

As today many networking functions are only usable within one vendor domain, the implementation of standard compliant inter-domain NNI interface functions will be required. To continuously ensure a high interoperability level of innovative networking solutions, especially in multi-domain environments, the GSN+ Demonstrator is based on OIF World Interoperability Tests results and will be interconnected to nationwide German and European-wide test beds. In this way a further step has been taken towards ASON/GMPLS and Ethernet network evaluations, and extensive experiences are gained in a realistic multi-domain network environment, covering all network hierarchies from the access to the backbone. On this basis, the evaluation and demonstration of new applications and services can be performed.

The GSN Demonstrator purely uses the overlay approach, in which the IP/MPLS and Ethernet client networks are strictly separated from the ASON/GMPLS domain, and interconnected via UNI. The overlay network architecture seems straightforward for short and medium term network evolution of incumbent operators' networks, such as Deutsche Telekom, as profitable network services today are

still widely based on traditional SDH networks. However, the basic underlying network architecture with its two extremes overlay and peer concepts is of profound importance in future transport networks. Its pros and cons are highly dependable for every operator on their main market segment. The more the network and the services converge towards IP centric Next Generation Networks (NGN) and IP based services, the more the peer approach with integrated network layer control seems advantageous. The migration from an overlay network concept with segregated and independently managed network levels, towards a peer network concept with an integrated CP, however, has to be seamless. The evolution from today's transport network towards future ASON/GMPLS based transport networks and the convergence of OIF/ITU-T and IETF approaches, as well as the interworking of the IP/MPLS and ASON/GMPLS network domains need further studies and will be conceptually investigated and practically evaluated in future project phases of the GSN+ Demonstrator.

References

- 1 ITU. *Architecture for the Automatically Switched Optical Network (ASON)*. ITU-T G.8080/Y.1304. Geneva, Nov 2001. (www.itu.int/ITU-T)
- 2 IETF. *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*. RFC 3471. Jan 2003.
- 3 OIF. *User Network Interface (UNI) 1.0 Signaling Specification, Release 2: Common Part*. OIF-UNI-01.0-R2-Common. The Optical Internetworking Forum, Fremont, CA, Feb 2004. *and* OIF. *RSVP Extensions for User Network Interface (UNI) 1.0 Signaling, Release 2*. OIF-UNI-01.0-R2-RSVP. The Optical Internetworking Forum, Fremont, CA, Feb 2004. *and* OIF. *Intra-Carrier E-NNI Signaling Specification*. OIF-E-NNI-Sig-01.0. The Optical Internetworking Forum, Fremont, CA, Feb 2004. (www.oiforum.com)
- 4 Gladisch, A et al. *Global Seamless Network Demonstrator – A Proof of Concept*. *ECOC2003*, Rimini, Italy, Sep 2003.
- 5 Foisel, H-M. *Optical Internetworking Forum: World Interoperability Tests and Demonstrations*. Post deadline paper, *ECOC2004*, Stockholm, Sweden, Sep 2004.
- 6 Foisel, H-M. *Global Seamless Network Demonstrator*. *Broadband Europe 2004*, Brugge, Belgium, Dec 2004.
- 7 *VIOLA – Vertically Integrated Optical Testbed for Large Application in DFN*. April 26, 2005. [online] – URL: www.viola-testbed.de
- 8 *MUPBED – Multi-Partner European Testbeds for Research Networking*. April 26, 2005. [online] – URL: www.ist-mupbed.org
- 9 Hanik, N et al. *Extension of All-Optical Network-Transparent Domains based on Normalized Transmission Sections*. *Journal of Lightwave Technology*, 22 (6), 1439–1452, 2004.

Hans-Martin Foisel received his Dipl.Ing. degree in electrical engineering from the University of Kassel and Technical University Berlin, Germany. He worked as student research staff member and later as member of the scientific staff at the Heinrich-Hertz-Institute in Berlin for 19 years in the area of optical transmission systems. Since 1998 he has been member of the Photonic Networks and Network Architecture Department of T-Systems/Technology Centre (division of Deutsche Telekom) in Berlin. He is currently involved in several projects in the area of IP transport networks and in related standardisation activities, e.g. he is the T-Systems/Deutsche Telekom representative at the Optical Internetworking Forum (OIF), where he is actually chairing the Carrier Working Group.

email: Hans-Martin.Foisel@t-systems.com

Andreas Gladisch received his Dipl.Ing. Degree in Theoretical Electrotechnics from the Technical University of Ilmenau in 1986. From 1986 to 1990 he was engaged in research on coherent optical communication and optical frequency control at Humboldt University Berlin, where he received his PhD in Optical Communication. He joined the Research Institute of Deutsche Telekom in 1991 where he was involved in projects on coherent optics, wavelength division multiplexing systems, wavelength control and frequency stabilisation. From 1996 to 1998 he was head of a research group working in the field of design and management of optical networks, and in 1999 he became head of the department "network architecture and system concepts" of T-Nova, the research and development subsidiary of Deutsche Telekom. The development of migration concepts to introduce additional optical functionality into the networks is the main responsibility of the department and moreover, the department represents Deutsche Telekom in OIF and IEEE and other fora, dealing with transmission networks and IP-WDM integration. Dr. Gladisch is member of ITG and IEEE and has authored or co-authored more than 80 national and international technical conference or journal papers, and holds several patents.

email: Andreas.Gladisch@t-systems.com

Monika Jäger has been working as senior scientist in the department of Photonic Networks and Network Architecture at T-Systems, Technologiezentrum in Berlin since 1998. She previously held positions with the Fraunhofer Institute for Open Communication Systems and DeTeWe. Monika Jäger graduated in Electrical Engineering from Munich University of Technology, Germany in 1992. Her current research interests are in the area of optical transport network design. She is involved in the Deutsche Telekom project GSN+ and in the European IST project NOBEL.

email: Monika.Jaeger@t-systems.com

Sabine Szuppa has been working a professional Research and Development Engineer in the department of Photonic Networks and Network Architecture at T-Systems, Berlin since 1989. Sabine Szuppa graduated in Teaching for Mathematics and Physics from Humboldt University of Berlin, Germany in 1985. Her current research interests are in the area of optical Metro Ethernet Network. She is involved in the Deutsche Telekom project GSN+.

email: Sabine.Szuppa@t-systems.com

Armin Ehrhardt received his Dipl.Phys. degree from the Kharkov State University, Ukraine, in 1987 and his PhD degree in physics from the Humboldt University Berlin in 1991 for research on WDM-systems with heterodyne receiver schemes. From 1991 to 1996 he worked on photonic components, generation of short optical pulses and high speed optical transmission systems at the Heinrich-Hertz-Institute Berlin. In 1996 he joined the Technology Center of Deutsche Telekom. Currently he is with the Photonic Networks and Network Architecture Department of T-Systems where he is involved as a senior expert in investigations on high capacity WDM transmission and development of extended SAN solutions based on optical networks.

email: Armin.Erhardt@t-systems.com

Optical transmission systems in the Telenor network – In retrospect

STEINAR M. SVENDSEN



Steinar M. Svendsen is Chief Engineer at Telenor Nordic – Fixed

For more than two decades we have been witnessing the establishment and expansion of an optical transmission network in Norway. Initially international research and development was mainly focusing on high-capacity and long-haul systems. New “world records” were frequently reported, e.g. with respect to bit rates and reach. The most remarkable milestones were achieved when introducing optical amplifiers and wavelength division multiplexing (WDM). In recent years we have seen a notable slow-down in business, while focus has shifted to metropolitan networks and cost reduction.

Introduction

The intention of this article is to present an overview of the development of optical transmission systems that are being used in the Telenor network. Starting with the first system in 1983, in the following I will attempt to report in chronological order all the major evolutionary steps that have taken place, as well as to review important decisions that needed to be made along the way.

A slow start during the multimode era

The first optical transmission system in Norway was established in 1983, on a 9 km cable route between Aurland and Flåm in Sogn, Western Norway. The cable was partly submarine, containing ITU-T G.651 multimode fibres. Ericsson supplied the optical equipment, a laser-based 34 Mbit/s system operating at 850 nm wavelength – also referred to as the first window. Following a three months test period, the system was carrying regular traffic before the end of the year.

More multimode systems followed in the years from 1983 to 1988. These included:

- Different capacities: 2, 8, 34 and 140 Mbit/s;
- Different manufacturers: Ericsson, NEC, Philips, SEL, Siemens and Telettra (although the number was soon reduced to three);
- Two transmission windows: 850 and 1300 nm;
- Two types of light sources: Laser and LED.

Likewise, the quality of fibre was not unique. G.651 allows for a wide variety of quality classes. After some considerations, we standardized on a specific type with the following main characteristics:

1300 nm (specified): Bandwidth > 1000 MHz · km, attenuation < 1 dB/km

850 nm (not specified): Bandwidth > 200 MHz · km, attenuation < 4 dB/km

850 nm equipment was approximately 30 % cheaper than equivalent 1300 nm equipment – partly depending on system bit rate. But because of much shorter span lengths for 850 nm, we soon decided to restrict ourselves to 1300 nm systems.

Table 1 presents an overview of the total number of multimode systems to be acquired by Telenor 1983–88. Span lengths for the different sub-categories are indicated. These are based on the characteristics of our “standard” multimode fibre, and also interface characteristics of optical line terminals from the most relevant suppliers. As one will notice, a few of the systems – mostly the 850 nm systems – are bandwidth-limited, while the remaining are attenuation-limited.

The total number of multimode systems is not very substantial, and a significant number are short, rural connections and not main routes.

An important switch to single-mode fibres

A highly important decision was made in the middle of 1985: As quickly as practically possible, we would switch from multimode to single-mode fibres on all new installations. The goal was that from the first half of 1987 all new cables should contain single-mode fibres. The arguments for this decision were the following:

- Single-mode fibres were superior with respect to transmission characteristics (signal capacity and reach).
- Single-mode fibres had recently been introduced to the market, and in the USA sales figures increased rapidly, while price levels showed a corresponding decrease. In Europe the market would still experi-

ence a significant price level difference between multi- and single-mode fibres in many countries. However, in Norway the price of single-mode fibres was just negligibly higher, and seemed to continue to decrease below the price of multimode fibres.

- With respect to optical equipment, 140 Mbit/s line terminals for single-mode fibres were already available. A 1310 nm laser terminal would cost about 15 % more than a corresponding 140 Mbit/s multimode terminal, and lower order systems were not yet available. However, it was expected that this situation would change rather quickly.
- The rather limited amount of fibre-optic cables already installed in Norway made it easier to quickly switch to the new fibre type.

Single-mode transport systems are settling in

The first single-mode system in the Telenor transport network was ready for service on 18 June 1986. It was a 140 Mbit/s system, installed on the important route between Oslo and Drammen. Total cable length was 53 km, therefore one intermediate regenerator was needed (Asker). The system made use of lasers operating in the 1310 nm window. The manufacturer was AT&T Philips (Nederland).

From 1986 onwards, things were moving faster. In 1987 the first 565 Mbit/s system was established. At this time the full range of hierarchical bit rates from 2 to 565 Mbit/s were present in the network, all of them based on 1310 nm operation. The high capacity systems, 140 and 565 Mbit/s dominated the market. Table 2 contains an overview of systems ordered between 1985 and 1987. All laser systems were attenuation-limited, although the 565 Mbit/s was dispersion-limited to approximately the same reach. During these years NEC was our main supplier of 2, 8 and 34 Mbit/s systems, Siemens and Philips supplied the 140 Mbit/s systems, and Siemens most of the 565 Mbit/s systems.

Bit rate (Mbit/s)	Total no. of systems	Systems needing regenerator(s)	Light source	Wavelength (nm)	Span length for LED / laser systems (km)
2	3	0	1 laser, 2 LED	1310	30 / 63-80
8	16	0	5 laser, 11 LED	1310	23 / 60
34	30	0	14 laser, 16 LED	1310	15 / 46-58
140	179	6	All laser	1310	46
565	92	6	All laser	1310	33

Table 2 Overview of first generation of single-mode transport systems, 1985–87

Bit rate (Mbit/s)	Total no. of systems	Wavelength (nm)	Light source	No. of systems	Span length (km)
2	4	850	Laser	1	12
		1300	Laser	3	31
8	3	850	Laser	1	10
		1300	Laser	2	27
34	16	850	Laser	3	4
		1300	LED	10	13
		1300	Laser	3	27
140	45	850	Laser	7	2
		1300	LED	10	4
		1300	Laser	28	22

Table 1 Overview of multimode systems established

Although optical systems were installed in most parts of the country at this time, the optical network within the capital of Oslo dominated the domestic market, as more than 60 % of the 140 Mbit/s line terminals and more than 70 % of the 565 line terminals were to be installed here.

Table 2 indicates an interest for using LEDs in low bit rate systems. Conventional surface-emitting LEDs had too wide a bandwidth to be used for high bit rate systems. With the introduction of edge-emitting diodes – ELEDs – this problem could be overcome, thereby providing more economic and reliable 140 and 565 Mbit/s systems. However, another new component was to arise in 1987, the low-power laser (LP-LD). The LP-LD was able to compete with the ELED with respect to price and reliability – and at the same time offering longer span lengths.

The growing interest for more economic short-haul systems proved to be an attractive niche that could initiate local development and production in Norway. Telenor (at that time government-owned Televerket) wanted to support domestic activity in this field, and development contracts were signed with two compa-

nies: the Norwegian branches of Siemens for 2 and 34 Mbit/s LED systems and with Telettra for 140 Mbit/s LED systems. Before any delivery could take place, LED was changed to LP-LD for both contracts. Actual deliveries started in 1988.

As the laser was a cost-dominant component within the line terminal equipment, the price difference between conventional laser and low-power laser would constitute a 15–20 % reduction of equipment price for one complete line terminal. Within Oslo, it was calculated that optical systems of 20 km reach could cover 90 % of the routes to be established.

Opening the third window

The single-mode fibre cable we had adopted in Norway, was according to the ITU-T Recommendation G.652 standard single-mode fibre and cable. This fibre was originally intended to be used in the 1310 nm window, and a complete set of specified characteristics referred only to this window.

As the G.652 fibre has its minimum attenuation at about 1550 nm – referred to as the third window – trying to make use of this advantage was another challenge, both with respect to fibre, cable and equipment.

The cable specifications were now extended to include transmission parameters in the third window. From the middle of 1988 all cables to be installed in Norway contained double-window G.652 fibres. It may be noted that G.652 fibre produced from 1985 to 1988 may be used at 1550 nm at medium temperatures. However, at low temperatures bending loss will become an obstacle.

Equipment-wise, there were also problems to overcome. Not in the optical receivers, as the detectors would accept 1550 as well as 1310 nm. But because of the high value of dispersion at 1550 nm, it was necessary to develop a new type of laser with a very narrow spectral width to avoid the optical systems to be dispersion-limited. These lasers were named DFB (Distributed Feed-Back) lasers. Having a spectral width of about 0.3 nm, while the corresponding figure for conventional Fabry-Perot lasers normally would be ten times higher, they would constitute the real single-mode lasers. However, the price level for 1550 nm lasers was high. From an initial cost of about 100,000 NOK, during 1987 it was reduced to 30,000 – 60,000 NOK for a volume of 100 units.

In October 1987 we received the first quotations for 140 and 565 Mbit/s systems for the 1550 nm window. The first installations could start the following year.

However, although prices for 1550 nm lasers continued to decrease, they would still be more expensive than 1310 nm lasers. For this reason 1310 nm systems continued to dominate the market for the years to follow, 1550 nm was mainly chosen for long-haul systems where regenerators could be saved.

To illustrate the difference in span lengths, the following figures were typical for systems with standard lasers:

- 140 Mbit/s: 1310 nm: 52 km 1550 nm: 64 km
- 565 Mbit/s: 1310 nm: 45 km 1550 nm: 68 km

Birth of WDM

The idea of wavelength division multiplexing (WDM) originated in the early days of optical fibre transmission. WDM-systems were developed in laboratories. One manufacturer reported that three 565 Mbit/s systems in each of the two transmission windows could be realizable, the channels having a spacing of 20 nm. However, this concept did not prove economically advantageous.

A very simple alternative of WDM appeared in 1988 in Oslo. On a few short-haul routes one 565 Mbit/s channel was transmitted in each of the two windows 1310 and 1550 nm. Simple and cheap optical couplers were used to combine and separate the two signals at transmit and receive end respectively. Conventional line terminals could be used. Awaiting the 2.5 Gbit/s systems, this solution proved a good alternative when the number of fibre pairs was scarce.

A new single-mode fibre type

Before 1990 an alternative single-mode fibre type was introduced to the market. It was referred to as a “dispersion-shifted” fibre. ITU-T issued the first version of a recommendation for such in 1988 and named it G.653. The purpose was to shift the zero-dispersion wavelength from the 1310 to the 1550 nm window. As both dispersion and attenuation would have a minimum value around 1550 nm, this fibre type was therefore optimized for use in this region. However, fibre cost was high.

Several network operators saw this fibre type as a must for future high-capacity long-haul systems, and soon started deploying it on main routes. In Telenor we were more reluctant. We engaged our research department to carry out a thorough evaluation of G.653. The conclusion was reached in 1992. For different reasons, mainly because of high fibre price and the advantage of co-ordination of backbone and local systems, the dispersion-shifted fibre should not be introduced in

the network for the time being. It was realized that the standard fibre, G.652, had great potential. Time also showed that technical difficulties that seemed severe at the time, like high dispersion, could be overcome.

The decision not to switch to G.653 in the domestic network was a lucky one, as this fibre type would have caused severe problems (four-wave-mixing – a non-linear effect) when multi-channel WDM-systems were introduced some years later. It should be mentioned, though, that for two submarine cable systems between Norway and Denmark, G.653 was chosen.

From PDH to SDH

Starting from the late 1980s, a considerable amount of 140 and 565 Mbit/s systems were installed in Norway, most of them supplied by Siemens. From 1989 our main supplier of 565 Mbit/s equipment came to be the Danish company NKT (in 1994 to become DSC Communications, then Alcatel Denmark and today Tellabs Denmark). In particular, before the 1994 Winter Olympics in Lillehammer, there was some intensive activity establishing an extensive high-bit-rate network in Southern Norway. At this time, from 1992 onwards, the first pre-SDH 622 Mbit/s and 2.5 Gbit/s systems were available. During 1992 and 1993 NKT and PKI (German Philips) delivered a large number of these systems, as shown in detail in Table 3. The total value of this equipment was more than 73 million NOK.

The following nominal span lengths were recommended:

- 622 Mbit/s
1310 nm: 54–57 km 1550 nm: 90–96 km
- 2.5 Gbit/s
1310 nm: 46–48 km 1550 nm: 71–74 km

These figures assume an attenuation coefficient of 0.45 dB/km at 1310 nm and 0.31 dB/km at 1550 nm for planning purposes, including cable margins and installation splices, and are also inclusive of loss in the cable termination in the stations. A few years later the figure for 1550 nm was changed to 0.275 dB/km.

These values refer to buried cables and submarine cables. For aerial borne cables attenuation values were 0.05 dB/km higher, because of exposure to low temperatures.

In December 1994 a contract was signed with NKT/DSC, making the company our main supplier of SDH equipment, including optical transmission equipment. Following this contract, and continued international standardization activities, the first “real” SDH 622 Mbit/s and 2.5 Gbit/s systems followed thereafter.

WDM is back

In the mid 1990s a new fibre-optic component quickly rose to fame: the optical amplifier. It paved the way for the breakthrough of WDM. Where before it was necessary to equip one regenerator per optical channel in each repeater point, now one optical amplifier could amplify all channels within the same unit. In addition, the saving of optical fibres was also an important factor for total cost reduction. The number of fibre pairs in the cables was often quite scarce during the first years.

The first multi-channel WDM system in Norway was installed in 1997, equipment delivered by DSC. The system was Oslo–Bergen, with a total route length of 520 km, comprising six intermediate line repeaters; i.e. seven repeater sections. The system comprised four 2.5 Gbit/s channels, spaced 400 GHz apart (1549.32 nm, 1552.52 nm, 1555.75 nm, 1558.98 nm). No regeneration of signals was needed between terminals, also no extra optical equalization and no dispersion compensation were needed.

8-channel WDM systems followed in 1998, 16-channel systems in 1999 and 32-channel systems in 2000. The 32-channel system used 100 GHz channel spacing, practically filling the ITU-T defined C-band (1530–1565 nm). Later also 64-channel systems became available, making use also of the L-band (1565–1625 nm) with the same channel spacing; however, so far we have seen no need for that many channels.

System capacity	NKT		PKI	
	Line terminals	Line regenerators	Line terminals	Line regenerators
622 Mbit/s	33	16	14	4
2.5 Gbit/s	46	36	50	15

Table 3 Pre-SDH 622 Mbit/s and 2.5 Gbit/s systems delivered in 1992/93

The success of WDM managed to slow down the urgent need for optical channels of higher capacity. For a time this delay was fortunate, as technical challenges had to be overcome, like dispersion equalization accuracy, non-linear effects, PMD. However, as technical development progressed, 10 Gbit/s channels became available – although not needed in our network until recently. Today also 40 Gbit/s channels are available from a limited number of suppliers.

Another challenge to the standard single-mode fibre

In 1996 ITU-T standardized a new fibre type, G.655 “Characteristics of a non-zero dispersion-shifted single-mode optical fibre and cable”. For WDM-transmission the G.653 fibre had proven unsuitable, because too low dispersion would cause non-linear phenomena. Low dispersion was an important advantage for high-capacity channels, but dispersion should not fall below a certain minimum value. This was the idea of G.655. The fibre manufacturers were promoting the new fibre type quite extensively, and one could easily get the impression that G.655 fibres were a must when high-capacity systems of some length were to be established. The problem of G.655, however, was that the fibre price seemed to be two or three times the price of standard G.652 fibres.

At Telenor we carried out a study of the two fibre types G.652 and G.655 in 2000. At this time only technical data and price information on 10 Gbit/s systems were available to us. The conclusion was reached in October that year: At that point we saw no advantage, neither of technical nor economical nature, to switch to the new fibre type.

A more extensive evaluation took place in 2001/2002, now including also 40 Gbit/s transmission. Three manufacturers were invited to submit technical data and price information on 40 Gbit/s equipment, and our research department was engaged to perform an analysis of dispersion and non-linear effects based on simulations. The conclusion of our evaluation was that, based on technological investigations and economical analysis, the standard single-mode fibre G.652 was the preferred fibre type to be used in the Telenor transport network for the following years.

The new fibre type was favoured by many newcomers in the business, particularly in the USA. On the other hand, it seems that most of the traditional network operators reached the same conclusion as Telenor did.

10 Gbit/s WDM goes far north

In 2003 Telenor participated in the world’s northernmost submarine cable project, two cables (of 1436 km and 1413 km respectively) between Harstad on the Norwegian mainland and Longyearbyen on Spitsbergen (Svalbard) at about 78° north. Cables and equipment were delivered by Tyco Telecommunications. Channel capacity is 10 Gbit/s. The owner of the cable system is the Norwegian Space Centre; however, Telenor was responsible for the technical implementation of the project, and has also undertaken responsibility for operation and maintenance.

CWDM attracts interest

All WDM systems that have so far been installed in the Telenor network are of the conventional DWDM type (Dense WDM – with narrow channel spacing). As the WDM technique proved to be advantageous also in the metropolitan network, thereby saving fibres, cheaper technical solutions were sought. In December 2003 ITU-T standardized a CWDM wavelength grid (Coarse WDM), Recommendation G.694.2, followed by a recommendation on CWDM optical interfaces in 2004 (G.695). Channel spacing is 20 nm, about 25 times the spacing of a conventional 32-channel DWDM system. G.694.2 allows for 16 (18) CWDM-channels in a wide wavelength range from 1271 to 1611 nm. Optical amplifiers are in many cases not foreseen in the metro network, and because of the wide channel spacing the laser stability is not critical and coolers thereby not needed.

In 2004 Telenor started a project on CWDM, resulting in an RFQ and subsequent contract negotiations. A final decision has not yet been reached.

Towards a future all-optical transport network

During recent years ITU-T has developed a family of optical transport network (OTN) standards. These include recommendations for the physical/optical layers, Layer 1 standards and layers implemented in software. An optical transport hierarchy (OTH) has been defined.

At the start of the century hopes were high for a steady development into an all-optical network, which would offer transparent transport of client signals. Optical cross-connects were introduced to the market, in particular the MEMS technology (micro-electro-mechanical-system) seemed promising. 32x32 MEMS-based all-optical switch modules were commercially available in 2001. Lucent in particular was promoting the new technology. However, because of market conditions, Lucent had to bring an end to the

development of their MEMS cross-connect, the LambdaRouter, in July 2002. Nortel did a similar thing, and Cisco Systems before that. Attention was instead steered towards OEO switches. A less pragmatic view will allow signal processing in the electrical domain, inside the optical interfaces of a functional module. The all-optical network may be some future goal instead.

Telenor has not yet seen any need for DWDM-channels in the L-band, although they have been available for a few years. Likewise, 40 Gbit/s channels have not had any success so far. However, reports indicate that the situation is slowly improving. It is also interesting to note that optical interfaces today are conquering new grounds outside the traditional transport networks domain. To demonstrate this expansion, we will now see optical interfaces implemented in access networks, in IP- and DSL equipment.

Steinar M. Svendsen is Chief Engineer at Telenor Nordic – Fixed, Fornebu, Norway. Since 1983 he has been working with fibre-optic transmission systems for the transport network. During recent years he has been engaged in different projects mainly focusing on WDM transmission in transport and metropolitan networks.

email: steinar-moe.svendsen@telenor.com

Building a fibre-optic network in Norway

HARALD JANSEN



Harald Jansen works at Telenor Telecom Solutions

Televerket/Telenor has been using optical fibres in the transmission network for nearly 25 years. From a careful start, with strong elements of experimentation in both fibres and transmission, the network has grown to a rather finely meshed network (with some exceptions). The focus of this article is on planning and building of the core network. The article will describe the start of the use of optical fibres, how the work was organized, and briefly describe the trials and the technological choices that were made. Further, a chronological description is given of how the network expanded, with some comments on competition.

1 Start- and trial phase

The organized start of fibre optic activities in Televerket started in 1980, when Televerket “saw the light”. Activities had till then been limited to research and some co-operation with domestic cable producers on fibres (also domestic production), cable types and installation methods.

1.1 Organization

The organized activities started with the establishment of the FOTgroup in the Line section of Teledirektoratet (Norwegian Telecommunications Administration – NTA). This group was placed under the Line section’s Network planning office, and had participants from cable and transmission offices, as well as from the research department.

The group’s mandate was to guide the introduction of fibre-optics in the network through trials and planning. The trials were conducted with the aim to test interesting technologies in practical projects, covering fibre, cables and transmission.

1.2 Trials

In the first activities there was a close connection between transmission testing and fibre technology. Both fields were new, and soon proved to be developing at a very fast pace. This led to the situation that as soon as a new trial project was initiated the technology that was chosen in the planning was outdated and no longer relevant in further planning. Transmission capacity had increased, developments in fibre technology had made better fibres available, and new and better components were available (at acceptable prices).

This can be illustrated in the (incomplete) list of trial projects in the box below left.

1.3 Choices

1.3.1 Technology

The rapid changes in both fibre- and transmission technology made planning difficult. On this background it was an advantage for Televerket to be a bit of a late starter in this field. Televerket did not have time to invest large amounts of money in technologies that would become obsolete in a very short time.

- Multimode (MM) fibre soon showed the limitations in length-bandwidth product, and in practical use limited to 140 Mbit/s. MM fibres with graded index gave better performance, but could not deliver the needed capacities and distances.
- The 850 nm window was very soon closed as the 1300 nm window was opened (light sources became readily available).
- Availability of single mode fibre, single mode light sources and practical techniques for the use of these components made SM fibre the only choice for all levels of the network. First for 1300 nm, since this window had near zero chromatic disper-

Time-frame	Projects
1980	Oslo – Økern. Planned 140 Mbit/s on 6 km dual window MM fibre. Oslo – Røverkollen. Planned 140 Mbit/s on 16 km 1300 nm MM fibre. Bryne – Nærbø. Planned 34 Mbit/s on 9 km 850 nm MM fibre.
1982	More trial projects were selected, among them Aurland – Flåm (34 Mbit/s at 850 nm on 9 km sea cable with MM fibre). Testing of aerial cable in difficult weather conditions.
1983	Installation and testing on all the mentioned trial projects. Decision to stop planning with 850 nm equipment. Discussions with NSB on use of their masts for the electrical wires on electrified rail lines for aerial fibre-optic cables (non-metallic).
1984	MM fibre was standardized for 1300 nm. Trial projects for single mode (SM) fibre were selected. Further planning of cable projects was only for SM fibre.

sion, and long reach of transmission systems could be dispersion limited.

- Availability of light sources and other components in the 1550 nm window gave longer reach due to lower attenuation, but higher dispersion.

1.3.2 Installation methods

Installation of underground cables is expensive. The low-weight, relatively small dimensions of fibre cables make them easy to handle, and several methods of installation were investigated.

A natural method was to suspend the cable on poles, both Televerkets own and others. NSB has been mentioned, and some trials were made in co-operation with power companies. In these environments a non-metallic cable was necessary. One method was to “spin” the cable around the power line. This demanded both a special cable and installation equipment. The conclusion was that this method was not practical for Televerket. The power companies would of course not let us do this for free, and the practical problems were many, for example repairs. Shooting birds with a shotgun when they were sitting on the cable/power line was never good for the cable, and repairs demanded that there were no voltage on the line. As will be seen later the power companies came to a different conclusion.

Underwater cables should be a natural choice in Norway. Several trials were conducted, and a project was started to evaluate the possibilities of using fibre-optic underwater cables. This was called FOS (Fibre Optic Sea cables).

An important factor was the co-operation with NSB to lay the cables along railway lines, mostly in the ground, but also on poles. Right of way was in most instances paid for by some of the fibres in the cable.

2 Plans

Already when the FOTgroup was started it was clear that using optics would be the way to go for the future transport network. It was only a question of when components for practical use would be available at acceptable prices. The inherent bandwidth capacity in the fibre was there from the beginning; it was just a question of getting the components to utilize this bandwidth. The traffic demand was increasing rapidly and optical systems seemed to have the capability of meeting this demand in a cost-effective way.

In 1982 it was decided to stop planning with coaxial cables in the core network.

2.1 FOTplan

At the same time as the trial projects started, work also started on a plan to use fibre-optics in the core network. It was clear at an early stage that fibre-optics would be capable of delivering the necessary capacities, could be used over the necessary distances, and that the developments in this field made fibre optics the economical choice as well as the technical one.

The FOTplan focused on the core network, which then was the network between the main telephone exchanges (FSIIs). This excluded all parts north of Tromsø, which was a “white” area on the maps. In other areas of Norway the regional and local networks could benefit from the work in the core network, as cables also passed other stations and local exchanges. The plan provided a framework, and was continually adjusted on a more detailed level for changes in technology, local needs, etc.

3 Building a network with fibre-optic cables

3.1 The first projects

As mentioned earlier the very first cables had multi mode fibres. These proved to have a very limited useful life due to developments in technology, and had very little impact on the optical network.

The real network started with the trial projects for singlemode (SM) fibres in 1984, on the routes Drammen – Oslo and Gjøvik – Hamar. These projects also included testing of transmission equipment. In 1985 it was decided to plan with SM fibre cables only. It is important to note that these fibres at that time had a full temperature specification only in the 1300 nm window. The specification did not include the 1550 nm window until 1988.

3.2 Building the core network

Further building started very tentatively in 1986 with a few scattered short routes. The activity increased in 1987, when some of the first projects were extended, and new and longer projects started. The FOTplan was introduced, which was to establish the main routes from Oslo to Trondheim, Bergen, Stavanger and Kristiansand by the end of 1991. In addition a strategic plan in Televerket stated that all FSII exchanges should be connected together by fibre-optic cables by year 2000. The building of this network started in 1988.

Evaluation of the fibres continued parallel to the building of the network. It became clear that the 1550 nm window would be very interesting in the

Figure 1 Status by end of 1988




Figure 2 Status by end of 1990

core network due to lower attenuation and the possibility for longer system reach.

Components for this window began to appear, in particular lasers with the necessary characteristics in output power and line bandwidth. It was therefore important that the network would be able to handle 1550 nm signals. The standardized SM fibre was optimized for 1310 nm, and tests showed that cables with these fibres could not be used at very low temperatures due to bending of the fibres. This situation would be expected on several main routes, and especially where aerial cables were used. The routes through the Østerdalen and Gudbrandsdalen valleys are good examples of routes that would be exposed to low temperatures. This led to new specifications in 1988, which included the temperature range in the 1550 nm window.

In 1988 work was ongoing on several routes. In Østerdalen work was stopped at Rena to wait for the new cables. This marked the end of the old type SM fibres, no more cables were installed with that type of fibre.

By the end of 1988 the network was as shown on the map in Figure 1. Very scattered projects, but the start of some of the main routes can be seen.

The building activities between 1988 and 1990 resulted in a marked extension of the network. The status by end 1990 is shown in Figure 2. One route between Oslo and Trondheim was established, and the route from Bergen through Stavanger round the coast to Oslo was well under way.

The graph in Figure 3 shows the dramatic increase in cable kilometres per year from 1986 to 1989. Comparing it to the graph for fibre kilometres per year in Figure 4 there is no such dramatic increase. This illustrates that even if many cables were installed the number of fibres per cable was low. In fact, the typical cable size was with six fibres (three pairs) in the beginning, increasing to eight and 12 fibres over a period. Fibres were expensive, so the number of fibres in the cables was kept low to reduce costs. Even so, the cost of a new cable project between Oslo and Trondheim could not compete with a radio link solution when compared only for end-to-end transmission. An important advantage for cables was that they would also cover the needs for many nodes between the end-points.

This low number of fibres in the cables meant that transmission technology had to develop fast enough to meet the demands for capacity. This meant in practice that the next hierarchical system level had to be available before capacity demands were higher than




Figure 3 Cable kilometres per year




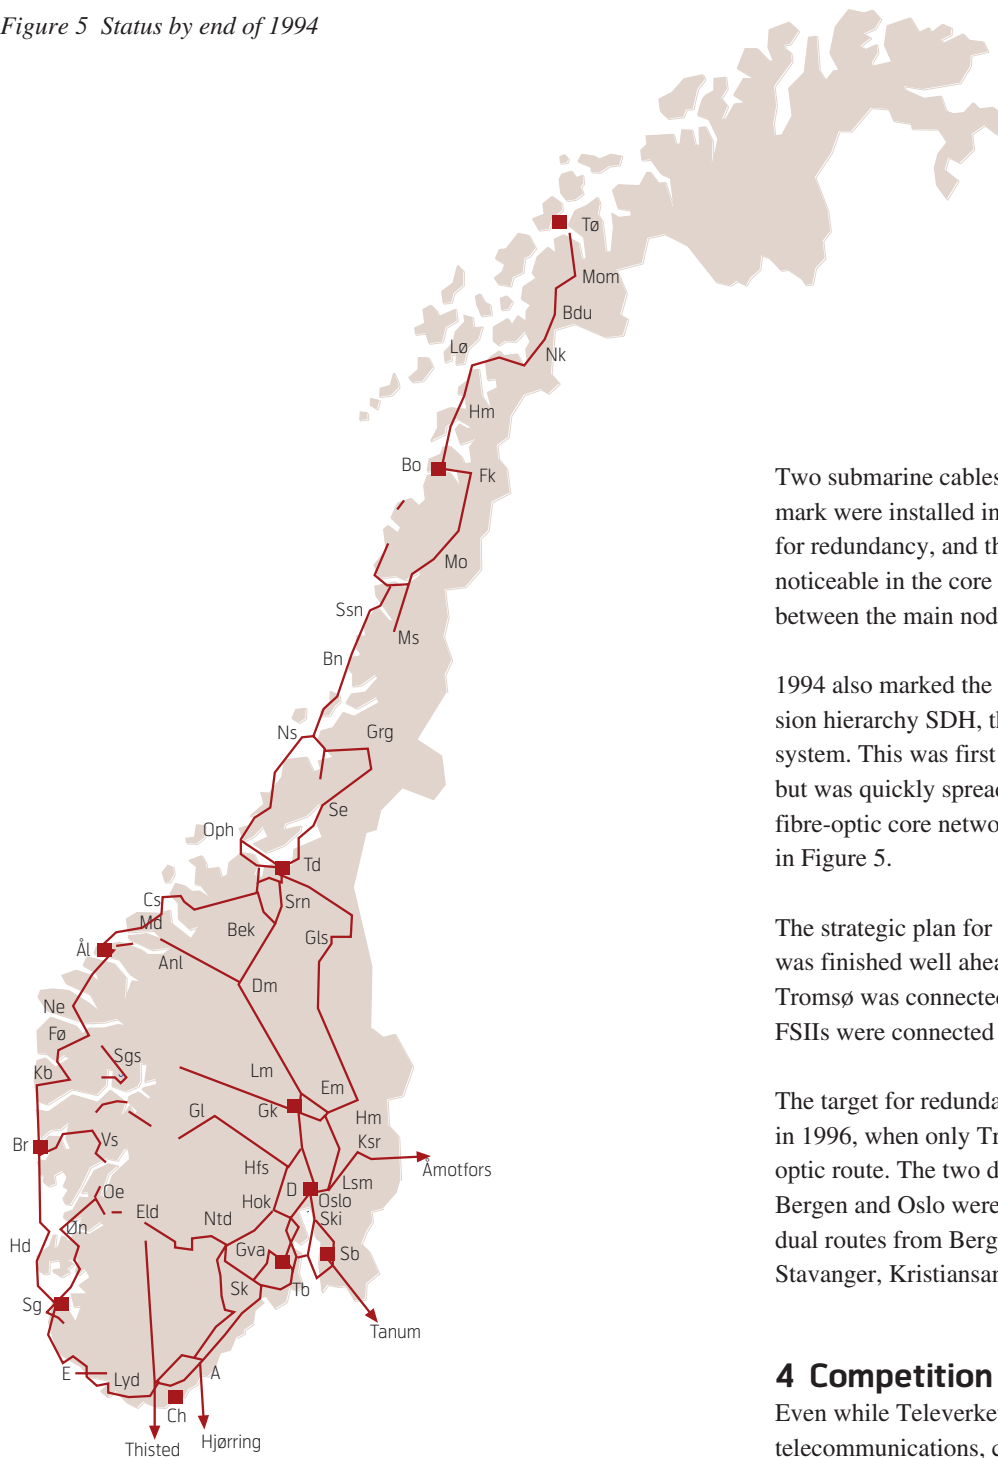
Figure 4 Fibre kilometres per year

what could be delivered with parallel systems at present level. With six fibres available, this meant that 565 Mbit/s systems had to be available before the demand was higher than what three parallel 140 Mbit/s systems could deliver, 2.5 Gbit/s systems before three parallel 565 Mbit/s were filled with traffic, and so on. This expectation was fulfilled up to 2.5 Gbit/s systems.

Several new types of fibres were developed, all with the aim of reducing the chromatic dispersion problem in the 1550 nm window. First was the dispersion shifted (G.653) fibres, with zero dispersion at 1550 nm. With these fibres the system reach could be extended, fewer regenerators would be needed, and consequently costs would be lower. However, the regenerator distance was in many cases dispersion limited.

These fibres were evaluated, but in 1992 the decision was made to continue with the “standard” SM fibre. The advantages of having only one type of fibre in

Figure 5 Status by end of 1994



Two submarine cables between Norway and Denmark were installed in 1992. Two cables were chosen for redundancy, and this redundancy thinking is noticeable in the core network. Redundant routes between the main nodes were the target.

1994 also marked the introduction of a new transmission hierarchy SDH, through the use of 2.5 Gbit/s system. This was first used in the Olympic network, but was quickly spread out in the core network. The fibre-optic core network at the end of 1994 is shown in Figure 5.

The strategic plan for connecting the FSII exchanges was finished well ahead of plan in 1994, when Tromsø was connected as the last FSII. All the other FSIIs were connected in 1993.

The target for redundant routes was nearly achieved in 1996, when only Tromsø was left with one fibre-optic route. The two direct connections between Bergen and Oslo were finished that year, as well as dual routes from Bergen around the coast through Stavanger, Kristiansand and Tønsberg, to Oslo.

