Highlighted pages from these sources are found starting with the 2nd page.

1. **Source #1:** S. Sarkar, S. Dixit, B. Mukherjee (IEEE Fellow), *"Hybrid Wireless-Optical Broadband-Access Network (WOBAN): A Review of Relevant Challenges"*, IEEE Journal of Lightwave Technology, vol. 25, no. 11, Nov 2007

http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=04357916

2. **Source #2:** H. Willebrand, B.S. Ghuman, *"Free space optics: enabling optical connectivity in today's networks"*, SAMS Publishing, 2002

http://books.google.co.uk/books?id=iSk7r67xyboC&lpg=PA97&ots=ovbDa7G7Ym&pg=PA96#v=snippet&q=%22as%20metropolitan%20DWDM%20systems%20migrate%20into%20the%20access%20arena%22&f=false

3. **Source #3:** R. Ramaswami, K.N. Sivarajan, G. Hajime Sasaki, *"Optical Networks: A Practical Perspective"*, 3rd Ed, Elsevier (Morgan Kaufmann Publishers), 2010

http://books.google.co.uk/books?id=WpByp4Ip0z8C&lpg=PA1&ots=7O5WNDrAL2&pg=PA1#v=onepage&q=%22relentless%20need%20for%20more%20capacity%22&f=false

http://books.google.co.uk/books?id=WpByp4Ip0z8C&lpg=PA1&ots=7O5WNDrAL2&pg=PA629#v=snippet&q=%22with%20fiber%20now%20directly%20available%22&f=false

http://books.google.co.uk/books?id=Apy17B5ZrsYC&lpg=PA593&ots=Xl0MPaQjoK&dq=%22an%20RN%20broadcasts%20the%20data%20it%20receives%22&pg=PA593#v=onepage&q=%22different%20combinations%20of%20services%22&f=false

4. **Source #4:** B. Mukherjee, *"Optical WDM networks"*, Springer Science, 2006

Ch1: http://www.springerlink.com/content/t16854p05h3x0615/fulltext.pdf

Ch12: http://www.springerlink.com/content/ku3x37248427683p/fulltext.pdf

Ch16: http://www.springerlink.com/content/g85n1827433q1184/fulltext.pdf

(Full book: http://www.springerlink.com/content/978-0-387-29188-8#section=472628&page=1&locus=0)

5. **- Source #5:** K. Zhu, H. Zhu, B. Mukherjee, *"Traffic grooming in optical WDM mesh networks"*, Springer Science, 2005

http://dutta.csc.ncsu.edu/csc778_fall07/wrap/student_presentations/groom_1.pdf

Hybrid Wireless-Optical Broadband-Access Network (WOBAN): A Review of Relevant Challenges

Suman Sarkar, *Student Member, IEEE*, Sudhir Dixit, *Senior Member, IEEE*,
and Biswanath Mukherjee, *Fellow, IEEE*

Source #1

(Invited Paper)

***Abstract*—The hybrid wireless-optical broadband-access network (WOBAN) is a promising architecture for future access networks. Recently, the wireless part of WOBAN has been gaining increasing attention, and early versions are being deployed as municipal access solutions to eliminate the wired drop to every wireless router at customer premises. This architecture saves on network deployment cost because the fiber need not penetrate each end-user, and it extends the reach of emerging optical-access solutions, such as passive optical networks. This paper first presents an architecture and a vision for the WOBAN and articulates why the combination of wireless and optical presents a compelling solution that optimizes the best of both worlds. While this discussion briefly touches upon the business drivers, the main arguments are based on technical and deployment considerations. Consequently, the rest of this paper reviews a variety of relevant research challenges, namely, network setup, network connectivity, and fault-tolerant behavior of the WOBAN. In the network setup, we review the design of a WOBAN where the back end is a wired optical network, the front end is managed by a wireless connectivity, and, in between, the tail ends of the optical part [known as optical network unit (ONU)] communicate directly with wireless base stations (known as "gateway routers"). We outline algorithms to optimize the placement of ONUs in a WOBAN and report on a survey that we conducted on the distribution and types of wireless routers in the Wildhorse residential neighborhood of North Davis, CA. Then, we examine the WOBAN's routing properties (network connectivity), discuss the pros and cons of various routing algorithms, and summarize the idea behind fault-tolerant design of such hybrid networks.**

***Index Terms*—Architecture, broadband access, fault tolerance, optical network, routing, wireless network.**

I. INTRODUCTION

THE DOMINANT broadband-access network that is emerging from today's research and development activities is a point-to-multipoint (P2MP) optical network known as passive optical network (PON). The basic configuration of a PON connects the telecom central office (CO) to businesses and residential users by using one wavelength channel in the downstream direction [from optical line terminal (OLT) at CO to optical network units (ONUs)] and another wavelength channel in the upstream direction [from ONUs to OLT]. A PON does not have any active element in the signal's path from source to destination; hence, it is robust. The only interior elements used in such a network are passive combiners, couplers, and splitters.

A PON provides much higher bandwidth for data applications [than current solutions such as digital subscriber line (DSL) and cable modem (CM)], as well as deeper fiber penetration. Based on current standards, a PON can cover a maximum distance of 20 km from the OLT to the ONU. While fiber-to-the-building, fiber-to-the-home (FTTH), or even fiber-to-the-PC solutions have the ultimate goal of fiber reaching all the way to end-user premises, fiber-to-the-curb may be a more economical deployment scenario today [1], [2].

The traditional single-wavelength PON (also known as the time-division-multiplexed PON or TDM-PON) combines the high capacity of optical fiber with the low installation and maintenance cost of a passive infrastructure. The optical carrier (OC) is shared by means of a passive splitter among all the users, so the PON topology is a tree, as in most other distribution networks, e.g., those for power, voice, video, etc. As a consequence, the number of ONUs is limited by the splitting loss and by the bit rate of the transceivers in the OLT and in the ONUs. Current specifications allow for 16 ONUs at a maximum distance of 20 km from the OLT and 32 ONUs at a maximum distance of 10 km from the OLT.

The per-user cost of such a network can be low as the bandwidth (typically up to 1 Gb/s in current practice and expected to increase to 10 Gb/s in the future) is shared among all the end users, but, as end users demand more bandwidth, the need to upgrade the existing PON architectures [viz., Ethernet PON (EPON), Broadband PON (BPON, based on ATM), Gigabit PON (GPON), Generic Framing Procedure PON (GFP-PON), etc.] to Wavelength-Division-Multiplexed PON (WDM-PON) is essential. A WDM-PON solution provides excellent scalability because it can support multiple wavelengths over the same fiber infrastructure, it is inherently transparent to the channel bit rate, and, depending on its architecture, it may not suffer power-splitting losses (see [3] for a review of WDM-PON architectures).

Manuscript received February 15, 2007; revised May 10, 2007.

S. Sarkar and B. Mukherjee are with the Department of Computer Science, University of California, Davis, CA 95616 USA (e-mail: sarkar@cs.ucdavis.edu; mukherje@cs.ucdavis.edu).

S. Dixit is with the Nokia Siemens Networks, Mountain View, CA 94043 USA (e-mail: sudhir.dixit@nsn.com).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2007.906804

The straightforward approach to build a WDM-PON is to employ a separate wavelength channel from the OLT to each ONU, both in the upstream and downstream directions. This approach creates a point-to-point (P2P) link between the OLT and each ONU, which differs from the P2MP topology of the traditional PON. In the WDM-PON, each ONU can operate at a rate up to the full bit rate of a wavelength channel. Moreover, different wavelengths may be operated at different bit rates, if necessary; hence, different types of services may be supported over the same network. This is clearly an advantage of WDM-PON over the traditional PON [4].

