EMERGING OPTICAL NETWORK TECHNOLOGIES

Recent Related Titles

Metropolitan Area WDM Networks: An AWG-Based Approach
Martin Maier
ISBN 1-4020-7574-X, 2003
http://www.wkap.nl/prod/b/1-4020-7574-X

Optical Networks: Architecture and Survivability
Hussein T. Mouftah and Pin-Han Ho
ISBN 1-4020-7196-5, 2002
http://www.wkap.nl/prod/b/1-4020-7196-5

WDM Mesh Networks: Management and Survivability
Hui Zang
ISBN 1-4020-7355-0, 2002
http://www.wkap.nl/prod/b/1-4020-7355-0

EMERGING OPTICAL NETWORK TECHNOLOGIES
Architectures, Protocols and Performance

Edited by
KRISHNA M. SIVALINGAM
University of Maryland, Baltimore County

SURESH SUBRAMANIAM
George Washington University

Springer

eBook ISBN: 0-387-22584-6
Print ISBN: 0-387-22582-X

Created in the United States of America

Visit Springer's eBookstore at: http://www.ebooks.kluweronline.com
and the Springer Global Website Online at: http://www.springeronline.com

This book is dedicated to our families.

Contents

Dedication v

Contributing Authors xi

Preface xiii

Part I NETWORK ARCHITECTURES

1
Enabling Architectures for Next Generation Optical Networks 3
Linda Cline, Christian Maciocco and Manav Mishra

2
Hybrid Hierarchical Optical Networks 23
Samrat Ganguly, Rauf Izmailov and Imrich Chlamtac

3
Advances in Passive Optical Networks (PONs) 51
Amitabha Banerjee, Glen Kramer, Yinghua Ye, Sudhir Dixit and Biswanath Mukherjee

4
Regional-Metro Optical Networks 75
Nasir Ghani

Part II SWITCHING

5
Optical Packet Switching 111
George N. Rouskas and Lisong Xu

6
Waveband Switching: A New Frontier in Optical WDM Networks 129
Xiaojun Cao, Vishal Anand, Yizhi Xiong, and Chunming Qiao

7
Optical Burst Switching 155
Hakki Candan Cankaya and Myoungki Jeong

Part III SIGNALING PROTOCOLS AND NETWORK OPERATION

8
GMPLS-based Exchange Points: Architecture and Functionality 179
Slobodanka Tomic and Admela Jukan

9
The GMPLS Control Plane Architecture for Optical Networks 193
David Griffith

10
Operational Aspects of Mesh Networking in WDM Optical Networks 219
Jean-Francois Labourdette, Eric Bouillet and Chris Olszewski

Part IV TRAFFIC GROOMING

11
Traffic Grooming in WDM Networks 245
Jian-Qiang Hu and Eytan Modiano

12
Efficient Traffic Grooming in WDM Mesh Networks 265
Harsha V. Madhyastha and C. Siva Ram Murthy

Part V PROTECTION AND RESTORATION

13
A Survey of Survivability Techniques for Optical WDM Networks 297
Mahesh Sivakumar, Rama K. Shenai and Krishna M. Sivalingam