4 Competition

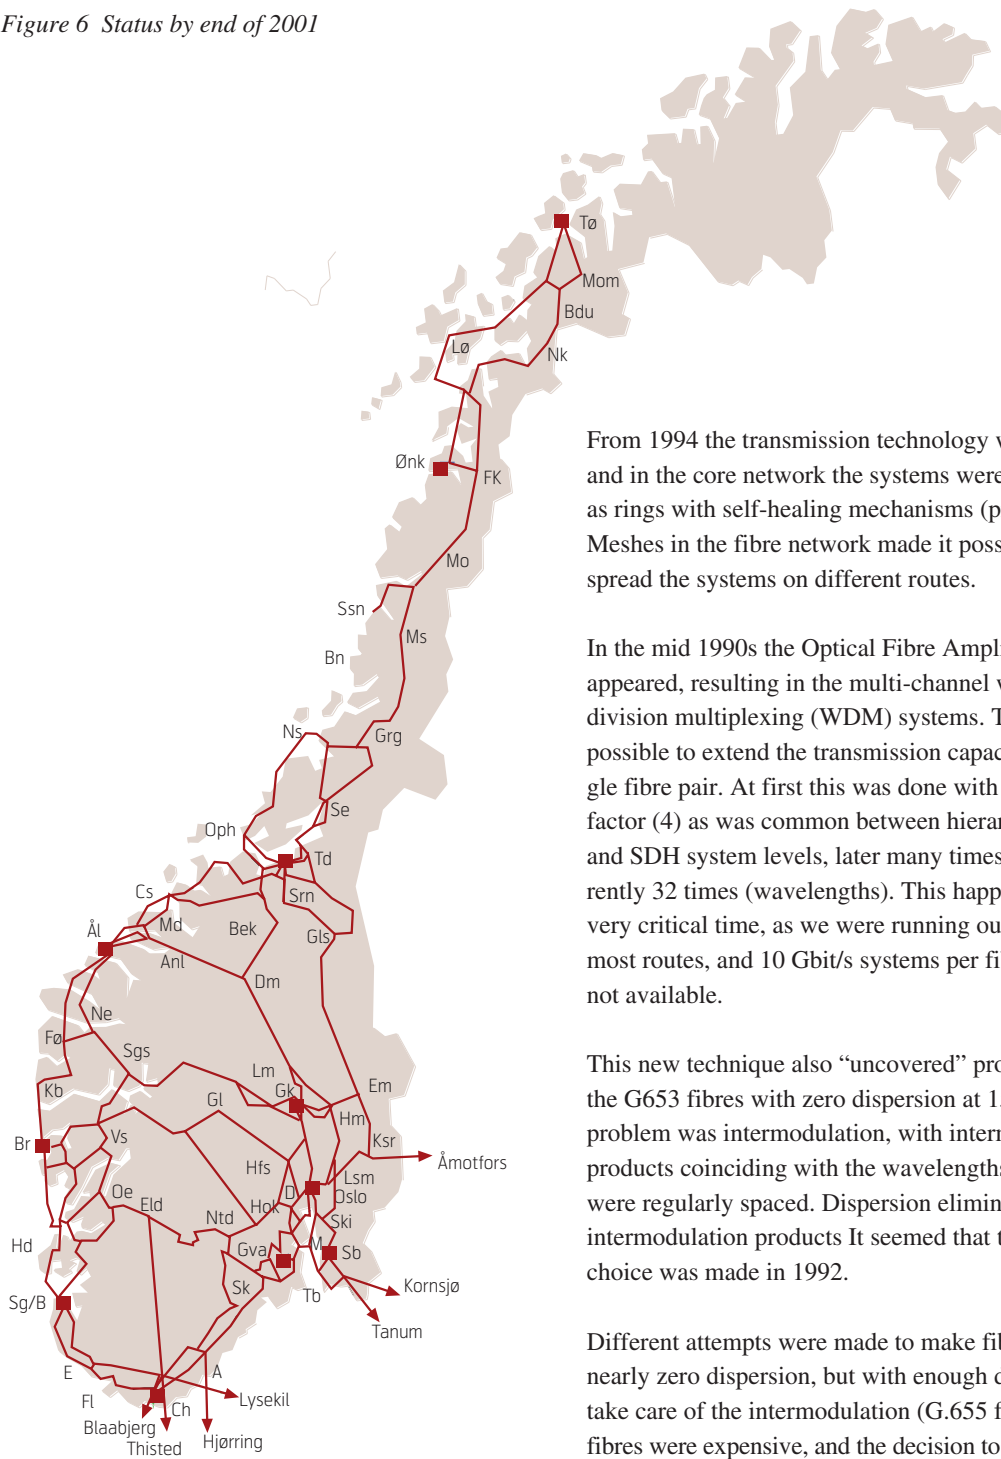
Even while Televerket had a monopoly on public telecommunications, companies and organizations could build their own private networks. NSB is a good example. They had a railway network that covered large parts of the country, and soon they also had an internal network with fibres, partly through the co-operation with Televerket.

The power companies also had a need for communications and found the method of spinning cables on the power lines more practical for their own use. Over some time quite a network has been built by this method. Mergers and takeovers of the different companies have led to a countrywide network that can compete with Telenor in delivering telecommunications solutions, even if we feel that they do not have quite the same diversity in their network.

the network were great. Components were developed which would counter the problems of standard SM fibres; the most important being dispersion compensating elements.

The Winter Olympics in Lillehammer in 1994 led to a lot of network building, especially in the local area, but also other parts of the network needed strengthening. This included connections abroad, and to the two domestic satellite stations.

Figure 6 Status by end of 2001



From 1994 the transmission technology was SDH, and in the core network the systems were constructed as rings with self-healing mechanisms (protection). Meshes in the fibre network made it possible to spread the systems on different routes.

In the mid 1990s the Optical Fibre Amplifier appeared, resulting in the multi-channel wavelength division multiplexing (WDM) systems. This made it possible to extend the transmission capacity of a single fibre pair. At first this was done with the same factor (4) as was common between hierarchical PDH and SDH system levels, later many times that, currently 32 times (wavelengths). This happened at a very critical time, as we were running out of fibres on most routes, and 10 Gbit/s systems per fibre pair were not available.

This new technique also “uncovered” problems with the G653 fibres with zero dispersion at 1550 nm. The problem was intermodulation, with intermodulation products coinciding with the wavelengths when they were regularly spaced. Dispersion eliminated these intermodulation products. It seemed that the right choice was made in 1992.

Different attempts were made to make fibres with nearly zero dispersion, but with enough dispersion to take care of the intermodulation (G.655 fibres). These fibres were expensive, and the decision to stay with the standard G.652 fibre was renewed as late as in 2002.

5 Further refinements in the fibre-optic network

When the foundation network with dual routes was in place in 1998, the work continued to build flexibility into the network. Diversity and redundancy are key words in this respect. Even if most of the diversity in the transmission network was taken care of in the transmission equipment, it was necessary to have diversity also in the fibre network. The network changed into a meshed network. In the south of Norway several east-west routes were established. This is pictured in Figure 6, which shows the network in 2001.

Protecting the fibre cables also became important, as some incidents uncovered weaknesses in the network. Even if most of these incidents were a result of several simultaneous faults, they resulted in heavy activity in the years 2000–2001. This is clearly seen in the graphs in Figures 3 and 4. Parallel routes were established, several independent cable intakes to the stations were made, and so on. The aim was to make the network more robust. The network is normally dimensioned to handle one fault, but since work in the network can have the same effect as a fault (the protection route is used during work), a “real” fault can have a dramatic effect.

In the core network the focus is now on diversity, with parallel routes, more meshes, and careful “placing” of the transmission systems in relation to each other. New transmission methods (ATM, IP, Ethernet) have taken the protection up a level from the transmission equipment. This places more focus on the fibre network (and the WDM systems) and the diversity there.

An extension of the fibre network north of Tromsø is going on, but slowly. The traffic demand is still within what can be delivered with radio links, but demand is increasing. Activities in the oil sector is one driving factor in this.

Fibre cables are spreading out in all parts of the network, even if the “fibre to the home” is not yet a reality. However, activities are going on, but that is another story with other challenges.

Harald Jansen was educated at the Norwegian Institute of Technology and started work in 1972 in the Line Section of the Norwegian Telecommunications Administration. His initial work was in the planning of line systems on metallic pairs, and moved on to systems on fibre-optic pairs, a field he is still working in.

email: harald.jansen@telenor.com

Cost and benefits of survivability in an optical transport network

GUIDO MAIER, MASSIMO TORNATORE AND ACHILLE PATTAVINA



Guido Maier is Researcher at CoreCom, Milan, Italy



Massimo Tornatore is a PhD candidate at Politecnico di Milano, Italy



Achille Pattavina is Professor at Politecnico di Milano, Italy

In optical networks a link failure may cause a huge data loss due to the ever-increasing capacity of WDM links. Survivability to failures in the optical layer is thus of great importance. This paper presents the most common protection techniques for optical mesh networks and introduces the reader to the approaches that can be used to design the network minimizing the excess cost due to survivability. On the other hand, we will show how the effectiveness of different protection mechanisms can be compared in terms of lightpath availability, a quality-of-service parameter that gives a measure of the degree of network resilience.

1 Introduction

The Optical Transport Network (OTN) has today become the key to high-capacity network infrastructures. The use of optical-fibre technology, Wavelength Division Multiplexing (WDM) and the now ample set of well-established control and management protocols, allow for high-capacity connections on-demand. By employing advanced photonic technology, optical networks can provide switching and routing of optical circuits in the space and wavelength switching domains. On the switching side, Optical Cross Connect (OXC) systems have recently become available in addition to the more mature Optical Add-Drop Multiplexers. This opens the possibility of deploying complex WDM networks based on mesh topology, while in the past single ring or overlaid multi-ring have been the most used architectures for WDM networking. Mesh topologies are preferred to rings because they attain a considerably better use of the available bandwidth as well as provide better traffic engineering and efficient $M : N$ restoration schemes (that is where M working paths share the same N protection paths).

These past years have also seen a continuous growth of system aggregated bitrate. Today WDM transmission systems allow the multiplexing of 160 distinct optical channels on a single fibre, while recent ex-

perimental systems support up to 256 channels [1]. Given the high bitrate carried by a single WDM channel, e.g. 2.5 to 40 Gbit/s [2], the outage of a high-speed connection operating at such bitrates, even for few seconds, means huge loss of data. The increase in WDM complexity associated with the evolution from ring to mesh architectures, together with the tremendous bandwidth carried by each fibre, brought the need for suitable protection strategies into the foreground.

Survivability, i.e. the capability of keeping services active even in the presence of failures, is obviously a general property that applies not only to optical networks but to networks in general. Resilience strategies have been developed in the past for a range of network architectures and at many protocol layers. For example, in the case of the IP protocol, survivability is achieved essentially by routing the packets (*datagrams*) through the network dynamically, keeping the network-element state into account. In IP, routing is distributed, i.e. any IP router takes the routing decisions applying the same algorithm on its own "image" of the network. Each router has a direct knowledge of only a small part of the network: its neighborhood. In order to create its network "image" it has to receive and gather information from its peer routers. Information-exchange between routers occurs according to a dynamic routing protocol, the best-known and widest spread being Open Shortest-Path First (OSPF). The basic IP resilience mechanism, then, works as follows. When a failure occurs, some routers detect it and inform the other routers by sending OSPF signaling messages. Meanwhile, they modify the routing and direct packets to bypass the failed elements. When all the routers are aware of the failure, the new routing that skips the failures is consistent in the whole network: in this way, traffic protection is automatically achieved.

IP has not been casually mentioned: given the predominance of TCP/IP as the protocol architecture that supports the great majority of telecom applications

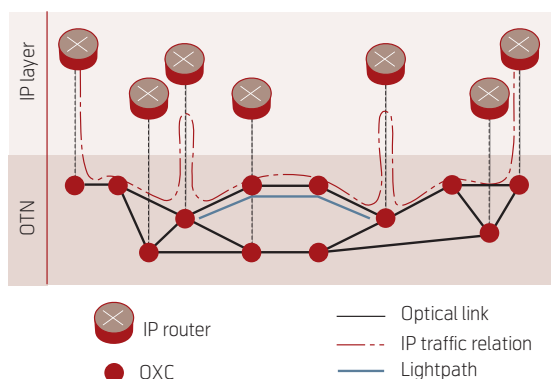


Figure 1 Layer structure of an IP-over-OTN network

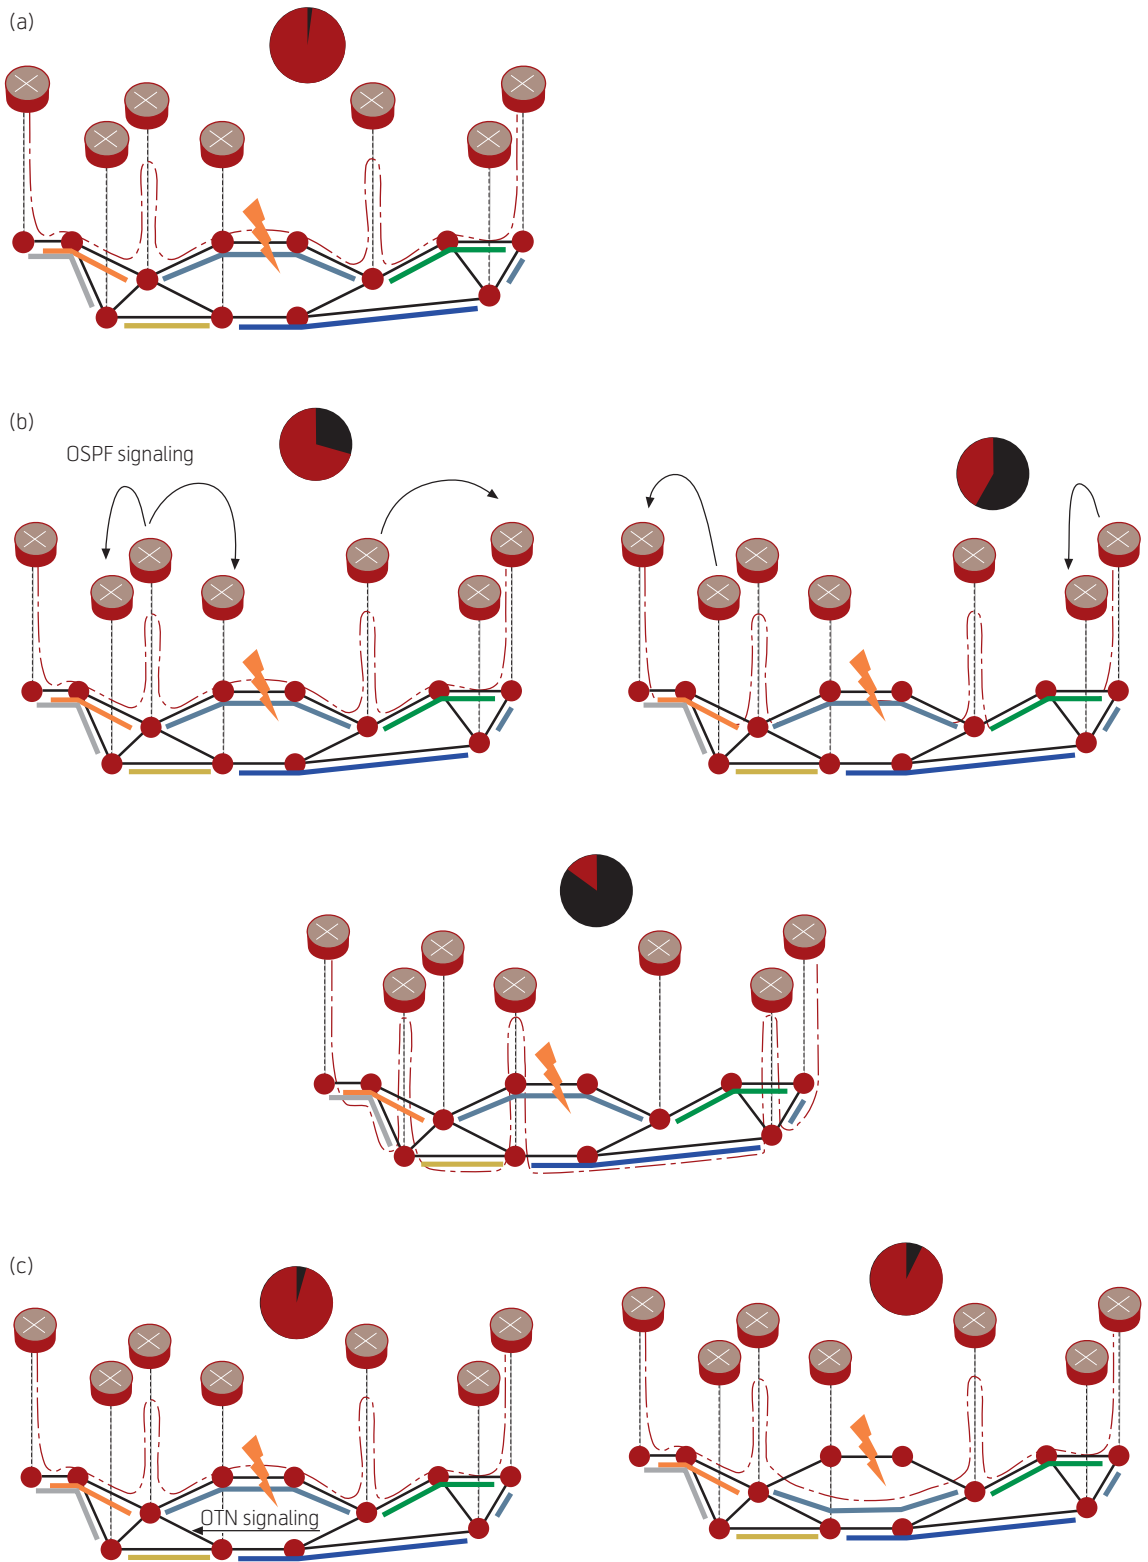


Figure 2 (a) A failure occurs on a fibre link of an IP-over-OTN network. (b) The failure may be recovered at the IP layer, e.g. by OSPF: the procedure takes time in the order of tens of second. (c) Optical protection is instead much faster and reacts in a few milliseconds

today, IP has become the most important and frequently adopted client for the optical layer. Figure 1 shows schematically the IP-over-OTN architecture.

After our brief digression of IP (OSPF-based) protection, the reader may wonder why in an IP-over-OTN

network resilience could not be provided at the IP layer, making protection in the optical layer obsolete. The main reason to implement resilience in the optical layer with its own protection mechanisms, is the failure response time. Let us assume that at a given time a failure struck a physical link of an IP-over-

WDM network, as represented in Figure 2(a). As we have explained above, IP traffic is recovered by a dynamic change in routing, via OSPF protocol (see Figure 2(b)): this implies delay for signaling propagation and processing and delay for router reconfiguration. It should be noted that OSPF messages are sent within IP datagrams and thus require complex layer-3 packet processing. Figure 2(c) shows the case in which protection is performed at the OTN layer. Still we have delay for signaling propagation and reconfiguration. However, signaling is sent at this layer exploiting optical control circuits directly from the final node to the ingress node of the protected lightpath, without the need of processing in intermediate OXCs. Reconfiguration is also fast.

In conclusion, the main reason for implementing protection at the optical layer is to achieve a fast recovery of faulted connections: optical protection mechanisms at the layer are able to restore connectivity in less than 100 ms (typically, well below 50 ms). OSPF-based traffic recovery requires tens of seconds to be carried out completely. The difference is of several orders of magnitude¹⁾. Such difference allows to recover connectivity so fast in the optical layer that OSPF is not even able to detect the failure.

Let us go back now to the OTN, explaining optical resilience in more detail. In the traditional ring-based networks the protection requirements are satisfied by well-known and tested solutions existing for quite a long time. For their simplicity and ease of integration with SDH structures, WDM ring topologies can be considered historically the second stage in the evolution of optical networking and represent the environment in which WDM protection techniques have come to be standardized.

In recent years the issue of survivability of optical connections has become of outstanding importance also in mesh WDM networks and has raised much interest in the research community. Undoubtedly, the adoption of protection techniques is traded off by a more complex network design; this has to include a further aspect of dimensioning and handling of the additional resources required to face the link failure, for example for the rerouting of lightpaths involved in a failure. These problems can no longer be manually solved in complex network architectures, as usually happened in the earlier experimental WDM system deployments. Computer-aided planning tools and procedures are needed in order to achieve an efficient utilization of network resources. Research on optical

networking has recently been investigating design and optimization techniques in order to provide operators with the most efficient and flexible procedures to solve the network design problem.

The improvement of the network performance attainable by introducing protection can be quantitatively measured. Generally speaking, availability and reliability are therefore parameters to be used for both repairable and non-repairable systems. Given that OTN is clearly repairable, the most important feature is connection availability. By this parameter the operator is able to quantify the quality of service that is offered to the user in terms of maximum downtime percentage.

Clearly any protection technique requires additional network costs to deploy spare resources that are traded off by the network operator's capability of guaranteeing agreed levels of connection availability to customers. While methods aimed at planning survivable networks have been extensively studied in the last decade and have resulted in a number of protection methods, the related topic of how these affect availability is receiving growing interest today. In particular, the definition of a standard model of service level agreement for the optical layer (O-SLA) is today largely debated. A service level agreement is a formal contract between a service provider and a subscriber containing detailed technical specifications called service level specifications (SLSs). An SLS is a set of parameters and their values that together define the service offered to a traffic stream in a network. Until now, no standards for the contents of an SLS have been finalized, but interesting proposals have been published as Internet drafts by the Internet Engineering Task Force (IETF) [3]. A recent proposal [4] identifies the service unavailability as a key parameter to define a class of service distinction for optical circuits (see Table I).

Availability and Reliability (A&R) analysis is a fundamental tool for the operators to understand the relations between the protection mechanisms they install and the performance of connection integrity of their network. The final goal is to optimize the trade-off

CoS	Premium	Gold	Silver	Bronze
Service Unavailability	10 ⁻⁵	10 ⁻⁴	10 ⁻³	10 ⁻²

Table I Optical circuits class of service

¹⁾ Often MPLS is adopted as intermediate layer between IP and OTN. Several protection mechanisms have been proposed for MPLS. These mechanisms are faster than OSPF, but still in the range of seconds.

between extra deployment costs and higher revenues from more advantageous service level agreements.

The first aim of this paper is to compare the performance of some protection techniques that have been largely discussed in previous literature in terms of the number of fibres required to support a given offered traffic. We will show how to obtain optimal solutions by exploiting exact methods in order to guarantee comparison between optimal results. Using heuristic approaches to accomplish network dimensioning would imply an uncertainty due to the approximations and/or sub-optimality inherent in such methods. In particular we focus on Integer Linear Programming (ILP), a widespread technique to solve exact optimization. ILP formulations used to carry on this comparative study are based on the universally accepted flow and route paradigm [5] that we will explain in the following.

In the second part of this article we focus our attention on the analysis and comparison of the availability performance of protected OTNs. In particular, we will consider any possible end-to-end protection technique: each dedicated and shared configuration will be analyzed by a combinatorial approach, providing a closed-form algebraic equation (sometimes by introducing approximations). These simple back-of-the-envelope equations are, however, sufficient to reveal useful properties of end-to-end protection that are in turn presented later on.

The rest of the paper is organized as follows. Section II describes the features of the protection strategies that will be analyzed and compared. Section III briefly introduces the most common approaches to model protected network design, focusing on the ILP-based method, where some consideration on the advantages and drawbacks of exact vs. heuristic methods are also given. To conclude the first part of this work, in section IV results obtained by means of the ILP formulations to a case-study network are shown; this allows us to point out the network cost implied by the adoption of the different protection techniques. Section V opens the availability-focused part of the paper illustrating the assumptions and basic principles on which our analytical model is based; in Section VI we present the derivation of the algebraic relations that evaluate the availability performance of the dedicated and shared $N : M$ end-to-end protection schemes. In Section VII we report some numerical examples to compare the availability degree provided by the different protection techniques and highlight dependencies of A&R on some network parameters.

II Protection techniques in WDM network

After the introductory discussion on WDM networks and the drivers for WDM survivability, let us review the details of the protection techniques that have been taken into account in this comparison study. In the rest of the paper we will assume a mesh network as the reference topology. Although the ring is the most common physical topology today, WDM mesh networks are gradually attaining growing importance, especially thanks to the development and improvement of the OXC. In a mesh network, survivability is a more complex problem than in a ring topology because of the greater number of routing and design decisions that need to be made [10]–[12].

Two general and “orthogonal” criteria can be assumed in order to classify these techniques. A first classification criterion regards the entity to be protected, so that protection can be applied directly on the single optical link or on a whole lightpath connecting two end-nodes. Actually, this simple distinction reflects the particular sublayer of the WDM layer [ASON] in which a given protection mechanism operates. Two alternatives exist: Optical Channel (OCh) sublayer or Optical Multiplex Section (OMS) sublayer. In the former case the lightpath is the entity to be protected, so that OCh-protection is also called *path protection*. In case of failure each single interrupted lightpath is switched on its protection path [6]. Recovery operations are activated by the OCh equipment hosted in the end-nodes (source and destination) of the lightpath. These systems also have the duty of monitoring lightpaths for failure detection. The protected entity is called *working lightpath*, while after the failure the optical circuit is switched over to a *protection lightpath*. This lightpath can be pre-allocated or dynamically established.

On the other hand, the OMS-sublayer managed entity is the multiplex of WDM channels transmitted on a fibre. Thus at this sublayer fault recovery regards each network link individually, so that this approach is also called *link protection* [7]. The OMS equipment in the terminations of the fibres composing a single link locally manages fault-detection and protection switching. The protection mechanism reacts to a failure by diverting the interrupted WDM multiplex to an alternative path, thus bypassing the damaged components. The main difference from path protection is that all the lightpaths travelling along a broken fibre are simultaneously re-routed. Link protection is commonly implemented adopting one of two alternative modes: depending on signalling capabilities, either all the fibres belonging to a failed link must be jointly re-routed, or the protection scheme can be applied at

the level of the single channel, setting an alternative path for each failed wavelength.

A second classification criterion distinguishes between *dedicated protection* and *shared protection*. The simplest and most conservative procedure is the reservation of a set of spare resources exclusively to one working entity (a lightpath in OCh protection or a link in OMS protection). This is the so-called dedicated protection: it reduces the complexity of failure recovery, but requires that at least 50 % WDM channels cannot be used by the (non-preemptive) working traffic. Since pre-planned protection is based on the assumption that a multiple failure is a very unlikely event, two or more protection entities (lightpaths or fibre sequences for OCh and OMS protection, respectively) can actually share some resources (WDM channels or a fibre, respectively). This is possible provided that the corresponding working entities cannot be simultaneously involved by a single failure event, i.e. they cannot belong to the same Shared Risk Link Group (SRLG), a concept introduced in recent literature [8], [9]. In this case all the fibres in the same link (bundle) form an SRLG²). Shared-protection strategy exploits this property by preplanning the network so that some WDM channels or fibres are shared by more protection entities. Shared protection allows to sensibly reduce the amount of spare resources and to improve network utilization for working traffic, at the cost of increasing the recovery procedure complexity (this point will be discussed later).

A Path protection

Path protection at the OCh layer is obviously well applicable to mesh networks. To satisfy each connection request a pair composed of a working and a protection lightpath has to be established (Figure 3). For the protection mechanism to be effective against link failures, the links of the working and protection lightpaths must be independent in the sense of failure occurrence. In our analysis, this condition is satisfied by setting up the two lightpaths in physical-route diversity: the primary and backup paths cannot share any link (link disjointness³).

Care must be taken when imposing physical route diversity. A network topology simply representing fibres or cables as separated arcs may be misleading. Ref. [8], [13] discuss cases in which distinct arcs

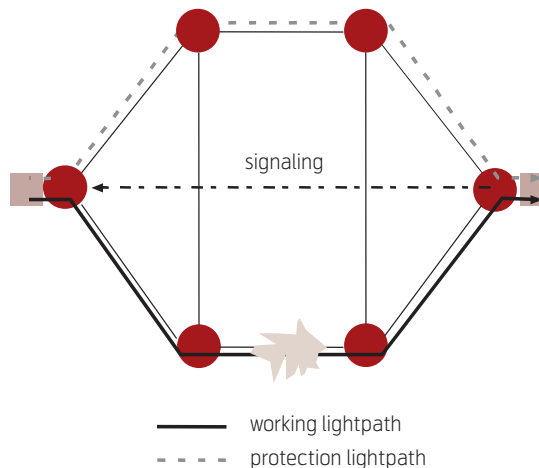


Figure 3 Path protection in a mesh network: in 1 : 1 dedicated protection, signaling is required

of the physical topology share the same infrastructure (e.g. two different fibre cables crossing a river on the same bridge). Two dedicated path protections are defined, 1 + 1 and 1 : 1. In the former case the same signal is transmitted on two diverse paths by the transmitter node, while the receiver node is in charge of choosing the signal with the higher SNR (or, more generally, with better characteristics). A link failure event can be bypassed without signalling exchange. In the second case (also called *protection transferring*), low priority traffic can be transmitted on the protection lightpath in absence of failure, but end-to-end signalling becomes necessary (Figure 3).

Dedicated path protection (DPP) is quite resource consuming in mesh networks because of the physical route diversity constraint. Sharing of WDM channels among protection paths may reduce the physical resources employed for protection. Shared protection may be applied in an end-to-end sense using a single protection lightpath for N working lightpaths with the same source-destination node pair. This technique is a special case of sharing in which N protection lightpaths share all their WDM channels (known as 1 : N protection). Obviously 1 : N protection requires that $N + 1$ link-disjoint paths are available between the source and the destination nodes of the connection. So this protection strategy implies a high connectivity degree in the source and destination nodes that exploit it, but a realistic scenario of WDM network deployment, especially in wide-area application,

2) Let us observe that when applying a link protection strategy, two or more protected entities (the link) cannot be involved in a single failure event, under the hypothesis of failures affecting links but not nodes. So the dedicated case provides a large redundancy of backup capacity that will improve the survivability of the network against multiple failure events.

3) The term "link disjoint paths" has entered the common usage in literature to indicate the condition of preventing physical resource sharing (see [7]). The term disjoint is not entirely appropriate, since in probability theory it refers to events not happening at the same time: independent should be used instead. We will however follow the common convention in this paper.

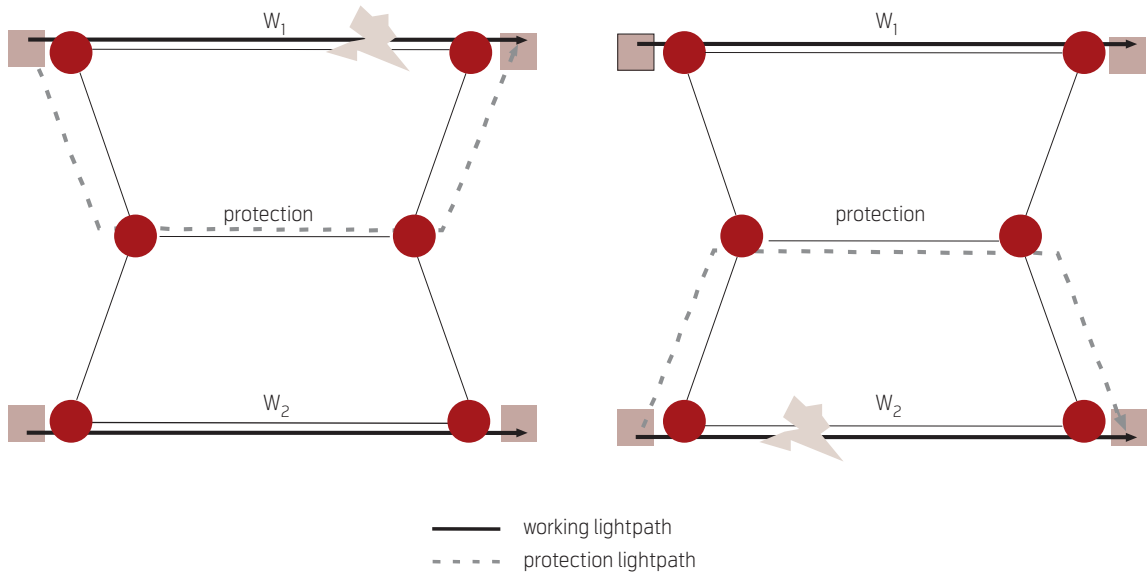


Figure 4 Shared path protection in a mesh network. Network configurations when a failure affects the lightpath w_1 (a) or the lightpath w_2 (b), whose protection-lightpaths share a common fibre

would probably be characterized by low values of connectivity index.

The end-to-end shared protection can be generalized by adopting more than one – e.g. M – protection paths to backup N working lightpaths. This protection technique, indicated as $M : N$ can achieve higher reliability compared to $1 : N$, as we will show later. It is worth mentioning that with $M : N$ we need a total of $M + N$ disjoint paths between the two end-points.

The *shared path protection (SPP)* scheme is implemented in a wider sense on a mesh network by allowing partial sharing among the protection lightpaths. In this case an additional constraint must be taken into account: protection lightpaths sharing WDM channels must be associated to working lightpaths that are mutually link disjoint [6], [11]. It is important to notice that sharing allows savings in terms of transmission resources, but it also increases control plane complexity. In $1 : 1$ and in $M : N$ protections, when a failure occurs, only the end-nodes are involved in the recovery process, because the protection lightpaths are completely set-up in advance. When shared-path protection is adopted in the wide sense in a mesh network, the fault event activates a more complex recovery procedure that requires a lot of signalling among several network elements. It is in fact necessary to reconfigure all the OXCs that are terminations of shared WDM channels (see Figure 4) according to which particular working lightpath needs to be recovered [14]. These operations increase the recovery delay, which will be limited by the time taken by the signalling messages to reach all the involved elements plus the time taken to reconfigure all the OXCs.

Since shared protection is a pre-planned strategy, the recovery operation could be controlled in a distributed rather than in a centralized way, thus eliminating the intervention of the network management system and reducing the amount of signalling. In this case the OXCs must be able to autonomously identify the faulty working lightpath in order to switch accordingly. The first operation requires real-time detection of the lightpath identity and it is one of the main motivations that fostered the definition of an OCh identifier in the framework of the standardization of the OCh supervisory channel (ITU-T G.872, G.709, G.798 recommendations).

B Link protection

In WDM mesh networks, *link protection* at the OMS sublayer under some aspects can be preferable to path protection. In a complex topology, a local recovery mechanism, more suitable to distributed than to centralized control, is easier to manage than an end-to-end mechanism. The present-day realizations of this protection technology are implemented by means of self-healing rings that provide a local (along the ring) shared utilization of backup resources. Link protection on a mesh network can be realized in various ways [15]. Basically two approaches can be followed to accomplish link protection; following a link failure either all the fibres crossing the link are rerouted on a common protection route, or each channel is rerouted independently on different paths (Figure 5).

In our approach link protection consists in providing a single alternative path to each link in the network. In other words, given the number of fibres on a link needed to support offered traffic, an equivalent number of fibres has to be planned along an alternative




Figure 5 Link protection in a mesh network. When a link fails: (a) all the fibres crossing a link are rerouted on a common protection route, or (b) each channel is re-routed independently on different paths

route, by-passing the link to be protected. This can be done reserving distinct backup capacity for each link (*Dedicated Link Protection, DLP*). Clearly, in order to avoid an excessive waste of spare fibre capacity, a shared strategy is preferable, also considering that a single failure may not affect more than one protected entity (link) (*Shared Link Protection, SLP*). In this latter case of SLP we will consider the two different protection approaches represented in Figure 5: protection is guaranteed altogether for the whole fibre (SLP-F), or independently for each channel supported by the fibre (SLP-C). This latter strategy (applied in a shared scenario) is expected to provide a more efficient utilization of spare resources, while it implies a more complex switching architecture to process failures and route each channel separately at termination nodes. This approach to reduce resource over-provisioning can be effectively implemented thanks to the new capabilities provided by (G-)MPLS protocol.

III Fibre number estimation

Solving the routing and wavelength assignment problem in WDM networks has been proven to be an NP-hard problem [16]. Our objective involves the introduction of other two terms of complexity in the problem: the models of protection techniques and the evaluation of the minimum number of fibres on each link to support a given traffic matrix. So the Routing and Wavelength Assignment (RWA) problem scales to a more computational intensive Routing, Fibre and Wavelength Assignment (RFWA) problem with protection objectives. In order to solve the problem in a reasonable computational time, in some cases we have introduced some simplifications. These approximations will not affect the validity of the comparison between the different protection techniques under analysis [11]. According to many studies that show

the marginal effect wavelength converters have on the global amount of required transmission resources, we have decided to solve the case of networks with all nodes equipped with wavelength converters; these networks are usually referred to as Virtual Wavelength Path (VWP) networks. This assumption allows us to neglect the problem of wavelength assignment (wavelength continuity constraint), keeping the other constraints unchanged [5].

Of course, ILP represents a flexible mathematical tool to model graph problems, such as those arising from network routing and design when protection requirements are introduced. The application of LP to solve the design problem in optical networks is a mature problem and a very rich literature exists on this topic. The basic analysis has regarded the single-fibre case, in which the RWA problem has been studied [11], [17]. In the multifibre scenario, the problem scales to the more complex RFWA problem: formulations to model and solve it can be found in [5], [10], [18]–[20]. All of these studies are based on two traditional approaches: the flow formulation and the route formulation. In the former the basic variables are the flows on each link relative to each source-destination node pair; in the latter the basic variables are the paths connecting each source-destination pair.

ILP models to solve the RFWA problem are characterized by a well-defined set of constraints:

- solenoidality constraint;
- capacity constraint;
- integrality constraint.

First of all, the network flow problem requires a basic constraint to guarantee that the traffic offered by a source node reaches its destination node. The so-

called solenoidality constraint sets the flow conservation condition; in other words, for each node and for each connection request in the network, this condition states that the total flow leaving a node must be equal to the total flow incident on that node. This equation is slightly modified in the source (destination) node, where the outgoing (incoming) flow must be equal to the required traffic.

Secondarily, the capacity constraint allows us to dimension the physical network capacity. In order to ensure a feasible resource allocation, it ensures that on each link the sum of flows generated by all the nodes is smaller than the product of the number of fibres by the number of wavelengths per fibre (i.e. the capacity of the link expressed in terms of λ -channel).

Let us observe that, in the following comparison, only VWP networks have been investigated, where one has only to deal with capacities, reducing the formalization of the RFWA problem to the capacitated network design problem [21]. When the nodes have no wavelength conversion capabilities, every path and protection structure becomes coloured, so that the problem has to also consider the wavelength constraint needed to impose the same wavelength along a path. Therefore, the number of variables is multiplied by a factor $|W|$, when W is the set of available wavelengths per fibre. In today's WDM transmission systems, realistic values of W are in the order of tens of λ -channel (typical values are 20, 40, 64, 128 or 160). This makes the ILP approach even for small networks infeasible and one has to rely on heuristics with lower complexity. Anyway our assumption of VWP network does not affect the objective of the proposed comparison: different studies have highlighted the marginal role of wavelength conversion under static traffic showing that in this case the two scenarios lead to very similar results. We can thus argue that the efficiency of the different protection strategies in terms of required additional resources is not significantly affected by this assumption.

The integrality constraint has to be applied on flow (or route) and capacity variables. Actually, these two groups of variables play completely different roles. Flow variables are related to the routing and multi-commodity flow problems and in these fields good results have been obtained by relaxing the integrality constraints. It has been proven that for a single flow unit the previous constraint is superfluous, while, in the generic n -connection case, techniques such as randomized rounding based on LP relaxations have shown some merits. On the other hand, the introduction of the capacity variables implies that RFWA scales from a multicommodity flow problem to a more complex localization problem. The application

of a relaxation on these last variables does not often allow us to obtain a significantly lower bound.

Beside these basic conditions, additional constraints must be introduced in the formulations to model the different protection techniques. Actually, this additional set of constraints can be imposed in different ways with respect to the choice of flow or route variables and to the detail in the description of the problem (e.g. taking into account further circumstances such as node failures, partial wavelength conversion, cost function typology would require a different structure of the ILP formulation). In any case it is possible to identify some common conditions to be satisfied. As far as the DPP (Dedicated Path Protection) case is concerned, the main constraint stems from the link disjointness condition: no more than one lightpath associated to a connection request can coexist in the same link (or more generally in the same SRLG) [5], [10]. This check could be avoided only if we exploit as basic variable a diverse path routed pair, composed by a link-disjoint couple of a working and a spare connection.

In the shared case a (pre-determined) protection path is set up only if the corresponding working path fails due to a network failure that occurs in any location. To handle such a mechanism in mathematical programming we have to introduce new indicator variables that imply a large increase in the number of variables. More generally, the huge complexity involved with shared mesh protection exact models is due to the following control: an optical channel can be shared between several spare lightpaths, only if their associated working lightpaths are link-disjoint. In other words, if some working lightpaths are routed on a common link, their corresponding spare lightpaths cannot share an optical channel. When this condition is fulfilled then if a link fails, it will always be possible to reroute traffic on spare paths because the two connections will be utilising different channels.

In order to deal with the complex SPP management, the set of basic variables and constraints of dedicated case must be extended to store such kind of information: the working lightpath associated to a given connection crosses the link i and the associated spare lightpath crosses link j . The increase in complexity due to the collection of this network knowledge makes the ILP infeasible also in very small networks. Our optimizations have failed in finding optimal solution starting from simple low-connected six-node topologies [20]. So an approximate route-based approach has been carried out, reducing the field of admissible paths. Anyway, a large body of previous studies have confirmed that approximate solutions are sufficiently close to optimum solutions.

The link protection has been subject to several modelling approaches, too [6], [15]. To the basic constraints in this case we have added a new condition that must be applied to each single link: all the flows on a link need an alternative and link-disjoint path to reach the opposite end-node. A route-based approach will pre-compute all (or just a subset of) the admissible paths that circumvent a given link; in a flow-based case, we could impose additional solenoidality constraints on the end-nodes of each link to reroute all the traffic flowing on the link (paying attention to reroute the traffic on the same network excluding the link in object). The basic variable that models the entity to be re-routed (fibre or channel) will be different in the *SLP-F* and the *SLP-C* scenario.

IV Comparison on a case-study network

After having presented the formulation for each protection strategy, we compare now their performance. We have set as objective function the number of fibres needed to support a given static traffic. This coarse cost function has some merits: while minimizing the fibre number, the objective function includes also the cost of transmission equipment associated to each fibre and tries to minimize the global amount of switching fibre port in the network. Clearly we are referring to a simplified estimation with respect to the actual amount of network resources: on the other hand, a more complex description of network cost would increase the number of variables and constraints, leading to computational infeasibility.