There are various industry efforts to build PON architecture for commercial deployment. In the United States, Verizon has introduced its "Fiber-to-the-Premises" architecture, called FiOS, to deliver high-speed voice and data services to the home. FiOS service consists of three consumer broadband speeds: up to 5 Mb/s downstream and up to 2 Mb/s upstream (5 Mb/s/ 2 Mb/s), 15 Mb/s/2 Mb/s, and 30 Mb/s/5 Mb/s. The FiOS network is migrating from current BPON to future GPON architecture, thus moving toward higher upstream/downstream speed and eliminating ATM [5]. Among other efforts, Novera Optics has launched TurboLIGHT, a dense-WDM fiber-to-the-X optical-access technology, which allows flexible multimode-transport capabilities at different bit rates (125 Mb/s–1.25 Gb/s) [6]. In Asia, a similar effort can be found in WE-PON, which has a combined architecture of WDM (from CO to WDM device) and TDM (from WDM device to ONU through splitters) with bit rates on the order of 100 Mb/s [7].

Another promising access solution is a wireless network. Recently, we have seen tremendous growth in the research and deployment of various wireless technologies. There are three major techniques that have been employed for wireless-access networks worldwide, viz., "Wireless Fidelity" (known as WiFi), "Worldwide Interoperability for Microwave Access" (known as WiMax), and "Cellular Network." These technologies have their own advantages and disadvantages.

WiFi is one of the most popular wireless technology (standards: IEEE 802.11a/b/g), and it is mainly used for wireless local-area networks. WiFi can operate in both the "Infrastructure" and "*Ad Hoc*" modes. In infrastructure mode, a central authority, known as access point, is required to manage the network. But, in *ad hoc* mode, the users are self-managed, and there is no concept of an administrator. WiFi technology can exploit the flexibility of "multihopping." WiFi offers low bit rate (max 54/11/54 Mb/s for 802.11a/b/g, respectively) and limited range (typically 100 m).

WiMax (standard: IEEE 802.16) is gaining rapid popularity. It is essentially a P2MP broadband wireless-access service. WiMax can be used efficiently for single-hop communication (for multihop, WiMax suffers from higher delay and lower throughput). It provides high bandwidth and uses less-crowded spectrum. Thus, WiMax is particularly suitable for wireless metropolitan-area networks because of its high bit rate and long range. It can support data rates up to 75 Mb/s in a range of 3–5 km and, typically, 20–30 Mb/s in longer ranges. Transmission over longer distances significantly reduces bit rates due to the fact that WiMax does not work efficiently for nonline-of-sight communications. WiMax base stations (BSs) can be placed indoor (installed by customer) or outdoor (installed by network operator) to manage the wireless network. Recently, WiMax is being examined as an alternative for fixed-wired infrastructures, viz., DSL and CM, to deliver "last mile" broadband access to users.

Cellular technology is used for low-bit-rate applications (maximum of 2 Mb/s). A cellular network is mainly used to carry voice traffic and is unoptimized for data traffic. In addition, the data component of the cellular network, such as the high-speed downlink packet access and high-speed uplink packet access, jointly known as high-speed packet access (HSPA) in the third-generation (3G) evolution, can deliver a downstream bandwidth of up to 14 Mb/s and upstream bandwidth of 5 Mb/s. A more advanced version, namely, HSPA+, will offer a downlink speed of up to 40 Mb/s and up to 10 Mb/s in upstream direction. They use Federal Communications Commission regulated expensive spectrum (licensed band) with 3G, beyond-third-generation (B3G), and fourth-generation (4G) standards. WiFi technology, on the other hand, uses the free industrial, scientific, and medical (ISM) band, while WiMax uses both licensed and ISM bands.

There are several industry efforts to build WiMax architecture for commercial deployment, and a few examples are stated as follows. In the U.S., Sprint Nextel holds the license in 2.5-GHz band to build a nationwide wireless-access network, which is expected to cover 100 million U.S. customers in 2008 [8]. Towerstream has deployed wireless networks, which have bit rates of tens of megabits per second, in several locations in the U.S. [9]. Among other regions, Intel WiMax trials have been launched in several locations in Europe and India in collaborations with local service providers [10].

The growing customer demands for bandwidth-intensive services (such as "Quad-play," which refers to voice, video, Internet, and wireless—all are delivered over IP whether on a fixed, mobile, or a hybrid access infrastructure to bring operational efficiencies and convenience to end-users) are accelerating the research efforts needed to design an efficient "last mile" access network in a cost-effective manner. Thus, the radio-on-fiber (ROF) technology has gained momentum, where radio signals can be effectively carried over an existing optical-fiber infrastructure (saving "last mile" costs) by means of the "hybrid fiber radio" (HFR) enabling technology. Recent research works propose ROF-based technologies in millimeter-waveband [11], [12] and demonstrate integrated broadband services in a ROF downstream link [13]. HFR helps to reduce the design complexity at the remote antenna units (RAU) (consequently leading to cheap and simple RAUs), because up/down-conversion, multiplexing/demultiplexing, modulation/demodulation, etc. can be performed at a CO (also known as HFR head end). It is also possible to transmit multiple radio signals over the same fiber. The ROF-enabled access network may have different topologies, such as "optical star–radio P2P," "optical tree–radio star," "optical star–radio cellular," etc. Among various research efforts, Lin [14] proposes a dynamic wavelength-allocation scheme at the bursty traffic load for WDM fiber-radio ring access networks. Reference [15] demonstrates simultaneous wireline (600 MHz) and wireless (5.5 GHz) data transmission in a hybrid fiber-radio access network over cable-service

TABLE I
SAMPLE OF MUNICIPAL MESH NETWORKS

Area/Location	Architecture	Compatibility		Configuration	Operating Range	Player
		Present	Future			
Akron, OH	Flat AP infrastructure	WiFi	WiFi	Multiple radio	2.4 GHz	MobilePro
Athens, GA	Multi-layered MP2P	WiFi	WiFi, WiMax	Multi-radio	2.4, 5 GHz	Belair
Bristol, UK	Multi-layered deployment	WiFi	WiFi, WiMax	Multi-radio multi-antenna	2.4, 5 GHz	Belair
Chaska, MN	Flat deployment	WiFi	WiFi, WiMax	Single radio (omni-directional)	2.4 GHz	Tropos, Pronto
Corpus Christi, TX	Flat deployment (GPS-compatible)	WiFi	WiFi	Single radio (omni-directional)	2.4 GHz	Tropos, Pronto
Culver City, CA	Flat (intermesh capable)	WiFi	WiFi	Multi-radio omni-directional	2.4, 5 GHz	Firetide
Farmers Branch, TX	Gateways with OC-3 ingress	WiFi	WiFi, WiMax	Multiple radio	2.4, GHz	NeoReach, Pronto
Galt, CA	Multi-layered deployment	WiFi	WiFi, WiMax	Multi-radio multi-antenna	2.4, 5 GHz	Belair
Gilbert, AZ	Flat AP infrastructure	WiFi	WiFi	Multiple radio	2.4 GHz	MobilePro
Gordes, France	Flat with intermesh	WiFi	WiFi	Multi-radio omni-directional	2.4, 5 GHz	Firetide
Isla Vista, CA	Flat deployment	WiFi	WiFi	Multi-radio	2.4, 5 GHz	Firetide
Islington, UK	3-tier hierarchical deployment	WiFi	WiFi	Multi-radio multi-antenna	2.4, 5 GHz	Belair
Moorehead, MN	P2MP with fiber optic backbone	WiFi	WiFi	Single radio (omni-directional)	2.4 GHz	Tropos
New Orleans, LA	WiFi routers with digital IP cameras attached to IP-backbone for video surveillance system					Tropos
Philadelphia, PA	Currently being deployed					Earthlink
San Francisco, CA	Currently being deployed					Earthlink, Google
Springfiled, MO	Hierarchical (L2 VLAN capable)	WiFi	WiFi	Multi-radio	2.4, 5 GHz	Belair
St. Maarten, Carribean	Hierarchical deployment	WiFi	WiFi	Multi-radio	2.4, 5 GHz	Belair, Lucent
Tempe, AZ	Gateways with OC-3 ingress	WiFi	WiFi, WiMax	Multi-radio multi-antenna	2.4, 5 GHz	Strix, NeoReach
Wavion, Inc. is a new player with their "spatially adaptive" MIMO-based routers having an antenna array and six radio transreceivers.						