14
Tradeoffs and Comparison of Restoration Strategies in Optical WDM Networks 333
Arun K. Somani

15
Facilitating Service Level Agreements with Restoration Speed Requirements 357
Gokhan Sahin and Suresh Subramaniam

16
Failure Location in WDM Networks 379
Carmen Mas, Hung X. Nguyen and Patrick Thiran

Part VI TESTBEDS

17
A Multi-Layer Switched GMPLS Optical Network 403
Aihua Guo, Zhonghua Zhu and Yung J. (Ray) Chen

18
HORNET: A Packet Switched WDM Metropolitan Network 423
Kapil Shrikhande, Ian White, Matt Rogge and Leonid G. Kazovsky

Index 449

Contributing Authors

Vishal Anand, *State University of New York, College at Brockport*

Amitabha Banerjee, *University of California, Davis*

Eric Bouillet, *Tellium, Inc.*

Hakki Candan Cankaya, *Alcatel USA*

Xiaojun Cao, *State University of New York at Buffalo*

Yung Chen, *University of Maryland, Baltimore County*

Imrich Chlamtac, *University of Texas at Dallas*

Linda Cline, *Intel Corporation*

Sudhir Dixit, *Nokia Research Center*

Samrat Ganguly, *NEC Labs America Inc.*

Nasir Ghani, *Tennessee Technological University*

David Griffith, *NIST*

Aihua Guo, *University of Maryland, Baltimore County*

Jian-Qiang Hu, *Boston University*

Rauf Izmailov, *NEC Labs America Inc.*

Myoungki Jeong, *Alcatel USA*

Admela Jukan, *Vienna University of Technology*

Leonid Kazovsky, *Stanford University*

Glen Kramer, *Teknovus, Inc.*

Jean-Francois Labourdette, *Tellium, Inc.*

Christian Maciocco, *Intel Corporation*

Harsha Madhyastha, *University of Washington*

Carmen Mas, *AIT, Greece*

Manav Mishra, *Intel Corporation*

Eytan Modiano, *MIT*

Biswanath Mukherjee, *University of California, Davis*

Hung Nguyen, *EPFL, Switzerland*

Chris Olszewski, *Tellium, Inc.*

Chunming Qiao, *State University of New York at Buffalo*

Matt Rogge, *Stanford University*

George Rouskas, *North Carolina State University*

Gokhan Sahin

Ramakrishna Shenai, *University of Maryland, Baltimore County*

Kapil Shrikhande, *Stanford University*

Mahesh Sivakumar, *University of Maryland, Baltimore County*

Krishna Sivalingam, *University of Maryland, Baltimore County*

C. Siva Ram Murthy, *Indian Institute of Technology, Chennai (Madras)*

Arun Somani, *Iowa State University*

Suresh Subramaniam, *George Washington University*

Patrick Thiran, *EPFL, Switzerland*

Slobodanka Tomic, *Vienna University of Technology*

Ian White, *Stanford University / Sprint Advanced Technology Laboratories*

Yizhi Xiong, *State University of New York at Buffalo*

Lisong Xu, *North Carolina State University*

Yinghua Ye, *Nokia Research Center*

Zhonghua Zhu, *University of Maryland, Baltimore County*

Preface

Optical networks have moved from laboratory settings and theoretical research to real-world deployment and service-oriented explorations. New technologies such as Ethernet PON and optical packet switching are being explored, and the landscape is continuously and rapidly evolving. Some of the key issues involving these new technologies are the architectural, protocol, and performance aspects.

The objective of this book is to present a collection of chapters from leading researchers in the field covering the above-mentioned aspects. Articles on various topics, spanning a variety of technologies, were solicited from active researchers in both academia and industry. In any book on such a quickly growing field, it is nearly impossible to do full justice to all of the important aspects. Here, rather than attempting to cover a large ground with a limited treatment of each topic, we focus on a few key challenges and present a set of papers addressing each of them in detail. It is our hope that the papers will be found to have sufficient detail for the new entrant to the field, and at the same time be a reference book for the experienced researcher.

This book is aimed at a wide variety of readers. The potential audience includes those who are interested in a summary of recent research work that cannot be found in a single location; those interested in survey and tutorial articles on specific topics; and graduate students and others who want to start research in optical networking. We hope that readers gain insight into the ideas behind the new technologies presented herein, and are inspired to conduct their own research and aid in further advancing the field.

Organization of the book

The book is divided into six parts, each dealing with a different aspect: network architectures, switching, signaling protocols, traffic grooming, protection and restoration, and testbeds. At least two chapters have been selected for each part, with three or more chapters for most parts.

Part I is on network architectures and contains four chapters. The first chapter by Cline, Maciocco and Mishra from Intel Labs takes a look into the services and architectures for next generation optical networks. The second chapter by researchers from NEC Labs and UT Dallas presents a hybrid hierarchical network architecture wherein both all-optical and OEO switching co-exist within a cross-connect. Chapter 3 summarizes recent developments in passive optical network (PON) architectures. This chapter is written by researchers from UC Davis, Teknovus, and Nokia Research. Chapter 4, by Nasir Ghani of Tennessee Technological University, presents a detailed survey of the recent activities in regional and metro network architectures.