We present and discuss the results obtained by performing dimensioning on a case-study network, the (United States) National Science Foundation Network (NSFNET) that includes 14 nodes and 22 links. Its physical topology is shown in Figure 6; the offered traffic matrix (360 connection requests distributed on 108 node couples) is taken from Ref. [10].

The mathematical details of ILP formulations exploited to obtain the following results can be found in [19] for the unprotected case, in [20] for SPP and in [22] for SLP and DPP.

All the obtained results are the optimum of the problem, except for the shared protection cases, which anyway are proven to be close to the optimal ones. The computation time spreads from a few seconds to a maximum exceeding one day.

Figure 7 shows the total network fibre requirements M associated to each protection strategy. The most expensive technique is the link protection in the dedicated case; there is no advantage from a backup-

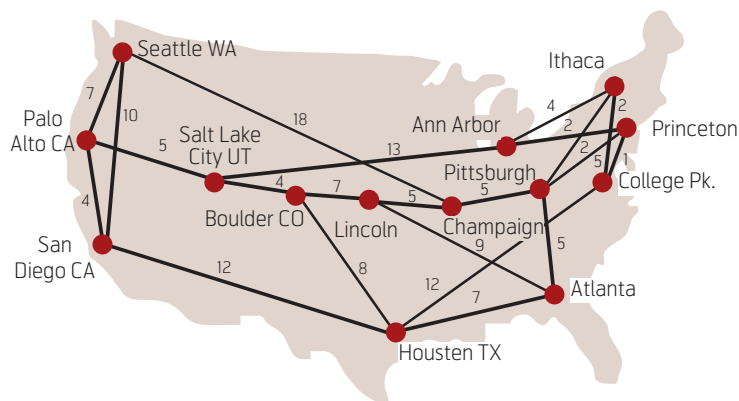


Figure 6 NSFNET network physical topology




Figure 7 Total fibre number on NSFNET exploiting different protection techniques

capacity planning point of view in reserving the protection resources separately to protect single failure events. The positive effect on survivability of this large capacity redundancy emerges when multiple failures occur. DPP returns a more efficient result than dedicated link protection (DLP), but it still requires more resources than shared strategies: both shared path protection (SPP) and shared link protection (SLP) show a better utilization of fibres; in particular, the increase in fibre number with respect to unprotected case is always lower than 100 %.

Table II numerically reports the additional amount of physical capacity needed to support the different protection techniques. We express the percentage extra cost with respect to the unprotected case (un) by defining the parameter Add_{pt} for each protection technique ($pt = \{DLP, DPP, SLP, SPP\}$):

$$Add_{pt} = \left(\frac{M_{pt}}{M_{un} - 1} \right) \cdot 100$$

W	Add_{DLP}	Add_{DPP}	Add_{SLP-F}	Add_{SPP}
2	327 %	162 %	52 %	48 %
4	327 %	160 %	58 %	50 %
8	330 %	158 %	65 %	52 %
16	330 %	145 %	82 %	50 %
32	343 %	116 %	95 %	56 %

Table II Percentage extra cost Add_{pt} with respect to unprotected case for different protection techniques

Considering any specific protection strategy, the additional term of capacity shows a small variation for all the values of W (number of wavelengths per fibre) that we have analyzed. Only the SLP-F case seems to require a larger number of fibres for increasing values of W .

Figure 8, comparing the two protection techniques SLP-C and SLP-F, shows that for a small number of wavelengths there is no significant difference. However, for fibres supporting a larger number of wavelengths individually rerouting the single channels of the failed link appears to be more efficient. This gain on the number of fibres is paid by a more complex management of the switching activity in the nodes. By increasing the complexity of the switching equipment, it can be verified that link protection is able to achieve the same performance as path protection.

V Assumptions and fundamentals of the WDM-network availability model

The following A&R analysis is developed according to the following classical scheme: a) system identification and decomposition in functional elements; b)

characterization of each element in terms of its A&R parameters; c) development of an A&R mathematical model taking into account the relations among the elements within each subsystem and among the subsystems within the system; d) A&R evaluation of each subsystem and of the whole system.

Since this paper will provide a comparison of different end-to-end protection mechanisms, the system that we are going to study for each case of protection is the set of optical connections that may be involved by common protection actions. We call this set of connections a *protection group* (PG). We will see that, according to their various implementations, the protection mechanisms can create interdependency between connections that have the same source and destination ($M : N$ case) or even connections among different couples of nodes of a network (mesh shared-protection). We will assume that routing and wavelength assignment have already been solved for the working and protection lightpaths of all the connections of the PG under study. This means that a WDM channel has been reserved and is in use for every WDM link of the network crossed by a Working Lightpath (WL) of the PG. On the other hand, a WDM channel has been assigned for every WDM link of the network on which a Protection Lightpath (PL) of the PG will be routed in case of failure.

Each connection of the PG is a subsystem of our model. The functional elements should comprise all the transmission and switching equipment crossed by each lightpath. In this work we have, however, considered ideal WDM switching devices, i.e. perfectly reliable and free from any kind of failure (assumption not far from reality, according to Ref. [23]). This ideal-behaviour assumption extends also to any device providing switching of the optical signals of a connection from working to protection paths in case of failure. Thus only *WDM channels* have to be taken into account as functional blocks. A WDM channel is part of a WDM link, composed of the fibre cable installed between two adjacent nodes and equipped by a set of line devices (e.g. optical amplifiers). The A&R parameters of a WDM channel can be obtained by suitably combining those of the line devices plus those of other possible devices such as transponders, transmitters, receivers, WDM multi-demultiplexers, etc. Such parameters are commonly specified by technology vendors. The details of the reliability description of a WDM channel (see for example Ref. [24]) are not of interest in this paper and will be omitted. We shall only say that the model is based on the usual approximation of considering a constant rate of failure $z(t) = \eta$, corresponding to a negative exponential reliability function $R(t) = e^{-\eta t}$. According to such an approximation, the Mean Time To Failure (MTTF)




Figure 8 Total fibre number on NSFNET with different granularity of rerouted entities with shared link protection

of a WDM channel is independent of the components' age. Moreover, the WDM channels of a given optical connection are mutually failure-independent [25]–[27]. This assumption allows us to exploit the theory of Lee on the analysis of switched networks [28] for all the $M : N$ cases. The same cannot be applied, instead, to the mesh shared protection case, as explained later on.

WDM links can be realistically considered repairable systems: we thus assume the MTTF of a WDM channel to be equal to its *Mean Time Between Failures* (MTBF); thus: $MTBF = 1 / \eta$. The *Mean Time To Repair* (MTTR) of a WDM channel is also assumed to be constant in time. Eventually, for the purpose of this paper, we will assume each functional element of our system (i.e. each WDM channel assigned to any lightpath of the PG) characterized by a known average steady-state availability $A = MTBF / (MTBF + MTTR)$ or by a known MTTF (the mean value of the reliability distribution).

All the components included in our model have been characterized in terms of their intrinsic availability. Externally-provoked failures are not considered⁴⁾. In the examples reported in the following we will assume WDM channels assigned to PLs to have the same A&R parameters as those assigned to WLs. It should be considered that a common routing method is to route the WL on the first shortest path between source and destination and the PL on the second link-disjoint shortest path: the total A&R of the standby path can be even worse than that of the primary path, the former usually being longer than the latter. Finally, let us specify that in this work we are not considering for simplicity the presence of disjoint links belonging to the same shared risk link group (e.g. passing through the same conduit), nor protection or restoration errors.

VI Availability of WDM path-protection schemes

In this section, we provide algebraic equations to evaluate the availability of the single optical connections (subsystems) and of the entire PG (system). We will start from the simple dedicated 1 : 1 schemes, then we will increase first the number of working lightpaths in the PG (1 : N) and then the number of spare lightpaths ($M : N$). We will conclude with the mesh shared-cases for which we introduce a simple approximation that has been shown to provide very good results. All these schemes may be of practical interest in WDM network planning. It should be

noted however that due to the fundamental requirement of path-protection (see Section I), at least all the WLs of a PG are link-disjoint. The increase of the number of mutually link-disjointness constraints in the same PG makes the most complex schemes applicable only in extremely highly-connected network topologies.

The following notation will be used, also in the figures. Events, negated events and availability are identified by E , \bar{E} and A , respectively. These symbols always appear with a subscript, the first letter of which indicates what the symbol refers to: the whole PG system (s), a connection (k), a working (w) or a protection (p) lightpath, a working (λ) or a spare (π) WDM channel. Except for the whole PG, a second letter of the subscript identifies the particular element in the considered system: e.g. A_{w1} is the availability of working lightpath number 1. Each connection obviously corresponds to one and only one WL. Therefore a connection always has the same identifier of its WL. The same does not apply to PLs when they are shared.

The equations are obtained by a combinatorial method [29], enumerating all the favourable cases and summing their probabilities. The well-known formulas of the availability of parallel and series systems [30] are often applied. For instance, a WL w_i is a series of WDM channels. Thus its availability is the product of the availability of all the elements λ_j of the set Λ_{wi} of WDM channels assigned to it

$$A_{wi} = \prod_{\forall \lambda_j \in \Lambda_{wi}} A_{\lambda_j}$$

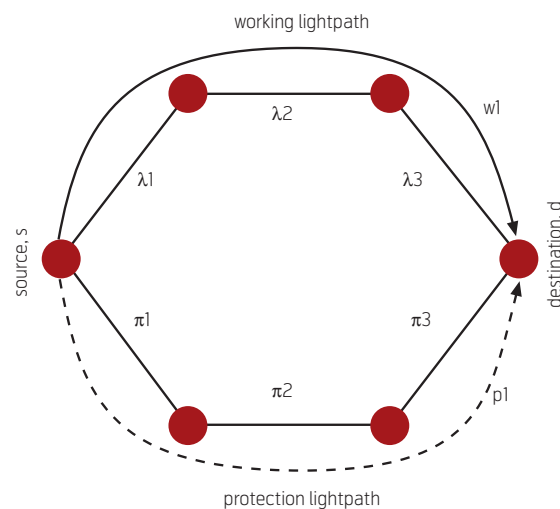


Figure 9 Protection group of 1 : 1 dedicated protection

⁴⁾ Statistically modelling external failure agents is generally difficult: often, intrinsic availability only appears in system specifications.

A 1 : 1 dedicated protection

In the 1 : 1 technique (Figure 9) the PG is simply composed of one connection (connection $k1$), which is coincident with the entire system and comprises a working ($w1$) and a link-disjoint protection lightpath ($p1$). The backup path, which is used when a failure occurs on the working lightpath, is in this case dedicated to one single connection.

The system availability is given by the union of two disjoint events: the WL is available (E_{w1}); the WL is not available ($\overline{E_{w1}}$), but the PL is available and can be used (E_{p1})

$$P\{E_s\} = P\{E_{k1}\} = P\{E_{w1} \cup (\overline{E_{w1}} \cap E_{p1})\}$$

The connection (and PG) availability is given by

$$A_s = A_{k1} = A_{w1} + A_{p1} - A_{w1}A_{p1} \quad (1)$$

B 1 : N protection

The PG is composed of N connections with the same source and destination, sharing a single PL (Figure 10).

We can similarly extend the 1 : 1 case to the general case of N connections (N WLs plus one PL, all mutually link-disjoint). The system availability is expressed by:

$$P\{E_s\} = P\left\{\bigcap_{j=1}^N E_{w_j} \cup \left(\overline{E_{w_h}} \cap E_{p1} \cap \bigcap_{j=1}^N E_{w(j \neq h)}\right)\right\}$$

$$A_s = (1 - NA_{p1}) \prod_{j=1}^N A_{w_j} + \sum_{h=1}^N \left[A_{p1} \prod_{j=1}^N A_{w(j \neq h)} \right]$$

C M : 1 protection

In this scheme the PG comprises one single connection $k1$ (Figure 11). Its WL $w1$ is protected by multiple link-disjoint PLs $p1 \dots pM$. p_i is used when $w1$ and all the PLs from $p1$ to $p(i-1)$ are unavailable. Up to M failures can be recovered.




Figure 11 Protection group of $M : 1$ protection

Eq. (2) expresses system availability in the general $M : 1$ case.

$$P\{E_s\} = P\left\{E_{w1} \cup \bigcup_{h=1}^M \left[\overline{E_{w1}} \cap E_{p_h} \bigcap_{j=1}^{h-1} \overline{E_{p_j}}\right]\right\} \quad (2)$$

D M : N protection

The most general path-protection configuration involving connections between the same end nodes is obtained by combining 1 : N and $M : 1$ in the $M : N$ case. Unfortunately, a general equation for the $M : N$ availability cannot be written in a closed form, since its algebraic form changes with M and N .

E Mesh shared-protection

Let us start with the sample PG of Figure 12, composed of only two connections: the $2 \times (1 : 1)$ case. This simple layout will help understand both the availability evaluation mechanism and the approximation that we are going to introduce to make this evaluation feasible under more complex scenarios.

The two WLs $w1$ and $w2$ are protected by two PLs ($p1$ and $p2$) that share the WDM channel $\pi 5$. The system availability is the probability that both connections are routed successfully and is obtained in




Figure 10 Protection group of $1 : N$ protection




Figure 12 PG of the $2 \times (1 : 1)$ mesh shared-protection

Eq. (3) and Eq. (4) by the union of three disjoint events.

$$P\{E_s\} = P\{(E_{w1} \cap E_{w2}) \cup E_a \cup E_b\} \quad (3)$$

where, keeping in mind that $E_{p2} = E_{\pi2} \cap E_{\pi5} \cap E_{\pi4}$ ($A_{p2} = A_{\pi1} \cdot A_{\pi5} \cdot A_{\pi3}$) and $E_{p1} = E_{\pi1} \cap E_{\pi5} \cap E_{\pi3}$ ($A_{p1} = A_{\pi2} \cdot A_{\pi5} \cdot A_{\pi4}$), we set

$$E_a = E_{w1} \cap \overline{E_{w2}} \cap E_{p2}$$

$$E_b = \overline{E_{w1}} \cap E_{w2} \cap E_{p2}$$

Thus

$$A_s = A_{w1}A_{w2} + A_{w1}(1 - A_{w2})A_{p2} - (1 - A_{w1})A_{w2}A_{p1} \quad (4)$$

To evaluate the availability of a single connection, we have to distinguish different double-link failure scenarios. For instance, even if lightpath $w2$ and WDM channel $\pi2$ fail, connection $k1$ can be routed successfully. So the first subsystem (protected connection $k1$) is characterized by the following availability:

$$P\{E_{k1}\} = P\{E_{w1} \cup E_a \cup E_b \cap E_\gamma\}$$

where

$$E_a = \overline{E_{w1}} \cap E_{p1} \cap E_{w2}$$

$$E_b = \overline{E_{w1}} \cap E_{p1} \cap \overline{E_{w2}} \cap \overline{E_{\pi2}}$$

$$E_\gamma = \overline{E_{w1}} \cap E_{p1} \cap \overline{E_{w2}} \cap E_{\pi2} \cap \overline{E_{\pi4}}$$

Thus

$$A_{k1} = A_{w1} + (1 - A_{w1})A_{p1}A_{w2} + (1 - A_{w1})A_{p1}(1 - A_{w2})(1 - A_{\pi2}) + (1 - A_{w1})A_{p1}(1 - A_{w2})A_{\pi2}(1 - A_{\pi4}) \quad (5)$$

The need to consider all the possible multiple-failure combinations makes the problem intractable for larger PGs. We introduce an approximation by neglecting multiple failure scenarios. This is equivalent to considering only terms in which $(1 - A)$ appears at the first order, neglecting higher-order terms. It can be proven that the second order terms are always absent even without the approximation, except when the spare path is totally shared (but this case coincides with the $1 : N$ case). In the next section we will show by numerical examples that the approximated formula converges to the real availability values for highly available components (rare-event approximation). The approximated availability of connection $k1$ is calculated in Eq. (6) and Eq. (7).




Figure 13 PG of the $m \times (1:1)$ mesh shared-protection

$$P\{E_{k1}\} \approx P\{E_{w1} \cup [\overline{E_{w1}} \cap E_{p1} \cap E_{w2}]\} \quad (6)$$

$$A_{k1} \approx A_{w1} + (1 - A_{w1})A_{p1}A_{w2} \quad (7)$$

We extend now our analysis to a PG comprising the m protected working connections whose protection lightpaths share some optical channels ($m \times (1 : 1)$ scheme, Figure 13).

The system availability formulas Eq. (8) and Eq. (9) are obtained neglecting multiple-failure cases.

$$P\{E_s\} \approx P\left\{\bigcap_{j=1}^m E_{w_j} \bigcup_{h=1}^m \left[\overline{E_{wh}} \cap E_{ph} \bigcap_{k=1}^m E_{w(k \neq h)}\right]\right\} \quad (8)$$

$$A_s \approx \prod_{j=1}^m A_{w_j} + \sum_{h=1}^m (1 - A_{wh})A_{ph} \prod_{k=1}^m A_{w(k \neq h)} \quad (9)$$

VII Availability numerical examples

In this section we analyze the protection techniques through numerical examples. We assume that each working lightpath w_i is composed of a single hop (channel) with availability $A_{wi} = 1 - U$. Each protection lightpath p_x has the length $L_{px} = 3$, being the avail-

Protection technique	Unavailability $U_{k1} = 1 - A_{k1}$
2 : 1	8.99825×10^{-12}
3 : 1	2.77556×10^{-15}
4 : 1	1.11022×10^{-16}

TABLE III Connection unavailability in $M : 1$ protection

ability of each of its 3 WDM channels $A_{wj} = 1 - U$. The total spare path availability is $A_{px} = (1 - U)^3$. The reported numerical values refer to the availability of a single protected connection.

A $M : N$ protection

For all the results in this section: $U = 10^{-4}$. The connection unavailability values of $1 : N$ are plotted in Figure 14 as a function of N . The plot shows that unavailability grows for increasing values of N with a linear slope of about 10^{-8} per $\Delta N = 1$.

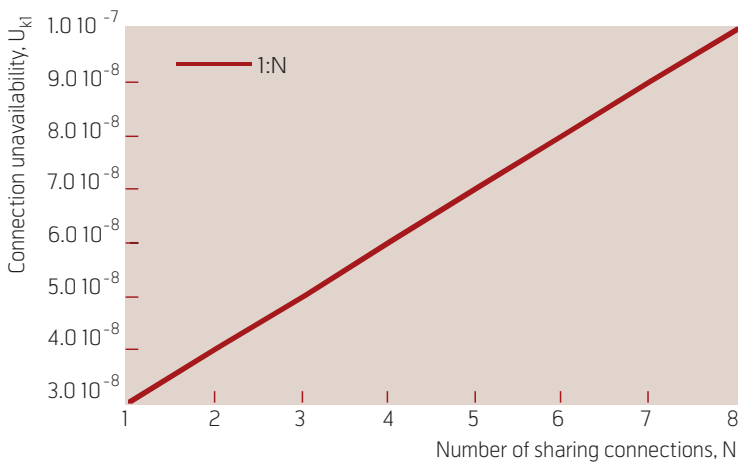


Figure 14 Connection unavailability of protection scheme 1:N




Figure 15 Connection unavailability of $m \times (1:1)$ mesh shared-protection

U	U_{k1} exact	U_{k1} approx	% error
10^{-1}	3.30049×10^{-2}	3.439×10^{-2}	4.2
10^{-4}	3.9992×10^{-8}	3.9994×10^{-8}	5×10^{-3}

Table IV Unavailability of connection $k1$ in the pg of Figure 12

Table III refers to $M : 1$ protection: unavailability decrease of orders of magnitude by adding protection lightpaths, since a higher number of connection failures can be recovered.

We can conclude that availability in $M : N$ protection is primarily determined by M , corresponding to the number of simultaneously recoverable failures. The number N of working paths that share the backup paths has instead a marginal effect compared to M . For example, from Table III and from Figure 14 we see that 2 : 1 unavailability is $\cong 9 \times 10^{-12}$ with $M / N = 2$, while in the 3 : 4 case unavailability is $\cong 1.2 \times 10^{-14}$ with $M / N = 0.75$. Actually, 3 : 4 provides protection against any three link failures, achieving a higher level of availability.

B Mesh shared-protection

In Sec. VI-E we have obtained the single connection availability using either the exact Eq. (5) or the approximated Eq. (7). Table IV shows the accuracy of our approximations considering the network of Figure 12. In the first row ($U = 0.1$) unavailability is selected on purpose with values unrealistic for optical networks. We can observe that even in these extreme conditions the percentage of error of the approximated result is quite small, while it is almost negligible with realistic unavailability ($U = 10^{-4}$).

The values in Figure 15 refer to the general $m \times (1 : 1)$ protection. The graph displays the unavailability of a generic connection k_i . We consider the two cases when its spare path has length $L_{pi} = 5$ and $L_{pi} = 7$, respectively. The number m_{ki} of connections sharing backup channels with p_1 varies from 2 to 6. As already observed for 1 : N (single fault recovery), Figure 15 shows a linear increase of connection unavailability, associated to the increase of m_{ki} . In terms of availability performance, increasing the length L_{pi} of the protection lightpath by x hops is equivalent to increasing the number of sharing connections m_{ki} by x .

C Final comparison

In the previous sections we have separately studied the different protection approaches. Now we can jointly compare the performances of the various approaches (Figure 16).

As already explained, unavailability improves considerably when the protection scheme is able to recover multiple failures: this behaviour is apparent in Figure 16. Mesh shared-protection and 1 : N, recovering single failures, give similar unavailability results.

VIII Conclusions

In this paper we have dealt with four protection strategies that are candidates to be the best choice for the next-generation WDM network: path and link protection in both the dedicated and the shared case. After outlining the schemes and their technological requirements, we have described the mathematical formulations that model these protection techniques. Using a case-study network, we have compared the resource requirements of each scheme by exploiting ILP. The shared protections (path or link) provide very good results, as they require about 50 % of added capacity with respect to the unprotected case. The link protection in dedicated configuration needs a huge backup-capacity increase (330 %), while dedicated path protection achieves more efficient results (150 %). Finally, we have shown that shared link protection needs less capacity when failed channels are rerouted individually.

In the second part of the paper, we provided formulas to evaluate connection availability under several protection schemes. In treating shared protection we have introduced an approximation that allows us to analyze complex topologies. The formulae have been used in a comparative analysis of the different resilience mechanisms, leading us to the following interesting general finding: the number of simultaneous failures a protection scheme can recover sets the order of magnitude to the availability of its protected connections.

References

- 1 Bigo, S. Transmission of 256 wavelength-division and polarization-division-multiplexed channels at 42.7 Gb/s (10.2 Tb/s capacity) over 3x100km of TeraLight fibre. In: *Proceedings OFC'02*, 2, 2002, FC5-1–FC5-3.
- 2 ITU. *Network node interface (NNI) for the optical transport network (OTN)*. ITU-T International Communication Union, Feb 2001. (G.709/Y.1331, Amendment 1)
- 3 Goderis, D et al. *Service level specification semantics and parameters*. Internet Draft, draft-tequila-sls-01.txt, Tech. Rep., June 2001.




Figure 16 Connection-unavailability comparison of various protection schemes

- 4 Fawaz, W, Daheb, B, Audouin, O, Du-Pond, M, Pujolle, G. Service level agreement and provisioning in optical networks. *IEEE Communications Magazine*, 42 (1), 34–35, 2004.
- 5 Caenegem, B V, Parys, W V, Turck, F D, Deemester, P M. Dimensioning of survivable WDM networks. *IEEE Journal on Selected Areas in Communications*, 16 (7), 1146–1157, Sept 1998.
- 6 Ramamurthy, S, Mukherjee, B. Survivable WDM Mesh Networks, part. I - Protection. In: *Proceedings, IEEE INFOCOM'99*, 2, 1999, 744–751.
- 7 Stern, T E, Bala, K. *Multiwavelength Optical Networks: A Layered Approach*. Reading, MA, Addison Wesley, 1999.
- 8 Strand, J, Chiu, A, Tkach, R. Issues for routing in the optical layer. *IEEE Communications Magazine*, 39 (2), 2001.
- 9 Doverspike, R, Yates, J. Challenges for mpls in optical network restoration. *IEEE Communications Magazine*, 39 (2), 2001.
- 10 Miyao, Y, Saito, H. Optimal design and evaluation of survivable WDM transport networks. *IEEE Journal on Selected Areas in Communications*, 16 (7), 1190–1198, Sept 1999.
- 11 Baroni, S, Bayvel, P, Gibbens, R J, Korotky, S K. Analysis and design of resilient multifibre wave-

- length-routed optical transport networks. *Journal of Lightwave Technology*, 17 (5), 743–758, 1999.
- 12 Anand, V, Qiao, C. Static versus dynamic establishment paths in WDM networks. Part i. In: *Proceedings ICC '00*, 2000, 198–204.
 - 13 Zang, H, Ou, C, Mukherjee, B. Path-protection routing and wavelength assignment (RWA) in WDM mesh networks under duct-layer constraints. *IEEE/ACM Transactions on Networking*, 11 (2), 248–258, 2003.
 - 14 Xiong, Y, Xu, D, Qiao, C. Achieving fast and bandwidth-efficient shared-path protection. *Journal of Lightwave Technology*, 21 (2), 2003.
 - 15 Lumetta, S, Medard, M, Tseng, Y. Capacity versus robustness: A tradeoff for link restoration in mesh networks. *Journal of Lightwave Technology*, 18 (12), 2000.
 - 16 Chamtlac, I, Ganz, A, Karmi, G. Lightpath communications: an approach to high-bandwidth optical WAN's. *IEEE/ACM Transactions on Networking*, 40 (7), 1172–1182, 1992.
 - 17 Ramaswami, R, Sivarajan, K N. Routing and wavelength assignment in all-optical networks. *IEEE/ACM Transactions on Networking*, 3 (5), 489–500, 1995.
 - 18 Banerjee, D, Mukherjee, B. Wavelength-routed optical networks: linear formulation, resource budgeting tradeoffs and a reconfiguration study. *IEEE/ACM Transactions on Networking*, 8 (5), 598–607, 2000.
 - 19 Tornatore, M, Maier, G, Pattavina, A. WDM Network Optimization by ILP Based on Source Formulation. *Proceedings, IEEE INFOCOM '01*, June 2002.
 - 20 Concaro, A, Maier, G, Martinelli, M, Pattavina, A, Tornatore, M. QoS Provision in Optical Networks by Shared Protection: An Exact Approach. In: *Quality of service in multiservice IP Networks*, ser. Lectures Notes on Computer Sciences, 2601. Springer, Feb. 2003, 419–432.
 - 21 Bienstock, D, Gonluk, O. Computational experience with a difficult mixed-integer multi-commodity flow problem. *Mathematical Programming*, 68 (32), 213–237, 1995.
 - 22 Tornatore, M. *Modelli matematici di programmazione lineare a numeri interi per l'ottimizzazione delle reti ottiche wdm*. Politecnico di Milano, 2001. (Master's thesis)
 - 23 Jereb, L, Jakab, T, Unghvary, F. Availability analysis of multi-layer optical networks. *Optical Networks Magazine*, March/April 2002.
 - 24 Tornatore, M, Maier, G, Pattavina, A, Villa, M, Righetti, A, Clemente, R, Martinelli, M. Availability optimization of static path-protected wdm networks. In: *Proceedings, OFC 2003*, Mar. 2003.
 - 25 Antonopoulos, A, O'Reilly, J J, Lane, P. A framework for the availability assessment of sdh transport networks. *Proceedings, Second IEEE Symposium on Computers and Communications*, 666–670, July 1997.
 - 26 Inkret, R, Lackovic, M, Mikac, B. WDM network availability performance analysis for the COST 266 case study topologies. In: *Proceedings, Optical Network Design & Modelling*, Feb. 2003.
 - 27 Clouqueur, M, Grover, W. Availability analysis of span-restorable mesh networks. *IEEE Journal on Selected Areas in Communications*, 20 (4), May 2002.
 - 28 Lee, C Y. Analysis of switching networks. *Bell Systems Technical Journal*, 34 (6), 1287– 1315, 1955.
 - 29 Mood, A M, Boes, D C, Graybill, F A. *Introduction to the Theory of Statistics*, 3rd ed. New York, McGraw-Hill, 1974.
 - 30 Lewis, E E. *Introduction to Reliability Engineering*. New York, John Wiley, 1987.

Guido Maier received his Laurea degree in Electronic Engineering at Politecnico di Milano (Italy) in 1995 and his PhD degree in Telecommunication Engineering at the same university in 2000. He is researcher at CoreCom, where he has the position of Head of the Optical Networking Laboratory. His main areas of interest are optical network modelling, design and optimization, ASON/GMPLS architecture and WDM switching systems. He has authored more than 30 papers in the area of Optical Networks published in international journals and conference proceedings. He is currently involved in industrial and European research projects.

email: maier@corecom.it

Massimo Tornatore received his Laurea degree in Telecommunications Engineering from Politecnico di Milano in 2001. He is currently a PhD candidate in the Electronics and Information Department of Politecnico di Milano under the supervision of Prof. Pattavina. His research interests include Design, Protection Strategies, Traffic Grooming in Optical WDM Networks and Group Communication Security.

email: tornatore@elet.polimi.it

*Achille Pattavina received his DrEng degree in Electronic Engineering from La Sapienza University of Rome (Italy) in 1977. He was with the same University until 1991 when he moved to Politecnico di Milano, Milan (Italy), where he is now Full Professor. He has authored more than 100 papers in the area of Communications Networks published in leading international journals and conference proceedings. He has been author of the book *Switching Theory, Architectures and Performance in Broadband ATM Networks* (John Wiley & Sons). He has been Editor of *Switching Architecture Performance of the IEEE Transactions on Communications* since 1994 and Editor-in-Chief of the *European Transactions on Telecommunications* since 2001. He is a Senior Member of the IEEE Communications Society. His current main research interests are in the area of optical networks and switching theory.*

email: pattavina@elet.polimi.it

Why bother with optical packets? An evaluation of the viability of optical packet/burst switching

EVI ZOUGANELI, RAGNAR Ø. ANDREASSEN, BONING FENG, ASTRID SOLEM, NORVALD STOL, HEIDI KJØNSBERG, AASMUND SUDBØ, BJARNE E. HELVIK AND ROLF B. HAUGEN



Evi Zouganeli is Senior Research Scientist at Telenor R&D



Ragnar Ø. Andreassen is Research Scientist at Telenor R&D



Boning Feng is Associate Professor at Oslo University College (OUC)



Astrid Solem is Research Scientist at Telenor R&D

Optical packet/burst switching (OPS/OBS) combines the high throughput and scalability of optical technologies with the granularity and flexibility of packet switching. The question is whether the gains OPS/OBS brings outweigh the extra complexity and justify development costs. In this paper we quantify the gains that can be obtained by OPS/OBS using network simulations combined with a theoretical analysis. We also carry out a direct techno-economic comparison between OBS, IP/MPLS and a circuit switched optical network (OCS) in order to reveal the viability of optical packet/burst switching.

The advantages of packet switching are by and large indisputable: packet based services are increasingly dominating in the network, the internet protocol (IP) has prevailed, and there is wide consensus that circuits do not fit the bill for datacom. But given the inherent poor processing capabilities of optics and the maturity of electrical IP routers, why bother with optical packets? Optical packet switching (OPS) is not about processing packets optically. An optical packet or optical burst comprises a series of IP packets, typically 10–100, i.e. it is a sort of mega-packet. The characteristic of an optical packet switched or optical burst switched (OBS) network is that contrary to electrical (IP) packet switched networks the optical packets in OPS/OBS remain in the optical domain throughout the network, all the way between the ingress and egress point. The label, on the other hand, is typically detected at each node and processed electronically. This way, decisions about packet forwarding are made for a bunch of packets rather than on a packet-to-packet basis, hence achieving considerably better efficiency and enabling higher throughput and overall scalability compared with IP/MPLS. This is further enhanced by the fact that since the packets remain in the optical domain the node is transparent to bit-rate and format and therefore scales transparently with increasing interface bit-rate. Optical packet switching combines the advantages of two worlds: it retains the high switching efficiency high throughput and scalability of optical technologies, and combines these with the granularity and flexibility of packet switching. The price to pay is higher complexity – in any case when compared with optical circuit switching. So the question is whether the gains OPS/OBS brings outweigh the extra complexity and justify development costs.

Considerable research effort has been invested in optical packet/burst switched networks in recent years, yet it has remained by and large unclear whether there is a realistic case for a paradigm shift to OPS/OBS networks. One main concern has earlier been whether such solutions are technologically fea-

sible in the first place – the switching speed and synchronisation requirements some solutions impose are not always easy to meet. For this reason OPS has often been regarded as an interesting topic of academic value, however, of limited relevance for real applications in the foreseeable future. The switching of a series of packets – a burst – instead of individual IP packets, leads to relaxed requirements in terms of both switching speed and synchronisation and is a more feasible version. However, the technology that actually put OPS/OBS back on the table as a possible candidate for a real implementation, is relatively new: namely the widely tunable laser. This is a laser module whose output wavelength can be tuned continuously and reproducibly over a wide spectrum (30 nm) using a software control and it has become commercially available in the past five years. The advent of widely tunable lasers has enabled an OBS node architecture based on Array Waveguide Gratings (AWG) and all-optical Tunable Wavelength Converters (TWC) that can be realised today and that has been found to be both technologically feasible and relatively low-cost. This architecture has been employed within the European Research Project IST STOLAS and studied both theoretically and experimentally [1]. In STOLAS the OBS network is optically label switched with the label and the payload co-propagating. As a result, the optical burst or packet resembles a traditional packet the way we know it from IP so that well-established IP/MPLS processes can be applied here with little – if any – modification.

In this paper we will use the node architecture employed in STOLAS to evaluate the feasibility and viability of OPS/OBS. The main question we want to shed light onto is whether there is a case for a paradigm shift to OPS/OBS – and if so, when and how. To this end we firstly attempt to quantify the gains that can be obtained by OBS using a theoretical study combined with network simulations. We then proceed to a direct techno-economic comparison between the capital investment required to build an OBS network, contra the investment required in the



Norvald Stol is Associate Professor at the Norwegian University of Science and Technology

case of an IP/MPLS network and an optical circuit switched (OCS) network of similar performance.

The STOLAS concept

The STOLAS network is an optically labeled burst switched network where the optical label is encoded orthogonally to the payload. Hence the label and payload signals can be encoded in parallel independently of each other.

This is achieved e.g. by modulating the phase of the optical carrier (DPSK) or by modulating the frequency of the optical carrier itself (FSK) at low speed (155 Mbit/s) in order to encode the label, whereas the payload is intensity modulated at a higher speed (10 Gbit/s) [1] [2]. IP packets are assembled in optical bursts of 10–100 packets at the edge of this network and are then forwarded to the destination node according to QoS and traffic engineering rules, as schematically shown in

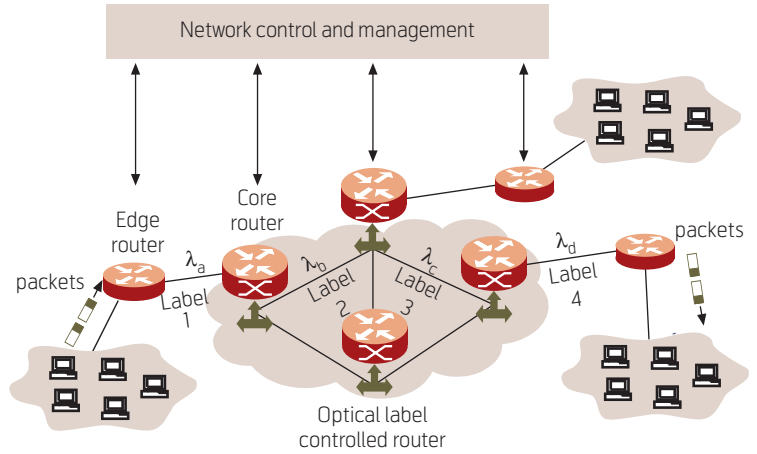


Figure 1 An optically burst switched network uses optical label switching to forward optical bursts of packets to their destination node

Figure 1. The signal remains in the optical domain throughout its journey from the ingress to the egress node and the wavelength in each link is a free resource for the bursts to contend for, subject to class of service. The generic STOLAS node is shown in Figure 2. It comprises a label reading and processing part, an all-optical label swapping part, and an optical



Heidi Kjønsgberg is Research Scientist at the Norwegian Computing Center



Aasmund Sudbø is Professor at the University of Oslo and Unik

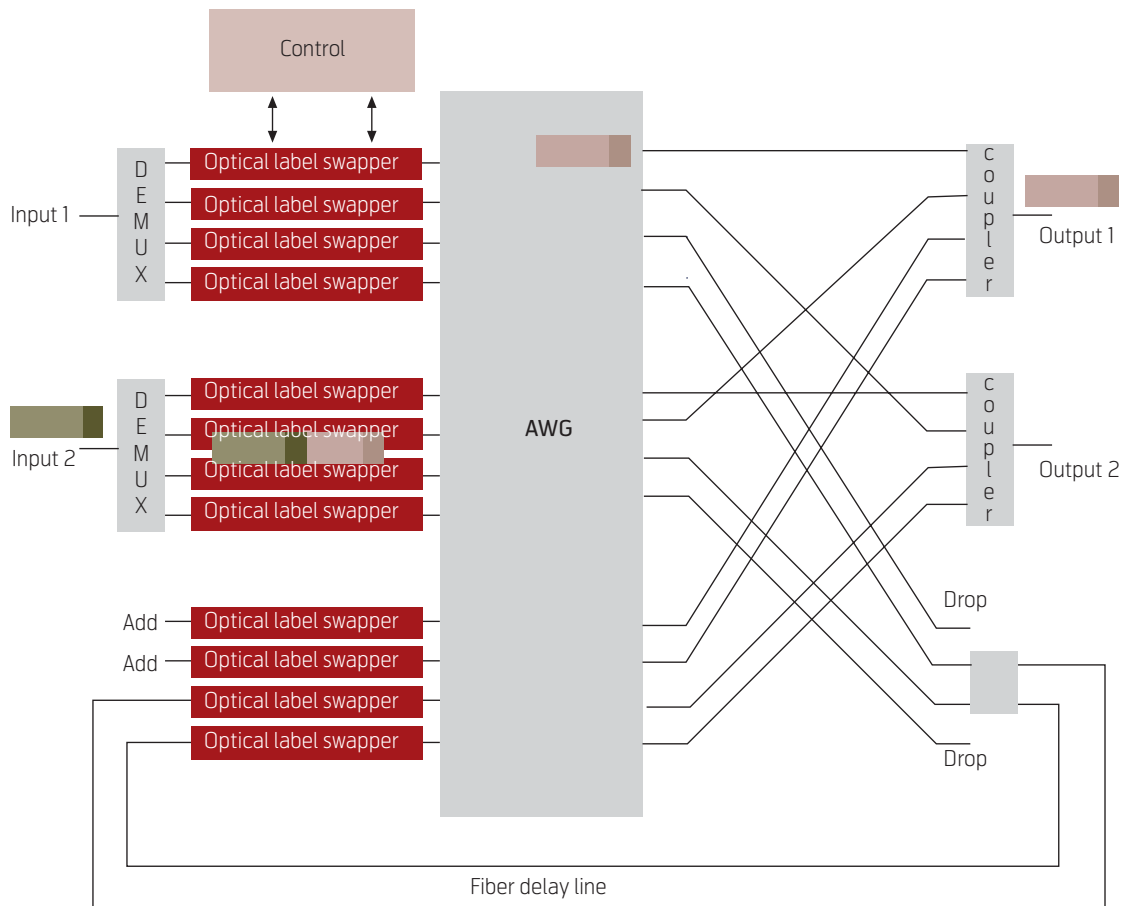


Figure 2 Schematic diagram of the generic STOLAS node comprising wavelength division de-multiplexers, Optical Label Swappers, an Array Waveguide Grating and multiplexers to the fibre outputs



Bjarne E. Helvik is Professor at the Norwegian University of Science and Technology



Rolf B. Haugen is in charge of coordinating external activities at Telenor R&D

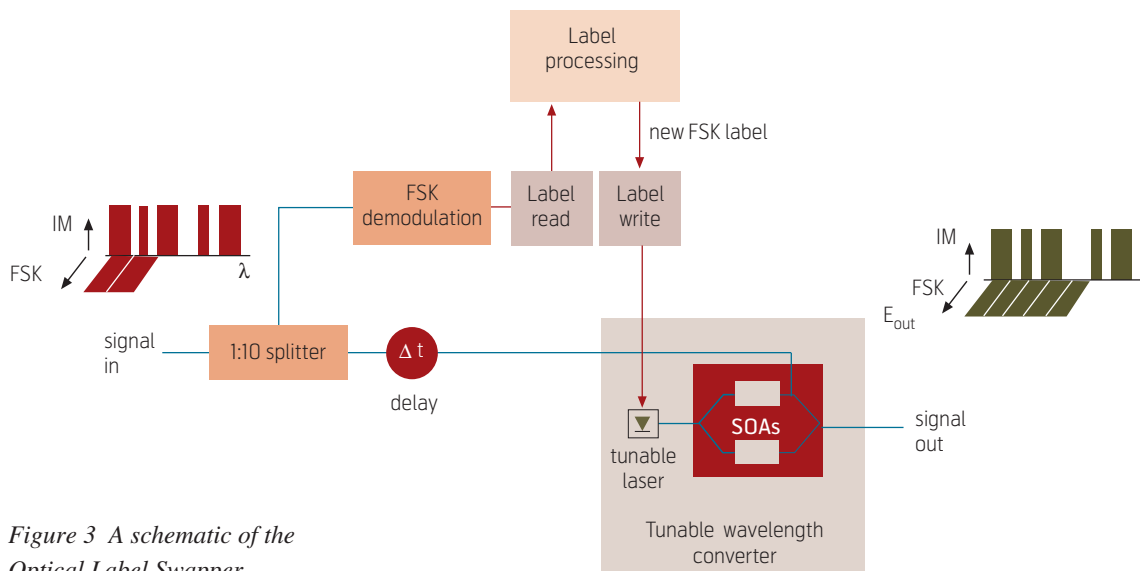


Figure 3 A schematic of the Optical Label Swapper

cross-connecting/forwarding part. The way this node attains forwarding of bursts to the right output is as follows. All fibre inputs – here depicted at the left hand side of the node – use the same optical carrier wavelengths. These are de-multiplexed via a standard passive optical WDM de-multiplexer and subsequently each of these signals goes through a processing card that comprises an Optical Label Swapper (OLS). Here, a small part of the signal is tapped out such that the label is read and processed electronically while the main part of the signal is delayed and remains in the optical domain. Label swapping takes place all-optically using a tunable laser and a label swapper arrangement that makes use of a non-linear effect that is called cross-phase modulation and which here takes place in an integrated Mach-Zehnder interferometer. At the output of the OLS, the burst is carried by a new wavelength and a new optical label has been encoded on the FSK signal. The output the burst is subsequently forwarded to, depends on this new carrier wavelength. Indeed, the Array Waveguide Grating is a passive component that acts as a “hardwired” wavelength dependent switch, where the hardwired connections depend on the input the signal enters the AWG from and the signal’s wavelength. Hence, the fibre output the burst is forwarded to depends on the wavelength that is assigned to it through the preceding optical label swapper – which is then controlled by standard (G)MPLS processes. It should therefore be noted that when this node architecture is implemented, *the wavelength is not available as a label*. The wavelength is in this configuration rather used to carry out the cross-connection function in the node instead.