interface specification, and a scheme for quantizing radio signals over fiber is investigated in [16]. A good overview of cost-effective wireless-over-fiber technology is provided in [17].

The rest of this paper is organized as follows. Section II reviews a novel architecture for broadband-access solution [called "hybrid wireless-optical broadband-access network (WOBAN)"], which captures the best of both the optical and wireless worlds and articulates the motivation behind WOBAN. It also summarizes (in Table I) the business drivers deploying an early incarnation of this network all over the world. In Section III, we briefly discuss and evaluate the algorithms for WOBAN deployment (network setup). In addition, some representative data from our survey of locations and types of wireless users in the Wildhorse residential neighborhood of North Davis, CA, are also examined. In Section IV, we discuss the routing characteristics of a WOBAN and study the pros and cons of various routing algorithms. Section V discusses the fault-tolerant behavior of a WOBAN, and Section VI concludes this paper.

This paper reviews in brief our research works on WOBANs (for more details, see the following papers: [18] and [19] for details on the WOBAN architecture presented in Section II; [18]–[20] for details on the WOBAN's network setup problem discussed in Section III; [21] for details on the WOBAN's routing problems and algorithms studied in Section IV; and [22] for details on the WOBAN's fault-tolerant properties outlined in Section V.

II. NOVEL WOBAN ARCHITECTURE

The concept of a hybrid WOBAN is a very attractive one. This is because it may be costly in several situations to run fiber to every home (or equivalent end-user premises) from the telecom CO; in addition, providing wireless access from the CO to every end-user may not be possible because of limited spectrum. Thus, running fiber as far as possible from the CO toward the end-user and then having wireless-access technologies take over may be an excellent compromise. How far should fiber penetrate before wireless takes over is an interesting engineering design and optimization problem.

The WOBAN architecture can be employed to capture the best of both worlds: 1) the reliability, robustness, and high capacity of wireline optical communication and 2) the flexibility ("anytime–anywhere" approach) and cost savings of a wireless network. A WOBAN consists of a wireless network at the front end, and it is supported by an optical network at the back end (see Fig. 1). Noting that the dominant optical-access technology today is the PON, different PON segments can be supported by a telecom CO, with each PON segment radiating away from the CO. Note that the head end of each PON segment is driven by an OLT, which is located at the CO. The tail end of each PON segment will contain a number of ONUs, which typically serve end-users in a standard PON architecture. However, for the proposed hybrid WOBAN, the ONUs will connect to wireless BSs for the wireless portion of the WOBAN. The wireless BSs that are directly connected to the ONUs are known as wireless "gateway routers," because they are the gateways of both the optical and the wireless worlds. Besides these gateways, the wireless front end of a WOBAN consists of other wireless routers/BSs to efficiently manage the network. Thus, the front end of a WOBAN is essentially a multihop wireless mesh network with several wireless routers and a few gateways (to connect to the ONUs and, consequently, to the rest of the Internet through OLTs/CO). The wireless portion of the WOBAN may employ standard technologies such as WiFi or WiMax. Since the ONUs will be located far away from the CO, efficient spectrum reuse can be expected across the BSs with much smaller range but with much higher bandwidth; thus, this WOBAN can potentially support a much larger user base with high bandwidth needs.

In a typical WOBAN, end-users, e.g., subscribers with wireless devices at individual homes, are scattered over a geographic area. An end-user sends a data packet to one of its neighborhood wireless routers. This router then injects the packet into the wireless mesh of the WOBAN. The packet

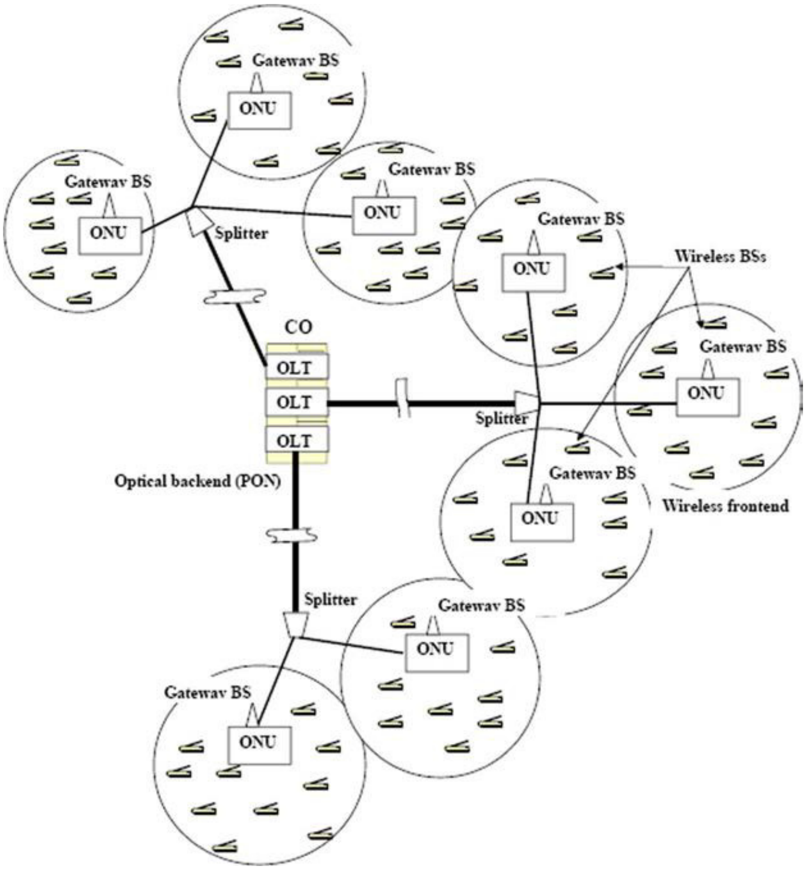

Fig. 1. Hybrid WOBAN architecture.

travels through the mesh, possibly over multiple hops, to one of the gateways (and to the ONU) and is finally sent through the optical part of the WOBAN to the OLT/CO. In the upstream direction of the wireless front end (from a wireless user to a gateway/ONU), the WOBAN is an anycast network, i.e., an end-user can try to deliver its packet(s) to any one of the gateways (from which the packet will find its way to the rest of the Internet). In the optical back end, the upstream (from an ONU to an OLT/CO) of a WOBAN is a multipoint media-access network, where ONUs are deployed in a tree network with respect to their OLT, and they contend for a shared upstream resource (or bandwidth), but in the downstream direction of the wireless front end (from a gateway/ONU to a wireless user), this network is a unicast network, i.e., a gateway will send a packet to only its specific destination (or user). In the optical back end, the downstream (from an OLT/CO to an ONU) of a WOBAN is a broadcast network, where a packet, destined for a particular ONU, is broadcast to all ONUs in the tree and processed selectively only by the destination ONU (all other ONUs discard the packet), as in a standard PON [1]. Fig. 2 captures a WOBAN's upstream- and downstream-transmit modes. A research proposal has been made for a bandwidth-allocation algorithm for an interactive video-on-demand system over a hybrid optical-wireless network in [23].