This node architecture makes use of the fast wavelength switching properties of tunable lasers (wavelength switching at 50 ns was realized in STOLAS)

and avoids large and complex switching matrices in order to realize cross-connection functionality. AWGs are inherently cheap components and currently broadly used as passive WDM multiplexers/de-multiplexers in existing networks. One of the prime cost determining factors of the node is the number of tunable wavelength converters comprising the optical label swapper, which needs to be minimized. At the same time, internal blocking at the node also needs to be minimized. The trade-offs involved in node design are discussed in [3].

Estimated OBS performance gains compared with OCS

Analytical study

The increased flexibility of OPS/OBS in relation to OCS should lead to better network traffic transport efficiency. Two main contributions to such increased efficiency will come from increased statistical resource sharing (i.e. multiplexing gain) and finer upgrade granularity. We will briefly consider both these effects in the following [4].

In optical circuit switching, packets share the resource that corresponds to a Label Switched Path (LSP). An LSP in this case may comprise one or more optical circuits with the same source-destination pair (optionally also class of service or other differentiation factors). In optical packet or burst switching on the other hand, all packets in all LSPs share the link resource on a given internal network link. As the link resources are generally much larger than LSP resources, statistical resource sharing in the OPS/OBS case will be better than for OCS. For a given quality requirement (i.e. in terms of packet loss or network overload), an OPS/OBS based network can




Figure 4 Link capacity increase factor (κ) for OCS versus OPS as a function of overall link capacity C/h or C/m for m/h -values of 0.05, 0.2 and 0.4. The left panel shows results for C/h variations while the right panel shows results for C/m variations. In all calculation the quality requirement was $P(\text{overload}) = 10^{-6}$ and the network size was 81 nodes

therefore carry the same traffic with less dimensioned capacity compared to an OCS-based network.

As end-to-end traffic increases, more wavelength channels need to be added to the network. In an OCS network, the capacity of each LSP will need to be upgraded incrementally by a whole wavelength channel to accommodate any traffic increase within this LSP. In the case of OBS, on the other hand, incremental capacity increase can be made at the internal link level. The incremental increase relative to LSP capacity will then be much smaller, as an internal link carries many LSPs. On average, this will give better utilisation in the network, particularly in the case where end-to-end LSP demand is in the range of one or very few optical channel capacities.

In the following we focus on a simple study of these effects, and attempt to quantify the OPS/OBS gains that may be expected under various circumstances. In brief, in the calculations we assume a simplified symmetric case. We consider a symmetric $\mu \times \mu$ torus¹⁾ network where all sub-links have the same capacity serving the same number of traffic streams. A given number of continuously active identical and independent variable bit-streams are considered, where individual streams have known peak rates h and mean rates m . As a simple approximation to a variable bitrate stream, we consider an on/off source model. We calculate the probability for a single stream to experience overload, and pose the requirement that this probability must be less than 10^{-6} . We then calculate the resources required in order to meet this requirement in the case of an OPS and an OCS network for different types of sources.

In Figure 4 we show results for sources of different sizes and a range of m/h values. The left panel shows

results when we consider sources of known peak-rates h relative to network link rate C , while the right panel shows the corresponding results if we assume we know the relative mean rates of sources. In both cases, we assume the network, source properties and quality requirement as given, and calculate on this basis the permissible loads and dimensioning required for the two different network technologies. The gain factor κ is given as the ratio between the OCS required resources and the OPS/OBS required resources.

In order to relate the figures to everyday life, we may for the left figure consider that traffic originate in xDSL-based networks of e.g. 10 Mbit/s maximum peak rate. A C/h ratio of 10^5 would then correspond to link transmission rates of 1 Tbit/s in the network. This would typically be a hundred 10 Gbit/s channels, i.e. one or two fibres. Our model predicts around 20 % extra capacity for circuit switching in this case. If there are no efficient constraints in peak rates of originating traffic sources, the right figure is more relevant. This may be the case if traffic originates in corporate networks, which may well have transmission capacities comparable to transport networks. If we for example then consider a backup application, generating a mean rate of 10 Mbit/s, a C/m ratio corresponds to the same network capacities as in the previous case, but results in an OCS increase factor of two (i.e. a 100 % increase) for the more bursty sources. The gains decrease with decreasing network size (number of nodes). For example, for a network of 25 nodes the gains are reduced by an estimated 12 % in relation to the values presented in Figure 4.

Granularity gains become significant in the case of networks with a relatively low number of channels per LSP. An example calculation is shown in Figure

¹⁾ A torus network is a regular grid network of $\mu \times \mu$ nodes where node (i, μ) is connected to $(i, 1)$ and node $(1, i)$ is connected to (μ, i) , for all $i = 1 \dots \mu$.




Figure 5 Granularity gains. End-to-end dimensioned sub-link capacities as a function of # of sources per link (left) and corresponding OCS versus OPS/OBS increase factor κ (right) for the source parameters $h = 50$ Mbit/s and $m/h = 0.05$ for a medium sized network of 25 nodes

5 where in effect the end-to-end traffic is shown to increase with smaller increments in the case of OBS. The OBS gain factor κ is depicted on the right hand side of the figure where high values of κ may be obtained at certain upgrade instances. Note that this is also an indication of the flexibility that can be obtained in case of failure or general cases of traffic engineering. Indeed, the OBS network can easier accommodate unforeseen traffic changes since streams can be broken down at sub-wavelength level and accommodated in different routes through the network.

On the whole this simplified theoretical analysis has pointed out that multiplexing and granularity gains with OBS are highly dependent upon traffic characteristics. Gains of significance are obtained first for relatively large sources (high mean value m above 2 Mbit/s) of high burstiness $h/m > 5$ and become more significant for larger networks (number of nodes).

Network simulations

Based on the insights that were gained by the analytical study we have carried out detailed network simulations to quantify the gains that may be obtained by OBS. The analytical study pointed out the significance of the assumptions with regard to traffic characteristics. Therefore a lot of effort was invested in the description of realistic traffic scenarios for the near to medium future and the subsequent derivation of realistic traffic characteristics for our network simulation studies [5][6]. The method for deriving a realistic input traffic that corresponds to around year 2009 is described in Appendix 1. In the same appendix we describe the simulator details whereas in the following we focus on an outline of the main assumptions and the obtained results.

In brief, we consider a mix of private and business clients and traffic is calculated as originating from a mix of nine application components for each of these

market segments. Each of these application components is modelled using a parameterisation of three generic source types, namely voice, video and data. Finally, we use estimated service penetration values for each of these applications in each market segment as well as an evolution of these with time. The resulting traffic is believed to be a good approximation of realistic traffic in a typical European region by year 2009. The network we use in our simulations is a 4 x 4 Manhattan Street network, or regular grid, as shown in Figure 6.

At the ingress node of the OBS network IP packets are assembled into optical bursts according to destination node and class of service (if applicable). We assume synchronous slotted operation, i.e. bursts of fixed length are created and released at given intervals that are identical for the whole network. Further, a thorough investigation using our network simulations [7] has revealed that best performance is obtained for relatively large bursts of about 100 μ s, or around 80 IP packets per optical burst, and this burst length is subsequently used in the simulations.

Limited buffers of about 17 Mbytes are assumed at the ingress. We set a requirement for the total packet

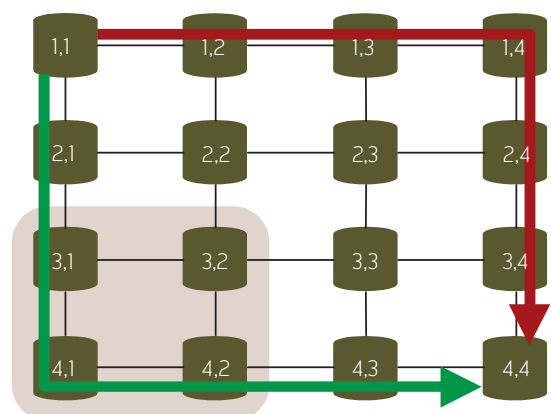


Figure 6 The 4 x 4 node grid network

From node \ To node	1,1	1,2	1,3	1,4	2,1	2,2	2,3	2,4	3,1	3,2	3,3	3,4	4,1	4,2	4,3	4,4
1,1	-	30.3	19.7	61.2	19.7	15.5	15.5	30.3	30.3	15.5	15.5	19.7	61.2	19.7	30.3	61.2
1,2	31.9	-	10.2	31.9	10.2	8.0	8.0	15.6	15.6	8.0	8.0	10.2	31.9	10.2	15.6	31.9
1,3	21.2	10.5	-	21.2	7.1	5.3	5.3	10.5	10.5	5.3	5.3	7.1	21.2	7.1	10.5	21.2
1,4	61.2	30.3	19.7	-	19.7	15.5	15.5	30.3	30.3	15.5	15.5	19.7	61.2	19.7	30.3	61.2
2,1	21.2	10.5	7.1	21.2	-	5.3	5.3	10.5	10.5	5.3	5.3	7.1	21.2	7.1	10.5	21.2
2,2	16.6	8.5	5.3	16.6	5.3	-	4.0	8.5	8.5	4.0	4.0	5.3	16.6	5.3	8.5	16.6
2,3	16.6	8.5	5.3	16.6	5.3	4.0	-	8.5	8.5	4.0	4.0	5.3	16.6	5.3	8.5	16.6
2,4	31.9	15.6	10.2	31.9	10.2	8.0	8.0	-	15.6	8.0	8.0	10.2	31.9	10.2	15.6	31.9
3,1	31.9	15.6	10.2	31.9	10.2	8.0	8.0	15.6	-	8.0	8.0	10.2	31.9	10.2	15.6	31.9
3,2	16.6	8.5	5.3	16.6	5.3	4.0	4.0	8.5	8.5	-	4.0	5.3	16.6	5.3	8.5	16.6
3,3	16.6	8.5	5.3	16.6	5.3	4.0	4.0	8.5	8.5	4.0	-	5.3	16.6	5.3	8.5	16.6
3,4	21.2	10.5	7.1	21.2	7.1	5.3	5.3	10.5	10.5	5.3	5.3	-	21.2	7.1	10.5	21.2
4,1	61.2	30.3	19.7	61.2	19.7	15.5	15.5	30.3	30.3	15.5	15.5	19.7	-	19.7	30.3	61.2
4,2	21.2	10.5	7.1	21.2	7.1	5.3	5.3	10.5	10.5	5.3	5.3	7.1	21.2	-	10.5	21.2
4,3	31.9	15.6	10.2	31.9	10.2	8.0	8.0	15.6	15.6	8.0	8.0	10.2	31.9	10.2	-	31.9
4,4	61.2	30.3	19.7	61.2	19.7	15.5	15.5	30.3	30.3	15.5	15.5	19.7	61.2	19.7	30.3	-

Table 1 Traffic interest in 4 x 4 matrix (in Gbit/s)

loss $Loss < 10^{-5}$ and calculate the resources that are required to meet this requirement for the given offered traffic. We then compare the resources required by the two technologies under consideration, namely OBS and OCS.

Note that load balancing is implemented in the network. We do not assume any over-dimensioning of OCS for protection purposes. However, traffic between two end-to-end node pairs is divided into two diverse “shortest path” LSPs. This means that in effect ASON (Automatically Switched Optical Network) functionality and QoS is assumed in OCS. In case of failure high priority traffic from the failed link will be directed over to the second LSP path where low priority traffic will have to be dropped to make space for the re-directed traffic from the failed link. Any other assumption would have required an increase of the OCS resources by about 30 %, which may not be a fair basis for a comparison with OBS. In other words, we have assumed a best-case scenario for OCS.

The traffic interest between the nodes is shown in Table I. Note the significantly diverse traffic levels between different node pairs. This is important in order to ensure the generality of the results.

Our results are shown in Figure 7. Different traffic levels have been simulated. The total offered traffic




Figure 7 Comparison of OBS and OCS in the 4 x 4 grid network. A burst length of 100 μs and a maximum burst assembly time of 300 μs have been assumed in the simulations

in the first scenario is 1581 Gbit/s. Due to the large total number (134) of LSPs, however, the average traffic per LSP is 11.8 Gbit/s. In the other two scenarios the offered traffic is halved to 790 and 395 Gbit/s, and the average traffic per LSP is 5.9 and 2.95 Gbit/s, respectively. According to the results, a relatively high degree of statistical multiplexing in OBS results in considerably better performance for OBS when compared with OCS. We also see that the lower the traffic per LSP, the higher is the gain obtained with OBS. The main reason for this is the granularity of the optical light path that results in reduced utilization of the available bandwidth in the case of OCS.

Our results indicate that an OBS network may typically require close to half the wavelength resources required by a corresponding OCS network. This can also be interpreted as follows: OBS node gear may cost roughly twice as much as OCS gear and still be a cost efficient solution.

Traffic burstiness

Given the importance of the input traffic, it is important to reflect on how well the assumed traffic burstiness corresponds to a likely scenario in 2009 and what may be expected further ahead. Our traffic model (Appendix 1) builds upon more or less existing applications whereas new applications may appear that affect traffic burstiness. Virtual reality type applications with extreme real-time demands may appear in the relatively near future, where information exchange of very large bursts of data is necessary. Such applications cannot allow the smoothing of traffic that we have introduced in our simulations to reflect bottlenecks that are inherent in the network today and in the near future. Hence such traffic would become a serious challenge for metro and core networks and would give rise to higher burstiness. The appearance of access networks with much higher capacity may in itself contribute to an increase of the burstiness of the traffic streams. However, the effect of mobile networks is difficult to predict in this respect. Even though mobile applications are also bursty, the peak rate of these services is low compared to similar services in the fixed network, due to limited radio resources. If video type mobile based services are a big success, a certain percentage of traffic carried in the metro and core network will also be of this type. This may actually contribute to the opposite effect, i.e. it may reduce the total burstiness of the composite traffic.

On the whole, we believe that neither access networks with larger capacities, extremely demanding virtual types of applications, or an explosion of video type mobile services will happen in the near future. Therefore our assumptions regarding traffic burstiness should hold for the time horizon of the current

simulation study where we have been aiming at about year 2009. However, it should be pointed out that higher burstiness may be expected in the future and this would lead to higher gains from OBS.

Conclusions of the theoretical and simulation study

According to the theoretical evaluation gains of any significance may only be obtained for:

- Traffic of relatively large mean value and high burstiness, i.e. of a different type than what is typical in networks today;
- Relatively large networks;
- Relatively low traffic volume per traffic stream.

The results of the network simulations have been in accordance with the analytical studies. The number of wavelength resources required in the OBS network may be of the order of 50 % of the number required in an OCS network. However, this is obtained for relatively low traffic per end-to-end traffic stream, whereas modest gains are obtained in the case of small networks with relatively high volumes of aggregated traffic – i.e. with several optical channels per end-to-end stream. The introduction of QoS differentiation will lead to an increase in the number of LSPs combined with a reduction of the average traffic per LSP and hence increase the gains obtained from OBS.

Techno-economic evaluation

So far we have discussed the gains that can be attained by OBS. Performance improvements alone are however by no means adequate – OBS solutions will have to pass the acid test of economic viability. For OBS to become a competitive solution, it is required that the same performance as competing technologies is achieved for lower total investment, and/or that attractive performance improvements are attained for lower investment than with competing technologies, and/or that OBS leads to attractive performance improvements that cannot be provided by other technologies. Note that in the latter case OBS should bring about an added feature that can increase revenues above any required excess investment, either by providing a competitive edge or by providing the means for the introduction of new services.

In this section we carry out a techno-economic analysis where we compare OBS directly with two commercial technologies, namely IP/MPLS and OCS. The current analysis is however limited to a comparison of the required investment costs, whereas revenues are not included.

For this study we implemented an adaptation of a techno-economic tool that has been developed within the IST project TONIC [8] [9] [10]. The tool we have implemented carries out a calculation of (among others) capital investment for an assumed ingress traffic volume – given a certain network topology, routing assumptions, technology and equipment choices. The network nodes are dimensioned according to the technology assumed and equipment components are chosen from a “shopping list” of viable types for each technology.

We have used this tool to compare a STOLAS OBS network with two commercially available technologies, namely OCS and IP/MPLS, where OCS is at a somewhat earlier stage of maturity than IP/MPLS. The equipment price implemented in our tool for these two technologies is based on typical market prices today. A STOLAS OBS network is based on primarily existing optical hardware; however, the technology as such is not commercially available and requires a good deal of development. Comparing commercially available technologies with emerging technologies is a rather challenging task that introduces a fair amount of uncertainty. We have chosen to employ a kind of scenario analysis in order to shed light and quantify the conditions and requirements that can lead to different technologies prevailing.

One of the main challenges of the current study lies with deducing a (future) sales price for the OBS technology when the available information today is primarily the current cost of the constituent components and subsystems. Clearly we cannot forecast the future price of OBS – this is determined by market forces and pricing strategy and is not directly foreseeable. What we can do, however, is find out what would be a reasonable price for this technology that would allow a manufacturer to make a profit and break even within a reasonable amount of time. In essence the assumption here is that cost-based pricing is employed. On this assumption we have derived possible price levels for STOLAS OBS as dictated by current empirical rules within the telecom equipment industry. The sales price, P , may be estimated from the cost of optical components (sub-systems) C_{HW} , according to the following empirical formula:

$$P = 2.5 \times C_{production} = 2.5 \times [3 \times [1.3 \times C_{HW}]],$$

where $C_{production}$ is the cost of production. The factor 1.3 is used to account for software costs of around 30 % of the cost of hardware. Note that such a large percentage of software costs corresponds to high-end of IP routers whereas in the case of optical equipment one estimates about 10 % ratio of software over hardware cost. However, we expect OBS to comply with

the packet switched paradigm and have therefore assumed high software costs. The factor 3 in the price equation shall account for non-optical hardware at the node, such as chassis, cables, power supplies, cooling fans etc. Typically, optical equipment accounts for about one fourth of the total cost in relatively mature systems, while in just developed systems it may dominate costs entirely. Hence a factor 3 is rather conservative and shall lead to an overestimation of the costs. In order to estimate the node price for future years, a price evolution with time is assumed for each component comprising the node. From these the hardware cost C_{HW} and hence the final price P may be estimated as a function of time. The evolution of constituent components is based on fitting historical price evolution for each component the past couple of years into a projection using learning curves [9].

The assumed cost of the main OBS subsystems and components is summarised in Table II.

Ingress interface modelling, traffic input, network topology and routing

Results are obtained from an analysis carried out on a 4 x 4 grid network with homogeneous traffic volumes at all ingress nodes and homogeneous traffic interest between node pairs across the network. In our analysis it is assumed that the traffic entering the network is in all cases IP packets. These are processed at an IP/MPLS edge router, the cost of which is not included in the comparison since it is identical for all cases. Then the packets enter our grid network through Layer 2 interfaces (Figure 8). The reason for this choice is that in order to be able to carry out a fair comparison between the different solutions we needed to define an interface that is identical in all cases. The three technologies we are comparing (OBS, OCS, IP/MPLS) are implemented rather dif-

	Price per unit US\$ year 2005
Tunable laser, CW	3500
Tunable laser, modulated at 2.5 Gbit/s	5000
Tunable laser, modulated at 10 Gbit/s	5000
Tunable laser, modulated at 40 Gbit/s	7000
Synch: 14 x (ms range 2 x 2 switch) + (µs range 2 x 2 switch)	35,000
Semiconductor Optical Amplifier (SOA)	300
Array waveguide grating, per port	40
Processing electronics (Optical Label Swapper)	100

Table II OBS hardware costs: list of main components and subsystem prices

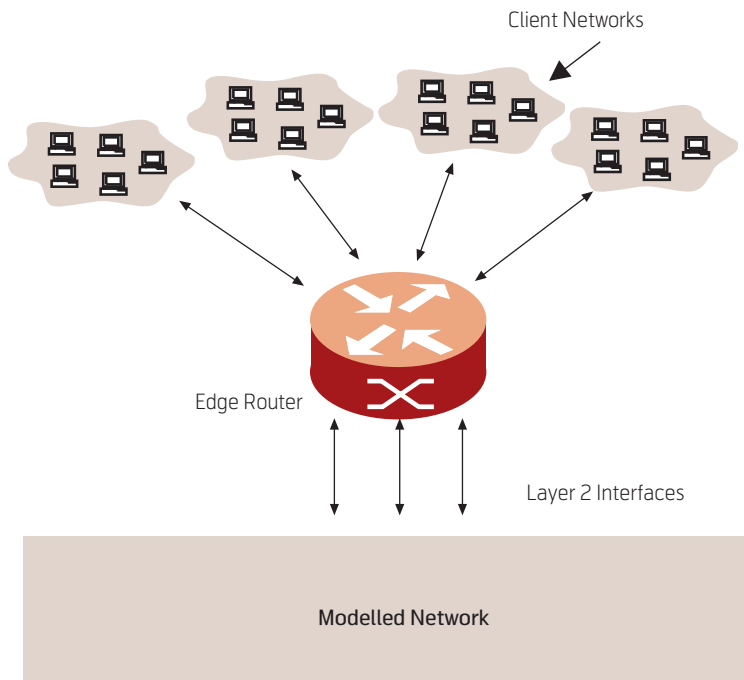


Figure 8 Schematic of the assumed ingress architecture, applicable to all technologies considered

ferently in real networks and have traditionally had, in any case to an extent, somewhat different areas of application.

We have assumed shortest path routing in our study. End-to-end traffic from each ingress node to each destination node is split to two equal parts and directed to the destination node following the (two) shortest paths. In the general case these are: a) one connection following the same row as the ingress node until the column of egress is reached where the path bends in the direction of the destination node, b) the remaining half of the end-to-end stream follows first the same column as the ingress in the direction of the row of the egress until it reaches this row where it bends towards the egress. In the case of nodes in the same row or column, two half-stream connections are assumed for each end-to-end stream; however, these now follow the same shortest path. In addition to balancing the traffic load, this routing scheme will allow for restoration of high priority traffic in case of failures in the network.

The same restoration scheme has been assumed for all technologies, where a failure situation is taken care of by dropping low priority traffic in order to accommodate rerouted high priority traffic. The required functionality is assumed to be included at the edge of the network, where similar Ethernet-based functionality will be used to differentiate between low and high priority traffic in case of node or link failure. Hence, for the studied technologies,

the network is not over-dimensioned with any additional capacity in order to take care of protection aspects. This implicitly assumes that the OCS network is automatically switched.

QoS differentiation has not been explicitly implemented in the comparison presented here.

IP/MPLS

The evolution of IP/MPLS has been used as a reference for scenario analysis in the rest of the study and hence no scenarios have been created for IP/MPLS. The evolution parameters of all other technologies have been defined in relation to the evolution assumed for IP/MPLS, hence the latter can be regarded in relative rather than absolute terms.

IP technology is broadly used and can be considered relatively mature. At the same time there are continuous improvements of the technology on both the switching and interface sides. We have considered the price reductions of the IP-MPLS switch itself to be in the order of 9–10 % the first coming years. This is reasonably well founded in historical data. Even though IP switches have been around for some time, the total world market is expected to be rather large, and thus only a small proportion of the total sales is bound to have been realised so far. In addition, constant developments with respect to MPLS, GMPLS as well as the incorporation of advanced optical components, may also pull in the same direction.

For the Ethernet interfaces, on both the client and network sides, sales are only in their infancy. Ethernet technology is rather new for such network node applications such that the expected learning curve coefficient may be in the range of 0.75. This is further supported by the price reduction that has been observed in the past couple of years. These lead to a price reduction in the order of 17 % in the first coming years.

OCS price evolution scenarios

OCS optical cross connects (OXC)s are commercially available and their price is given per port for a range of sizes. A typical commercial price per (10 Gbit/s) port has been used for OCS in reference year 2005, where a 10 % cost of software has been included.

Two scenarios are considered for the price evolution of OXC)s. In the “low-price” scenario it is assumed that the price evolution will be dictated primarily by the price fall of the constituent technology (i.e. hardware). Based on estimates of the market for OXC)s, the low volume today and the relative market maturity of the technology, we estimated a price reduction of 14 % for OXC)s for the first coming years in this scenario.

In the second scenario (called OCS ASON) it is assumed that the price of OCS switches will decrease with time at a rate lower than that of IP/MPLS due to the fact that significant development is required in the OCS case to achieve functionality comparable to IP/MPLS networks. The exact evolution is determined based on the following hypothesis. A very viable and cost-efficient solution to realise an all-optical OXC, is to base it on an AWG node architecture with tunable laser transmitters and tunable wavelength converters. The hardware is then rather similar to that of the OBS node – the only difference is that in the case of OCS there is no need for burstification processes, synchronisation, and label processing. Hence the software is simpler in the OCS case. In addition all hardware associated to label extraction and rewriting is not required; however, the corresponding costs are negligible. The cost associated with the pure OBS processes is on the whole not estimated to be very high so that in principle, there is no inherent reason why OBS should cost dramatically more than OCS when both technologies reach comparable levels of maturity.

To account for the above, we have defined the more pessimistic OCS ASON scenario as follows. Since burstification processes take place only at the ingress, we keep ingress costs separate and estimate these independently for OCS and OBS. The rest of the node cost, i.e. total switch costs including network interfaces but excluding ingress costs, are in this scenario considered to be comparable for OCS and OBS by year 2015 – albeit with OBS still being more expensive. More specifically, the OBS switch costs 33 % more than the OCS switch. Note that in the cost of the OBS switch in this case we include the cost of the optical label swapper unit. This OCS ASON scenario leads to a modest price reduction of about 5 % for the OCS switch the first coming years.

We consider both OCS scenarios as realistic. Note that both scenarios assume ASON functionality; however, they differ in their estimate of the development costs associated for introducing such functionality in an OCS network. OCS ingress interfaces are considered identical to IP/MPLS ingress interfaces, the only difference being that the transmission terminal costs

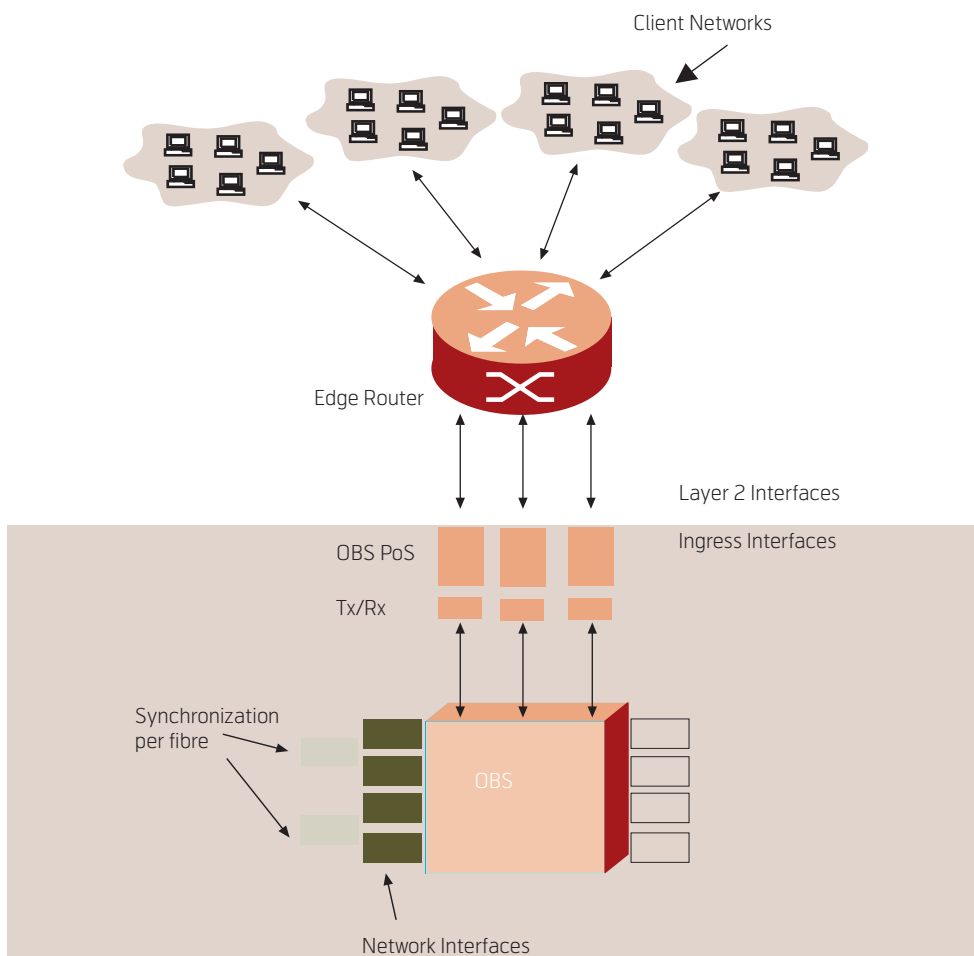


Figure 9 A schematic of the OBS node model used in the current study. Coarse synchronisation devices are included in each input fibre and the corresponding costs are included in the switch cost. The Network interfaces comprise Optical Label Swappers. Ingress interface costs simulate costs corresponding to an OBS ingress router

	OBS Interface cost in 2005	Price reduction per year	
Scenario A	3x IP/MPLS Ingress interface	22 %	realistic
Scenario B	3x IP/MPLS Ingress interface	17 %	unrealistically pessimistic
Scenario C	3x IP/MPLS Ingress interface	22 %	realistic
Scenario D	2x IP/MPLS Ingress interface	22 %	somewhat optimistic?
IP/MPLS Ingress	N/A	17 %	

TABLE III OBS scenarios considered in the analysis

twice as much as for the IP/MPLS case to account for the more stringent requirements for all-optical transmission.

OBS cost modelling

A schematic of the OBS node model we used in our study is shown in Figure 9. At the input of the modelled network there are Layer 2 interfaces whereas within the OBS network optical bursts are transmitted comprising a series of IP packets. These bursts are formed at the OBS-ingress and are then transmitted all optically through the OBS network until they reach their destination. For this purpose some sort of framing of the signal will be required for pure transmission quality purposes. This interface simulates an OBS edge router, where IP/MPLS packets are gathered in buffers according to class of service and destination, bursts are created and sent out in the OBS network. This function is quite similar to any IP/MPLS edge router and should in principle cost as much as an IP/MPLS edge router. It requires little hardware development, however, software development will be required in order to realise the burst management functions that are OBS specific. These are primarily related to the formation of bursts – such as determining the time-out factor, releasing or delaying bursts, prioritising, etc.

We account for OBS ingress costs by including an OBS Packet-over-SONET (PoS) type of interface at the ingress side that is dimensioned according to the number of bursts that are generated at each node. The idea is that packets will need to be encapsulated in some sort of transmission frame and here we can use current PoS interfaces as the starting point to estimate these costs. There are more stringent requirements on the transmission part in the OBS case due to the transparent by-passing of intermediate nodes and this is taken into account. In addition the OBS interface

will need to employ tunable lasers in this configuration instead of the conventional lasers currently used, so that their price will need to be adjusted upwards accordingly.

The question is still though how different should the price of the OBS ingress interface be from a conventional IP/MPLS ingress, both initially and in following years. As our study revealed, interface costs are rather dominant compared with the cost of the rest of the node hardware after the empirical formula has been applied. Therefore we created scenarios based on interface cost.

The scenarios we have chosen consider that the OBS ingress will cost somewhere between 2 and 3 times as much as an IP/MPLS ingress (PoS). This assumption is not doing OBS any favours. The scenarios are summarised in Table III. Since the OBS edge router has been modelled using the dedicated OBS PoS interfaces, we have considered four likely scenarios for the price and price evolution of these interfaces. In Scenario A, C and D the OBS ingress price falls faster than the IP/MPLS. This is a reasonable assumption since the much higher price of the OBS ingress is a result of its relative immaturity rather than inherently higher costs than IP/MPLS, as discussed above. Scenario B assumes that OBS ingress costs three times as much as IP ingress interfaces and at the same time these two follow the same learning curve, i.e. the same price evolution. This scenario is clearly unrealistically pessimistic and is used for sensitivity purposes only.

Results

Results were calculated for a number of end-to-end traffic levels. Note that both IP/MPLS and OBS scale nicely with input traffic, whereas in the case of OCS the results are very dependent upon the chosen input traffic level – as a result of the coarse granularity of OCS. Indeed, if the input traffic is close to an integer multiple of the traffic per channel, then the OCS channels are well filled and relatively good results are obtained from OCS. On the other hand, when the traffic level is such that one deviates significantly from an integer multiple of the traffic per channel, then less favourable results are obtained. This is more the case when the traffic per end-to-end stream is less than one channel. In a real network, some links will comprise well-filled channels while others will be less well filled, so that this effect is evened out. In our calculations, however, we assume a homogeneous network and equal traffic between all nodes, hence the choice of input traffic level will be decisive for the obtained results. In order to be able to accommodate traffic variations, a good margin over the mean traffic value needs to be allowed in the case of OCS,

also in real networks. In order to take these into account, we have carried out our study for input traffic levels that correspond to over 60 % filled OCS channels in the first reference year of an investment period. This is, however, also evened out over the whole five-year period of investment considered here, since we have assumed a yearly traffic increase of 25 %. We have thus taken care to attain a realistic representation of OCS.

In Figure 10 we present a comparison of the Net Present Value (NPV) of investment for IP/MPLS and the four OBS scenarios for an end-to-end traffic per node pair $T = 10$ Gbit/s. The NPV is presented for three five year investment periods. It may be seen that also the unrealistically pessimistic scenario for OBS (Scenario B) seems to give comparable results with IP/MPLS in 2015. The other three scenarios lead to lower investments for OBS compared to IP/MPLS.

In the following, we include scenario A and scenario C in our comparisons as we consider these two as the most likely scenarios, while we discard the worst case scenario (B) as well as the most optimistic scenario (D). These were used primarily in order to act as a sensitivity study and to illustrate the range within which we consider it most likely that scenarios may vary.

In Figure 11 we present a comparison of the chosen OBS scenarios, IP/MPLS, and the two OCS scenarios for a Core case with PoS interfaces and $T = 26$ Gbit/s. The two OBS scenarios come out positively in year 2015 while they are still too expensive as compared with OCS in year 2010. Note, however, that 10 Gbit/s channels have been assumed for all technologies to this point.

These results reflect the conclusion of the theoretical and simulation study, i.e. that multiplexing gains alone may not be sufficient justification for moving to an OPS/OBS paradigm. Indeed, although optical technologies appear to have an advantage by year




Figure 10 NPV of investment per five-year period for IP/MPLS and the four OBS scenarios. The results shown correspond to a Core case with PoS interfaces and an end-to-end traffic stream $T = 10$ Gbit/s per node pair. 10 Gbit/s channels have been assumed for all technologies




Figure 11 NPV of investment per five-year period for IP/MPLS, the two OCS scenarios and two of the OBS scenarios. The results shown correspond to a Core case with PoS interfaces and an end-to-end traffic stream $T = 26$ Gbit/s per node pair. 10 Gbit/s channels have been assumed for all technologies

2015, there is no clear advantage from OBS compared with OCS in terms of cost efficiency.

Thus far we have, however, forced all technologies to use 10 Gbit/s channels so that the obtained picture is




Figure 12 This figure shows the NPV of investment for two OBS scenarios, two OCS scenarios, and IP/MPLS for year 2005 and 2015. On the left hand side all technologies use 10 Gbit/s transmission channels. On the right hand side of the figure the choice between 10 Gbit/s and 40 Gbit/s channels is free: here the two OCS scenarios still use 10 Gbit/s, while IP/MPLS and OBS use 40 Gbit/s channels




Figure 13 Cost split per technology and scenario for 2005 and 2015. $T = 26$ Gbit/s




Figure 14 Total number of interfaces required per technology. IP/MPLS and OBS will choose high channel bandwidths (40 Gbit/s) to optimise cost, whereas for OCS this is not a viable option since an integer number of wavelengths needs to be assigned to each end-to-end (half) stream

in fact misleading. OBS and IP/MPLS can use higher channel bit rates since all connections on the link can share the available bandwidth. This is not the case for OCS where dedicated channels need to be allocated to each end-to-end (half) stream. In Figure 12 we present a comparison between technologies for the case




Figure 15 Comparison between technologies for a metro case (Ethernet interfaces) with relatively low traffic per end-to-end stream ($T = 3.5$ Gbit/s)

where 10 Gbit/s per channel are used for all technologies and for the case where the choice of bit rate may be optimised per technology. In this case IP/MPLS and OBS use 40 Gbit/s channels, leading to considerable savings. It should be noted that we have assumed that the cost of 40 Gbit/s interfaces is 2.5 times that of 10 Gbit/s interfaces – as indeed traditionally quadrupling of interface bit-rate has led to a cost increase by a factor 2.5.

In Figure 13 we see a comparison between yearly investment per technology, showing the detailed cost split per technology in year 2005 and 2015. It may be noted that as the total number of interfaces is reduced with the choice of higher bit-rate channels in the case of OBS, switching costs are also reduced. This is because switches are dimensioned according to the total number of interfaces in the case of OBS while IP/MPLS switches scale according to the total amount of traffic that is handled at the node. OBS is a clear winner in year 2015. Even when the network interface costs are included, the total cost of the node is primarily determined by the cost of the ingress interfaces. The same conclusions as for Figure 12 with $T = 26$ Gbit/s could be drawn also for $T = 43$ Gbit/s. In that case IP/MPLS attained even less favourable results.

In Figure 14 we see the number of interfaces per technology. The difference between OBS and IP/MPLS is due to the fact that we have assumed an increase of the traffic handled by OBS by 5 % to account for synchronisation guard bands. The main source of gains obtained with OBS is illustrated in this figure, where we can see how the number of interfaces may be reduced in the OBS case because higher bit rate transmission channels can be used. This is also the case for IP/MPLS; however, the cost of processing all through-passing traffic at each node pays its toll such that IP/MPLS is not a good solution for relatively high capacity per end-to-end stream.

For comparison purposes, we also present a study for a relatively low traffic metro network in Figure 15. Here we have included IP/MPLS equipment (IP/ MPLS Light) that is more suitable for lower traffic and that is cheaper per bit of traffic handled than the IP/MPLS Core equipment. The latter has been used in our study thus far since we have been considering relatively high total traffic at each node, around 1 Tbit/s. In addition we consider here OCS with 2.5 Gbit/s channels and have reduced the cost of the cross connect accordingly (assuming low price opto-electronic OXC). It is interesting to see that IP/MPLS appears to be a competitive solution also in year 2015 for this amount of traffic. However, note that OBS costs lie at the same level.

Summary and conclusion

The purpose of this study has been to evaluate the viability of optical packet or optical burst switching. To that end we first carried out an analytical evaluation of the gains that may be obtained with OPS/OBS combined with network simulations, where we addressed primarily two aspects: statistical multiplexing gains and granularity gains. The analytical study implied that gains of any significance may only be obtained for large networks and traffic of high mean value and large burstiness, or for relatively low traffic volumes. A thorough network simulation study was carried out both in order to quantify OBS gains and to give some insight as to where and when these become significant. The results were by and large in accordance with the analytical study. These have indicated that the number of wavelength resources required in the OBS network may be of the order of 50 % the number required in an OCS network. However, this is obtained for relatively low traffic per end-to-end traffic stream, whereas modest gains are obtained in the case of small networks with relatively high volumes of aggregated traffic – i.e. with several optical channels per end-to-end stream. The introduction of QoS differentiation will lead to an increase of the number of LSPs combined with a reduction of the average traffic per LSP and hence increase the gains obtained from OBS.

On the whole, the study of multiplexing and granularity gains of OBS as compared with OCS has pointed out that relatively modest multiplexing gains may be expected. Hence statistical multiplexing gains alone may not provide enough justification for a paradigm shift to OPS/OBS.

A techno-economic comparison of STOLAS OBS and two competing commercially available technologies – IP/MPLS and OCS – was carried out. The analysis was scenario based. The main cost for OBS is actually associated with the Ingress Interfaces so that OBS scenarios were created for the price and

price evolution of this interface. OCS scenarios were created according to whether the introduction of ASON functionality will require considerable development work or not. Our results ought to be regarded in relative rather than absolute terms and the study was carried out with reasonable sensitivity margins for a fair scenario assessment.

Our results indicate the following: For networks with several channels per link, optical technologies are preferable to IP/MPLS in the near future. If all technologies are forced to use the same kind of transport interfaces – 10 Gbit/s channels – then OCS achieves a good performance and there does not seem to be any clear advantage in moving to an OBS paradigm instead of OCS, in any case not in terms of clear cost savings. However, this conclusion is in fact misleading: when each technology optimises the type of transmission interfaces it uses, then 40 Gbit/s interfaces are employed in OBS and this can lead to considerable savings compared to OCS by year 2015.

It is, of course, not possible to predict with certainty the future evolution of the commercial technologies, let alone that of optical packet switching. Neither is it possible to know with certainty whether the particular technology we considered here will in fact become commercial. However, this exercise has made a couple of things rather clear. Firstly, it appears perfectly possible to build a cost-efficient OBS node and there is no component that makes this technology inherently prohibitively expensive. Secondly, OBS is more efficient than OCS also for relatively high end-to-end traffic between nodes, due to the fact that link resource sharing in OBS allows the use of higher bitrate interfaces whereas this would deteriorate granularity in the case of OCS. Thirdly, IP/MPLS ceases to be an efficient solution at relatively low traffic density levels and the introduction of optical interfaces brings costs down.

This study has focused on comparisons between networks with given traffic interests and has not taken into account traffic dynamics for example throughout the day or in case of failure. When the advantages of OBS that are related to flexibility and handling of dynamic changes are taken into account in addition to the indicated cost savings due to efficient use of the link capacity, the message is rather clear:

IP/MPLS as we know it is a great solution for low traffic density networks. However, for higher traffic density optical technologies take over because they eliminate excessive processing of through-passing traffic and provide efficient switching for high-throughput nodes. OCS networks can provide this functionality in the near future. In the longer term,

OBS will be developed to provide a combination of efficient switching/processing of large volumes of traffic – a typical characteristic of optical technologies – with the low granularity and flexibility of packet based networks.

References

- 1 *IST STOLAS*. 26 October 2005 [online] – URL: <http://www.ist-stolas.org/>
- 2 Koonen, T, Morthier, G, Jennen, J, de Waardt, H, Demeester, P. Optical packet routing in IP-over-WDM networks deploying two-level optical labelling. *ECOC 2001*, Amsterdam, The Netherlands, October 2001.
- 3 Cheyns, J, Develder, C, Colle, D, Van Breusegem, E, Demester, P. Cost efficient upgrading of OPS nodes. *Proc. OpNeTec*, Pisa (Italy), October 2004, 71–78.
- 4 Andreassen, R Ø, Helvik, B E. A note on traffic performance gain for optical burst/packet switching. *Proc. Optical Network Design & Modelling Conference, ONDM 2004*, Ghent, Belgium, 2–4 Feb 2004.
- 5 Zouganeli, E, Andreassen, R Ø, Haugen, R B, Solem, A, Feng, B, Sudbø, Aa, Stol, N, Helvik, B E, Emstad, P J. Performance evaluation of a mesh optical burst switched network with tunable lasers and tunable wavelength converters. *ConTel 2003*, Zagreb, Croatia, June 11–13, 2003.
- 6 Zouganeli, E, Andreassen, R Ø, Sudbø, Aa, Emstad, P J, Helvik, B E, Stol, N, Haugen, R B. Stacked label optical burst switched networks using tunable lasers: Prospects, Viability and Architectures. *Optical Network Design & Modelling Conference, ONDM'03*, Budapest, Hungary, Feb 2003.
- 7 Andreassen, R Ø, Stol, N, Zouganeli, E, Haugen, R B, Helvik, B E, Emstad, P J, Solem, A, Feng, B, Sudbø, Aa. Application dependent performance of a mesh optical burst switched network and effect of burst size. *European Conference on Optical Communications, ECOC'03*, Rimini, Italy, Sept 2003.
- 8 *TONIC (Techno-Economics of IP Optimised Networks and Services)*. 26 October 2005 [online] – URL: <http://www-nrc.nokia.com/tonic/>
- 9 Olsen, B T et al. RACE 2087/TITAN: Tool for Introduction scenarios and Techno-economic studies for the Access Network. *Proc. RACE Open Workshop on Broadband Access*, Nijmegen, The Netherlands, June 7–8, 1993.
- 10 Olsen, B T et al. Techno-economic evaluation of narrowband and broadband access network alternatives and evolution scenario assessment. *IEEE J. of Selected Areas in Communications*, 14 (6), 1184–1203, 1996.
- 11 Feng, B, Stol, N, Zouganeli, E, Andreassen, R Ø, Solem, A, Haugen, R B, Sudbø, Aa, Helvik, B E. Direct comparison between optical circuit and optical packet/burst switching using simulations. *OECC/COIN 2004*, Yokohama, Japan, July 2004.
- 12 Zouganeli, E, Andreassen, R Ø, Feng, B, Stol, N, Solem, A, Helvik, B E, Sudbø, Aa, Haugen, R B. *Optical packets, circuits, or packets and circuits: Viability of optical packet/burst switching*. Invited paper at OECC 2005, Seoul, July 2005.
- 13 Zouganeli, E, Solem, A, Kjøsberg, H, Andreassen, R, Feng, B, Stol, N, Sudbø, Aa, Helvik, B E. Techno-economic evaluation of optical packet/burst switching. *ECOC 2005*, Glasgow, Sept 2005.
- 14 *EURESCOM P918*. 26 October 2005 [online] – URL: <http://www.euresom.de>

Appendix 1 Traffic input and method for network simulations

Traffic input for the network simulation

The aim of this section is to produce (realistic) traffic matrices for input to the STOLAS network. To that end we follow the methodology reported in EURESCOM project P918 [14], albeit in a slightly simplified way. This methodology takes as input a number of geographical regions, i , with populations N_i , where certain fractions of the populations employ certain application services. The penetration of each individual service may differ from region to region, depending on the mix of user categories (residential users, large/small business users etc). From the knowledge of the number of users, penetration of the services and traffic generated per user (for the respective applications), we are able to establish the necessary traffic parameters. Specification of the latter is always related to the busy period.

The busy period for traffic originating from the business segment is between 10h00 and 15h00 whereas the busy period for the private segment is between 21h00 and 22h30. In order to dimension the network a rule of thumb for the calculation of the composite

maximum traffic volume is to add 15 % of the maximum volume generated by one segment to the maximum volume generated by the other segment. The maximum traffic is here calculated separately for the private and the business sector and the composite maximum is calculated using the above mentioned rule of thumb. The traffic parameters will all be related to year 2009 where we assume the core networks to be all-IP. The traffic predictions are derived from IST project TONIC [8] as well as internal work in Telenor.

Specification of the network

Traffic at the ingress is generated from “client networks” where this term is used in a wide sense, to be understood as the sum (or aggregation) of all possible client networks that generate traffic to the same ingress node. That is, if there are competing service providers in a region, with their respective client networks, we treat this multitude as one big network with the same traffic characteristics as the sum of the individual ones.

We consider sixteen geographical regions different in population size as well as composition (i.e. different mix of user categories). These are distributed evenly to different network nodes. Hence, with our client network definition above, there will be four types of client networks with N_i ($i = 1, \dots, 16$) users. Note that we assume most traffic to be packet based by 2009.

Table IV gives the number of Internet users (residential and corporate) in region i . A user is here defined as a user-line. In the private sector, a user is thus a household. In the business sector, a user is a user-line within a company.

The relative numbers of residential and corporate users are based primarily on empirical values combined with an additional assumption of our own that bigger regions will comprise a larger percentage of corporate users than smaller regions.

Application penetration

Nine application components have been considered here as summarized in Table V. The first three of these are Real-Time applications such that there is a maximum delay and jitter that can be tolerated by these. The remaining six applications are not Real-Time such that relaxed delay and jitter requirements are posed for these.

The penetration of application k ($k = 1, \dots, 9$) in region i is denoted by Q_i^{pk} and Q_i^{ck} for private and corporate users respectively. For simplicity we will, however, assume the penetration of applications to be the same in all the regions, i.e. $Q_i^{pk} \approx Q_k^p \forall i$ and

Residential and corporate users in Region i	1	2	3	4
N_i^p (millions)	0.7	1.0	1.5	2.0
N_i^c (millions)	0.8	1.0	1.5	3.0

Table IV Number of Internet users in Region i

$Q_i^{ck} \approx Q_k^c \forall i$. Note, however, that the relative mix of corporate and private users vary between regions. Penetration rates for each application are given in Table V separately for corporate and private users. These values are based on material from EURES-COM P918 that is slightly modified here in order to accommodate the somewhat different scenario assumptions that apply in our case. It is assumed here, for example, that all telephony is transported over our network, whereas only IP-telephony is considered in P918.

The number of potential users in region i of a given application k is then (using the “rule of thumb” mentioned earlier) given by

$$U_{ki}^p = N_i^p Q_k^p + \gamma N_i^c Q_k^c \quad k = 1, \dots, 9 \quad i = 1, \dots, 5$$

or alternatively

$$U_{ki}^c = \gamma N_i^p Q_k^p + N_i^c Q_k^c \quad k = 1, \dots, 9 \quad i = 1, \dots, 5$$

Application component	Penetration rates	
	Corporate users, Q_k^c	Residential users, Q_k^p
1 Telephony	1.0	1.0
2 Interactive video (HQ)	0.1 *	0.0
3 Interactive video (LQ)	0.2	0.2
4 Video/audio streaming (LQ)	0.3	0.05
5 Video/audio streaming (HQ)	0.0	0.6
6 Web browsing	1.0	1.0
7 File transfer (machine-human)	1.0	1.0
8 File transfer (machine-machine)	0.04 *	0
9 Messaging	1.0	1.0

* The penetration rates above are given with reference to the total number of business users. Services like Videoconferencing and Back-up file transfer are, however, carried out centrally at each company and not by each corporate user. The estimated penetration rate per company is divided by a factor of 10 in order to correct for this for these two services, where this factor reflects an estimate of the average number of employees per company. Note that this may vary from segment to segment and from country to country.

Table V Penetration per application for corporate and residential users (presumed to correspond to year 2009)

Application component	Continuous stream		Block transfer Residential users			Block transfer Corporate users		
	r_{up} bit/s	r_{down} bit/s	λ msg/s	v_{up} byte	v_{down} byte	λ msg/s	v_{up} byte	v_{down} byte
1 Telephony	64	64						
2 Interactive video (HQ)	400	400						
3 Interactive video (LQ)	100	100						
4 Video/au stream (LQ)	0	1500						
5 Video/au stream (HQ)	0	4000						
6 Web browsing			0.25	0	32k	0.25	0	500k
7 File transf (mach-hum)			0.025	0	30k	0.025	0	1.5M
8 File transf (mach-mach)			0.01	0	10M	0.01	0	10M
9 Messaging			0.002	20k	0	0.003	100k	0

Table VI Basic traffic parameters for applications

The first equation describes a scenario with focus on the private segment, while the last one describes a scenario with focus on the corporate segment. The factor γ in the above expressions denotes the fraction of corporate busy-hour traffic present in the private-segment busy hour, and vice versa (we assume symmetry in the rule of thumb).

Traffic parameters

In telephony, we know that the traffic generated by each user is $a = 0.1$ erlang in the busy period. This is the same as saying that he spends 10 % of the time on the phone, or, alternatively, 10 % of the users are active on the phone during the busy period. The bit rate traffic, \hat{a} , is defined as $\hat{a} \equiv a * \text{bit rate} = 6.4 \text{ kbit/s}$ for a telephone channel with 64 kbit/s bandwidth.

For Web browsing the ‘similar’ number is 0.4; that is, 40 % of the Web-users are connected to the Internet during the busy period. But in this case the 0.4 is not the (erlang) traffic, since the user is not generating messages continuously. And, in addition, the traffic a Web-user generates is basically in the downstream direction (towards the user). Hence, for Web-users, we also need to know λ , the number of messages generated per second, and v , the size of an average message. Observations give the following numbers for λ and v , $\lambda = 0.25$, $v_{up} \approx 0$ and $v_{down} = 32 \text{ byte}$. Here ‘up’ and ‘down’ refer to upstream and downstream traffic respectively. This gives a downstream bit rate traffic of $\hat{a} = a * \lambda * v * 8 = 0.4 * 0.25 * 32 * 8 \text{ kbit/s} = 25.6 \text{ kbit/s}$. This is four times the bit rate traffic of the telephone user, but the Web-user is also spending four times as long at his terminal.

Application component	Residential users			Corporate users		
	a	\hat{a} (kbit/s)		a	\hat{a} (kbit/s)	
		up	down		up	down
1 Telephony	0.1	6.4	6.4	0.3	6.4	6.4
2 Interactive video (HQ)	0.01	4	4	0.01	4	4
3 Interactive video (LQ)	0.04	4	4	0.07	7	7
4 Video/audio streaming (LQ)	0.01	0	2	0.01	0	2
5 Video/audio streaming (HQ)	0.1	0	200	0.1	0	200
6 Web browsing	0.4	0	25.6	0.4	0	400
7 File transfer (machine-human)	0.05	0	0.3	0.05	0	15
8 File transfer (machine-machine)	0.01	0	8	0.01	0	8
9 Messaging	0.25	0.39	0	0.25	0.6	0

Table VII User segment dependent traffic parameters

In order to have a unified notation, we define a as the fraction of time a user is engaged at his terminal during the busy period. For the continuous stream applications (application 1–5) a is simply the traffic per user. For applications 6–9 the traffic per user is most conveniently expressed through the bit rate traffic: $\hat{a} = a * \lambda * v * 8$. The total generated bit rate traffic for application k in region i is thus given by

$$\hat{A}_{ki}^x = \hat{a}_k^x U_{ki}^x \quad x = c, p$$

As mentioned above, the first five activity components in Table VI are called continuous stream applications. The last five are more conveniently looked upon as transfer of blocks of information (bytes) and are thus called block transfer. Table VII lists the essential parameters for our basic applications. From these parameters the (bit rate) traffic,

$$\hat{A}_{ki}^x, \quad x = c, p$$

can easily be calculated.

Traffic matrix

The traffic matrices specify the traffic between any two regions i and j . Routing will in general depend on application, but we will here consider as an approximation that the geographical distribution of traffic

is the same for all applications. Let $m_{i,j}$ then be a matrix which gives relative numbers for the traffic flow, i.e. the matrix elements denote the fraction of the total traffic generated in region i that is destined for region j . Some applications are bi-directional while others are essentially unidirectional. For the bi-directional applications, traffic originating in one direction will create traffic also in the opposite direction. For the unidirectional applications, the direction of generated traffic is given in Table VII, with the convention that upstream is the direction from i to j – in other words there is no traffic in the direction where n or r are zero. The resulting traffic matrix is included in the main article (Table I).

Network simulations

Simulation model

Figure 16 is a schematic of our generic simulator model. Slotted operation is assumed, which implies that a synchronization unit would be required at the input to the node. Bursts are created per destination at each ingress node, and each burst will in general contain packets from many different applications. Bursts are not allowed to be longer than one time slot. In addition, a time-out value has been defined such that ingress delay cannot exceed a certain maximum value.

Application		Parameters	Mean traffic [kbit/s]
1	Telephony	Silence detection variant of 1	25.57
2	File-transfer (Machine-Machine)	8	0.01 msgs/s, mean length 10x8 Mbit
3	Video/audio streaming (HQ)	5	Fixed rate
4	Interactive video (HQ)	2	Fixed rate
5	Interactive video (LQ)	3	Fixed rate
6	Video/audio streaming (LQ)	4	Fixed rate
7	Web browsing	6 residential	0.25 msgs/s, mean length 32x8 kbit
8	File-transfer (Man-Machine)	7 residential	0.025 msgs/s, mean length 30x8 kbit
9	Messaging	9 residential	0.002 msgs/s, mean length 20x8 kbit
10	Web browsing	6 corporate	0.25 msgs/s, mean length 500x8 kbit
11	File-transfer (Man-Machine)	7 corporate	0.025 msgs/s, mean length 1.5x8 Mbit
12	Messaging	9 corporate	0.003msgs/s, mean length 100x8 kbit

Table VIII Mapping of Applications to simulator source types




Figure 16 Simulator structure with one slot feedback fiber (generic model)

A scheduler is defined to take care that bursts are sent out synchronized to the switching node as well as delay bursts at a local queue when there is no free wavelength with which to enter the network. We assume pure ingress switching nodes in the network, i.e. no control signaling is needed from the optical part of the switching node to the scheduler. The core nodes are assumed to have no internal congestion [3].

Bursts that cannot be served may overflow to a one slot delay feedback fibre link with a given number of available wavelengths (32). This feedback fibre is shared by all directions/outputs from a switching node. At the core nodes, the arbitration is fair in the sense that it starts serving incoming bursts (of same priority) available at a certain slot period with equal probability and then serves the input fibres in a round robin fashion. Bursts originating from the feedback loop are however served first.

The simulator is implemented in the programming language Simula, based on the DEMOS (“Discrete Event Modelling On Simula”) package developed by Graham Birtwistle. The freeware conversion tool and script system “Cim”, which is developed at the University of Oslo, is used to run the simulations. Cim converts the Simula code to language C and runs it using the C run-time system available on the used hardware platform. The results that are presented in the paper have been obtained from a series of runs. The load was varied by taking a fraction of the traffic that corresponds to the scenario in Section IV. In addition, the slot size has been varied.

Traffic sources

We have chosen to use three basic source types to model all the different service components described in Section IV. The three basic source types used are as follows:

- “Speech”: Speech codec with silence detection. Active (ON) period negative exponentially distributed with mean value 0.96 seconds, and silent (OFF) period negative exponentially distributed with mean value 1.69 seconds. Active rate is 64 kbit/s, passive rate is 3.73 kbit/s. This source type is used in the simulator to model the “Telephony” source type.
- “Data”: Negative exponentially distributed time between time instants where information units become available. Length of information unit (to be divided into IP packets for transport) is assumed Pareto distributed. Seven different parameter sets are used to model the source types “Web browsing – residential”, “Web-browsing – corporate”, “File-transfer (machine-human) – residential”, “File-transfer (machine-human) – corporate”, “File-transfer (machine-machine)”, “Messaging – residential” and “Messaging – corporate”.
- “MPEG-compressed Video”: Deterministic time between generation of frames. Length of a frame is taken from a Gamma distribution, but is also correlated with the previous frame, if in the same “Group-of-pictures” (GOP). GOP lengths are different and out-of-phase for each source, as are the time instants where IP packets are sent towards the ingress nodes. Four different parameter sets are

used to model the source types “Interactive video (HQ)”, “Interactive video (LQ)”, Video/audio streaming (LQ)”, and “Video/audio streaming (HQ)”.

Note that this leads to a “streaming fashion” behaviour for “speech”- and “video”-based sources, but to very heavy offered traffic bursts generated from all “data”-based source types. We make no assumptions on any available transport level flow control protocol for “data”-based source types (i.e. no TCP or TCP-like protocol is in use); overflow from these traffic bursts must be handled by the network itself (e.g. by smoothing of traffic streams in electronic buffers and/or deliberate burst dropping – both in the ingress nodes).

The main limitation on the size of the network scenarios that could be studied in our initial simulation runs was found to be RAM of the computers used. A main contributor to RAM processes is the source processes, especially the data sources which must be modeled as one process for each actual source. However, the traffic scenario we used in D121 leads to a large number of active sources of especially certain types that generate very small amounts of traffic. These can safely be excluded from the scenario without influencing our results in any measurable way. Control simulation runs were carried out to verify this. Therefore the following source types were excluded from the our final simulation runs:

- “File Transfer (Machine-Machine) – source type 8 corporate
- “File-Transfer Man-Machine) – source type 7 residential
- “Messaging” – source type 9 residential
- “Messaging” – source type 9 corporate.

Introduction of bottlenecks in the network

If a direct access to an ingress node of the optical network is assumed for all source types, then for data type sources, this means that a (Pareto distributed) information unit is divided into IP packets and offered to the network at once, regardless of length. This is an assumption that easily leads to unrealistically high peak rates. In our simulations we assume that IP packets generated from sources will pass through an access network with limited capacity before arriving at the ingress node of the optical network. This also means that traffic streams from traffic from data type sources will be manifested at the edge of the optical network as considerably less bursty.

- For “residential data” type sources, an access network capacity of 10 Mbit/s has been assumed, implying fast Ethernet connections or similar.

- For “corporate data” type sources, an access network capacity of 1 Gbit/s has been assumed, implying Gigabit Ethernet or similar.

In practice in the simulator this has been implemented directly into the data source processes. IP packets are then sent out with the proper spacing in accordance with this.

Explicit modeling of limited buffer space in the ingress

In our study an amount of RAM is modeled for each ingress node. The buffer is a shared resource for all packets arriving at the ingress node. It is, however, possible to define limits of buffer usage depending on priority class of the arriving packet, i.e. highest priority class will always be able to use all free buffer space, the next priority class may be given access to the buffer only if (e.g.) less than 90 % of the buffer is occupied, and so on.

OCS simulations

In the current study the focus has been to quantify the gains of OBS compared with OCS, hence OCS networks had to be simulated. It is important to clarify the method used here.

The first point to make is that we focus on core (or metro) networks and simulate only the actual transportation of information through the core (or metro) network. Higher protocol level functionality is modeled only in as far as it has implications for our simulations, i.e. in order to make sure that our results are realistic in a certain context. One such higher-level functionality is the realization of an “emulated” end-to-end “circuit-switched” connection type by applying some type of priority mechanism to information units in a packet based network. If the optical core network is also packet (or burst) based and statistical multiplexing is used to arbitrate resources (in this case wavelengths in a certain direction at a certain time), such a mechanism must also be present in the core (or metro) network. Packets or bursts must then be divided into priority classes and differentiation with regard to access to resources must be implemented.

It is, however, still an open question whether an optical core network should be based on statistical multiplexing or not. An alternative is to implement transport through the core network as pre-allocated light-paths (a series of wavelengths through the core network) from the given ingress node to an egress node. This may be implemented for example as a GMPLS label switched path (LSP) where the actual wavelength is a label between two nodes of the core net-

work. The wavelengths are then part of an OCS connection through the core network and not available as a resource for statistical arbitration.

It is thus important to underline that it is only the optical core network that is circuit switched, not the access networks. The actual usage of lightpaths will still be statistical as seen from higher protocol level.

In our simulations we assume that one efficient way of doing this will be to organize packets into bursts before using the lightpaths to transport these bursts through the core network. An alternative would have been to transport the packets directly. However, it would in this case be required that packets are sorted into different priority classes at higher protocol levels, so that organizing packets into bursts of different priorities may be an efficient way of achieving this.

For a presentation of Evi Zouganeli, please turn to page 2.

Ragnar Ø. Andreassen received his MSc degree in Physics from the University of Oslo in 1988, and his PhD from the Norwegian University of Science and Technology, Trondheim, in 1997. Since 1988 Andreassen has been with Telenor R&D, where he has worked in various fields associated with telecommunication services, network technologies and business aspects. His current research interests include network architectures, network traffic and dependability performance, and charging and pricing of telecommunication services.

email: ragnar-oivind.andreassen@telenor.com

Boning Feng received his ME and PhD degrees in telematics from the Norwegian Institute of Technology (NTH), Trondheim, Norway, in 1985 and 1990, respectively. Since 2005 he has been Associate Professor at Oslo University College (OUC), Faculty of Engineering. Prior to that he was Senior Research Scientist at Telenor R&D (1998–2004) and Associate Professor at Norwegian University of Science and Technology (1990–1997). From 1996 to 1997 Boning Feng was Visiting Scholar at University of California at Davis (UCD), Davis, CA. He is also Adjunct Associate Professor at University Graduation Center at Kjeller, Norway. His field of interests include internet protocols, traffic modelling, control aspects of networks, and discrete event simulation.

email: Boning.Feng@iu.hio.no

Astrid Solem graduated from the Norwegian University of Science and Technology in 1988. She has since then been working at Telenor, the first four years in the Network division and then at Telenor R&D where she has been working with technology and business aspects of network architectures for fixed and mobile networks on contract from Telenor business units and European Research projects. She is currently in charge of monitoring the work in Study Group 13 on behalf of Telenor.

email: astrid.solem@telenor.com

Norvald Stol received his Siv.Ing. degree (MSEE) in 1983 from the Norwegian Institute of Technology (NTH), Trondheim, Norway. He was awarded the degree Dr.Ing. in telematics from the Norwegian University of Science and Technology (NTNU) in 1997. Since 1998 Norvald Stol has been Associate Professor at NTNU, Department of Telematics. Prior to that he was Research Scientist at SINTEF (1984–1991 and 1994–1998) and Research Scholar at NTH/NTNU (1991–1994). His interests include high capacity access and transport networks, mobile communication, Quality of Service aspects of networking, traffic modelling, and discrete event simulation. His current research focus is on different aspects of optical networking. He has managed a number of research projects and been active in EU initiated research collaborations. Norvald Stol is author/co-author of a number of papers and reports. He has held lectures and presentations in Norwegian and international fora. He is former Chairman of the NIK foundation (1999–2003).

email: Norvald.Stol@item.ntnu.no

Heidi Kjønberg received the degree Dr.Scient (PhD) from the University of Oslo in 1999. She has since 2005 been employed as a research scientist at the Norwegian Computing Center. Prior to that she was a research scientist at Telenor R&D (2000–2005) and has also had a post doc position financed by the Norwegian Research Council at Indiana University, Bloomington, IN, USA. At Telenor R&D she mainly worked on internet protocols, traffic engineering and traffic modelling. She now works on stochastic modelling in connection to the oil industry.

email: Heidi.Kjonsberg@nr.no

For a presentation of Aasmund Sudbø, please turn to page 33.

Bjarne E. Helvik is Professor at the Norwegian University of Science and Technology (NTNU), Department of Telematics. He is also with the Norwegian Centre of Excellence (CoE) in Quantitative Quality of Service in Communication Systems. He is Siv.ing. (MSc) from the Norwegian Institute of Technology (NTH) in 1975 and was awarded the degree Dr. Techn. in 1982. He previously held various positions with ELAB, SINTEF Telecom and Informatics and as Adjunct Professor at NTH.

His field of interests includes QoS, dependability modelling, measurements, analysis and simulation, fault-tolerant computing systems and survivable networks. His current research focus is on distributed, autonomous and adaptive fault-management in telecommunication systems, networks and services. He also deals with teletraffic issues and the architecture of networks and network based systems.

email: Bjarne.E.Helvik@item.ntnu.no

Rolf B. Haugen received his PhD in Physics from the University of Colorado in 1972. After one year as visiting scientist at Centre for Theoretical Physics in Trieste, Italy, he joined Telenor Research in 1973 and has since then been with the company. Haugen started his research work in the field of teletraffic theory and performance analysis. He has later, as Chief of Research, been in charge of several research groups within Telenor R&D. He is currently in charge of coordinating Telenor R&D's external activities, which include standardization, EU and EURESCOM projects and collaboration with Norwegian universities.

email: rolf-bjorn.haugen@telenor.com

Burst, packet and hybrid switching in the optical core network

STEINAR BJØRNSTAD, MARTIN NORD, TORODD OLSEN, DAG-ROAR HJELME AND NORVALD STOL



Steinar Bjørnstad is Senior System Architect at Canal Digital



Martin Nord is Research Scientist at Telenor R&D



Torodd Olsen is research scientist at Telenor R&D

Dag-Roar Hjelme is Chief Scientist at Optomed



Norvald Stol is Associate Professor at NTNU

The cost of transmission in fibre-optical networks has been much reduced in recent years. With increasing bandwidth demand, the cost and speed limitations of electronic switching are expected to become the limiting factors in future telecommunication core networks. Optical packet switching (OPS) and optical burst switching (OBS) have the potential to remove this bottleneck, by introducing statistical multiplexing in optical switched networks. Optical hybrid switching combines optical packet or burst switching with Optical Circuit Switching (OCS), and, thus offering both high throughput and a circuit switched quality. In this paper we give the rationale for introducing OPS, OBS or a hybrid network and describe the required switching functionality in each case. Then we discuss and analyse different packet/burst handling schemes and make a detailed simulation analysis of an optical packet switch, identifying the most promising design choices. Furthermore we describe a hybrid switch and present simulations showing its basic properties.

Introduction

Improved residential access and corporate communication needs may trigger an increased demand for core networks supporting multimedia applications such as video-conferences, online gaming, video-on-demand and web-television streaming. Supporting these services requires a network with a better performance than today's IP network. The future "layer 3" network will benefit from being served by a high-capacity optical layer network that complies with the following requirements:

- 1 Support high utilisation of resources; i.e. handle high capacity and link loads in a cost efficient way;
- 2 Support fine granularity that can be tailored to the applications' needs;
- 3 Support the quality needed by strictly real time services; i.e. low packet loss and low delay;
- 4 Support variable length IP packets, as used in the Internet today.

We will start by giving a background for the current developments in technologies for core networks, before discussing how the different switching techniques have the potential to meet these requirements.

In the mid-1990s, the introduction of the erbium doped fibre amplifier (EDFA) enabled the more bandwidth efficient wavelength division multiplexing (WDM) transmission systems that resulted in reduced core transmission cost. However, cost-effective core networks also require cost-effective switching and routing. Today, the switching is performed electronically; optical to electrical to optical (O/E/O) conversions are hence required in the switching nodes. With

increasing bandwidth demand, the cost of these conversions and the limited capacity of electronic switching may become a network capacity bottleneck. Implementing optical switches with ms-range switching time, e.g. based on micro electro-mechanical systems (MEMS) technology, is nowadays becoming feasible. Currently, operators and vendors are working on control plane architectures for resulting optical circuit switched (OCS) networks. OCS networks offer explicit transfer guarantees. This leads however to set-up delays of at least several milliseconds, due to the propagation delay induced by awaiting the set-up confirmation.

Future optical networks should be able to serve a client layer that includes packet-based networks, such as the Internet. These may have a highly dynamic connection pattern, with a significant portion of bursty traffic between the communicating pairs. In this case, OCS transport at a granularity of 2.5 Gbit/s, 10 Gbit/s or 40 Gbit/s may lead to low resource utilisation. It would require over-dimensioning of the number of connections as well as of the bandwidth reservation of each connection, in order to avoid excessive delay and extensive buffering at the ingress router. Optical packet switching (OPS) and optical burst switching (OBS) have the potential to overcome these problems by introducing statistical multiplexing (SM) at the optical layer. Several optical switching technologies have demonstrated nanosecond to microsecond range switching times, opening up for these new switching schemes. Telenor R&D and NTNU participate actively in European based studies of both concepts, through the ongoing IST projects NOBEL [1] and STOLAS [2], the recently concluded COST 266 action "Advanced Infrastructure for Photonic Networks" [3], as well as the ongoing COST 291 action "Towards digital optical networks". Fur-

thermore, a hybrid concept combining packet switching, optical or electronic, with optical circuit switching, is studied in the Optical Packet switched Migration capable Network with Service Guarantees (OpMiGua) project. The circuit switched part of the network offers a Guaranteed Service Transport (GST), while the Statistically Multiplexed (SM) part of the network ensures high throughput for a Best Effort (BE) type of traffic. The OpMiGua project receives parts of its funding from the Norwegian national research council (NFR). Project partners are Telenor, NTNU and Network Electronics, a Norwegian company developing and manufacturing WDM systems for broadcasting applications.

OPS and OBS concepts

Inevitably, there are some differences in terminology within the research community. We start this section by explicitly describing some concepts and terms used in this article. Both *OPS* and *OBS* are based on the idea of separating forwarding from switching in the network nodes. Forwarding decisions are based on the information contained in the optically transmitted control information. As illustrated in Figure 1, in the OPS case this is realised as a packet header while in the OBS case it can be a separate packet preceding the burst (Burst Control Packet, BCP). The BCP may travel together with the burst at a separate wavelength, alternatively it is encoded on the burst using a different modulation format than used for the payload, e.g. as in IST STOLAS [2]. The BCP and OPS header are converted to the electronic domain at the packet/burst switch input interface. It is subsequently electronically processed in the control unit, which then configures the node resources for switching, wavelength conversion or buffering of the data unit. On the other hand, the burst or packet payload is optically switched, thus avoiding the costly O/E/O conversion, and in principle simplifying the interface cards to the optical fibres.

OPS definition

OPS assume in-band encoding of control information (same wavelength, simultaneous transmission). Since the header follows the rest of the packet closely, there is no reservation possible. Reading and reinsertion of packet headers with strict timing requirements are required, due to the short packet duration – typically around 1 μ s. Typically, OPS proposals assume packet sizes from around 40 to 1500 B, representative of today's Internet traffic [4]. Some aggregation of client packets to form an optical packet may be required, since the smallest client packet sizes only correspond to 32 ns payload at 10 Gbit/s bitrates. The relation between packet format, switching technology and overhead is studied in more detail in [5].

OBS definition

OBS was introduced quite recently [6] and aims at offering granularity in between optical wavelength channels and optical packets, at the same time minimising optical complexity. OBS assumes more extensive burst aggregation to form bursts with tens of kB of payload, thereby relaxing the requirement to optical switches, from the nanosecond to the microsecond range. Client layer packets are assembled in edge nodes, and transported through the optical network in optical bursts. Another key concept of OBS is *one-pass reservation*; i.e. burst transmission is initiated without awaiting a set-up confirmation. As illustrated in Figure 1, BCPs are transmitted ahead of the corresponding burst. The relative delay is termed offset, and as a minimum it equals the time required to settle the switches and perform control processing in the network nodes. The reservation scheme depends on the amount of information that the BCP contains on its corresponding burst. *Reserve a Fixed Duration* (RFD) type BCPs include both the start and end times of bursts. The combination of RFD scheme, offset and the principle of delayed reservation (DR), enables advanced burst scheduling. For example, a newly arriving burst can be reserved in a gap left by already reserved bursts. This can optimise bandwidth usage and enable offset based QoS differentiation methods. A more detailed overview of OBS is provided in [7].

Optical Circuit and Packet switched Hybrid (OpMiGua)

When SM is efficiently exploited, constant delay and zero packet loss is not guaranteed. Combining properties of Optical Circuit Switches (OCS) with those of optical packet switches (OPS), may offer both better cost and performance [8]. Recently, a hybrid network concept was proposed [9, 10], combining packet switches with a Wavelength Routed Optical Network (WRON); i.e. an optically circuit switched network as shown in Figure 2. In this concept, a frac-




Figure 1 Main differences for transferring control information in OBS and OPS. a) illustrates that OBS typically has an “out-of-band” BCP, transmitted on a different wavelength and with a time-offset compared to the burst payload. b) illustrates that OPS typically has an “in-band” header, transmitted at the same wavelength and simultaneously as the packet payload. A number of variations on these main principles have been proposed

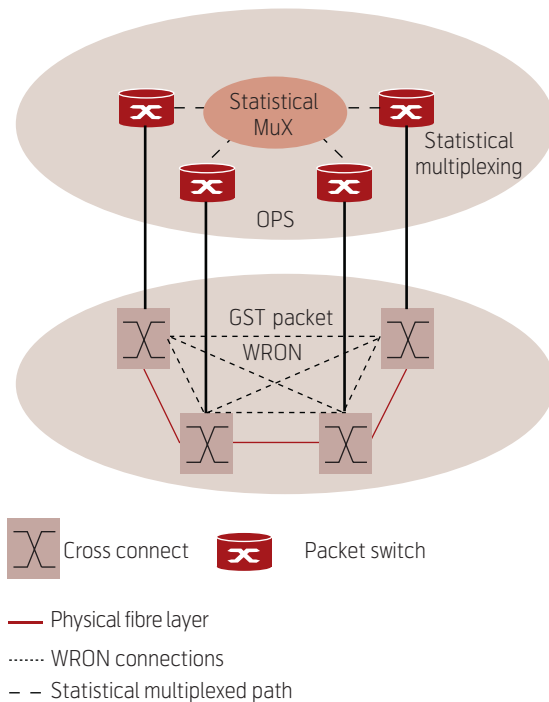


Figure 2 Hybrid network model illustrating the efficient sharing of the physical layer. If the WRON network is an S-WRON, the cross connect can be a matrix, manually configurable. If a D-WRON network is preferred, the cross connect should be configured by a control plane. The connected cross-connects and packet switches are physically co-located. These can be separate units with a common control unit, or they can be integrated, sharing physical resources in the node

tion of the packets, the Guaranteed Service Transport (GST) packets, follow the wavelength paths defined by the configuration of the optical cross connects (OXC). The remaining packets (SM packets) are switched in packet switches according to their header information. The GST packets are not subject to packet loss, and delay is fixed. On the other hand, SM packets have a BE quality. They will be dropped during congestion, and the use of buffering implicates a variable delay.

In the hybrid network concept the link resources are divided in the time domain with packet length granularity. By sending GST packets through the WRON, and sending lower quality packets through the packet switches on link bandwidth not consumed by GST traffic, both high throughput efficiency and the guaranteed service are achieved [8]. The GST class is able to accommodate services requiring circuit switched quality, while lower quality traffic is statistically multiplexed and experience only a moderate penalty from the GST traffic. We term this network concept Optical

packet switched Migration capable network with Guaranteed service (OpMiGua). If the WRON network is used for transit traffic, studies show that hybrid networks achieve potentially large savings compared to both packet- and circuit switched networks [11]: Compared to a circuit switched network, one can save wavelengths because SM improves the utilization of the wavelength paths. Compared to a packet switched network, reduction in packet switch sizes can be achieved, because the packet switches do not process packets that are sent through the WRON. Assuming a reference network applied within the COST 266 project with associated traffic demand forecast for 2008, it is found that the hybrid network only needs 1/3 of the wavelengths needed in a pure circuit switched network. Compared to a pure packet switched network, the amount of packet processing in the nodes is reduced corresponding to a 30 % reduction in number of wavelengths.

Comparison of optical switching schemes

In Table 1 we compare basic properties of OCS, OPS and OBS, aiming at identifying main properties of each switching paradigm. Note that these concepts are not standardised, and no universally accepted definition exists.

Recent papers on OPS/OBS show a converging trend: First, the packet/burst handling schemes discussed below are becoming more similar. Second, in contrast to conventional OBS, buffers are now being implemented in new OBS designs, similar to OPS designs. Note that implementing buffers imply that packets may arrive out-of-order even when packets in the same stream follow the same network path. Such out-of-order delivery is avoided in OCS. However, some important differences remain. These include finer granularity, but also higher overhead in OPS. OBS still assumes out-of-band control encoding and more advanced scheduling. Furthermore, its coarser granularity gives more tolerance to the switching time of switch matrices.

In OCS, since each connection has reserved dedicated resources, it can offer explicit transfer guarantees. This means that once a connection is set up, there will be no loss of packets inserted on that circuit. In contrast, such explicit/absolute transfer guarantees require retransmission protocols for OPS and OBS networks. However, properly dimensioned networks, possibly combined with "MPLS type" circuit oriented traffic engineering, can offer statistical quality of service guarantees which may serve the purpose, e.g. average packet loss rate (PLR) of 10^{-6} . Moreover, employing QoS differentiation gives the network designer the possibility of offering two or more PLR levels to the clients.

	OCS (including set-up)	OBS	OPS	Hybrid
Recommended size of transfer unit	> GB	~ tens of kB	~ 40 – 1500 B	~ 40 B – > GB
Explicit transfer guarantee	Yes	No	No	Yes (GST part)
Statistical QoS guarantees possible	N/A	Yes	Yes	N/A (GST part) Yes (BE part)
Loss type	Set-up request rejection	Loss of burst	Loss of packet	Loss of packet (BE part)
Control	Out-of-band	In-band or Out-of-band	In-band	Out-of-band (GST part) In-band (BE part)
Latency mainly given by ¹⁾	~ 3 x prop. delay	~ Prop. delay + burst assembly delay, offset	~ Prop. delay	~ Prop. delay
Control overhead given by	Connection set-up	BLP	Packet header	Packet header (BE part) and Connection Set-up part if D-WRON for GST

Table 1 Main properties of Optical Circuit-, Burst- and Packet switched networks

¹⁾ Latency is end-to-end delay. This row is valid in the case of no set-up rejections, packet or burst loss

The hybrid scheme combines absolute and statistical QoS guarantees. Headers are only needed on the BE type of traffic. If a Static WRON (S-WRON) is employed, overhead for setting up connections is avoided for the GST packets. GST connections may be of the duration of a packet, burst or dynamic circuit. In this paper we consider the connection duration to correspond to that of a burst.

Packet/burst handling schemes

Both in OBS and in OPS the switching matrix is required to be re-configurable on the burst/packet time scale, to allow the packets/bursts to efficiently share node and fibre resources. The switching operation is hence more demanding in OPS than in OBS. The basic principles for switching architectures and functionalities are in principle independent of the packet/burst-handling scheme. In the hybrid scheme, the BE traffic may be burst switched or packet switched, thereby having the same requirements to the switching architectures as OBS and OPS, respectively. The GST packets follow circuit switched paths, hence demands to the switching matrix for these packets are equal to those for circuit switching. In Figure 3 we show the four potential handling schemes for OPS and OBS, classified according to synchronisation and data unit type.