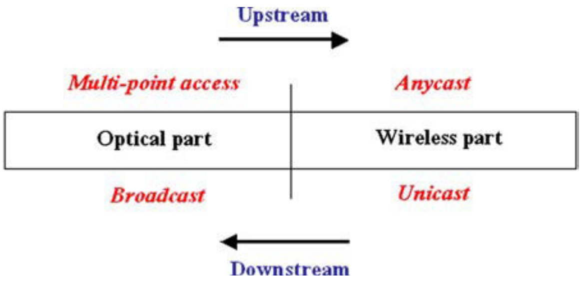

Fig. 2. WOBAN's upstream and downstream protocols.

The WOBAN architecture assumes that an OLT is placed in a telecom CO and that it feeds several ONUs. Thus, from ONU to the CO, we have a traditional fiber network; moreover, from ONUs, end-users are wirelessly connected (in single-hop or multihop fashion).

A common vision of a next-generation converged (fixed and wireless) network is that of the IP-based end-to-end (between the end nodes) network, which enables devices to access common services over one or more networks seamlessly. In a WOBAN, end terminal mobility can be supported at the IP layer by one of the three dominant approaches developed at the Internet Engineering Task Force (IETF), namely, mobile IP, migrate, and host-identity protocol. Mobile IP has unquestionably received the most attention and has already been demonstrated to work well in large networks [24]. Since mobility at the IP layer is an overlay protocol and can be easily supported on a WOBAN, we do not cover it in this paper.

Source #2

Most metropolitan DWDM systems available today depend on the connected SONET equipment to provide protection against node failures or fiber cuts, or they provide optional automatic protection switching. The disadvantages of this approach include the following:

- WDM system duplication, nearly doubling the cost in many cases
- Separate protection system required for each optical channel
- Non-SONET elements not being protected (unless optional APS is employed)

As metropolitan DWDM systems migrate into the access arena, they will be supporting both SONET and native data services, increasing the requirement for protection and restoration in the optical domain. Simple APS is available today on most vendors' equipment, whereas others (Nortel) are beginning to employ electrical crossconnects at the core to provide selectable wavelength protection.

Protection and restoration are not synonymous. In today's network, these functions represent two distinct functions of fiber-optic equipment. Protection refers to the simple, fast (< 50 ms) switching of traffic from one optical route to another predetermined route in the event of a detected failure. SONET equipment performs protection switching today at acceptable rates. Thus, optical layer equipment must perform at least as well to be justified in the network. Protection switching on SONET routes today typically requires 100% excess bandwidth on a given route, which creates additional demand for fiber. Optical line protection will enhance quality of service (QoS) levels for non-SONET traffic—such as ATM and IP—by providing faster restoration than possible in those protocols.

Restoration is a secondary mechanism that can be much slower than protection because it determines routes on-the-fly as nodes fail or become saturated. In optical networks, restoration will be performed by optical crossconnects, most often in mesh topologies. Crossconnects will have the embedded intelligence to select available paths on the network to route wavelengths or entire fibers around saturated or failed nodes. This can lead to more efficient and cost-effective networks as the need for SONET equipment diminishes. Optical layer restoration will be needed for events such as optical amplifier failures, fiber cuts, transponder faults, and SONET LT protection. This will, however, require sophisticated software to compute the efficient alternative routes.

Eventually, restoration will evolve to full wavelength restoration, in which each wavelength will be able to be restored separately. This will require wavelength translation in most cases, but adds the benefit of the most efficient method of utilization of fiber resources. This capability is being realized at the core of networks and not in the metropolitan area at this time.

As noted, most tiered optical bandwidth services proposed today tend to be associated with long-haul network operators. These services usually come in the form of leased OC-n circuits across the wide area and are often wavelength services. In the long-haul network, the deployment of optical switching systems enables this capability, whereas in the metropolitan network, optical switching systems or optical edge systems can provide this capability either at the optical layer (for tiered wave services) or at Layer 1 and 2 for tiered leased line or data services. This kind of flexibility will be appealing to metropolitan carriers that serve ASPs and broadband ISPs because each benefit from the flexible pricing that is associated with tiered bandwidth services and the high degree of customer network management.

Wavelength on Demand

Wavelength on demand is probably the "hottest" MAN service offering from the new breed of metropolitan area carriers with new names like Yipes and Telseon leading the way. The allure is not only cost oriented, which is the intended benefit. The allure can be compared to the exuberance felt by users a couple decades ago as they abandoned the mainframe in favor of doing spreadsheets on their own PC. It definitely connotes "power to the people."

In the long-haul world, wavelength on demand is most often found in the literature of national wholesale network operators. In these networks, idle wavelengths on a backbone trunk can be quickly allocated to other carriers or service providers through the implementation of optical switching systems. These systems allow an operator to treat the optical layer of its network much like it treats the ATM layer: as a pool of available bandwidth within a "cloud" to be quickly allocated in virtual circuits. In the case of optical networks, these virtual circuits are now optical circuits that are managed by optical switching systems using constraint-based routing algorithms. If vendors can develop optical edge equipment that can be agile enough with wavelengths, carriers might find it cost effective in certain instances to offer service providers or major corporate users the opportunity to purchase wavelength services not as a fixed lease or IRU, but as a flexible service. This would require a fully distributed metropolitan DWDM network in which a large percentage of the available interfaces on network equipment were installed and ready to be called into service by the network operator. Although this scenario is feasible in long-haul networks, it might not be in many metropolitan networks with limited DWDM deployments.

The class of optical-edge network gear that includes integrated DWDM functionality might allow some carriers to begin offering this service, although again it would require a widescale deployment of DWDM interfaces throughout a network. Today this comes at a cost of roughly $20,000 per DWDM interface, which is clearly cost prohibitive unless that interface is supporting a revenue-generating service from its initial implementation.

It is clear from all of this that multiple applications at the edge will drive a need for high bandwidth. If the needs for high bandwidth are not addressed, it will eventually lead to a total connectivity bottleneck. FSO can help service providers address this proactively. [2]

FSO in Metropolitan Optical Networks

Now that you understand the overall MAN picture, you need to know how FSO fits into this overall hierarchy. The answer is simple. FSO is an optical technology that can address connectivity needs at any point in the network, be it core, access, or edge. FSO, with its capability to be Layer 1 and protocol transparent, is able to integrate with and interoperate with a variety of network elements and interfaces. This allows it to seamlessly be a part of the growing optical networking family.