The following summarises main points to consider when making a choice concerning operation mode and data unit type (fixed or variable length).

In asynchronous optical networks, each transmitter/output port in the nodes emits optical packets asynchronously, so that packets arrive at random moments at the OPS nodes' interfaces. Synchronous operation requires optical synchronisers for packet alignment at the switch interfaces and a global network clock for practical realisation.

Fixed length packets (FLP) schemes require fragmentation of client packets and padding to fill the optical packets. Both factors increase the *overhead*; i.e. bandwidth consumption for successful transmission through the OPS network [5], and thus the blocking ratio for the same client load. In addition, fragmentation calls for reassembly of client packets, which increases egress node complexity. For fragmented packets, it can in some cases be hard to tell during reassembly whether a missing fragment is lost or simply delayed. Fibre delay line (FDL) based buffer design and management is simpler for FLP than for variable length packets (VLP), and it is simpler to maintain the packet order. The complexity of other contention resolution methods may be independent of the packet length. In general, a switch matrix operating in synchronous, FLP mode will have less contention than when operating in asynchronous mode [12]. Furthermore, in this mode, a re-arrangeable non-blocking switch matrix may have equal performance to a strictly non-blocking switch matrix [13]. For optimum scheduling, variable packet length requires coding of packet duration in the packet header.

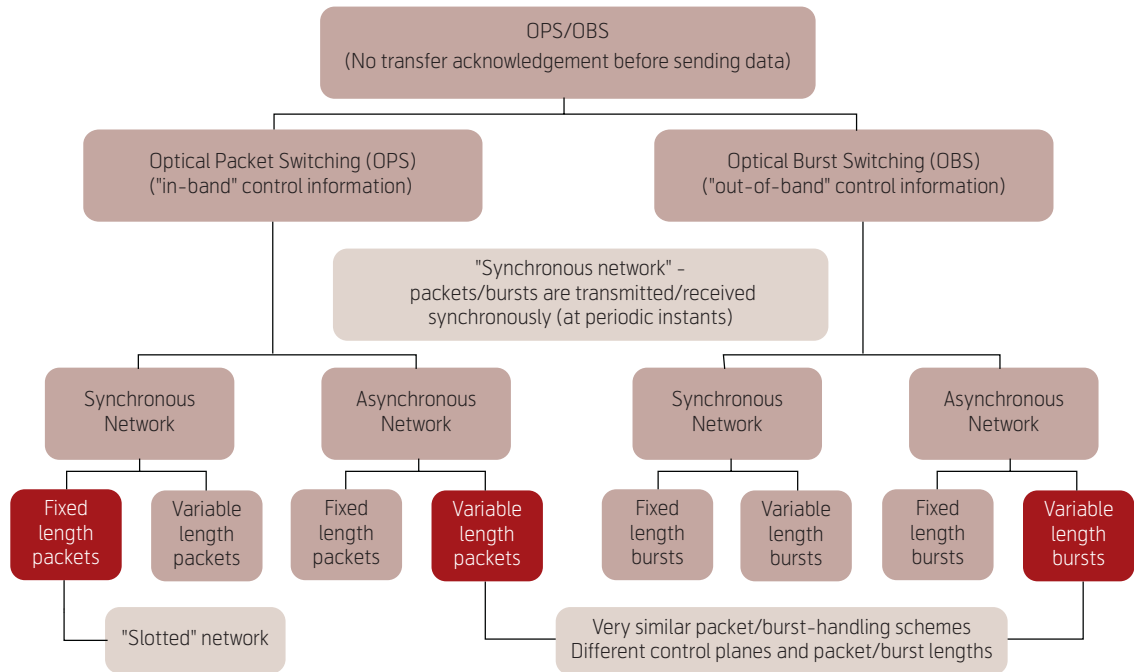


Figure 3 Potential and viable (dark red boxes) OBS and OPS packet/burst handling schemes. As mentioned earlier, control information in OBS may also co-propagate in-band with the burst

A main motivation for OBS is to reduce the optical technology complexity. Therefore, avoidance of burst alignment is essential, and asynchronous operation required. Burst assembly mechanism typically includes timers to prevent excessive delay, and variable burst length gives the assembler more freedom to optimise burst lengths. Furthermore, in asynchronous operation, fixed burst length does not in general reduce blocking probability. Hence, we consider in this study that OBS networks operate asynchronously, with variable length bursts, as is the custom in most work on OBS.

On the other hand, as further discussed in [7] OPS has two schemes that in general are more attractive, namely synchronous operation with FLP and asynchronous operation with VLPs. The two schemes' main advantages are summarised in Table 2.

Contention resolution

In the hybrid scheme, contention resolution is not needed for the GST packets since they have an absolute guarantee and follow reserved wavelength paths through the network. For the BE traffic, requirements to contention resolution will be as for OPS or OBS, depending on the scheme chosen.

Time domain contention resolution

Buffering enables so-called store-and-forward networks, in which contending packets are stored until they can be forwarded. For unlimited buffer capacities, the network can guarantee that all packets reach their destination. With optimum routing schemes, this scheme also guarantees that the packet follows the shortest possible path, which maximises the network overall capacity.

Optical buffering

Most optical buffers proposed rely on Fibre Delay Lines (FDL), which are known to be bulky and to

OPS packet handling scheme	Synchronous FLP	Asynchronous VLP
Synchronisers / network clock	Required	Not required
Edge node assembly	More complex	Less complex
Architecture required to avoid internal blocking	Reconfigurably non-blocking	Strictly non-blocking
Requirements to contention resolution	Lower	Higher
Overhead	Higher	Lower

Table 2 Main properties of attractive OPS packet handling scheme

require stable external temperature to maintain a stable propagation delay. The time spent in the FDL is determined by the fibre length – it is not a RAM. More complex buffer designs, achieved by interconnecting re-circulating FDLs or by broadcast-and-select architectures [14] can increase the accessibility after insertion. However, both the increased loss, which must be compensated for by EDFAs, and the switch cross-talk introduced, decrease signal quality. Variable length packets further increase buffer size and buffer design complexity. Complex FDL multi-stage buffers may be required, and available size in OPS nodes is in practice limited. Alternatively, [15] has demonstrated a variable optical delay circuit for use in optical buffers, suitable for variable length packets. However, this requires extensive hardware such as re-circulating loop, filters, circulators and λ -converters.

Electronic buffering

Most OPS designs assume an electronic control unit, since electronic processing is superior to optical processing. Similarly, electronic memory is inexpensive compared to FDLs. However, electronic buffering requires opto-electronic interfaces. Avoiding such interfaces is one of the main targets in OPS. We believe that buffering in OPS core routers should be used in combination with other contention resolution schemes. Therefore, only a small fraction of the packets undergo buffering, thus only a few buffer input interfaces are needed. An OPS design partly based on electronic buffering is proposed in [16], and a design that relies on electronic buffering only is proposed in [17].

Wavelength dimension contention resolution

By sharing the total traffic on several wavelengths, i.e. decreasing the load on each wavelength, the contention problem is minimized, and with a sufficient number of wavelengths a sufficient packet loss ratio can be achieved [18]. However, because efficient utilisation of the link-capacity is important for achieving cost-effective transport networks, a relatively high load on each of the wavelengths is desirable. Contention can be resolved by using λ -converters to convert packets to vacant wavelengths at the same fibre without increasing the number of wavelengths. Due to properties of statistical multiplexing, this method improves with the number of wavelengths per fibre. Furthermore, the signal quality aspects are potentially excellent, since the non-linear response of some wavelength converters has a re-shaping effect (2R). Adding a clock-recovery and a pulsed local source, full regeneration (3R) can be achieved [19], and the integration potential opens up for compact devices. This must be exploited to achieve cost-effective contention resolution due to the high number of wavelength converters needed in this approach. The tuning speed is an important parameter for tuneable wave-

length converters since it determines the optical guard band needed between subsequent packets.

Contention resolution by deflection routing

When two packets contend for the same output fibre/wavelength pair, the conflict can be solved by simply forwarding one of the packets on another fibre, and let an alternative OPS node be responsible for forwarding and switching the packet towards its destination. This contention resolution scheme is called “deflection routing”, and can be seen as spatial contention resolution. In a sense the whole network is used as a buffer. The great advantage of this approach is its simplicity, since buffers and wavelength converters may be avoided. An important drawback is the significant loss of packet order that follows, which increases the edge node complexity. Furthermore, the overall network capacity decreases with the increased average hop-count, so that deflection routing does not work well for high loads, especially in asynchronous networks [20].

Comparison of contention resolution schemes

The main advantages of electronic buffering are compatibility with asynchronous scheme, low footprint, excellent signal quality aspects and the commercial availability. The main drawback is the limited transparency and the potentially higher cost of the O/E interfaces.

Wavelength conversion has the same advantages, and somewhat better transparency properties. Furthermore, the performance improves with the number of wavelengths. Demonstrations of wavelength conversion at high bit rates prove their feasibility.

Optical buffering is transparent, but has no random access in the time domain. FDLs have negative impact on the signal quality. Asynchronous schemes require more advanced scheduling processing than slotted schemes. Relying purely on FDLs in this scheme seems unrealistic due to the large footprint required for extensive FDL arrays or use of complex multi-stage buffers.

Deflection routing’s greatest advantage is the low hardware complexity in the core nodes and its transparency. However, the increased hop count, which limits network capacity, especially for asynchronous networks, is a drawback. Furthermore, deflection routing in slotted networks requires complex edge nodes, since the need for fragmentation in combination with the possibility of loss of packets, will require larger edge buffers and mechanisms capable of reassembling fragments and suitable time-out values to decide when a fragment is lost or simply delayed.

	Electronic buffering	Optical buffering	Wavelength conversion	Deflection routing
Random access	Yes	No	N/A	N/A
Packet sequence	Modified (typically)	Modified (typically)	Maintained	Modified
Hop count	Minimum	Minimum	Minimum	Increased
Additional delay	~ μ s	~ μ s	~ ns	~ ms
Payload Clock Recovery	Required	Not required	Only for 3R	Not required
Transparency	Limited	Yes	Possible	Yes
Scheme control	Simple	May be complex	Simple	Simple
SNR	Very good (3R)	Degraded	Good (2/3 R)	Degraded
HW complexity	Intermediate	Complex	Intermediate	Low

Table 3 Comparison of four buffering schemes

Combining contention resolution and packet handling schemes

A combination of contention resolution schemes should be used for optimum performance/complexity ratio. [21] demonstrates the technical feasibility by implementing contention resolution in an OPS node using wavelength conversion, optical buffering and deflection routing. Wavelength conversion has good signal quality aspects and is effective in resolving contention. To achieve very low PLRs it should be used with buffering, e.g. FDLs as in [22, 23]. Here, the wavelength conversion minimises the buffering, and in slotted scheme a small amount of single unit length FDLs can be used. This combination will not work as well for asynchronous packet handling scheme, which requires more buffers and more complex FDL buffer design. We believe that asynchronous mode will benefit from the random access and simple design of electronic buffers in combination with wavelength conversion.

Node design in OPS

A scalable packet switch design

A modular and scalable design, scaling to a very high number of wavelengths, and a high node degree is shown in Figure 4. The design is based on using λ -converters and Array Waveguide Gratings (AWG), thereby avoiding the use of optical switches. Benefits of this design are that the signal path does not introduce large attenuation, contributing to the scalability of the switch, and that no additional λ -converters are needed when using the wavelength dimension for contention resolution. When a signal is sent into an AWG, the wavelength of the signal will decide at which output on the AWG the signal will occur. This principle is exploited for routing the packets to the desired output. The input WDM signal of each fibre

is demultiplexed to its corresponding wavelengths and fed to the input of the λ -converters. The outputs of each λ -converter are then fed to the AWG inputs. By tuning the output wavelength, packets can be sent to any of the AWG outputs. The packet will be sent to the scheduled output if a vacant wavelength can be found. If no output with correct destination is available, the packet will be sent to one of the buffer inputs, if a vacant buffer input can be found. If not, the packet will be dropped. Buffered packets are clocked out of the buffer and sent back to an AWG input as soon as a wavelength output to the destination becomes available. At the buffer output, the wavelength, and thus the output of the first AWG, is decided by tuning a tuneable laser. This type of architecture is called a feedback design, and has the benefit of supporting packet priority also when FDLs are used for buffering [24].

In this design, each of the input fibres is coupled to a corresponding input plane. Hence, the AWG size in each plane equals the sum of the number of wavelengths, 'W', and buffer inputs, 'B', which is independent of the number of fibres and node degree. Each plane has 'W' outputs. The outputs are fed to a second set of 'W' AWGs with size $N \times N$, making the maximum allowable number of fibres depend on the AWG size.

The buffer is shared among all the planes. Total number of buffer inputs corresponds to the number of buffer output ports in each plane. The buffer ports from each of the planes are passively coupled together using 1 : N couplers so that port '1' in plane '1' are coupled to port '1' in plane '2' through 'N', and the same goes for the rest of the 'B' ports. At each of the buffer outputs, wavelengths are set using tuneable lasers. By setting the wavelength, the packet is forwarded to the decided output.




Figure 4 The scalable packet switch design, with two switching stages, using a fully shared buffer. In the circle is shown the principle of the second switching stage

Forwarding of a packet undergoes two stages. In the first stage, the packet's wavelength at the output fibre is decided. In the second stage, the packet is forwarded to the destined fibre. A packet occurring at one of the input fibres is forwarded to its destination as follows: In each plane, which is the first stage in the forwarding, the wavelengths in each fibre are demultiplexed to 'W' corresponding fibre-lines. Each line is then fed to a λ -converter, setting the wavelength so that the packet is forwarded to any of the desired outputs of the AWG, making it possible to freely choose at which AWG at the second stage the packet will occur. This makes the first stage wide-sense non-blocking as described for the basic design.

After passing the first AWG, the packet is sent through a λ -converter before occurring at the input of the chosen AWG in the second stage. By setting the output wavelength of the λ -converter, forwarding to any output fibre can be chosen, thereby making the whole switch wide-sense non-blocking. The packet then undergoes a third wavelength conversion, converting to a vacant wavelength at the output fibre before the packet is sent through a passive coupler on to the output fibre.

Hybrid node design

For separating SM and GST packets in the node, the State of Polarization (SOP) may be used [9]. Figure 5 illustrates the node design with polarization separation and combination mechanisms as well as an OPS

with an electronic packet switch buffer [10]. Polarization Beam Splitters (PBSs) physically separate SM and GST packets at the input interface. Neither guard band between GST and SM packets nor headers on GST packets and processing in the detecting node, are required. Because of polarization variations in the transmission fiber, Automatic Polarization Control (APC) at the node inputs is necessary. It will however operate on a relatively slow timescale of milliseconds to seconds, corresponding to the frequency of variations in the fibre's physical environments [25]. The viability of a transmission path through a model-node is experimentally verified in [9]. The physical viability of a node input interface, with combined wavelength conversion and header payload separation is proposed and experimentally verified in [26]. An alternative to segregating by polarization is to replace the PBS by a fast switch and send the class information in a packet header. The switch setting is then based on the header information [10].

The cross-connect in the node may be static (S-WRON) or dynamic (D-WRON). An S-WRON cross-connect is simple and reliable and does not need remote control, avoiding a control plane. Secondly, in S-WRONs, path setup delay is avoided enabling a GST without waiting time. For D-WRONs, however, connections can be established dynamically for relatively short time intervals, supporting traffic patterns that vary with a moderate frequency.

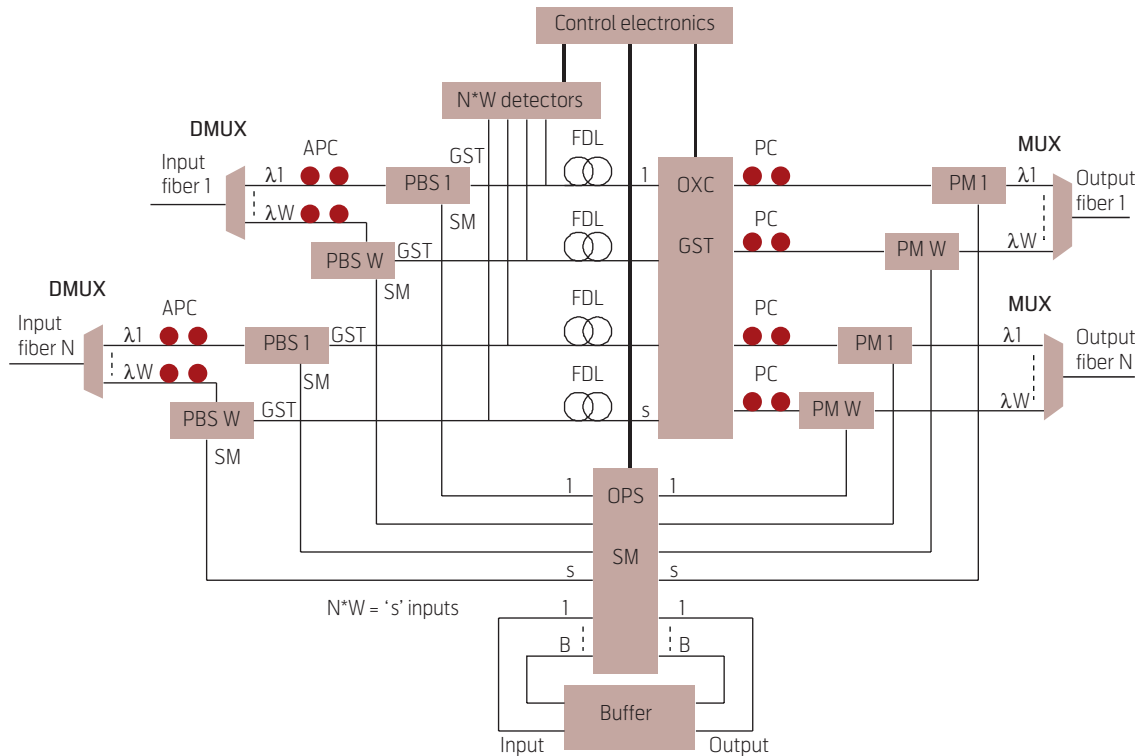


Figure 5 Node design with an optical cross connect (OXC) for the GST packets and an optical packet switch with shared buffering for the SM packets. The GST packets are separated using a PBS for each of the wavelength inputs, before they are delayed through FDLs, cross connected and scheduled at the output. The OXC may be static (hardwired) or dynamic, and may contain wavelength converters at the output for reducing blocking probability. FDL = Fibre Delay Line, W = number of link-wavelength. N = Number of link inputs, B = number of buffer inputs, $s = W * n$ = number of switch inputs, PM = Polarization Maintaining coupler, APC = Automatic Polarisation controller, PC = Polarisation Controller

Simulations

The aim of our simulations is to compare the performance of synchronous and asynchronous packet scheduling and fixed/variable length packet schemes. Furthermore, we analyse the basic characteristics of a hybrid node. We do this by characterising packet delay and packet loss. In our simulations we assume an optical packet switch using the wavelength dimension for contention resolution in combination with a minimised number of electronic buffer inputs, as described above and in [17]. Since the wavelength dimension can be used for contention resolution, the switch must be capable of doing wavelength-conversion. In this switch, full wavelength conversion is assumed, thus packets arriving at any input with any wavelength can be switched to any output and converted to any wavelength.

Since electronic buffering provides random access in the time domain, packets that are buffered are allowed clocking in and out of the buffer at random times. The switching matrix is assumed to be strictly non-blocking, thus it avoids internal blocking.

Simulation parameters

We assume an optical switch with ' N ' fibre inputs, each with ' W ' wavelengths. We have used $N * W$ independent traffic generators, generating fixed or variable length packets according to a Poisson arrival process, corresponding to a load of 0.8 normalised load. The traffic generators generate packets asynchronously. When simulating synchronous packet arrivals, packets are synchronised to discrete time slots at the input, thereby causing a slotted operation of the packet switch.

Both the input and the output traffic are uniformly distributed at the inputs and outputs, respectively. Since electronic buffering is assumed, the packets can stay in the buffer for an arbitrary period of time. The buffer size is set to 100 kB, which in our simulations is sufficient to avoid additional packet loss due to buffer overflow. Packets in the buffer are handled with FIFO priority. We expect that the future transport network will employ links with high channel (wavelength) counts, but have approximately the same node degree as we see in the networks of today. Therefore we have chosen to evaluate the performance of the switch in a backbone network with a

node-degree of '8' and a high link channel count, starting at 32. Performance is analysed varying both number of buffer interfaces and channel count. In the case of asynchronous operation, both fixed and variable packet length is evaluated. The fixed packet length is 500 bytes, while the variable packet length distribution is based on measurements of packet lengths in Internet [4]. It has the following distribution, expressed in number of packet occurrences: 40–44 bytes 62 %, 45–552 13%, 553–576 8 %, 577–1500 17 %.

In [24] the mean delay is shown to decrease quickly as the number of wavelengths increases. We find the same tendency. The longest mean delay observed is with 32 link channels, asynchronous packet arrival and fixed packet length, where a delay as low as $6.8 \cdot 10^{-4}$ ($\pm 7 \cdot 10^{-5}$) of the duration of a packet is found. With moderate packet lengths of 500 bytes, and a bit rate of 10 Gbit/s, this corresponds to a 270 ps mean delay. Transmission delay is therefore dominant and will be increasingly dominant as the link channel count increases. Packet loss is therefore the main topic in this analysis.

PLR performance electronic buffering; Fixed versus variable packet length, asynchronous operation

In our simulations, we have chosen to vary the number of buffer interfaces. For a given link channel count, this is the parameter to minimise since it is an important cost factor that directly governs the PLR. In Figure 6 we have compared the performance between using fixed and variable length packets when switching asynchronously, assuming an electronic Random Access in Time Memory (RATM) buffer. Results when varying both the number of wavelengths and the number of buffer interfaces are shown.

Results for both fixed and variable length packets (VLP) are shown for 32 and 256 wavelengths. For both wavelength counts there is a marginal performance difference between performance of fixed and variable packet lengths. The difference is close to or within the confidence interval. Since there is no synchronization mechanism at the inputs of the switch, packets do not arrive in distinct timeslots, causing little or no difference in performance when switching fixed or variable length packets. Since the performance difference is small, we see it as beneficial to use VLPs, and therefore for the rest of this paper VLPs are assumed when switching is asynchronous.

As expected, packet loss is shown to decrease as the number of channels is increasing when keeping the number of buffer inputs at a constant level. This is due to the wavelength dimension's impact on con-

tention resolution. This implies that as the number of channels increases, the number of buffer inputs can be decreased while maintaining an acceptable packet loss probability. When the number of link-wavelengths reaches 256, 15 buffer inputs are sufficient to reach a packet loss probability better than 10^{-6} . At 256 wavelengths, a PLR of $3.2 \cdot 10^{-5}$ is achieved without using buffers, hence with a sufficiently high number of wavelengths, an acceptable PLR can be reached without using buffering. This is confirmed in Figure 7, where results from simulations without buffering are plotted. We see clearly that the slotted scheme achieves a better PLR. We also observe that when the link channel count increases, the difference in packet loss performance between the two schemes increases. At the most demanding PLR threshold of 10^{-6} , buffering very few packets is necessary. At 128 channels per link, the highest packet loss probability is that of the asynchronous scheme, with a PLR of 10^{-3} . This implies that for this high channel count system, even this scheme will only need very limited buffering resources.

In Figure 8, the PLRs for the slotted and asynchronous scheme, for link channel counts of 32 and 128, are shown. When having 32 wavelengths in each link, a PLR of 10^{-6} can be achieved by using approximately 50 and 24 buffer interfaces in the asynchronous and slotted schemes, respectively. Thus, the slotted scheme requires less than half as many buffer inputs as the asynchronous scheme. Assuming 128 wavelengths in each link, these numbers decrease to 24 and 14, respectively. The drastic decrease in required buffer inputs, especially for the asynchronous scheme, indicates that the asynchronous approach becomes increasingly attractive with the increasing number of link wavelengths.




Figure 6 Packet loss ratio (Y-axis) as a function of number of buffer inputs (X-axis). Performance using fixed (broken lines), and variable length packets (continuous lines) are compared. 'W' is the number of wavelengths in each fibre, 'F' indicates fixed packet length. The error bars mark the limits of a 95 % confidence interval. Where only the upper limit is given, the lower limit is lacking. Higher precision can be achieved, making simulation time excessively long




Figure 7 Packet loss as a function of number of link channel counts. No buffering has been used in the simulations. Results are shown for the asynchronous and the slotted schemes. The error bars mark the limits within a 95 % confidence interval. When only the upper limit is given, lower limit is inadequate because it is less than 0. Higher precision can be achieved, making simulation time unduly long




Figure 8 PLR as a function of number of buffer interfaces. Results are shown for asynchronous variable packet length and slotted scheme with 32 and 128 wavelengths (W) per link




Figure 9 PLR and buffered packets delay of SM traffic as a function of GST traffic share. $B = 20$ buffer inputs; GST relative packet length = 100 times mean SM packet length; $1 \cdot 10^8$ packets simulated for each of the plotted values. The error bars indicate the 95 % confidence interval. Delay values are given relative to the duration of an SM packet with mean packet length

Performance of a hybrid node

The GST class is given absolute priority so that its quality is given solely from physical parameters in the WRON network, like transmission delay and signal quality. The quality of the SM (= BE) traffic does however depend on characteristics of the GST traffic, especially on the GST traffic share, which may typically vary. To illustrate how the GST traffic share in an OpMiGua node influences the performance of the SM traffic, we have simulated Packet Loss Ratio (PLR) and buffered delay for the SM traffic. We assume the node design shown in Figure 5 with a passive OXC and an OPS consisting of a strictly non-blocking switching matrix and a limited number B buffer inputs. Parameters are as for the previous simulations. The number of link wavelengths is $W = 32$. The switch experiences a uniform traffic distribution pattern from both the traffic classes. The GST packets are forwarded to a fixed destination and wavelength, thereby avoiding congestion with other GST packets.

We have chosen to evaluate the performance of the system when the GST packet length is fixed. The GST packet length is set to 100 times the mean SM packet length, making overhead caused by the GST packet's output reservation negligible. The long fixed-length GST packets correspond to an aggregation of several, e.g. IP-packets in a fixed length container.

We assume the buffering to be randomly accessible in the time domain, i.e. electronic buffering within a short time frame, and potentially optical memory within a longer time frame. Consequently, the packets can stay in the buffer for an arbitrary period of time. However they are scheduled to an output as soon as a wavelength to the destination becomes vacant. The Packet Loss Ratio (PLR) and delay results are plotted in Figure 9.

There are two counteracting effects impacting on the PLR. First, the GST packets are given absolute priority and as the GST traffic share increases, the lower priority SM traffic is increasingly penalized, resulting in an increase in PLR. On the other hand, as the GST traffic share increases, the remaining SM traffic share gets a relatively higher share of buffering resources available, resulting in a decrease in PLR. In Figure 9 we find that for low GST traffic shares of up to 30 %, the first effect is dominating, resulting in an increase in PLR. As the GST traffic share increases, the second effect becomes dominant and the PLR decreases significantly. For 90 % GST, no packet loss was observed for the $1 \cdot 10^8$ packets that were simulated. This indicates that given a specific PLR requirement, packet switch resources may be reduced if GST traffic shares are high.

The GST induced contention on SM packets increases with increasing GST traffic share. Consequently, the delay increases because buffered packets will wait longer before a wavelength becomes vacant. In the delay values in Figure 9, the clocking of the packets into the buffer is not accounted for. This corresponds to a cut-through solution, where packets can be clocked out of the buffer while they are still being clocked in. If such a solution is not available, one unit of delay must be added. The magnitude of the delay values indicates that packet reordering within a wavelength channel will occur. However, the bit-rate of a wavelength channel is typically in the order of 10 Gb/s. If the packets of an application are multiplexed evenly into the wavelength channel, e.g. for a 100 Mb/s application, every 100 packet in the wavelength will belong to the application. With mean delay values below 1.5 mean SM packet duration, packet reordering will therefore be of low probability.

Conclusion

There seems to be a convergence of the optical packet switching and optical burst switching concepts when it comes to packet/burst handling and contention resolution. Acceptable loss rates as well as QoS differentiation can be achieved in both concepts. In general, we see OBS as a first step towards a real dynamic network, with OPS as the ultimate goal. The hybrid network combines circuit switching with either packet- or burst switching. This concept allows guaranteed service with no packet loss and fixed delay, throughput efficiency of statistical multiplexing, and bypassing of transit traffic through a WRON network allowing smaller packet switches. A hybrid network can be implemented today, combining electronic packet switches with optical circuit switches. Furthermore it enables a migration path towards the technology found most beneficial for the future, be it a pure OPS network or the hybrid combination including an OPS as part of the hybrid node.

The optical packet switch designs discussed in this paper all rely heavily on wavelength conversion, and also on relatively large AWGs. Hence, for these designs to become attractive it is crucial that the mentioned components become cost effective.

Since the complexity of optical buffering is regarded as the chief argument against optical packet switching, we have compared buffering techniques and evaluated their performance in slotted and asynchronous mode. We conclude that optical packet switches using FDLs is an attractive solution if the number of wavelengths in each link is low and the bitrate at each wavelength is too high for electronic processing/buffering. In this case, the slotted scheme

appears to be advantageous, as it reduces the complexity of FDL based buffers.

However, if extensive use of wavelength division multiplexing is preferred, and the bitrate at each wavelength is sufficiently low to be processed electronically, electronic buffering appears advantageous. Only a very small number of electronic buffer interfaces are necessary to obtain sufficient performance, even when using the asynchronous scheme.

For high capacity optical packet switches, capable of handling several Tbit/s, it is our view that the limited increase in buffering required for the asynchronous scheme is less complex than the slotted scheme's requirements for network synchronisation, optical packet alignment, and electronic edge reassembly.

Therefore, the most viable optical packet switched network design could employ an asynchronous, variable packet length handling scheme, with packet switches using a combination of wavelength conversion and electronic random access memory for contention resolution. For even higher link wavelength count, we expect that buffering will not be required, and the network will become transparent.

Acknowledgments

Thanks to Telenor R&D and the Research Council of Norway (NFR) for financial support for this work.

References

- 1 *Next Generation Optical Networks for Broadband European Leadership*. March 15, 2005 [online] – URL: http://www.ist-nobel.org/Nobel/servlet/Nobel.Main?seccio=3_1
- 2 Vlachos, K et al. An Optical IM/FSK Coding Technique for the Implementation of a Label-Controlled, Arrayed Waveguide Packet Router. In: *IEEE Journal of Lightwave Technology*, 21 (11), 2617–2628, 2003.
- 3 Bjørnstad, W, Gauger, C, Nord, M (eds). Chapter 4, Optical Burst Switching and Optical Packet Switching. In: Extended Final Report of Cost Action 266, *Advanced Infrastructure for Photonic Networks*. Inkret, R, Kuchar, A, Mikac, B (eds). Croatia, University of Zagreb, Faculty of Electrical Engineering and Computing, 2003. [URL: www.ure.cas.cz/dpt240/cost266/index.html]
- 4 Claffy, K, Miller, G, Thompson, K. The nature of the beast: Recent traffic measurements from an Internet backbone. *Proc. INET'98*, 1998.

- 5 Technical Committee Telecommunications COST 266, workgroup 2, *Advanced Infrastructure for Photonic Networks, Optical Packet and Burst switching Report. (ed. S Bjørnstad) Chapter 4. (M Nord)*. Analysis of required switching times and identification of suitable switching technologies in optical packet switching. <http://www.ure.cas.cz/dpt240/cost266/> or available on request.
- 6 Yoo, M, Jeong, M, Qiao, C. A High Speed Protocol for Bursty Traffic in Optical Networks. In: *SPIE's All-Optical Communication Systems*, 3230, 79–90, 1997.
- 7 Nord, M, Bjørnstad, S, Gauger, C. OPS or OBS in the Core Network? *Proc. of ONDM conference*, Jan 2003, Budapest.
- 8 Packet switched hybrid optical networks. Bjørnstad, S, invited paper. In: *Proceedings of International Conference on transparent optical networks 2004 (ICTON 2004)*, July 5-8, Wroclaw, Poland.
- 9 Bjørnstad, S, Nord, M, Hjelme, D R. QoS differentiation and header/payload separation in optical packet switching using polarisation multiplexing. *Proceedings of ECOC 2003*, Mo3.4.6, 28–29.
- 10 Bjørnstad, S, Hjelme, D R, Stol, N. A highly efficient optical packet switching node design supporting guaranteed service. *Proceedings of ECOC 2003*, Mo4.4.5, 110–111.
- 11 Van Breusegem, E, Cheyns, J, Colle, D, Pickavet, M, Demeester, P. Overspill Routing In Optical Networks: a new architecture for future-proof IP over WDM networks. *OptiComm 2003*, Dallas (TX, USA), Oct. 2003.
- 12 Hansen, P B, Danielsen, S L, Stubkjaer, K E. Optical packet switching without packet alignment. *Proc. ECOC'98*, Madrid, Spain, 1998, 591–592.
- 13 Haas, Z. The 'staggering switch': An electronically controlled optical packet switch. *JLT*, May 1993, 925–936.
- 14 Guillemot, C et al. Transparent Optical Packet Switching: The European ACTS KEOPS Project Approach. In: *IEEE Journal of Lightwave Technology*, 16 (12), 2117–2133, 1998.
- 15 Sakamoto, T et al. Variable optical delay circuit using wavelength converters. In: *Electronic Letters*, (37) 7, 454–455, 2001.
- 16 O'Mahony, M J et al. The application of optical packet switching in future communication networks. In: *IEEE Communication Magazine*, (39) 3, 128–135, 2001.
- 17 Bjørnstad, S, Hjelme, D R, Stol, N. An optical packet switch design with shared electronic buffering and low bit rate add/drop inputs. In: *Proc. of IEEE ICTON*, Mo.C.6, Warsaw, April 2002.
- 18 Danielsen, S L, Joergensen, C, Mikkelsen, B, Stubkjaer, K E. Optical Packet Switched Network Layer Without Optical Buffers. *Photonics Technology Letters*, 10 (6), 1998.
- 19 Lavigne, B. All-optical 3R regeneration based on semiconductor technology. In: *Proc. of PIS*, TuA1, Cheu Island, Korea, July 2002.
- 20 Yao, S, Mukherjee, B, Yoo, S J B. Electrical Ingress Buffering and Traffic Aggregation for Optical Packet Switching and Their Effect on TCP-Level Performance in Optical Mesh Networks. In: *IEEE Communications Magazine*, 66–72, Sep 2002.
- 21 Yoo, S J B et al. Optical-Label based Packet Routing System with Contention Resolution in Wavelength, Time and Space Domains. In: *Proc. of OFC'02*, WO2, Anaheim, March 2002.
- 22 Callegati, F, Corazza, G, Raffaelli, C. Exploitation of DWDM for Optical Packet Switching with Quality of Service Guarantees. In: *IEEE Journal On Selected Areas in Communications*, 20 (1), 190–201, 2002.
- 23 O'Mahoney, M J et al. An optical packet switched network (WASPNET)– concept and realization. In: *Optical Networks Magazine*, 46–53, Nov/Dec 2001.
- 24 Chia, M C et al. Packet Loss and Delay Performance of Feedback and Feed-Forward Arrayed-Waveguide Gratings-Based Optical Packet Switches With WDM Inputs–Outputs. In: *IEEE Journal of Lightwave Technology*, 19 (9), 1241–1244, 2001.
- 25 Nicholson, G et al. Polarization Fluctuation Measurements on Installed Single-Mode Optical Fiber Cables. In: *IEEE Journal of Lightwave Technology*, 7 (8), 1989.
- 26 Nord, M, Bjørnstad, S, Nielsen, M L, Dagens, B. Demonstration of optical packet switching scheme for header-payload separation and class-based forwarding. *OFC 2004*, TuQ2, March 22–27.

Steinar Bjørnstad received his MSc in electronics from the Institute of Physics, University of Oslo in 1991, and his PhD in telecommunication from the Institute of Physical Electronics at the Norwegian University of Science and Technology (NTNU) in 2004. He worked with electronic design at Ring-Master 1991–1995 and at NERA Contec Design 1995–1996. He was then with Telenor R&D as Research Scientist on optical communication until 2005, when he started work as Associate Professor at the Institute of Telematics at NTNU, and as Senior System Architect on transport networks at Canal Digital. His interests include optical communication systems and networks, systems for television broadcasting, quality of service aspects of networking, discrete event modelling, electronic design and electromagnetic compatibility. He has managed several research projects and been active in several projects funded by the EU within the ACTS, IST and COST programs. He has authored and co-authored a number of papers, and has currently three patents pending.

email: steinar.bjornstad@iet.ntnu.no

Martin Nord received MSc degrees from the University of Essex and Institut National des Sciences Appliquées de Toulouse in 2000. He subsequently started working as a research scientist for Telenor R&D on optical networking. From 2002 to 2005 he pursued PhD studies on performance aspects and physical realisation of optical packet switched networks at the Research Center COM, Technical University of Denmark. His current research for Telenor R&D targets broadband network architectures for core-, metro- and access networks.

email: martin.nord@telenor.com

Torodd Olsen is a Research Scientist at Telenor R&D and has been working in the field of optical communications since 1986, including participation in several projects of the EU-funded ACTS and IST programme. He received his higher national diploma in electronics from Østfold Technical College, Sarpsborg, Norway in 1976. He is currently involved in project management and support, and recent contributions feature techno-economical studies of optical networking technologies.

email: torodd.olsen@telenor.com

Dag-Roar Hjelme is Chief Scientist at Optomed and Adjunct Professor at the Department of Electronics and Telecommunication, NTNU, both Trondheim. He earned his Siv.Ing. degree in Electrical Engineering from the Norwegian Institute of Technology (NTH), Trondheim, and his PhD in Electrical Engineering from the University of Colorado, USA. Prior to joining Optomed he was Senior Research Scientist at the Department of Photonics, SINTEF Electronics and Cybernetics, where he managed research projects on optical sensor technology and optical communication systems. After receiving his PhD he was Post-Doctoral Research Associate and later Staff Scientist at the University of Colorado, where his research focused on semiconductor lasers, integrated optics, microwave optics, and fiber optic systems. He is the author and co-author of more than 30 refereed journal papers and 70 conference presentations. He is co-inventor of six international patents. Dr. Hjelme is a member of the Norwegian Academy of Technological Sciences.

email: Dag.R.Hjelme@optomed.no

For a presentation of Norvald Stol, please turn to page 146.

Revealing the wide spectrum installed fibre cable loss in Telenor's network

SVEND HOPLAND



Svend Hopland is Chief Engineer at Telenor Nordic Fixed

By performing extensive measurements on installed single mode fibre cables in Telenor's network, we have revealed statistical values for fibre attenuation and splice loss in cables with standard G.652 fibres from different time periods. Data are provided at six different wavelengths in the 1241–1642 nm wavelength range. A mathematical model has been derived for the continuous full spectrum fibre attenuation and used to fit the field data. As a result, total link loss statistics for the installed cables have been calculated at the 16 different CWDM wavelengths.

The main observations are that average fibre attenuation, splice loss and their standard deviations have decreased significantly at all wavelengths since the late 1980s. In particular, fibre attenuation at the water peak at 1383 nm has decreased dramatically. In cables installed after 2000, the fibre attenuation and splice loss are very low in the entire wavelength range 1260–1675 nm.

The present levels of hydrogen due to corrosion of armour wires in a limited number of old underwater cables without hydrogen barrier do not result in significant additional loss in major parts of the transmission bands. Long term field measurements and high temperature accelerated tests have shown that the hydrogen levels in these cables have stabilized after 8–10 years and will decrease during the next 20–30 years.

Extensive testing of duct cables have shown that temperature induced loss will not occur in installed cables for wavelengths up to 1675 nm. Additional loss due to tension and crush will not occur provided the cables have been installed properly. The field measurements have provided evidence for the test results.

1 Background

1.1 Rationale for undertaken work

The extended use of wavelengths in the range 1260 nm up to nearly 1700 nm for WDM optical transmission has received considerable attention in the later years. Responding to this interest, the leading standardisation body ITU-T has defined a continuous optical spectrum from 1260 nm to 1675 nm as listed in Table 1. Up to now, the ITU-T does not consider the U-band to be suitable for transmission, due to expected high losses.

In so-called Coarse Wavelength Division Multiplexing (CWDM), where the spacing between evenly distributed channels is approximately 20 nm, one certainly foresees to use the whole span approved by ITU.

As for Dense Wavelength Division Multiplexing (DWDM), wavelengths through the S-, C- and L-band are for the moment considered to be optimal, but clearly also other wavelength bands can be taken into use.

Figure 1 shows the intended wavelength allocation of CWDM and DWDM relative to a typical optical fibre attenuation spectrum.

Band designation	Wavelength [nm]
O	1260 – 1360
E	1360 – 1460
S	1460 – 1530
C	1530 – 1565
L	1565 – 1625
U	1625 – 1675

Table 1 ITU-T band designations for continuous optical spectrum from 1260 nm to 1675 nm

The majority of installed cables up to now, at least as far as the traditional operators including Telenor are concerned, have contained standard G.652-fibres. In the Telenor case 100 % have been standard G.652 fibres.