Following are some of the common applications using free-space optics in MANs:

- Metropolitan network extensions: FSO can be deployed to extend an existing metropolitan ring or to connect new networks. These links generally do not reach the ultimate end user, but are more an application for the core of the network.
- Enterprise: The flexibility of FSO allows it to be deployed in many enterprise applications, such as LAN-to-LAN connectivity, storage area networking, intracampus connectivities, and so on.
- Last-mile connectivity: These are the links that reach the end user. They can be deployed in PTP, point to multipoint, or mesh connections. Fiber deployment in urban areas could cost \$300,000–\$700,000 given the costs involved in digging tunnels and getting right-of-way. By contrast, a short FSO link of 155 Mbps might cost only \$10,000–\$18,000 or as little as \$166 per month (plus interest) on a 60-month amortization. This is a fairly monumental fact to grasp—the equivalent of three DS-3 lines for \$166 per month! The present cost for three DS3s as leased lines from an ILEC could run as high as \$10,000 or more per month!
- Fiber complement: FSO can also be deployed as a redundant link to back up fiber. Most operators who are deploying fiber for business applications connect two fibers to secure a reliable service plus backup in the event of outage. Instead of deploying two fiber links, operators could opt to deploy an FSO system as the redundant link.
- Access: FSO can also be deployed in access applications such as Gigabit Ethernet access. Service providers can use FSO to bypass local loop systems and to provide FSO-based high-capacity links to businesses.
- Backhaul: FSO can be used for backhaul such as LMDS or cellular backhaul, as well as for Gigabit Ethernet "off-net" to transport network backhaul.
- DWDM services: With the integration of WDM and FSO systems, independent players aim to build their own fiber rings, yet might own only part of the ring. Such a solution could save rental payment to ILECs, which are likely to take advantage of this situation.

Summary

POTS, SONET, wireless, first-generation optical networks, second-generation optical networks, and now free-space optics—this is quite a transition over a couple of decades. Although most of the other applications were new and disruptive changes in the telecommunications networks, free-space optics was not. Unknown to most people, free-space optics has been around for more than three decades, but interestingly enough, due to multiple market drivers, free-space optics has found a renewed value-added interest. It is fast becoming a value-added application for MANs that are enabling service providers to accelerate their deployment of optical networks, thus addressing the needs of their end users quickly and cost effectively.

With the evident growth in optical networks, it is clear that the connectivity bottleneck will continue to be shifting problems all across the optical networks. It is also clear that although innovation is key to such a growth, cost reduction is also a driving force. The all-optical network is focused on decreasing the cost per bit and making optical capacity available to the end users. Alas, some dreams are not easily realized, and the vision of the all-optical network finds itself in this dilemma of cost versus infrastructure.

To address and enable the acceleration of optical networks while addressing the need to be cost effective, free-space optics is presenting the users with an opportunity to do so. FSO is a perfect fit for the growing MANs fitting into multiple areas and not just last mile. Regardless of whether you use free-space optics in the core, access, or edge, one thing is clear: FSO addresses the connectivity bottleneck of today.

Sources

[1] These three paragraphs relating to storage area networks were taken from Chapter 3 of the report by Pioneer Consulting, LLC, "Optical Edge Networks: Market Opportunities for Integrated Optical Network Solutions in Metro Networks." August 2000. `http://www.pioneerconsulting.com/report.php3?report=13`

[2] Much of the material presented in the VPN Services section through the Wavelength on Demand section was taken from Chapter 3 of the report by Pioneer Consulting, LLC, "Optical Edge Networks: Market Opportunities for Integrated Optical Network Solutions in Metro Networks." August 2000. `http://www.pioneerconsulting.com/report.php3?report=13`

Source #3

1

chapter

Introduction to Optical Networks

As we begin the new millennium, we are seeing dramatic changes in the telecommunications industry that have far-reaching implications for our lifestyles. There are many drivers for these changes. First and foremost is the continuing, relentless need for more capacity in the network. This demand is fueled by many factors. The tremendous growth of the Internet and the World Wide Web, both in terms of number of users and the amount of time, and thus bandwidth taken by each user, is a major factor. Internet traffic has been growing rapidly for many years. Estimates of growth have varied considerably over the years, with some early growth estimates showing a doubling every four to six months. Despite the variations, these growth estimates are always high, with more recent estimates at about 50% annually. Meanwhile, broadband access technologies such as digital subscriber line (DSL) and cable modems, which provide bandwidths per user on the order of 1 Mb/s, has been deployed widely. For example, in 2008 about 55% of the adults in the United States had broadband access at home, while only 10% had access through dialup lines of 28–56 kb/s. Fiber to the home has shown steady growth with Asian markets showing the highest market penetration.

At the same time, businesses today rely on high-speed networks to conduct their businesses. These networks are used to interconnect multiple locations within a company as well as between companies for business-to-business transactions. Large corporations that used to lease 155 Mb/s lines to interconnect their internal sites are commonly leasing 1 Gb/s connections today.

There is also a strong correlation between the increase in demand and the cost of bandwidth. Technological advances have succeeded in continously reducing the cost of bandwidth. This reduced cost of bandwidth in turn spurs the development of a new set of applications that make use of more bandwidth and affects behavioral patterns. A simple example is that as phone calls get cheaper, people spend more time on the phone. This development in turn drives the need for more bandwidth in the network. This positive feedback cycle shows no sign of abating in the near future.

Another factor causing major changes in the industry is the deregulation of the telephone industry. It is a well-known fact that monopolies impede rapid progress. Monopolistic companies can take their time adapting to changes and have no incentive to reduce costs and provide new services. Deregulation of these monopolies has stimulated competition in the marketplace, which in turn has resulted in lower costs to end users and faster deployment of new technologies and services. Deregulation has also resulted in creating a number of new start-up service providers as well as start-up companies providing equipment to these service providers.

Also, traffic in a network is dominated by data as opposed to traditional voice traffic. In the past, the reverse was true, and so legacy networks were designed to efficiently support voice rather than data. Today, data transport services are pervasive and are capable of providing quality of service to carry performance sensitive applications such as real-time voice and video.

These factors have driven the development of high-capacity optical networks and their remarkably rapid transition from the research laboratories into commercial deployment. This book aims to cover optical network technologies, systems, and networking issues, as well as economic and other deployment considerations.

1.1 Telecommunications Network Architecture

Our focus in this book is primarily on the so-called *public* networks, which are networks operated by *service providers*, or *carriers*, as they are often called. Carriers use their network to provide a variety of services to their customers. Carriers used to be essentially telephone companies, but today there are many different breeds of carriers operating under different business models, many of whom do not even provide telephone service. In addition to the traditional carriers providing telephone and leased line services, today there are carriers who are dedicated to interconnecting Internet service providers (ISPs), carriers that are in the business of providing bulk bandwidth to other carriers, and even virtual carriers that provide services without owning any infrastructure.

In many cases, the carrier owns the facilities (for example, fiber links) and equipment deployed inside the network. Building fiber links requires right-of-way privileges. Not anybody can dig up streets! Fiber is deployed in many different ways

11 chapter

Access Networks

In previous chapters, we have explored the use of optical networks for metro and long-haul network applications. The *access network* is the "last leg" of the telecommunications network that runs from the service provider's facility to the home or business. With fiber now directly available to many office buildings in metropolitan areas, networks based on SONET/SDH or Ethernet-based technologies are being used to provide high-speed access to large business users. Business users are big consumers of data services, many of which are delivered in the form of leased lines at various speeds ranging from 1.5 Mb/s to several gigabits per second. While this is happening, the telephone and cable companies are also placing a significant emphasis on the development of networks that will allow them to provide a variety of services to individual homes and small to medium businesses. This is the focus of this chapter.

Today, homes get essentially two types of services: plain old telephone service (POTS) over the telephone network and broadcast analog video over the cable network. Recently added to this mix have been data services for Internet access using either digital subscriber line (DSL) technology over the telephone network or cable modem service over the cable network.