Optical fibre cables with standard G.652-fibres have cut-off values of approximately 1260 nm, and are thus able to transmit this large wavelength span. The fibres also have dispersion values of some positive magnitude from around 1300 nm and upwards, which

is necessary to reduce non-linear effects in DWDM-systems. It is thus possible to use this fibre type for DWDM in a very broad wavelength range.

Fibre production methods as well as fibre cable constructions and fibre splicing techniques have improved since the late 1980s, resulting in a historical change of the installed fibre cable loss spectrum. However, in practice, a link is only routinely measured shortly after installation at one specific wavelength, usually at 1550 nm using OTDR equipment. Hence the magnitude of fibre attenuation and splice loss are provided at reasonably low cost and gives also an indication of the link loss in the C-band (1535–1565 nm). Obviously, other parts of the spectrum remain unknown.

This is the reason why very little wide spectrum data exist worldwide from actual field installations. Telenor has recently performed extensive field measurements with corresponding data analysis and curve fitting procedures to provide wide spectrum data on installed cables from different time periods. This work has for the first time revealed accurate and reliable field data [1] [2] [3] [4] [5].

Supplementary to the field measurements, Telenor have performed extensive fibre cable testing to investigate possible additional losses due to installed cable environmental conditions such as temperature, tension and crush [6].

1.2 Fibre and cables used by Telenor

The fibres used in Telenor's cables are Standard G.652 with fibre cut-off values in the 1190 nm – 1330 nm range. The installed cables consist of fibres in stranded loose tubes or in a central steel tube surrounded by different kinds of protection and tension members to form duct, aerial and underwater cables, respectively.

1.3 Factors affecting the spectral loss

In principle, the total fibre attenuation in an installed cable link in the 1200–1750 nm wavelength region consists of contributions originating from fibre production as well as contributions originating from the cable itself. The intrinsic fibre loss is expected to dominate the overall fibre cable loss. After installation, however, environmental factors such as temperature, crush and tension, or hydrogen intrusion, may cause additional attenuation in the cabled fibres. Fibre splices, including the bending of fibres in organizers will also add to the total loss.

1.3.1 Fibre intrinsic attenuation

The fibre intrinsic attenuation is composed of three basic mechanisms:




Figure 1 Typical fibre attenuation spectrum and relative to the use of CWDM and DWDM

- Rayleigh scattering loss at the lower wavelength end;
- Water (OH^-) absorption loss with its peak at 1385 nm;
- Infrared (IR) absorption loss at longer wavelengths.

The Rayleigh loss is caused by scattering of light by the fibre core glass molecules. It is dominant at lower wavelengths and decreases with increasing wavelength. The so-called “water peak” is caused by OH^- hydroxyl ions originating from the fibre production process and has its maximum value at 1385 nm. The infrared absorption is caused by glass molecular vibrations within the fibre core. It increases with increasing wavelength and becomes dominant at longer wavelengths.

Each of these contributions exhibit variations due to variations in fibre production parameters. Specifically, the water peak may vary significantly from one fibre manufacturer to another, from production to production, and even during the same production. Moreover, it has shown a historical evolution towards generally lower values.

1.3.2 Cable induced attenuation

Bend loss

During installation and through the cable lifetime, environmental factors such as temperature, tension and crush may cause additional attenuation in the cabled fibre. The origin of this loss is different types of bending of the fibres inside the cable structure. Generally, all bending loss increases with increasing wavelength and moreover, the different bending types have different wavelength dependencies.

Hydrogen induced losses

In cables with metallic armouring such as underwater cables, hydrogen will be generated due to corrosion of the armouring wires and hydrogen molecules may diffuse into the cable core and eventually into the fibres. As a result, the fibre attenuation increases according to a specific wavelength dependent absorption pattern. This can be prevented if a hermetic hydrogen barrier is used between the armouring and the fibres. Except for a limited amount of underwater cables in the late 1980s, Telenor has installed underwater cables including hydrogen barriers.

Fibre splices in cable

Fibre fusion splices along the installed cable including the bending of fibres in splice organizers will cause additional wavelength dependent loss.

2 Measurement methods used for wide spectrum

2.1 Spectrum analyser

In the field measurements as well as in the cable test measurements, we have used a white light source and a spectrum analyser to measure the continuous spectral loss curve of individual fibre paths from 1200 nm to 1750 nm.

This method integrates all losses between input and output connector and it is not possible to distinguish between fibre attenuation and splice loss. However, the detailed form of absorption peaks such as OH and molecular hydrogen may be revealed. Spectral measurements are rather difficult and time consuming to perform. They require careful data analysis, and are not highly accurate. Both ends of the cable have to be accessed during measurement. In practice, the maximum fibre length that could be measured by this method is approximately 8–10 km.

2.2 Optical Time Domain Reflectometry (OTDR)

An OTDR measurement is a single wavelength measurement utilizing the signature of backscattered laser light along the fibre length. Only one end has to be accessed during measurements, and fibre lengths up to 100 km can easily be reached. Data analysis may be automated and fibre attenuation length distribution and splice losses can easily be separated and measured with high accuracy.

If multiple OTDR wavelengths are used to measure the same fibre length, it is possible to deduct the full fibre attenuation spectrum for individual fibre lengths between splices using a mathematical curve fitting method [2]. It is however required that the OTDR

wavelengths are distributed in a certain strategic pattern over the full optical spectrum.

3 The basic field measurements

3.1 Methodology

In the field we performed OTDR measurements on six different wavelengths: 1241 nm, 1310 nm, 1383 nm, 1550 nm, 1625 nm and 1642 nm. A coupling fibre of approximately 1 km was used at the measurement end. Normally, we would measure cable sections that could be reached end-to-end at all wavelengths. The sections were measured from both ends in order to calculate the splice losses. As a general rule we only used individual fibre attenuation data measured on fibre lengths longer than 1 km in order to obtain high measurement accuracy.

For each section, fibre attenuations for individual fibre lengths between splices as well as splice loss were analysed and calculated for all wavelengths. Subsequently, loss statistics for fibre attenuation and splice loss were produced.

The cables included in the measurements are basically duct cables and some underwater cables.

In the late 1980s, Telenor installed some underwater cables without hydrogen barrier, which resulted in some limited hydrogen ingress into the fibres. The hydrogen induced losses stabilized after a few years at low levels at 1310 nm and 1550 nm [6] [7] [8]. Using cables with hydrogen barriers since 1992, all H₂-induced peaks have been removed, and the loss values at all wavelengths are expected to be similar to duct cables. Investigation of hydrogen levels and field measurements of spectral loss in old underwater cables are treated separately in Chapter 5. We have not included any old underwater cables in the basic field measurements.

3.2 Results

3.2.1 Fibre attenuation

In Figure 2, the fibre attenuation average values and standard deviations are summarised. The most important observations are:

- Average values decrease at all wavelengths with time, in particular at the water peak.
- In 2003 cables the average value at 1383 nm is 0.346 dB/km, which is nearly as low as the average value at 1310 nm (0.336 dB/km).




Figure 2 Fibre attenuation and standard deviation for different time periods

- The average value at 1625 nm is 0.023 dB/km higher than at 1550 nm in cables from before 1990, and only 0.013 – 0.014 dB/km higher in the newer cables.
- The average value at 1642 nm is 0.016 dB/km higher than at 1625 nm and in cables from before 1990, and only 0.012 – 0.014 dB/km higher in the newer cables.
- Not surprisingly, the standard deviation is much larger at the water peak than at other wavelengths.




Figure 3 Splice losses and standard deviations versus wavelength for different time periods

3.2.2 Splice loss

Figure 3 shows the average splice loss and standard deviation versus wavelength for cables installed before 1990 and around 2000. Splice loss in cables installed in 2003 has not been measured yet, but we may assume that splice loss is of the same magnitude or lower than in 2000 cables.

Important observations are:

- Splice loss has generally decreased at all wavelengths from around 0.1 dB to less than 0.05 dB from 1990 to 2000.
- In newer cables, splice loss decreases with increasing wavelength up to 1625 nm. A very small and insignificant increase is noted at 1642 nm.

3.2.3 Total loss

The total link loss is calculated by adding fibre attenuation and splice loss. A typical distance between splices may be from 1 km up to 4–6 km, depending on the cable type and method of installation. Therefore, different distances between splices should be taken into account when calculating the total loss.




Figure 4 Fibre and splice loss for different time periods with 2 km between splices

Depending on the number of cable lengths spliced together, the fibre attenuation and splice loss in a link are averaged to a certain extent. Therefore, the predicted link loss value will be somewhere between the average value and the average value plus the standard deviation.

The upper limit (worst case) of the total loss at each wavelength, is calculated using the equation:

$$\text{Total loss } (\lambda) = \text{ave}(\text{fibre att.}) + \text{ave}(\text{spl.loss}) + \text{stdev}(\text{fibre att.}) + \text{stdev}(\text{spl.loss}) \quad (3.1)$$

The results based on equation (3.1) are illustrated in Figure 4.

4 Extension of basic field measurements

4.1 Methodology

The previous section have given statistical results for fibre attenuation at fixed wavelengths; however, fibre attenuation values in the gaps between the OTDR-wavelengths still remain unknown. To solve this problem we have tried to find a mathematical model for the full spectrum fibre attenuation in the field. We have investigated a number of spectral field measurements performed on our installed cable base; in particular, we have isolated and studied in detail the water peak absorption in order to reveal the detailed form of the water peak.

The derived mathematical model for total fibre attenuation has been applied on the 6-wavelength data sets and full spectrum fibre attenuation has been calculated for all measured fibre lengths using an automated mass curve fitting procedure. Accordingly, we have been able to calculate statistical parameters for fibre attenuation at any wavelength of interest for cables installed in different time periods. An empirical model has been used to estimate splice loss versus wavelength in order to include them in link loss calculations.

Using this method we are able to present statistical parameters for total link loss including fibre attenuation and splice loss at all the CWDM wavelengths on installed cables with standard G.652 fibres from different time periods. We can also present estimated fibre attenuation values in the region 1650–1675 nm based on the curve fittings. This method can be used to give information of fibre attenuation and link losses at any wavelength of interest in cables installed at different time periods. The method will be described in the following chapters.

4.2 Mathematical model for fibre attenuation

The total fibre attenuation in an installed cable length in the wavelength region 1200–1700 nm consists of contributions originating from fibre production as well as contributions originating from the cable itself.

The spectral loss of cabled fibre can be expressed as:

$$\alpha = \alpha_R + \alpha_{OH} + \alpha_{IR} + \alpha_{UV} + \alpha_{H2} + \alpha_{bend} \quad (4.1)$$

where α_R is the Rayleigh scattering loss, α_{OH} is the OH⁻ absorption loss, α_{IR} is the infrared absorption loss, α_{H2} is hydrogen induced loss caused by molecular hydrogen ingress from cable materials such as armouring, and α_{bend} is bend loss due to cable environmental conditions such as low temperature, fibre tension or impact/crush in the cable.

The Rayleigh loss is proportional to λ^{-4} and the following relationship is assumed:

$$\alpha_R = A \cdot \frac{1}{\lambda^4} \quad (4.2)$$

The infrared molecular vibration absorption is given by:

$$\alpha_{IR} = B \cdot e^{-\frac{C}{\lambda}} \quad (4.3)$$

where B and C are constants.

α_{UV} , the absorption loss by the UV tail of the germanium doping can be ignored because it is very small.

Hydrogen intrusion only exists in a limited number of Telenor's old underwater cables, and the OTDR measurements used for curve fitting have not included such cables. Therefore, α_{H2} can be ignored.

Bend loss increases with increasing wavelength and is not easy to distinguish from other fibre loss. However, field measurements by OTDR at 1625 nm and 1642 nm have indicated that bend losses are very small in Telenor installed cables [1]. In addition, extensive testing of the cables has confirmed that bend loss due to low temperature or fibre tension / cable crush is not likely to occur in Telenor's field installations [6]. Therefore, α_{bend} can be ignored.

4.3 Model for water peak derived from field measurements

The OH absorption spectrum is an important loss factor, especially in the older installations where the peak value is high. We have measured spectral loss curves from 1200 nm to 1750 nm in a number of installations of different age and different cable types and the detailed spectral characteristics of the water

peak have been investigated. An OH^- spectrum derived from a spectral measurement of a duct cable installed in 1988 is shown in Figure 5.

We have found that all OH^- absorptions revealed in this investigation can be adequately represented by three Gaussian absorption components, giving an accurate description of the total absorption.

The general expression for the OH^- absorption is given by:

$$\alpha_{OH} = \alpha_{gauss1} + \alpha_{gauss2} + \alpha_{gauss3} = A \cdot e^{-\frac{(\lambda-I)^2}{2B^2}} + C \cdot e^{-\frac{(\lambda-G)^2}{2D^2}} + E \cdot e^{-\frac{(\lambda-H)^2}{2F^2}} \quad (4.4)$$

where A, B, C, D, E, F, G, H and I are constants.

We have isolated a number of OH^- absorptions taken from cable installations from different time periods and normalized them in order to make a qualitative and quantitative comparison. As shown in Figure 6, different curve forms have been revealed. However, it is evident that the differences mainly occur at the edges of the absorption band in the wavelength region 1410–1480 nm. For small OH^- peaks, the effect of these variations on the total loss will be minor.

We have selected two models, which are representative for the outer limits of the variations. *OH-model 1* (= duct 1988) represents the lowest values all through the OH^- absorption, while *OH-model 2* (= duct 1999) represents the highest values. These two OH^- models will be referred to in later calculations.

4.4 Curve fitting method

The OTDR values at 1241 nm and 1310 nm are selected to make a fit to the Rayleigh scattering using equation (4.2). This is possible because the OH^- absorption does not contribute at 1241 nm and 1310 nm, and the IR absorption is so small that it can be ignored.

Subsequently, the Rayleigh scattering contributions are calculated for the remaining wavelengths 1383 nm, 1550 nm, 1625 nm and 1642 nm – and subtracted from the original values. From the difference values, the 1550 nm, 1625 nm and 1642 nm values are selected to make a fit to the IR-absorption using equation (4.3). This is possible because the Rayleigh contributions have been removed and the OH^- absorption does not contribute here.

Next the IR absorption contribution is calculated at 1383 nm, and subtracted from the result of the previous step. The remaining value is now clean and can be used to make a fit to the OH^- absorption using




Figure 5 Water peak in 1988 duct cables




Figure 6 Normalized field OH^- absorptions

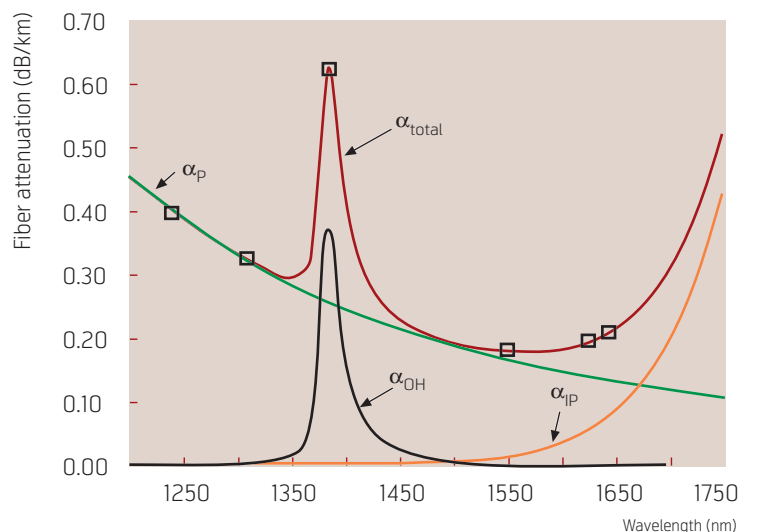


Figure 7 Single curve fit for fibre attenuation




Figure 8 Mass curve fits for fibre attenuations in cables installed before 1990 using OH-model 1




Figure 9 Mass curve fits for fibre attenuations in cables installed around 2000 using OH-model 1




Figure 10 Mass curve fits for fibre attenuations in cables installed 2003 using OH-model 1




Figure 11 Mass curve fits for cables installed around 2000 using OH-model 2

equation (4.4) and OH-model 1 or OH-model 2. Figure 7 illustrates the different steps in the curve fitting procedure.

4.5 Results

4.5.1 Mass curve fitting of fibre attenuation

Figure 8, Figure 9 and Figure 10 show the results of the mass fittings for cables installed before 1990, around 2000 and 2003, respectively. The calculated spectra have been plotted from 1200 nm to 1675 nm and we have chosen OH-model 1. Figure 11 shows the results for cables installed in 2000 when using OH-model 2. Comparing the spectra in Figure 9 and Figure 11, we observe only minor differences in the wavelength region 1410–1500 nm.

As a result of the curve fitting method described here, we are able to calculate statistical parameters for fibre attenuation in the field at any wavelength. A case of great interest is attenuation statistics for CWDM wavelengths.

Table 2 shows the calculated statistics at the CWDM-wavelengths for cables installed around 2000 using OH-model 1. For comparison, the results of calculations using OH-model 2 are also included. We note the small differences in the statistical values in the region 1410–1500 nm.

In order to compare the curve fitting results with the actual field measurements, we have calculated the fibre attenuation statistics for the OTDR wavelengths using OH-model 1. The results are shown for cables installed 2000 in Table 3. We note that the calculated values are in excellent agreement with the actual measured values at the OTDR wavelengths.

We also note that the estimated average attenuation values in the longer wavelength region from 1650–1675 nm are in the range 0.23–0.26 dB/km for cables installed 2000, which are very low values.

4.5.2 Splice loss

The measured average splice loss and standard deviation versus wavelength for cables installed before 1990 and around 2000 or later have been shown earlier in Figure 5. We have used this graph to find the splice loss average and standard deviation at wavelengths other than the OTDR wavelengths, assuming a linear relationship between the wavelength measurement points.

4.5.3 Total loss

Equation (3.1), shown earlier, is used to calculate total loss in a link at the CWDM wavelengths in the case of 2 km or 4 km between splices. Table 4 shows

Cables installed 2000				
CWDM edge wavelength (nm)	Fiber attenuation (dB/km) OH-model 1		Fiber attenuation (dB/km) OH-model 2	
	Average	Stdv	Average	Stdv
1261	0.396	0.018	0.396	0.018
1281	0.371	0.017	0.371	0.017
1301	0.349	0.016	0.350	0.016
1321	0.329	0.015	0.330	0.016
1341	0.313	0.01	0.315	0.016
1361	0.307	0.020	0.314	0.025
1381	0.499	0.196	0.500	0.200
1401	0.371	0.098	0.389	0.115
1421	0.285	0.037	0.305	0.054
1441	0.253	0.020	0.272	0.036
1461	0.229	0.012	0.249	0.026
1481	0.215	0.010	0.228	0.017
1501	0.206	0.010	0.211	0.012
1521	0.199	0.010	0.201	0.010
1541	0.194	0.010	0.195	0.010
1561	0.192	0.010	0.192	0.010
1581	0.192	0.010	0.192	0.010
1601	0.196	0.010	0.196	0.010
1621	0.204	0.011	0.205	0.010

Table 2 Fibre attenuation statistics at CWDM wavelengths for cables installed around 2000

Cables installed around 2000				
Wavelength (nm)	Fitted values (dB/km)		Measured values (dB/km)	
	Average	Stdev	Average	Stdev
1241	0.422	0.019	0.420	0.020
1310	0.340	0.016	0.342	0.017
1383	0.509	0.209	0.507	0.207
1550	0.193	0.010	0.194	0.010
1625	0.207	0.010	0.206	0.010
1642	0.221	0.011	0.221	0.012
1650	0.228	0.012	NA	NA
1660	0.240	0.014	NA	NA
1670	0.254	0.017	NA	NA
1675	0.263	0.020	NA	NA

Table 3 Fibre attenuation statistics at OTDR wavelengths for cables installed around 2000

Cables installed 2000						
CWDM center wavelength (nm)	Fiber attenuation (dB/km)		Total loss (dB/km) 2 km between splices		Total loss (dB/km) 4 km between splices	
	Average	Upper limit stdv	Average	Upper limit stdv	Average	Upper limit stdv
1271	0.396	0.414	0.421	0.467	0.409	0.441
1291	0.371	0.388	0.396	0.440	0.383	0.414
1311	0.349	0.365	0.373	0.416	0.361	0.391
1331	0.329	0.344	0.353	0.394	0.341	0.369
1351	0.313	0.329	0.337	0.377	0.325	0.353
1371	0.499	0.694	0.521	0.741	0.510	0.718
1391	0.509	0.718	0.531	0.765	0.520	0.742
1411	0.371	0.469	0.393	0.514	0.382	0.491
1431	0.285	0.322	0.307	0.366	0.296	0.344
1451	0.253	0.273	0.274	0.317	0.263	0.295
1471	0.229	0.242	0.250	0.284	0.240	0.263
1491	0.215	0.226	0.236	0.267	0.226	0.247
1511	0.206	0.216	0.226	0.256	0.216	0.236
1531	0.199	0.209	0.219	0.249	0.209	0.229
1551	0.194	0.205	0.214	0.243	0.204	0.224
1571	0.192	0.202	0.211	0.240	0.201	0.221
1591	0.196	0.206	0.214	0.243	0.205	0.225
1611	0.204	0.215	0.223	0.251	0.213	0.233

Table 4 Loss statistics at CWDM-wavelengths for cables installed 2000

the results for cables installed in 2000. The value representing each centre wavelength is selected as the highest value in the corresponding interval.

5 Hydrogen induced loss in underwater cables

5.1 Background

Norway has a very long coastline, which makes the installation of underwater fibre optic cables very attractive. Telenor installed the first fibre optic sea cable in 1983 in Sognefjorden on the west coast of Norway, and the first cable in a freshwater lake was deployed in 1985 in Norway's largest lake, Mjøsa near Lillehammer. Up to now, Telenor has installed more than 2500 km of underwater cables along the Norwegian coast and more than 1000 km in lakes.

During the early installation period from 1986 to 1990, we had no hydrogen barrier in the cables to prevent hydrogen ingress in the fibres due to corrosion of the armour wires. From the early 1990s, cost effective solutions with hydrogen barriers such as hermetic fibres and hydrogen absorbing materials were available, and were included as a standard in our cables. Later on, other methods were introduced,

and since 1998 we have used an underwater cable construction with fibres inside a hermetically sealed central steel tube.

5.1.1 Old underwater cable construction

Our old underwater cable construction is basically very simple, with a fibre core consisting of fibres in stranded loose tubes or a slotted core element, surrounded by an inner sheath of low-density polyethylene. Then follows the armour wire layer(s) with a rubber asphalt/bitumen filling compound, and finally an outer sheath of high-density polyethylene, as shown in Figure 12. Different types of armour are used for different environments. Typically, SA and DA armour are used for sea installations, while SSA armour is used in lakes.

The cable core was identical to our standard duct cable. In these underwater cables, no hydrogen barrier is included and hydrogen from the armour wires is allowed to diffuse freely in and out of the cable construction.

5.1.2 Newer underwater cable constructions

From 1991–92, Telenor introduced hermetically coated fibres in the old cable construction. The hermetic ceramic coating is applied directly to the fibre




Figure 12 Old underwater cable constructions

surface during fibre production and prevents hydrogen from entering into the fibre core. Hermetic coatings have proven to be a lifetime full proof protection against hydrogen.

Figure 13 shows the underwater cable construction that we have used since 1998. Here, the unprotected fibres are placed inside a hermetic steel tube in the centre of the cable. The steel tube has been longitudinally laser welded to provide hermetic seal. The steel tube also gives better mechanical protection to the fibres in case of external crush and tension to the cable. The steel tube has proven to give a lifetime full proof hydrogen protection to the fibres.




Figure 13 Underwater cable construction used since 1998

5.2 Investigating old underwater cables

Telenor, along with many other cable owners has been concerned about hydrogen ingress in their installed cables without hydrogen barrier. As a consequence, Telenor started late 1988 an extensive measuring program on installed submarine cables in order to reveal the hydrogen level present in the cables. The initial results indicated low levels of hydrogen during the first few years after installation [7] [8]. However, it was not possible at that time to evaluate the long-term development of the hydrogen level.

Telenor continued to perform hydrogen induced loss measurements on selected installed fibre optic underwater cables. In addition, we performed long-term accelerated tests on armoured cables.

Considering the large number of field measurements over more than 10 years, and the results of the accelerated tests, we are in a position to give reliable indications of the long term behaviour of hydrogen induced loss.

5.2.3 Installed cable routes

The cable network along the coast connects the main populated areas and consists partly of underwater cable and partly of land cable. The overall average

underwater route length is 6.9 km, and the maximum sea depths typically vary between 50 and 800 m, with the deepest route of 1300 m. The sea bottom topographies along our submarine cable routes are large mountain formations, sharp and high ridges of several hundred metres, leads and troughs, as well as flat parts. Nearly 100 % of the underwater cables are surface laid due to the high cost of burial.

5.2.4 Field measurements

We decided early to measure mainly underwater cables installed at sea, since they are heavier armoured, situated in the harshest environments and also prone to mechanical damage due to fishing activities etc. The underwater cables were selected from different parts of the coast, with different maximum depths and route topographies. Also, we selected single armoured cables (SA) and double armoured cables (DA) in representative proportions. In total, we have measured 35 cables from the early installation period, which constitute nearly 30 % of all our submarine cables without hydrogen barrier. The selection is therefore in all respects representative of our marine installations.

For these early installed cables we have measured the hydrogen level 1–2 times a year over a period of more than ten years. In addition, from the later installation period we have measured ten cables with




Figure 14 Long term loss in old single armoured (SA) and double armoured (DA) underwater cables measured with OTDR 1 at 1238 nm

hydrogen barriers such as hydrogen absorbing materials and hermetic fibres for comparison. Some of these cables have been regularly measured over a period of five years.

Hydrogen molecules may diffuse in and out of the cable construction and the observed hydrogen level in the cabled fibre is determined by the equilibrium condition between inward and outward diffusion. This level may change with time if the hydrogen generation rate in the armour layers changes.

Hydrogen molecules interact with the glass molecules and a number of well-defined absorption peaks appears in the fibre attenuation spectrum. The main peak appear at 1244 nm and several smaller peaks may be observed at different longer wavelengths.

To monitor the hydrogen level over time we have used two different special OTDR-modules with slightly different centre wavelengths, which both are close to the peak wavelength of the hydrogen-induced loss at 1244 nm. With the reference module (OTDR 1 = 1238 nm), we are able to detect approximately 50 % of the 1244 nm peak height, while the other module (OTDR 2 = 1241 nm) detects approximately 70–80 % of the peak height. They are used simultaneously in order to compare and verify changes in fibre

losses. Both modules are checked for stability of centre wavelengths and half widths prior to all field measurements. We have measured several fibres in each cable and average values are used to characterise the hydrogen level in the cable.

In all field measurements the fibres were also measured at the 1310 nm and 1550 nm wavelengths.

Figure 14 shows the long-term fibre loss in our DA and SA underwater cables. Each cable shows an individual time dependent pattern. After 6–7 years, the hydrogen induced loss has more or less stabilised. The average value for the DA cables is clearly higher than the corresponding level for the SA cables. The ratio between average hydrogen induced loss in SA and DA cables after 6.5 years is found to be 1 : 2.58. This is in good agreement with the ratio of the surface areas of the armour wires for the SA cables and the DA cables, which is calculated to be 1 : 2.34. Thus, the hydrogen generation rate seems to be proportional to the surface area of the armour wires.

Hermetic fibres are considered to be a full proof lifetime protection against hydrogen levels of considerably higher magnitude than what is present in our cables. For comparison, measurements on these

cables are included in Figure 14. No sign of loss due to hydrogen absorption is visible in the hermetic fibres.

Figure 15 shows the corresponding long term loss at 1550 nm measured in SA cables and DA cables. Here the hydrogen-induced loss has been mapped into the 1550 nm wavelength. It should be noted that normal fibre loss in old cables was not always close to 0.20 dB/km and also other minor loss could be present at 1550 nm. However, we can deduce from Figure 20 that the hydrogen induced loss 6–7 years after installation at 1550 nm is in the region 0.01 – 0.02 dB/km and 0.02–0.04 dB/km for SA cables and DA cables, respectively.

5.2.5 Accelerated temperature tests

In order to study long-term hydrogen effects more in detail, a temperature controlled seawater bath was constructed. Cable samples were deployed in seawater and subjected to a temperature of + 60° Celsius. Accordingly, the hydrogen level rises to much higher values than in the field measurements. Some selected results are shown in Figure 16. Each sample shows an initial rise of the hydrogen level to a maximum value, followed by a long decaying period. SSA cables have less armour than the SA cables.

It is difficult to estimate the exact time before maximum value is reached because of a temporary variation in temperature in this period. However, a rise




Figure 15 Long-term loss in old single armoured (SA) and double armoured (DA) underwater cables measured at 1550 nm

time of 50–60 days for both SA and SSA cables is a reasonable approximation. The rise time does not seem to be strongly dependent on amount of armour.

The time needed to reduce the hydrogen level to half of its maximum value is approximately 300 days on average for the SA cables and approximately 100 days on average for the SSA cables. Clearly, the decay time increases with increasing amount of armour.




Figure 16 Fibre loss and temperature versus time in SA and SSA cables heated in seawater




Figure 17 Spectral loss in old underwater cables

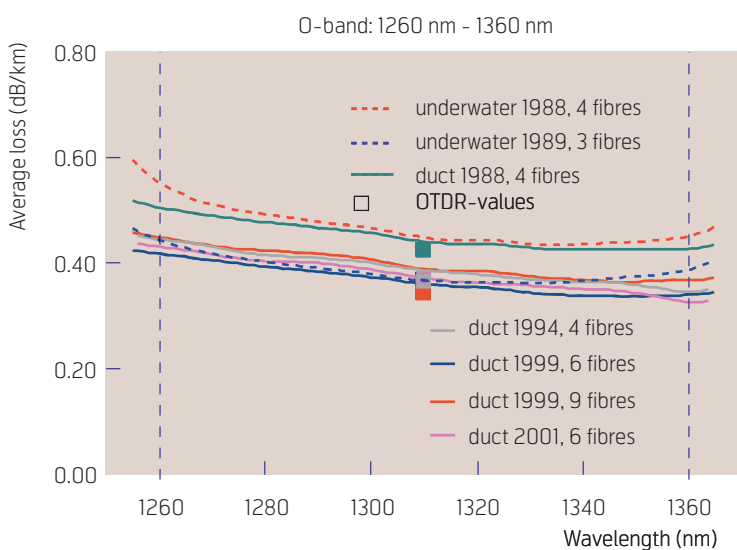


Figure 18 Spectral loss of fibre cables in the O-band




Figure 19 Spectral loss of fibre cables in the S-, C-, L- and U-band

We observe that the decay time is typically 5–6 times longer than the rise time for SA cables. For DA cables a higher ratio would be expected.

5.2.6 Expected long term field development

The high temperature experiment results are in agreement with the field measurements, which have shown increasing but stabilised hydrogen levels after 6–8 years, corresponding to the equivalent rise time of 50–60 days at the elevated temperatures. We can thus calculate a time reduction factor of 0.10 – 0.16 year/day for the high temperature experiment.

As a consequence, the time required for hydrogen levels in the field to decrease to half their maximum values, is estimated at more than 30 years for SA cables (decay-time · time reduction factor), and even longer times would be expected for DA cables.

In view of the test results, the installed cables will after 6–8 years go through a period of stable and decreasing hydrogen levels during the following 10–20 years. Some measurements in 2002 of spectral loss in 14 year old underwater cables (shown in Figure 17) have confirmed this prediction. Further field measurements will be performed.

5.3 Spectral loss signatures in old underwater cables

Figure 17 shows recently measured spectral loss (2002) for an underwater cable installed in 1988. We register the characteristic main molecular hydrogen peak at 1244 nm as well as two smaller H₂ related peaks at 1590 nm and 1700 nm, respectively. We also observe the OH⁻ peak at 1385 nm, which is not linked to molecular hydrogen as stated earlier. All these peaks will affect the different bands as outlined in more detail in the following.

Figure 18 shows the spectral loss in the O-band (1260 – 1360 nm). Due to the main H₂ absorption at 1244 nm and the rather high OH⁻ absorption at 1385 nm, the attenuation values for the old underwater cables increase slightly near the band edges in the O-band compared to duct cables, while in the central region of the band, the spectral loss is very similar to the duct cables.

In the E-band the loss in the old underwater cables as well as the old duct cables is generally high due to the peak OH⁻ absorption.

Figure 19 shows the spectral loss of old underwater cables and duct cables in the S-, C-, L- and U-band (1460 – 1675 nm). Due to the high OH⁻ absorption at 1385 nm, the attenuation for underwater cables is increasing towards the S-band edge at 1460 nm,

while the H₂ absorption at 1700 nm leads to an increase towards the U-band edge at 1675 nm. The small peak at 1590 nm is observed.

In addition to the peaks at 1700 nm and 1590 nm, it is a fact that the fibre attenuation is slightly higher than normal at 1550 nm and at 1625 nm due to the molecular hydrogen induced loss [7] [10].

6 Cable testing

6.1 Background

During installation and throughout the lifetime of the cable, it may be subjected to environmental conditions such as high/low temperatures, tension and crush, which may cause additional wavelength dependent loss in the fibres. At the time of a field measurement, this loss may not be present. In order to get the full picture of possible added loss during their lifetime, the different cable types have to be tested up to and beyond their specified limits.

For reasons outlined in Chapter 5, one crucial test to withstand for armoured underwater cables is the hydrogen test, which is performed by the cable manufacturer. The Telenor hydrogen test specification requires zero hydrogen ingress throughout the cable lifetime. Mechanical testing of armoured cables requires heavy equipment and is also performed by the manufacturer.

Telenor has its own facilities for testing land (duct) cables. This gives Telenor the possibility to examine and characterize the wavelength dependence of additional loss due to temperature, crush and tension on the cable. Such information is of crucial importance considering the extended use of wavelengths. All the test equipment and procedures are defined in detail by the standardisation body IEC for reproducibility and comparison between users.

Such data is not available from the cable manufacturer since their testing is performed at one wavelength only, usually 1550 nm.

In our investigation different duct cable samples from different time periods were subjected to extensive testing, during which the added spectral loss was recorded and analysed. In the following, the main results are presented.

6.2 Temperature test

Figure 20 shows the wavelength dependence of induced loss due to low temperatures in a duct cable from 2001. The additional loss is less than the specified limit of 0.1 dB/km at temperatures below -40° C,

for wavelengths up to 1700 nm. Similar results were obtained for duct cables produced in 1994–95. We should note that our aerial cables have better temperature performance than our duct cables due to the suspended part of the figure eight construction. Therefore, these results indicate that Telenor's aerial and duct cables are temperature resistant throughout the U-band (1625–1675 nm), even if they are specified only for 1550 nm operation.

6.3 Tension test

Figure 21 shows additional loss due to fibre elongation at cable tensions well above the specified limit, below which no fiber elongation is allowed, for a new duct cable. We note that induced loss is present at all wavelengths. Furthermore, they increase with increasing wavelength and elongation. Similar wavelength dependence, but with different loss increase




Figure 20 Temperature induced loss increase in a duct cable of recent date




Figure 21 Loss increase due to fibre elongation during tension test




Figure 22 Mandrel-plate crush induced loss in duct cable of recent date

magnitude has been observed in cables from other manufacturers.

Considering that it is unlikely that longer lengths of cable will remain tensioned above the limit after installation, we believe that loss due to excess cable tension should not be a problem if the cables are installed properly. This assumption has been confirmed in the field measurements results, where additional loss of this character has not been visible in the 1642 nm OTDR curves.

6.4 Crush test

Figure 22 shows the wavelength dependencies for loss increases for selected crushes at 15 min hold time during mandrel/plate crush test for a duct cable of recent date.

It should be noted that less than 5–10 % of the total number of crushes resulted in one single loss increase as shown in Figure 22. Thus the probability of loss increase is very low, and consequently a large number of crushes had to be performed in order to establish reliable statistics. The results are typical for new cables, while the statistics for older cables with other constructions may be poorer [10]. Added loss of this type has not been visible in the basic field measurements described in Chapter 3.

7 Summary

We have revealed statistical values for fibre attenuation and splice loss measured in the field on installed cables with standard G.652 fibres from different time periods, at six different wavelengths: 1242 nm, 1310 nm, 1383 nm, 1550 nm, 1625 nm and 1642 nm. The main observations are that the average fibre attenuation, splice loss and their standard deviations

have decreased significantly at all wavelengths since the 1980s. Especially the water peak attenuation at 1383 nm showed high values with considerable spread before 1990. In the 2003 cables, the average water peak attenuation is almost as low as the average attenuation at 1310 nm.

To study the fibre attenuation at wavelengths different from the measured ones, a mathematical model has been found to fit the fibre attenuation in the field by studying field measured spectra. The variations found in water peak form have been small and limited to the 1410 – 1480 nm region. The mathematical model has been applied to the field OTDR measurements at six wavelengths and the full fibre attenuation spectra have been calculated. The curve fitting results are in excellent agreement with the actual measured values at the OTDR wavelengths. We have calculated statistical parameters for fibre attenuation in the field at the CWDM wavelengths and also included the splice loss to form the total link loss at the CWDM wavelengths. The applied method gives extended and valuable information of fibre cable attenuation in installed cables from different time periods.

The levels of hydrogen present in our old underwater cables without hydrogen barrier do not result in significant additional loss in major parts of the transmission bands. More details with respect to the wide spectrum fibre attenuation in old underwater cables will be further investigated.

Extensive testing of Telenor's duct cables has revealed the detailed wavelength dependence of loss caused by excessive temperature, tension and crush. The test results have shown that temperature induced loss will not occur in our installed cables for wavelengths up to 1675 nm. Furthermore, the test results have shown that additional loss due to tension and crush is not likely to occur in cables if they are properly installed. Proof of this has been confirmed by the field measurements.

8 Acknowledgement

The author wishes to express great appreciation to Idar Gangsø at Telenor Networks who has performed the field measurements and the cable tests with great skills and much enthusiasm during many years.

9 References

- 1 Hopland, S. Loss Measurements in Installed Fibre Optic Cables At 1383 nm and 1642 nm. *Proceedings International Wire & Cable Symposium*

- 2003, Philadelphia, Pennsylvania, USA, 16–19 November 2003.
- 2 Hopland, S. A Full Spectrum Mathematical Model for Fibre Attenuation in the Field and its Application to OTDR Measurements at Six Selected Wavelengths. *Proceedings International Wire & Cable Symposium 2004*, Philadelphia, Pennsylvania, USA, 14–17 November 2004.
 - 3 Hopland, S. *Fibre cable field loss measurements in a broad wavelength range*. Contribution ITU-T meeting April 2004. Not official; presented as Working Document.
 - 4 Hopland, S. *Fibre cable field loss measurements in a broad wavelength range*. Contribution ITU-T meeting November 2004. White Paper COM15-12.
 - 5 Hopland, S. *Application of a mathematical model for fibre attenuation in the field to curve fit OTDR measurements in a broad wavelength range*. Contribution ITU-T meeting November 2004. Delayed document D-152.
 - 6 Hopland, S. Loss Performance of Standard Single Mode Fiber Cables in the wavelength Region 1200-1750 nm. *Proceedings International Wire & Cable Symposium*, Lake Buena Vista, Florida, USA, 18–21 November 2002.
 - 7 Hopland, S. Long Term Behaviour of Hydrogen Induced Losses in Installed Fibreoptic Submarine and Underwater Cables. *Proceedings International Wire and Cable Symposium*, Philadelphia, Pennsylvania, USA, 17–20 November 1997, 316–325.
 - 8 Hopland, S. Measurements of low hydrogen levels in installed open fibre optic submarine cables. *Proceedings International Wire & Cable Symposium*, St. Louis, Missouri, USA, 18–21 November 1991, 742–749.
 - 9 Hopland, S. Investigation of total and distributed hydrogen levels in installed fiberoptic submarine cables. *Proceedings International Wire & Cable Symposium*, Atlanta, Georgia, USA, 14–16 November 1989, 684–688.
 - 10 Hopland, S. Investigation of 1625 nm loss performance of fibre optic cables intended for use at 1550 nm. *Proceedings International Wire & Cable Symposium*, Atlantic City, NJ, USA, 15–18 November 1999.