Early efforts to develop high-capacity access networks were devoted to developing networks that would accommodate various forms of video, such as video-on-demand and high-definition television. However, the range of services that users are expected to demand in the future is vast and unpredictable. Today, end users

Table 11.1 Different types of services that must be supported by an access network. The bandwidth requirements are given for each individual stream.

Service	Type	Downstream Bandwidth	Upstream Bandwidth
Telephony	Switched	4 kHz	4 kHz
ISDN	Switched	144 kb/s	144 kb/s
Broadcast video	Broadcast	6 MHz or 19 Mb/s	0
Interactive video	Switched	6 Mb/s	Small
Internet access	Switched	A few Mb/s	A few Mb/s
IPTV	Switched	1–20 Mb/s	Small
Video-on-demand	Switched	1–20 Mb/s	Small
Videoconferencing	Switched	6 Mb/s	6 Mb/s
Business services	Switched	1.5 Mb/s–10 Gb/s	1.5 Mb/s–10 Gb/s

are interested in both Internet access and other high-speed data access services, for such applications as telecommuting, distance learning, entertainment video, and videoconferencing. Future, unforeseen applications are also sure to arise and make ever-increasing demands on the bandwidth available in the last mile. The term *full service* encompasses the variety of services that are expected to be delivered via access networks. A sampling of the different services and their characteristics is given in Table 11.1. Both telephone and cable companies are striving to become full-service providers.

At a broad level, these services can be classified based on three major criteria. The first is the bandwidth requirement, which can vary from a few kilohertz for telephony to tens of megabits per second per video stream or even tens of gigabits per second for high-speed leased lines. The second is whether this requirement is *symmetric* (two way), for example, videoconferencing, or *asymmetric* (one way), for example, broadcast video. Today, while most business services are symmetric, other services tend to be asymmetric, with more bandwidth needed from the service provider to the user (the downstream direction) than from the user to the service provider (the upstream direction). The last criterion is whether the service is inherently broadcast, where every user gets the same information, for example, broadcast video, or whether the service is switched, where different users get different information, as is the case with Internet access.

In the next section, we provide an overview of the different types of existing and emerging access network architectures. We then provide a more detailed description of the two most promising access architectures—the hybrid fiber coax (HFC) network and the fiber to the curb (FTTC) approach and its variants.

11.1 Network Architecture Overview

In broad terms, an access network consists of a hub, remote nodes (RNs), and network interface units (NIUs), as shown in Figure 11.1. In the case of a telephone company, the hub is a *central office* (also called a *local exchange* in many parts of the world), and in the case of a cable company, it is called a *head end*. Each hub serves several homes or businesses via the NIUs. An NIU either may be located in a subscriber location or may itself serve several subscribers. The hub itself may be part of a larger network, but for our purposes, we can think of the hub as being the source of data to the NIUs and the sink of data from the NIUs. In many cases, rather than running cables from the hub to each individual NIU, another hierarchical level is introduced between the hub and the NIUs. Each hub may be connected to several RNs deployed in the field, with each RN in turn serving a separate set of NIUs. The network between the hub and the RN is called the *feeder* network, and the network between the RN and the NIUs is called the *distribution* network.

We saw that services could be either broadcast or switched. In the same way, the distribution network could also be either broadcast or switched. Note that in the context of services, we are using the terms *broadcast* and *switched* to denote whether or not all users get the same information. In the context of the network, we are referring to the network topology. Different combinations of services and network topologies are possible—a broadcast service may be supported by a broadcast or a switched network, and a switched service may be supported by a broadcast or a switched network. In a broadcast network, an RN broadcasts the data it receives from the feeder network to all its NIUs. In a switched network, the RN processes the data coming in and sends possibly separate data streams to different NIUs. The telephone network that we will study later is a switched network, whereas the cable television network is a broadcast network. Broadcast networks may be cheaper than switched networks, are well suited for delivering broadcast services, and have the advantage that all the NIUs are identical, making them easier to deploy. (In some switched networks that we will study, different NIUs use different wavelengths, which makes it more complicated to manage and track the inventory of NIUs in the network.) Switched networks, as their name suggests, are well suited for delivering switched services and provide more security. For example, it is not possible for one subscriber to tap into another subscriber's data, and it is more difficult for one subscriber to

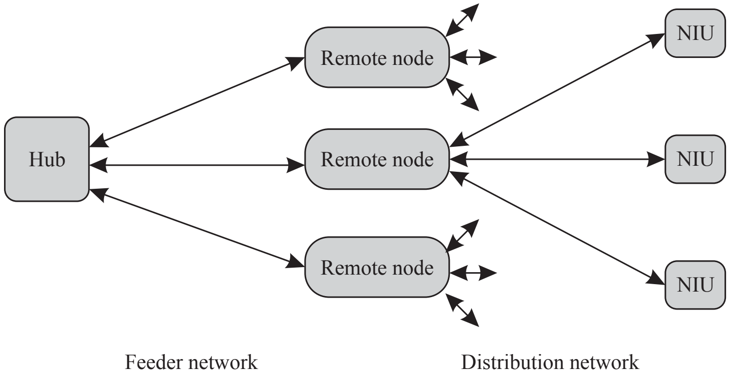

Figure 11.1 Architecture of an access network. It consists of a hub, which is a telephone company central office or cable company head end, remote nodes deployed in the field, and network interface units that serve one or more individual subscribers.

corrupt the entire network. Fault location is generally easier in a switched network than in a broadcast network. In broadcast networks, the "intelligence" is all at the NIUs, whereas in switched networks, it is in the network. Thus, NIUs may be simpler in switched networks than in broadcast networks.

Another way of classifying access networks is based on the type of feeder network, which is the network between the hub and the RN. In one scenario, the feeder network could assign each NIU its own *dedicated* bandwidth. By dedicated bandwidth, we mean that different NIUs are assigned different frequency (or wavelength) bands in the frequency (or wavelength) domain. In another scenario, the feeder network could have a total bandwidth that is *shared* by all the NIUs. By shared bandwidth, we mean that multiple NIUs share a given bandwidth in the time domain. In this case, each NIU could potentially access the entire bandwidth for short periods. For upstream transmission from the NIUs back to the hub, we will need some form of media access control to coordinate access to the shared bandwidth by the NIUs. If the traffic from/to the NIUs is bursty, it is more efficient to share a large total amount of bandwidth among many NIUs rather than assign each NIU its own dedicated bandwidth. On the other hand, with dedicated bandwidth, each NIU can be guaranteed a certain quality of service, which is more difficult to do with shared bandwidth. A disadvantage of the shared bandwidth approach is that each NIU must have optics/electronics that operate at the total bandwidth of the network as opposed to the bandwidth needed by the NIU.

The telephone and cable networks are vastly different. The telephone network provides very little bandwidth per home but incorporates sophisticated switching equipment and operations and management systems. The cable network provides a lot of bandwidth to each home, but it is all unidirectional and broadcast, with no switching and very simple management.

Several approaches have been used to upgrade the access network infrastructure to support the emerging set of new services. The *integrated services digital network* (ISDN) provides 144 kb/s of bandwidth over the existing twisted-pair infrastructure. The *digital subscriber line* (DSL) is another technique that works over the existing infrastructure but provides significantly more bandwidth than ISDN. It uses sophisticated modulation and coding techniques to realize a capacity of a few megabits per second over twisted pair, which is sufficient to transmit compressed video. This requires that the central office (CO) and the home each have a DSL modem. However, DSL has some limitations. The realizable bandwidth is inversely proportional to the distance between the CO and the home, and with today's technology, we can achieve several hundred kilobits per second to a few megabits per second over this infrastructure. The existing twisted-pair infrastructure incorporates several 4 kHz filters that must be removed. The bandwidth on the upstream (return) path is severely limited to a few hundred kilobits per second. Many variations and enhancements of DSL have been proposed. As in the conventional telephone network, ISDN and DSL can be classified as switched networks with dedicated bandwidth per NIU.