Svend Hopland is Chief Engineer on optical fibre cables and fibre transmission at Telenor Nordic Fixed. He graduated from the Norwegian Institute of Technology (NTH) in 1985 with a PhD on optical fibres and joined Telenor in 1986. He has been responsible for development projects on underwater fibre cables and fibre cable specifications at Telenor. He has pioneered and developed an advanced method for installing underwater cables in the extremely difficult sub sea terrain along the Norwegian coast, for which he received an award at the International Wire and Cable Symposium in 1993. He has done extensive investigations on hydrogen generation in installed underwater cables. Recently, he has pioneered field measurements on wide spectrum fibre cable link losses and done extensive cable testing to support his field investigations. He has published more than 25 papers in international conferences and journals, and given several contributions to ITU-T.

email: svend.hopland@telenor.com

Terms and acronyms in Optical Communications

16-QAM	16-level Quadrature Amplitude Modulation	A modulation method in which the amplitude of two quadrature waves is changed in order to represent the data signal. A constellation diagram is often used to represent this kind of modulation. In QAM the constellation points are usually arranged in a square grid with equal horizontal and vertical spacing, although other configurations are possible. The number of points in the grid is usually a power of 2 (2, 4, 8 etc.), and since QAM usually is square; the most common forms are 4-QAM (or QPSK), 16-QAM, 64-QAM, 128-QAM and 256-QAM. QPSK and 16-QAM are e.g. used in current and future mobile radio systems like 3G/HSDPA and WiMax/IEEE 802.16, while 64-QAM and 256-QAM are often used in digital cable television and cable modems.
ACTS	Advanced Communications Technologies and Services	Thematic Programme of the Fourth Framework Research Programme funded by the European Union. It ran from 1994 to 1998 (the last project ended 2000) with a total funding from EU of 681 million Euro. http://www.cordis.lu/acts
ADSL	Asymmetric Digital Subscriber Line	A data communications technology that enables faster data transmission over copper telephone lines than a conventional modem can provide. The access utilises the 1.1 MHz band and has the possibility to offer, dependent on subscriber line length, downstream rates of up to 8 Mb/s. Upstream rates start at 64 kb/s and typically reach 256 kb/s, but can go as high as 768 kb/s.
APON	ATM over PON	A Passive Optical Network (PON) that carries fixed length packets (cells) of 53 bytes. It is specified in the ITU-T G.983 standard and based on asynchronous transfer mode (ATM), and has therefore been referred to as APON (ATM PON). Further improvements to the original APON standard – as well as the gradual falling out of favor of ATM as a protocol – led to the final version of ITU-T G.983 being referred to more often as Broadband PON, or BPON. A typical APON/BPON provides 622 megabits per second (Mbit/s) of downstream bandwidth and 155 Mbit/s of upstream traffic, although the standard accommodates higher rates. www.ponforum.org
ASON	Automatically Switched Optical Network	ASON is a complete definition of the operation of automatically switched transport networks, management, control, data plane, and was developed by Study Group 15 of the ITU-T. ASON is an architecture that defines the components in an optical control plane and the interactions between those components. It also identifies which of those interactions will occur across a multi-vendor divide, and therefore require standardized protocols. Other areas are intentionally not standardized in order to allow vendors or operators to provide value adding.
ATM	Asynchronous Transfer Mode	A high bandwidth, low-delay, connection-oriented, packet-like switching and multiplexing technique. ATM allocates bandwidth on demand, making it suitable for high-speed connections of voice, data and video services. Access speeds are up to 622 Mb/s and backbone networks currently operate at speeds as high as 2.5 Gb/s. Standardised by ITU-T [Newton03].
	Availability	Fraction of the operational life of a component (a subsystem, a system) during which it is regularly functioning or providing its service.
	Back-up/spare capacity (lightpath)	Set of resources used to carry rerouted traffic in failure conditions.
BER	Bit Error Rate	The ratio of error bits to the total number of bits transmitted.
BPF	Band Pass Filter	A band pass filter is a device that passes frequencies within a certain range and rejects (attenuates) frequencies outside that range. An example of an analogue electronic band pass filter is an RLC circuit (a resistor-inductor-capacitor circuit).
CATV	Cable Television	Cable television or Community Antenna Television (CATV) (often shortened to cable) is a system of providing television, FM radio programming and other services to consumers via radio frequency signals transmitted directly to people's televisions through fixed optical fibres or coaxial cables as opposed to the over-the-air method used in traditional television broadcasting (via radio waves) in which a television antenna is required.
CP	Control Plane	
CWDM	Coarse Wavelength Division Multiplexing	A relatively low-cost WDM technology that uses wide spacing (about 10 – 20 nm) between wavelengths and un-cooled optics. CWDM is standardised by ITU-T G.694.1. www.itu.int

DA	Double Armour	
	Dedicated (shared) protection	Protection mechanism in which back-up resources are dedicated to a single (shared between many different) optical connection(s).
DGD	Differential Group Delay	Group delay is the rate of change of the total phase shift with respect to angular frequency through a device or transmission medium. In an optical fibre, it is the transit time required for optical power, traveling at a given mode's group velocity, to travel a given distance.
DLP	Dedicated Link Protection	A protection scheme where link protection is implemented, i.e. where in case of failure routes are diverted around the failed link or node but where the rest of the path remains unaltered; and where one protection link is dedicated to each working link.
DPP	Dedicated Path Protection	A protection scheme where path protection is implemented, i.e. where in case of failure the whole path is diverted through an alternative route; and where one protection path is dedicated to each working path.
DWDM	Dense Wavelength Division Multiplexing	A carrier class WDM technology that uses expensive, cooled optics and tight spacing between wavelengths of less than a nanometre based on specifications of the International Telecommunication Union (ITU). http://www.itu.int
EDFA	Erbium Doped Fibre Amplifier	Fibre amplifiers are optical amplifiers that use a doped optical fibre, which bears the communication signal, and is optically pumped with a laser having a high-powered continuous output at an optical frequency slightly higher than that of the communication signal. A fibre amplifier is capable of amplifying a complete wavelength-division multiplexed set of signals as a single wideband optical signal. The most common dopant used in fibre amplifiers is erbium, and such a fibre amplifier is known as an erbium-doped fibre amplifier (EDFA). Commercial EDFAs typically have a frequency of operation in the C and L band (approx. 1530 to 1605 nm), a low noise with noise figure of 3 – 6 dB, a high gain (20 – 40 dB) and less sensitivity to polarization of the light signal. Maximum optical output power is typically 14 – 25 dBm, with an internal gain of 25 – 50 dB. The gain variations are +/- 0.5 dB. The length of the active fibre is 10 – 60 m for C-band EDFAs and 50 – 300 m for L-band EDFAs. The number of pump lasers is 1 – 6, and the pump wavelength is 980 or 1480 nm.
EFM	Ethernet in the First Mile	The IEEE Std 802.3ah-2004 specification. This is an amendment to the Ethernet specification combining a minimal set of extensions to the Media Access Control (MAC) and MAC Control sublayers with a family of Physical (PHY) Layers. These Physical Layers include optical fibre and voice grade copper cable Physical Medium Dependent sublayers (PMDs) for point-to-point connections in subscriber access networks. It also introduces the concept of Ethernet Passive Optical Networks (EPONs). The specification has been superseded by the IEEE Std 802.3-2005 standard. www.ieee802.org/3/efm/
ELED	Edge Emitting Diodes	A light-emitting diode with output that emanates from between the heterogeneous layers. It has greater radiance and coupling efficiency to an optical fibre or integrated optical circuit than a surface-emitting LED, but not as great as the injection laser.
EPON	Ethernet over PON	A Passive Optical Network (PON) that carries variable length (up to 1518 bytes) packets. Promoted by the IEEE 802.3 Working Group www.ieee802.org/3/
FEC	Forward Error Correction	A technique of error detection and correction in which a transmitting host computer includes a number of redundant bits in the payload (data field) of a block or frame of data. The receiving device uses the extra bits to detect, isolate and correct any errors created in transmission.
FICON	Fibre Connection	
FR	Frame Relay	An efficient data transmission technique used to send digital information quickly and cheaply to one or many destinations from one point. It can be used for voice, data, local area network (LAN), and wide area network (WAN) traffic. Each frame relay end user gets a private line to a frame relay node. The frame relay network handles the transmission to its other end users over a path, which is always changing and is invisible to the end users. Frame relay is a packet switched Telecommunications network, commonly used at the data link layer (layer 2) of the Open Systems Interconnection reference model. Generally the concept of permanent virtual circuits (PVCs) is used to form logical end-to-end links mapped over a physical network. Switched virtual circuits (SVCs), analogous to circuit switching in the public switched telephone network, are also part of the Frame relay specification but are rarely applied in the real world. Frame relay was originally developed as a stripped-down version of the X.25 protocol.

FSAN	Full Service Access Network	Network concept defined by the FSAN Working Group in order to derive a common set of requirements among providers to give to equipment vendors in an attempt to accelerate the commercial deployment of optical access systems. The FSAN organization was established in 1995 by seven major telecommunications service providers. The work has included defining requirements for BPON, APON and GPON. The organization feeds their recommendations into the ITU-T, specifically Study Group 15, Question 2, and recommendations in the G.983.x and G.984.x family have been approved based on work done by the FSAN committee. www.fsanweb.org, www.ponforum.org
FTTB	Fibre To The Building	This is in reference to fibre optic cable, carrying network data, connected all the way from an Internet service provider to a customer's physical building.
FTTCab	Fibre to the Cabinet	Also known as FTTN (Fibre to the Neighbourhood). A Hybrid Fibre Coax (HFC) network architecture involving an optical fibre which terminates in either a street side or neighbourhood cabinet which converts the signal from optical to electrical. The subscriber connection is either over twisted pair cable or coaxial cable.
FTTH	Fibre To The Home	The installation of optical fibre from directly into the subscriber's home. FTTH is also referred to as fibre-to-the-building (FTTB), which includes optical fibre that is installed directly into a home or enterprise.
GFP	Generic framing procedure	A layer 2 telecommunications protocol. An encapsulating mechanism with a fixed amount of encapsulation overhead, which is independent from the binary content of the client. It is a frame based encapsulation for PDUs (Packet Data Units). GFP is a traffic adaptation protocol, providing convergence between packet-switched and transmission networks. GFP maps packet-based protocols such as Ethernet, Fibre Channel, FICON, ESCON, and various forms of digital video into SONET/SDH, typically using Virtual Concatenation (VCAT) to provide right-sized pipes for data services. Compared to older framing schemes such as PoS (Packet over SONET/SDH), ATM (Asynchronous Transfer Mode) and 8B/10B encoding, GFP offers significantly reduced latency and improved bandwidth utilization. GFP offers two different mapping modes, GFP-T and GFP-F. The ITU standard defining GFP is G.7041. GFP-F (GFP Framed) is normally used to encapsulate packet/frame-based protocols such as IP/PPP or Ethernet/ MAC. The frame is entirely assembled before transmission through the SONET/SDH network. GFP-T (GFP Transparent) offers direct transmission of data streams requiring low-latency, such as VoIP, digital video (e.g. DVB-ASI) and SAN (e.g. Fibre Channel) applications. GFP-T is optimized for protocols using 8B/10B encoding. www.itu.int
	Gigabit Ethernet	A term describing various technologies for implementing Ethernet networking at a nominal speed of 1 Gb/s. Gigabit Ethernet is supported over both optical fibre and twisted pair cable. Physical layer standards include 1000BASE-T, 1 Gb/s over cat-5e copper cabling and 1000BASE-SX for short to medium distances over optic fibre. While it is currently deployed in high-capacity backbone network links its speed is largely not yet required for small network installations. Gigabit Ethernet began to penetrate the desktop as of 2000.
GMPLS	Generalised Multi-Protocol Label Switching	Also known as Multi-protocol Lambda Switching, a technology that provides enhancements to Multi-protocol Label Switching (MPLS) to provide the control plane (signaling and routing) to support network switching for time, wavelength, and space switching as well as for packet switching. In particular, GMPLS will provide support for photonic networking, also known as optical communications. It is used not only to support packet based paths but other technologies such as Optical MUX (Multiplexer), ATM and Frame Relay switches. This common control plane promises to simplify network operation and management by automating end-to-end provisioning of connections, managing network resources, and providing the level of QoS that is expected in the new, sophisticated applications.
GPON	Gigabit PON	A PON (Passive Optical Network) which extends the BPON (Broadband PON) to support bitrates in excess of 1 Gb/s. It is specified by ITU-T G.984. www.ponforum.org, www.itu.int
GPRS	General Packet Radio Service	An enhancement to the GSM mobile communication system that supports data packets. GPRS enables continuous flows of IP data packets over the system for such applications as web browsing and file transfer. Supports up to 160 kb/s gross transfer rate. Practical rates are from 12 – 48 kb/s. http://www.etsi.org, http://www.3gpp.org
GSN	Global Seamless Network	An ASON/GMPLS network demonstrator built by Deutsche Telekom in a joint effort with system vendors in 2002. The focus of the GSN was the UNI (user network interface) and the ASON/GMPLS functionalities within one operator domain.

GSN+		The phase 2 of the GSN (Global Seamless Network) demonstrator ready in 2004, focusing on the NNI (Network-Network Interface) functionalities between domains and new services.
HDPE	High Density Polyethylene	A thermoplastic made from oil. It takes 1.75 kg of oil (in terms of energy and raw materials) to make 1 kg of HDPE. HDPE is resistant to many different solvents and is often used in food containers like containers for milk, liquid laundry detergent, etc., some plastic bags, containment of certain chemicals, chemical resistant piping systems like those used in high-tech industry, and geothermal heat transfer piping systems.
HDWDM	High Density Wavelength Division Multiplexing	
HFC	Hybrid Fibre Coax	A network which incorporates both optical fibre and coaxial cable to create a broadband network. It has been commonly employed by cable TV operators since the 1990s.
HIPER-ACCESS		A Point-to-MultiPoint (PMP) network architecture intended for high speed (up to 120 Mb/s) and high-QoS fixed wireless access. Applications include broadband access for residential and small business users to a wide variety of networks as a flexible and competitive alternative to wired access networks, however HiperACCESS is not an LMDS-type systems. HiperACCESS standardization focuses on solutions optimized for frequency bands above 11 GHz (e.g. 26, 28, 32, 42 GHz) with high spectral efficiency under LOS (Line Of Sight) conditions. For bandwidths of 28 MHz, both FDD and TDD channel arrangements as well as H-FDD terminals are supported. HiperACCESS is an interoperable standard, allowing for PMP systems with base station and terminal stations from different manufacturers. The HiperACCESS (HA) specifications are being developed by TC (Technical Committee) BRAN. www.etsi.org
HIPERLAN/2		A short-range variant of a broadband radio access network intended as a complementary access mechanism for UMTS systems as well as for private use as a wireless LAN type system. HiperLAN offers high speed (up to 54 Mb/s) access to a variety of networks including the UMTS core networks, ATM networks and IP-based networks. Basic applications include data, voice and video, with specific Quality of Service parameters taken into account. HiperLAN/2 systems can be deployed in offices, classrooms, homes, factories, hot spot areas such as exhibition halls and, more generally, where radio transmission is an efficient alternative or complements wired technology. The HiperLAN/2 specifications are being developed by ETSI Project BRAN. www.etsi.org
IEEE 802.11		Refers to a family of specifications developed by the IEEE for wireless local area networks. It also refers to the "Wireless LAN working group" of the IEEE 802 project. 802.11 specifies an over-the-air interface between a wireless client and a base station or between two wireless clients. The IEEE accepted the specification in 1997. There are several specifications in the 802.11 family, including i) 802.11 – provides 1 or 2 Mbit/s transmission in the 2.4 GHz band, ii) 802.11a – an extension that provides up to 54 Mbit/s in the 5 GHz band. It uses an orthogonal frequency division multiplexing encoding scheme rather than FHSS or DSSS, iii) 802.11b provides 11 Mbit/s transmission in the 2.4 GHz band and was ratified in 1999 allowing wireless functionality comparable to Ethernet, iv) 802.11g provides 20+ Mbit/s in the 2.4 GHz band, v) 802.11z is a method for transporting an authentication protocol between the client and access point, and the Transport Layer Security (TLS) protocol. More variants are also under preparation, including support of 100 Mbit/s traffic flows. http://www.ieee802.org/11/
IEEE 802.3		The permanent CSMA/CD (Ethernet) working group of the IEEE 802 project. http://www.ieee802.org/3/
IETF	Internet Engineering Task Force	A large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. The technical work of the IETF is done in its working groups, which are organized by topic into several areas (e.g. routing, transport, security, etc). Much of the work is handled via mailing lists. The IETF holds meetings three times per year. The IETF working groups are grouped into areas and managed by Area Directors (AD). The ADs are members of the Internet Engineering Steering Group (IESG). Providing architectural oversight is the Internet Architecture Board (IAB). The IAB also adjudicates appeals when someone complains that the IESG has failed. The IAB and IESG are chartered by the Internet Society (ISOC) for these purposes. The General Area Director also serves as the chair of the IESG and of the IETF, and is an ex-officio member of the IAB. The Internet Assigned Numbers Authority (IANA) is the central coordinator for the assignment of unique parameter values for Internet protocols. The IANA is chartered by the Internet Society (ISOC) to act as the clearinghouse to assign and coordinate the use of numerous Internet protocol parameters. www.ietf.org/rfc/rfc3935.txt

ILP	Integer Linear Programming	In mathematics, linear programming (LP) problems are optimization problems in which the objective function and the constraints are all linear. Linear programming is an important field of optimization. Many practical problems in operations research can be expressed as linear programming problems. Certain special cases of linear programming, such as network flow problems and multicommodity flow problems are considered important enough to have generated much research on specialized algorithms for their solution. A number of algorithms for other types of optimization problems work by solving LP problems as sub-problems. Historically, ideas from linear programming have inspired many of the central concepts of optimization theory, such as duality, decomposition, and the importance of convexity and its generalizations. If the unknown variables are all required to be integers, then the problem is called an integer programming (IP) or integer linear programming (ILP) problem. In contrast to linear programming, which can be solved efficiently in the worst case, integer programming problems are in the worst case undecidable.
IP	Internet Protocol	A protocol for communication between computers, used as a standard for transmitting data over networks and as the basis for standard Internet protocols. http://www.ietf.org
ISDN	Integrated Services Digital Network	A digital telecommunications network that provides end-to-end digital connectivity to support a wide range of services, including voice and non-voice services, to which users have access by a limited set of standard multi-purpose user-network interfaces. The user is offered one or more 64 kb/s channels. http://www.itu.int
ITU	International Tele-communication Union	On 17 May 1865, the first International Telegraph Convention was signed in Paris by the 20 founding members, and the International Telegraph Union (ITU) was established to facilitate subsequent amendments to this initial agreement. It changed name to the International Telecommunications Union in 1934. From 1948 a UN body with approx. 200 member countries. It is the top forum for discussion and management of technical and administrative aspects of international telecommunications. http://www.itu.int
ITU-T	International Tele-communication Union – Standardization Sector	http://www.itu.int/ITU-T/
LAN	Local Area Network	A network shared by communicating devices, usually in a small geographical area. A system that links together electronic office equipment, such as computers and word processors, and forms a network within an office or building.
LCAS	Link Capacity Adjustment Scheme	A mechanism to dynamically alter the SONET/SDH bandwidth. Together with Generic Framing Procedure (GFP) and Virtual Concatenation (VC) LCAS are key enablers of new and flexible data services over SONET/SDH. LCAS is specified by ITU-T G.7042. www.itu.int
LED	Light Emitting Diode	A semiconductor device that emits incoherent narrow-spectrum light when electrically biased in the forward direction. This effect is a form of electroluminescence. The color of the emitted light depends on the chemical composition of the semiconducting material used, and can be near-ultraviolet, visible or infrared. Nick Holonyak Jr. (born 1928) of the University of Illinois at Urbana-Champaign developed the first practical visible-spectrum LED in 1962.
	Lightpath	Optical circuit
	Link protection	A local protection mechanism protecting all the set of lightpaths crossing a and managed by the termination nodes of the link.
LPDE	Low Density Polyethylene	Low density polyethylene (LDPE) is a thermoplastic made from oil. It is unreactive at room temperatures, except by strong oxidizing agents, and some solvents cause its swelling. It can withstand temperatures of 80 °C continuously and 95 °C for a short time. It is translucent or opaque, quite flexible, and tough to the degree of being almost unbreakable. It is widely used for manufacturing various containers, dispensing bottles, wash bottles, tubing, and various molded laboratory equipment. Its most common use is in plastic bags.
LP-LD	Low-power Laser	
MAN	Metropolitan Area Network	Large computer networks usually spanning a campus or a city using wireless infrastructure or optical fibre connections to link their sites. Some technologies used for this purpose are Asynchronous Transfer Mode (ATM), Fibre Distributed Data Interface (FDDI) and Switched Multimegabit Data Service (SMDS). These older technologies are in the process of being displaced by Ethernet-based MANs in most areas. MAN links between LANs have been built without cables using either microwave, radio, or infra-red free-space optical communication links.

MEMS	Micro-Electro-mechanical Systems	Tiny mechanical devices that are built onto semiconductor chips and are measured in micrometres. In the research labs since the 1980s, MEMS devices began to appear as commercial products in the mid-1990s. They are used to make pressure, temperature, chemical and vibration sensors, light reflectors and switches as well as accelerometers for airbags, vehicle control, pacemakers and games. The technology is also used to make ink jet print heads, microactuators for read/write heads and all-optical switches that reflect light beams to the appropriate output port.
MM	Multimode	A type of optical fibre mostly used for shorter distances, e.g. with a building. It can carry 1 Gbit/s for typical building distances; the actual maximum speed (given the right electronics) depends upon the actual distance. The earliest fibre optic cables used a technique termed multi-mode transmission. The light signals from the laser are broken up into a number of paths along the length of the fibre and is reflected off the fibre wall. The amount of reflection that occurs dictates the quality of the signal. Typical transmission speeds/distances limits are 100 Mbit/s up to 2 km, 1 Gbit/s for distances up to 500–600 metres, and 10 Gbit/s for distances up to 300 metres. Multi-mode optical fibre has two categories: Step Index and Graded Index. The core is the light transmission element at the center of the optical fibre. All the light signals travel through the core. A core is typically glass made from a combination of silicon dioxide (silica) and other elements. Graded index multimode fibre uses a lower index of refraction towards the outer edge of the core. Therefore, the outer area of the core is less optically dense than the centre, and light can go faster in the outer part of the core. A light ray following a mode that goes straight down the centre of the core has a shorter path compared to a ray following a mode that bounces around in the fibre. All rays should arrive at the end of the fibre at the same time.
MPLS	Multi Protocol Label Switching	An IETF standard intended for Internet application. MPLS has been developed to speed up the transmission of IP-based communications over ATM networks. The system works by adding a much smaller "label" to a stream of IP datagrams allowing "MPLS"-enabled ATM switches to examine and switch the packet much faster. www.ietf.org
MSTP	Multi-Service Transport Platform	
MTTF	Mean Time To Failure	An expression of failure rate, which is the average frequency with which something fails. Failure rate is important in the fields of reliability theory and product warranties. The failure rate depends on the failure distribution, which describes the probability of failure prior to a specified time. The failure rate is not always constant, so the hazard function is used to describe the instantaneous failure rate at any point in time. A formal definition of failure rate can be: "The total number of failures within an item population, divided by the total time expended by that population, during a particular measurement interval under stated conditions." [MacDiarmid, P, Morris, S et al. (no date) <i>Reliability Toolkit: Commercial Practices Edition</i> , 35–39. Rome, NY, Reliability Analysis Center and Rome Laboratory]
MVDS	Multipoint Video Distribution Service	
MZM	Mach-Zehnder Modulator	A waveguide based optical modulator scheme based on a Mach-Zehnder (MZ) interferometer arrangement. Here the optical signal is divided into two equal parts that are subsequently sent to two waveguides of initially identical transmission properties and length such that the total phase shift on propagation through these is exactly equal. One of the waveguides includes a part where the phase shift can be varied. At the end of these two waveguides, the two signals are coupled back together and interfere with each other. By varying the phase change through the variable phase modulator, one can control the interference of the signals at the output and hence the amplitude of the output signal. An intensity modulator can be constructed this way.
NGN	Next Generation Network	A network concept that aims at providing a framework to encompass the large variety of existing and emerging protocols and services, facilitate a further evolution of these, decouple the evolution from the underlying network infrastructure and facilitate the interfacing of a plethora of available media. The rationale behind NGN lies founded in paradigm shifts that have been taking place within the technological solutions and the business models in the telecom industry as a whole. The concept is based on IP-technology and is being specified by ITU-T. www.itu.int
NMS	Network Management System	A combination of hardware and software used to monitor and administer a network.
NNI	Network to Network Interface	

NSFNET	National Science Foundation Network	A major part of early 1990s Internet backbone. Organisations connecting to the Internet in the early 1990s had to sign a usage agreement directly with NSFNET to gain access to large parts of the Public Internet, regardless of which Internet Service Provider they purchased Internet access from. From 1987 to 1995 the NSFNET was operated on behalf of the NSF by Merit Network, Inc, a non-profit corporation governed by public Universities. On 30 April 1995, the NSFNET Backbone Service was successfully transitioned to a new architecture, where traffic is exchanged at interconnection points called NAPs (Network Access Point).
OADM	Optical Add-Drop Multiplexing	A main element of optical fibre networks which combines or multiplexes several lower-speed streams of data into a single beam of light. OADMs can be used both in long-haul core networks and in shorter distance metro networks.
OCDMA	Optical Code Division Multiple Access	In optical communications different signature sequences of optical pulses (codes) are used to distinguish different data streams. The duration of the sequence must be at least equal to that of a data bit.
OCh	Optical Channel	
OEO	Optical-Electrical-Optical	A wavelength switch that terminates the optical signal at input and output and performs electrical switching.
OFDM	Orthogonal Frequency Division Multiplexing	A spread spectrum technique that distributes the data over a large number of carriers spaced apart at precise frequencies. This spacing provides the "orthogonality" in this technique, which prevents the demodulators from seeing frequencies other than their own. The benefits of OFDM are high spectral efficiency, resilience to RF interference, and lower multi-path distortion. This is useful because in a typical terrestrial wireless scenario there are multipath channels (i.e. the transmitted signal arrives at the receiver using various paths of different length). Since multiple versions of the signal interfere with each other (inter symbol interference (ISI)) it becomes very hard to extract the original information. OFDM is sometimes called multi-carrier or discrete multi-tone modulation. It is the modulation technique used for digital TV in Europe, Japan and Australia. It is used in DAB, ADSL and WLAN 802.11a and g and WMAN 802.16 standards.
OH	Hydroxyl	Hydroxide (OH ⁻) is a polyatomic ion consisting of oxygen (O) and hydrogen (H). It has a charge of -1. Hydroxide is one of the simplest of the polyatomic ions. The term hydroxyl group is used to describe the functional group OH when it is a substituent in an organic compound. Organic molecules containing a hydroxyl group are known as alcohols. The hydroxyl radical, OH, is the neutral form of the hydroxide ion. Hydroxyl radicals are highly reactive and consequently short lived.
OIF	Optical Internet working Forum	Industry group whose mission it is to foster the development and deployment of interoperable products and services for data switching and routing using optical networking technologies. OIF's purpose is to accelerate the deployment of interoperable, cost-effective and robust optical internetworks and their associated technologies. Optical internetworks are data networks composed of routers and data switches interconnected by optical networking elements. http://www.oiforum.org/
OLT	Optical Line Terminal	Optical termination in the local exchange.
OMS	Optical Multiplex Section - layer network	The optical multiplex section layer provides transport for the optical channels. The information contained in this layer is a data stream comprising a set of optical channels, which may have a defined aggregate bandwidth.
ONU	Optical Network Unit	A network element that is part of a fibre-in-the-loop system interfacing the customer analogue access cables and the fibre facilities.
OOO	Optical-optical-optical	Also called photonic switch or all-optical crossconnect. An optical wavelength switch that does not terminate the optical signal.
OPEX	Operations Expenditure	A company's operational cost in contrast to CAPEX – Capital Expenditure, which is the company's investment cost.
OSNR	Optical Signal-to-Noise Ratio	The measure of the ratio of signal power to noise power in an optical channel. For a typical optical communication system for which the OSNR is relevant, the signal consists usually of nearly monochromatic modulated light superimposed on a background comprised of (mostly unmodulated) optical power distributed over a broad wavelength range – a range including the signal wavelength. This noise arises typically in optical amplification and it is better thought of as a power density rather than a total power.
OSPF	Open Shortest-Path First	A link-state, hierarchical Interior Gateway Protocol (IGP) routing protocol. The shortest path tree is calculated using cost as its routing metric. A link state database is constructed of the network topology which is identical on all routers in the area.

OTDR	Optical Time Domain Reflectometry	An opto-electronic instrument used to characterize an optical fibre. An OTDR injects a series of optical pulses into the fibre under test. It also extracts, from the same end of the fibre, light that is scattered back and reflected back. The intensity of the return pulses is measured and integrated as a function of time, and is plotted as a function of fibre length. It may be used for estimating the fibre's length and overall attenuation, including splice and mated-connector losses. It may also be used to locate faults, such as breaks.
OTH	Optical Transport Hierarchy	www.itu.int
OTN	Optical Transport Network	ITU-T G.709 www.itu.int
OXC	Optical Cross-Connect	A space-division switch that can switch an optical data stream from an input port to an output port. Such a switch may utilise optical-electrical conversion at the input port and electrical-optical conversion at the output port or it may be all-optical. An OXC is assumed to have control plane processor that implements the signalling and routing protocols necessary for computing and instantiating optical channel connectivity in the optical domain. {based on [ID-ipofw]}.
	Path protection	An end-to-end protection mechanism protecting a point-to-point optical connection and managed by the source and destination nodes
PDH	Plesiochronous Digital Hierarchy	A technology used in telecommunications networks to transport large quantities of data over digital transport equipment such as fibre optic and microwave radio systems. The term plesiochronous is derived from Greek plesio, meaning near, and chronos, time, and refers to the fact that PDH networks run in a state where different parts of the network are almost, but not quite perfectly, synchronised. PDH is now being replaced by SDH (Synchronous Digital Hierarchy) equipment in most telecommunications networks. PDH allows transmission of data streams that are nominally running at the same rate, but allowing some variation on the speed around a nominal rate. The European and American versions of the PDH system differ slightly in the detail of their working, but the principles are the same. The European system is described. The basic data transfer rate is a data stream of 2.048 Mb/s. For speech transmission, this is broken down into 30 x 64 kb/s channels plus 2 x 64 kb/s channels used for signalling and synchronisation. Alternatively, the whole 2 Mb/s may be used for non speech purposes, for example data transmission.
PG	Protection group	A set of optical connections that may be involved by common protection actions.
PL	Protection Lightpath	An optical lightpath that is used as an alternative to a working lightpath in case the latter experiences failure, e.g. fibre cut or node failure.
PLANET	Photonic Local Access Network	Project AC050 in EUs 4th Framework Programme ACTS, lasting from 1996 to 2000. The objective of the project was to define a cost effective full-service access network for the year 2000 supporting a large splitting factor (up to 2000) and covering a wide range (up to 100 km). Especially the wide range is required for a future access network in order to anticipate the expected switching node consolidation. The studied access network, called a SuperPON, has to give the subscriber access to both narrowband/broadband, and distributive/interactive services.
PLOAM	Physical Layer Operation, Administration and Maintenance	
PMD	Polarisation Mode Dispersion	A form of modal dispersion where two different polarizations of light in a waveguide, which normally travel at the same speed, travel at different speeds due to random imperfections and asymmetries, causing random spreading of optical pulses.
PON	Passive Optical Network	A passive optical network (PON) is a system that brings optical fibre cabling and signals all or most of the way to the end user. Depending on where the PON terminates, the system can be described as fibre-to-the-curb (FTTC), fibre-to-the-building (FTTB), or fibre-to-the-home (FTTH). A PON consists of an Optical Line Termination (OLT) at the communication company's office and a number of Optical Network Units (ONUs) near end users. Typically, up to 32 ONUs can be connected to an OLT. The passive simply describes the fact that optical transmission has no power requirements or active electronic parts once the signal is going through the network.
POTS	Plain Old Telephone Service	A very general term used to describe an ordinary voice telephone service. See also PSTN.

PRISMA	Photonic Routing of Interactive Services for Mobile Applications	Project AC349 in EUs 4th Framework Programme ACTS. The main objective was to increase the flexibility of fibre-wireless networks, by introducing multiple wavelength channels in the fibre network and selecting one or more wavelength channels as needed at the wireless base transceiver stations (BTSs). http://www.cordis.lu/infowin/acts/rus/projects/ac349.htm
	Protection	A resilience mechanism in which the back-up capacity is pre-planned.
PSTN	Public Service Telephone Network	Common notation for the conventional analogue telephone network.
RAP	Radio Access Point	A device that "connects" wireless communication devices together to create a wireless network. It is usually connected to a wired network, and can relay data between devices on each side. Many RAPs can be connected together to create a larger network that allows roaming. In contrast, a network where the client devices manage themselves is called an ad-hoc network.
	Reliability	The probability of failure of a component (a subsystem, a system).
	Restoration	A resilience mechanism in which the back-up capacity is dynamically searched for after a failure.
RFWA	Routing, Fibre and Wavelength Assignment	The operation of assigning capacity to an optical connection in a mesh, multi-fibre and WDM network.
ROADM	Reconfigurable OADM	A new form of optical add-drop multiplexers (OADM) that adds the ability to switch individual networks. The key features of ROADMs is the ability to switch traffic and both the wavelength and SONET/SDH layers.
RWA	Routing and wave-length assignment	The operation of assigning capacity to an optical connection in a WDM network.
RZ	Return-to-Zero	A line code used in telecommunications signals in which the signal drops (returns) to zero between each pulse. This takes place even if a number of consecutive zeros or ones occur in the signal. This means that a separate clock does not need to be sent alongside the signal. The signal is self-clocking. A variant, Return-to-zero, inverted, swaps the signal values for one and zero.
SA	Single Armour	Optical fibre cable technology in which the core is protected by a single layer of armoured steel wires.
SCMA	SubCarrier Multiple Access	Variant of frequency division multiple access used in optics. Different packet streams are modulated onto different electrical carrier frequencies. They subsequently modulate the light intensity of the laser diode. The packet streams are thus put into different frequency bands. Each frequency band constitutes an independent communication channel.
SDH	Synchronous Digital Hierarchy	SDH is a standard technology for synchronous data transmission on optical media. It is the international equivalent of North American SONET. Both technologies provide faster and less expensive network interconnection than traditional PDH equipment. It is a method of transmitting digital information where the data is packed in containers that are synchronized in time enabling relatively simple modulation and demodulation at the transmitting and receiving ends. The technique is used to carry high capacity information over long distances. SDH uses the following Synchronous Transport Modules (STM) and rates: STM-1 (155 Mbit/s), STM-4 (622 Mb/s), STM-16 (2.5 Gb/s), and STM-64 (10 Gb/s). SDH is specified by ITU-T G.707. www.itu.int
SLA	Service Level Agreement	A contract between a provider and a customer that guarantees specific levels of performance and reliability at a certain cost. This contract should also define precisely what could be penalties and back-up solutions in case of problems. SLA is especially important to define when an important part of your system or activity relies on third party providers. SLA is also a very good approach for services provided internally to your organisation where you should also have a customer approach concern. [Awduche, D et al. <i>Overview and Principles of Internet Traffic Engineering</i> . May 2002. IETF RFC 3272]
SLP	Shared Link Protection	A protection scheme where link protection is implemented; i.e. where in case of failure routes are diverted around the failed link or node but where the rest of the path remains unaltered; and where the protection link provides protection to more than one other link that then share link protection.
SLS	Service Level Specification	

SM	Single Mode	An optical fibre in which only the lowest order bound mode can propagate at the wavelength of interest. Single mode fibres are best at retaining the fidelity of each light pulse over longer distances and exhibit no dispersion caused by multiple spatial modes; thus more information can be transmitted per unit time giving single mode fibres a higher bandwidth in comparison with multi-mode fibres. A typical single mode optical fibre has a core radius of 5 – 10 μm and a cladding radius of 120 μm . Currently, data rates of up to 10 Gb/s are possible at distances of over 60 km with commercially available transceivers.
SONET	Synchronous Optical Network	A standard for optical telecommunications transport formulated by the Exchange Carriers Standards Association (ECSA) for the American National Standards Institute (ANSI), in ANSI T1.105. http://www.ansi.org/ , http://www.iec.org/online/tutorials/sonet/
SPP	Shared Path Protection	A protection scheme where path protection is implemented; i.e. where in case of failure the whole path is diverted through an alternative route; and where the protection path provides protection to more than one path that then share path protection.
STM	Synchronous Transfer Mode	A transport level technique in which time-division multiplexing and switching is used across the user's network interface.
STM-16	Synchronous Transfer Mode transmission at 2.5 Gbit/s	
STM-64	Synchronous Transfer Mode transmission at 10 Gbit/s	
	Survivability (resilience)	The property of a system a network to completely or partly maintain its services in the presence of failures affecting some of its elements. This property is achieved by implementing in the system suitable mechanisms of failure reaction (resilience strategy or mechanism), usually based on signal duplication or traffic rerouting.
SW	Software	Essentially a computer program encoded in such a fashion that the program (the instruction set) contents can be changed with minimal effort. Computer software can have various functions such as controlling hardware, performing computations, communication with other software, human interaction, etc; all of which are prescribed in the program. The term "software" was first used in this sense by John W. Tukey in 1957; computer software is all information processed by computer system, programs and data. The concept of software was first proposed by Alan Turing in a report to National Physics Laboratory (UK) in 1946 entitled "Proposed Electronic Calculator".
TCP	Transport Control Protocol	Transport layer protocol defined for the Internet by Vint Cerf and Bob Kahn in 1974. A reliable octet streaming protocol used by the majority of applications on the Internet. It provides a connection-oriented, full-duplex, point-to-point service between hosts. http://www.ietf.org
TDMA	Time Division Multiple Access	A technology for shared medium (usually radio) networks. It allows several users to share the same frequency by dividing it into different time slots. The users transmit in rapid succession, one after the other, each using their own timeslot. This allows multiple users to share the same transmission medium (e.g. radio frequency) whilst using only the part of its bandwidth they require. Used in the GSM, PDC and iDEN digital cellular standards, among others. TDMA is also used extensively in satellite systems, local area networks, physical security systems, and combat-net radio systems.
TOBASCO	Towards Broadband Access Systems for CATV Optical Networks	Project AC028 in EUs 4th Framework Programme ACTS. The main objective of the project was to upgrade existing CATV networks having high splitting counts with broadband interactive services, by applying High-Density Wavelength Division Multiplexing (HDWDM). http://www.cordis.lu/infowin/acts/rus/projects/ac028.htm
TPON	Telephony over PON	The original proposal for a Fibre to the Home (FTTH) network, but superseded by BPON (Broadband PON), a much higher speed version capable of delivering multimedia high bandwidth services.
ULH	Ultra Long Haul	

UMTS	Universal Mobile Telecommunication System	The European member of the IMT 2000 family of 3G wireless standards. UMTS supports data rates of 144 kb/s for vehicular traffic, 384 kb/s for pedestrian traffic and up to 2 Mb/s in support of in-building services. The standardisation work began in 1991 by ETSI but was transferred in 1998 to 3GPP as a corporation between Japanese, Chinese, Korean and American organisations. It is based on the use of WCDMA technology and is currently deployed in many European countries. The first European service opened in 2003. In Japan NTT DoCoMo opened its "pre-UMTS" service FOMA (Freedom Of Mobile multimedia Access) in 2000. The system operates in the 2.1 GHz band and is capable of carrying multimedia traffic. http://www.3gpp.org/
	Unavailability	A fraction of the operational life of a component (a subsystem, a system) during which it is not functioning or it is out of service.
UNI	User Network Interface	A term used in ATM and Frame Relay networks, UNI is the interface between the ATM end user and a private ATM switch. It can also represent the interface between a private ATM switch and the public carrier ATM network.
VC	Virtual Container	Data unit on SDG layer 1
VCAT		VCAT enables transport pipes to be "right-sized" for various data payloads by allowing SONET/SDH channels to be multiplexed in arbitrary arrangements. VCAT breaks down data packets and maps them into the base units of TDM frames; e.g. STS-1 (51 Mb/s) for SONET, and AU4 (155 Mb/s) for SDH. This data is then grouped into multiple data flows of varying size to create larger, aggregate payloads optimally sized to match available SONET/SDH pipe capacity. VCAT is applied at the end-points of the connections, which permits each channel used to be independently transmitted through a legacy transport network. Data is typically encapsulated using GFP. The ITU standard for VCAT is G.707.
VDSL	Very high speed Digital Subscriber Line	VDSL transmits data up to 26 Mb/s over short distances of twisted pair copper wire. The shorter the distance, the faster the connection rate. Specified by ITU-T G.993.1. Second generation, VDSL2, supports up to 100 Mb/s data rate either symmetric or asymmetric in the 1 – 3 km range. Specified in ITU-T G.993.2. www.itu.int
VOA	Variable Optical Attenuator	
VoIP	Voice over Internet Protocol	Voice over Internet Protocol (also called VoIP, IP Telephony, Internet telephony, and Digital Phone) is the routing of voice conversations over the Internet or any other IP-based network. The voice data flows over a general-purpose packet-switched network, instead of traditional dedicated, circuit-switched voice transmission lines.
VWP	Virtual Wavelength Path	A method in which a signal path can travel on different wavelengths throughout the network. By avoiding a dedicated wavelength for an end-to-end connection, the network can reuse and optimize wavelengths to provide the greatest amount of capacity.
WDM	Wavelength division multiplexing	A means of data transmission that uses optical multiplexing to enable two or more wavelengths each with its own data source to share a common fibre optic medium. Wavelength Division Multiplexing is a technology that allows optical signals operating at different wavelengths to be multiplexed onto a single optical fibre and transported in parallel through the fibre. In general, each optical wavelength may carry digital client payloads at a different data rate and in a different format. For example, a mixture of SDH 2.5 Gbit/s and 10 Gbit/s can be carried over a single fibre. An optical system with WDM capability can achieve parallel transmission of multiple wavelengths gracefully while maintaining high system performance and reliability. Standardised WDM systems are specified by ITU as Coarse (CWDM) and Dense (DWDM). CWDM is standardised in ITU-T G.694.1 www.itu.int
WDMA	Wavelength Division Multiple Access	In optical communication different wavelengths are used to send packets from different sources. The wavelength constitutes independent communication channels and may carry different signal formats.
	Working capacity (lightpath)	A set of resources carrying traffic in normal network conditions (no failure).
XC	Crossconnect	