Satellites provide another way of delivering access services. The direct broadcast satellite system uses a geosynchronous satellite to broadcast a few hundred channels to individual homes. A satellite may provide more bandwidth than a terrestrial coaxial cable system. However, the main problem is that, unlike terrestrial systems, the amount of spatial reuse of bandwidth possible is quite limited, since a single satellite has a wide coverage area within which it broadcasts the signals. Also, there is no easy way to handle the upstream traffic. Today, it is possible to have high-speed Internet access delivered via satellite, with the upstream direction carried over a regular telephone line.

Wireless access is yet another viable option. Although it suffers from limited bandwidth and range, it can be deployed rapidly and allows providers without an existing infrastructure to enter the market. Among the variants are the *multichannel multipoint distribution service* (MMDS) and the *local multipoint distribution service* (LMDS), both of which are terrestrial line-of-sight systems. MMDS provides thirty-three 6 MHz channels in the 2–3 GHz band with a range of 15 to 55 km, depending on the transmit power. LMDS operates in the 28 GHz band with 1.3 GHz of bandwidth and is suitable for short-range (3–5 km) deployment in dense metropolitan areas (the distance is also dependent on the amount of rainfall, as rain attenuates signals in this band). LMDS is part of a family of wireless communication standards, IEEE 802.16 or commonly known as WiMAX. These standards can provide up to 70 Mb/s of symmetric bandwidth and up to a distance of 50 km. They have a variety of applications, including point-to-point links and portable Internet access. WiMAX can operate in a wide range of frequencies below 66 GHz, including 2.3 GHz to 3.5 GHz in the licensed spectrum and 5 GHz in the public spectrum.

A common wireless access technology to the Internet by laptop computers and other personal computing devices is the IEEE 802.11 wireless local-area network technology. It operates in the 2.5 and 5 GHz public spectrum and can provide data rates of about 50 Mb/s. They are limited by a very short range of tens of meters to an access point or "hot spot." These hot spots are often found in airports, coffee shops, restaurants, and hotels. They can be connected to the Internet in a number of ways including WiMAX.

Optical fiberless systems using lasers transmitting over free space into the home are also being developed as an alternative approach. These systems can provide about 622 Mb/s of capacity over a line-of-sight range of 200 to 500 m.

In the context of the next-generation access network, the two main architectures being considered today are the so-called hybrid fiber coax (HFC) approach and the fiber to the curb (FTTC) approach. The HFC approach is still a broadcast architecture, whereas the FTTC approach incorporates switching.

11.2 Enhanced HFC

Although we have used the term *HFC* to describe the existing cable infrastructure, this same term is used to describe an upgraded version of this architecture, which we will refer to as an *enhanced* HFC architecture. Since both the fiber and the coax cable carry multiple subcarrier modulated streams, and it is a broadcast network, a better term to describe the HFC architecture is *subcarrier modulated fiber coax bus* (SMFCB). The network architecture is essentially the same as that shown in Figure 11.3. In order to provide increased bandwidth per user, the network is being enhanced using a combination of several techniques. First, the transmitted frequency range can be increased, for example, up to 1 GHz from the 500 MHz in conventional HFC systems. Enhanced HFC systems deployed in larger metropolitan areas can deliver up to 862 MHz of bandwidth. Within each subcarrier channel, we can use spectrally efficient digital modulation techniques, such as 256 QAM (quadrature amplitude modulation), which provides a spectral efficiency of 8 bits/Hz. In addition, we can drive fiber deeper into the network and reduce the number of homes served by a remote node down to about 50 homes, from the 500 homes typically served by an HFC network. We can

Source #4

Chapter 1

Optical Networking: Principles and Challenges

1.1 Introduction

After experiencing rapid growth during the late 90s, the telecom industry in general, and optical networking in particular, has been experiencing some challenging times over the past several years. (An analysis of the underlying reasons will be offered later in this chapter.) Nevertheless, even though the telecom business market is unsettled today (but showing signs of improvement), we need to be ready with the appropriate technologies and engineering solutions to meet the growing bandwidth needs of our information society.

Optical networking using wavelength-division multiplexing (WDM) – the term WDM will be explained shortly in this chapter – is the technology of choice for meeting these growing demands [Mukh97, Mukh00]. While there may be an abundance of dark fiber and WDM transmission capacity today, we believe that there is – and there will continue to be – a tremendous need for optical switching equipment, namely high-capacity and high-density optical crossconnects (OXCs), for managing high-capacity optical signals. These technologies can be exploited by various categories of telecom businesses as outlined later in this chapter.

The rest of this chapter is organized as follows. Section 1.2 provides an overview of telecom networks. Section 1.3 discusses various categories of telecom business models. Section 1.4 makes the case for the important role software plays in bringing cost-effective and intelligent optical networking to the marketplace. Section 1.5 emphasizes the role of cross-layer design, analysis, and thinking for successful deployment of optical networks. Section 1.6 discusses the role of traffic engineering vs. network engineering vs. network planning in our networking investigations. Section 1.7 tries to clarify the question: "What is an Optical Network?" Section 1.8 starts with the basic characteristics of optics which can be exploited for optical networks. Section 1.9 clarifies the terms xDM vs xDMA. Section 1.10 introduces wavelength-division multiplexing (WDM). Section 1.11 outlines WDM networking evolution. Section 1.12 illustrates some WDM network constructions. Section 1.13 discusses some economic studies indicating the benefits of WDM. Section 1.14 outlines some important research problems and challenges in the WDM networking field today. Section 1.15 concludes this chapter with a "road map" of the rest of the book.

1.2 Telecom Network Overview

Figure 1.1 provides an overview of telecommunication networks. They consist of the access network, the metropolitan-area (or regional) network, and the backbone network.

The access network enables end-users (businesses and residential customers) to get connected to the rest of the network infrastructure. The access network spans a distance of a few kilometers (perhaps up to 20 km as some local exchange carriers (LECs) seem to prefer). Our current solutions for access are dial-up modems, higher-speed lines (such as T1/E1), digital subscriber line (DSL), and cable modem. However, the access network continues to be a bottleneck, and users require (and are demanding) higher bandwidth to be delivered to their machines. How to provide this high bandwidth in an inexpensive manner is a key R&D priority. Passive optical networks (PONs) based on inexpensive, proven, and ubiquitous Ethernet technology (and referred to as EPON) seem an attractive proposition for this market segment. PON technology in general, and EPON in particular, will be studied in Chapter 5.

The metro-area network typically spans a metropolitan region, cover-

Chapter

12

Light-Tree: Optical Multicasting

12.1 Introduction

This chapter studies architectures and approaches for establishing multicast connections in a WDM mesh network using "light-trees". It also discusses schemes for protecting multicast trees in a mesh network.

Advances in optical WDM networking have made bandwidth-intensive multicast applications such as HDTV, interactive distance learning, live auctions, distributed games, movie broadcasts from studios, etc., widely popular [Paul98, Mill99, MaZQ98, SuGT01]. These applications require *point-to-multipoint* (PtMP) connections from a source node to the destination nodes in a network. Multicasting provides an easy means to deliver messages to multiple destinations without requiring too much message replication.

An optical signal passing through an optical wavelength-routing switch (WRS) may be routed from an input fiber to an output fiber without undergoing opto-electronic conversion. As we already know, a *lightpath* is a end-to-end wavelength-routed channel connecting a transmitter at a source node to a receiver at a destination node, which may be used to carry circuit-switched traffic, and it may span multiple fiber links.

connection. The receipt of a signaling message usually triggers some local action, e.g., allocation of time slots or wavelengths when receiving a connection-setup message. Two protocols – resource reservation protocol (RSVP) and constraint-based routing label distribution protocol (CR-LDP) with traffic-engineering (TE) extensions – have been proposed as the standard signaling protocols in the GMPLS control plane. In a heterogeneous WDM network, the route computed for a connection request will be composed of a sequence of intermediate node identifiers as well as link bundle id. Since multiple candidate links may exist in a link bundle, an intermediate node needs to select one for the connection request when it needs to configure the optical crossconnect (OXC) and establish the connection. If the link is a bundled lightpath link, then, based on the available capacity of each lightpath in the bundle and the bandwidth requirement of the request, different link-selection schemes can be used, e.g., random selection, first-fit selection, best-fit selection, etc.

4. **Fault management**:
In an optical network, the high capacity of a link has the problem that a link failure can potentially lead to the loss of a large amount of data (and revenue). So, we need to develop appropriate protection and restoration schemes which minimize the data loss when a link failure occurs (see Chapter 11). Relative to the optical layer, upper layers of protocols (such as ATM, IP, and MPLS) have their own procedures to recover from link failures [ADDH94, Huit95, MSOH99]. However, the recovery time for upper layers is significantly larger (on the order of seconds), whereas we prefer that the fault-recovery times at the optical layer should be on the order of milliseconds in order to minimize data losses. Furthermore, it is beneficial to consider fault-recovery mechanisms in the optical layer for the following reasons [Gers98]: (a) the optical layer can efficiently multiplex protection resources (such as spare wavelengths and fibers) among several higher-layer network applications, and (b) survivability at the optical layer provides protection to higher-layer protocols which may not have built-in fault recovery.

Essentially, there are two types of fault-recovery mechanisms [Gers98, Wu92, Wu98] (see also Chapter 11). If backup resources (routes and wavelengths) are pre-computed and reserved in advance, we call it a *protection scheme* [RaMu99a, FCMJ99]. Otherwise, when a failure occurs, if another route and a free wavelength have to be discovsubprobl[redacted] solution for [redacted]ution to the whole problem. Moreover, this approach requires all the traffic requests to be known in advance, which cannot be satisfied in dynamic grooming.

Another approach is to solve the four subproblems as a whole. Since it can take into account all the constraints regarding the four subproblems simultaneously, this approach has a potential to achieve better performance. With static traffic, the traffic-grooming problem can be formulated as an integer linear program (ILP) [3], and an optimal solution can be obtained for some relatively small networks. However, an ILP is not scalable and cannot be directly applied to large networks. One way to make the problem tractable is to develop heuristic algorithms and jointly solve the grooming problem for one connection request at a time. To the best of our knowledge, no integrated heuristic algorithm for solving the traffic-grooming problem has been developed for wavelength-routed networks in previous work.

A. Previous Work

Traffic grooming is an important and practical problem for designing WDM networks and it is receiving increasing research attention both in academia and in industry. The work in [12] reviews most of the recent research work on traffic grooming in WDM ring and mesh networks.

Past research efforts on traffic grooming have mainly focused on SONET/WDM ring networks. The major cost of such a network is considered to be dominated by SONET add-drop multiplexers (ADMs). Therefore, minimizing the number of SONET ADMs has been the objective of static traffic grooming in recent research. The general traffic-grooming problem in a SONET/WDM ring network is proven to be NP-complete [13], [14]. An optimal algorithm for a single-hub ring is proposed in [13] and several optimal or near-optimal algorithms for traffic grooming and wavelength assignment to reduce the number of wavelengths and SONET ADMs are proposed in [15]. As a network design problem, the authors in [16] attempt to minimize the network cost, which is dominated by SONET ADMs, in an optical add-drop wavelength-division-multiplexed (OADM) ring network. Six optical WDM ring architectures are provided in [16] and the cost of different architectures, as well as the switching capabilities of different architectures under various traffic assumptions are compared. The maximum terminal-equipment savings using wavelength ADMs are quantified in [17] for WDM rings carrying uniform and distance-dependent traffic. Grooming with arbitrary traffic in bidirectional-line-switched rings (BLSRs) is addressed in [14]. In [18], based on a general formulation of the virtual-topology problem, a framework used to evaluate the performance of heuristics and requiring less computation than evaluating the optimal solution is presented. The authors in [19] formulate the grooming optimization problem as an ILP and compare single-hop grooming and multihop grooming. Instead of single-ring architectures, interconnected WDM rings are studied in [20] and several strategies for traffic grooming in such networks are compared. All the above references except [16] focus on static traffic only. The authors of [21] study the dynamic traffic-grooming problem in SONET/WDM rings and formulate it as a bipartite graph-matching problem.

As our fiber-optic backbone networks migrate from rings to mesh, traffic grooming on WDM mesh networks becomes an extremely important area of research. The work in [22] formulates the static traffic-grooming problem as an ILP and proposes a heuristic to minimize the number of transceivers. In [23], several lower bounds for regular topologies are presented and greedy and iterative greedy schemes are developed. However, in both [22] and [23], the authors relax the physical-topology constraints, assuming all the virtual topologies are implementable on the given physical topology, i.e., they do not consider lightpath routing and wavelength assignment. The authors in [3] propose several node architectures for supporting traffic grooming in WDM mesh networks and formulate the static traffic-grooming problem as an ILP. They present two heuristics and compare the performance with that of the ILP. The works in [24]–[29] consider a dynamic traffic pattern in WDM mesh networks. In [24], the authors propose a connection admission control scheme to ensure fairness in terms of connection blocking. A theoretical capacity correlation model is presented in [25] to compute the blocking probability for WDM networks with constrained grooming capability. In [26], two route-computation algorithms are proposed and compared, and the results indicate that, in order to achieve good performance in a dynamic environment, different grooming policies and route-computation algorithms need to be used under different network states. The work in [27] compares two schemes to dynamically establish reliable low-speed traffic in WDM mesh networks with traffic-grooming capability. In [28], the problem of planning and designing a WDM mesh network with certain forecast traffic demands, to satisfy all the connections as well as minimize the network cost, is studied. In [29], the authors investigate the design of multilayer mesh networks to satisfy each connection's bandwidth and protection requirements while minimizing the overall network cost.

B. Challenges of Traffic Grooming in a Heterogeneous WDM Mesh Network

The WDM backbone network is expected to emerge as a multivendor heterogeneous network. As WDM networks migrate from ring topologies to mesh topologies, it is very important to solve the traffic-grooming problem in a heterogeneous mesh network environment.

In terms of wavelength-conversion capability, heterogeneity means that some of the nodes in a network may have full wavelength-conversion capability (any incoming wavelength can be converted into any outgoing wavelength), some may have no wavelength-conversion capability (traffic must stay on the same wavelength when bypassing these nodes) [30], [31], and some may have partial wavelength-conversion capability (some wavelengths can be converted into some other wavelengths) [32]–[35]. In previous work, however, it was assumed that all the nodes in a network either have wavelength-conversion capability or none has wavelength-conversion capability. In addition, if a node has this capability, it always has full wavelength-conversion capability. This all-or-nothing assumption may not be practical or valid in the future WDM network. It is necessary to address the partial and sparse wavelength-conversion scenarios.