Part II focuses on switching and consists of three chapters. The first chapter presents an overview of optical packet switching and is written by Rouskas and Xu of North Carolina State University. Chapter 6, by researchers from the SUNY at Buffalo and Brockport, presents waveband switching OXC architectures, and algorithms for grouping wavelengths into wavebands. The last chapter of Part II is on the third main switching paradigm, namely optical burst switching (OBS). The article, written by researchers from Alcatel and Samsung, reviews OBS concepts and describes the work on OBS done at Alcatel USA.

Signaling protocols are the subject of Part III. The first chapter, by Tomic and Jukan of the Vienna University of Technology, discusses the architecture and functionality of GMPLS-enabled exchange points. The second chapter by David Griffith of NIST presents the GMPLS protocol framework including RSVP-TE, OSPF-TE, and LMP. Chapter 10, authored by three researchers from Tellium, explains the benefits and operational aspects of mesh optical networks.

Part IV contains two chapters on traffic grooming. The first chapter by Hu and Modiano introduces a simple traffic grooming problem and then presents various modifications and solution techniques. The next chapter by Madhyastha and Murthy presents a specific architectural solution for efficient traffic grooming.

Part V is dedicated to protection and restoration. The first chapter by Sivakumar, Shenai, and Sivalingam presents a survey of survivability techniques. The next chapter by Somani focuses on routing "dependable" connections and presents a novel solution. The following chapter by Sahin and Subramaniam presents a new strategy of scheduling restoration control messages to provide quality of protection in mesh networks using capacity sharing. The last chapter in this part, written by Mas, Nguyen, and Thiran, discusses methods to locate failures in WDM networks.

The final part of the book consists of two chapters describing the testbeds built at UMBC and Stanford. In Chapter 17, a multi-layered GMPLS optical network testbed is described and Chapter 18 describes the HORNET packet switched metro network developed at Stanford.

We invite you to sit back and read about the recent research in optical networking presented in these chapters and hope that it stirs your creativity and imagination leading to further innovations and advances in the field.

Acknowledgments

Naturally, this book would not have been possible without the time and effort of the contributing authors, and we are grateful to them. Each of the chapters selected were proofread by the editors and their graduate students who have also spent considerable time in taking care of the little details that make the book right. We also like to acknowledge the valuable assistance of Minal Mishra, Rama Shenai, Manoj Sivakumar, Mahesh Sivakumar and Sundar Subramani, graduate students at the University of Maryland, Baltimore County; and Tao Deng, Sunggy Koo, and Venkatraman Tamilraj at George Washington University.

We also gratefully acknowledge our research sponsors who provided partial support for this work. This includes DARPA under grant No. N66001-00-18949 (co-funded by NSA), National Science Foundation under grant Nos. ANI-0322959 and ANI-9973111, Cisco Systems and Intel Corporation.

We thank Kluwer Academic Publishers for the opportunity to publish this book. We are especially thankful to Alex Greene and Melissa Sullivan at Kluwer Academic Publishers for their constant help and patience, without which this book would not have been possible.

Krishna Sivalingam
Associate Professor
University of Maryland, Baltimore County
Email: krishna@umbc.edu

Suresh Subramaniam
Associate Professor
George Washington University
Email: suresh@gwu.edu

May 2004

Book Editor Biographies

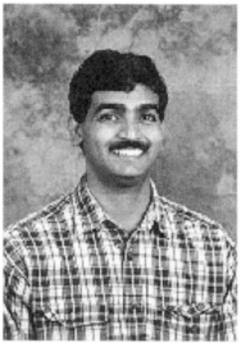

Krishna M. Sivalingam (ACM '93) is an Associate Professor in the Dept. of CSEE at University of Maryland, Baltimore County. Previously, he was with the School of EECS at Washington State University, Pullman from 1997 until 2002; and with the University of North Carolina Greensboro from 1994 until 1997. He has also conducted research at Lucent Technologies' Bell Labs in Murray Hill, NJ, and at AT&T Labs in Whippany, NJ. He received his Ph.D. and M.S. degrees in Computer Science from State University of New York at Buffalo in 1994 and 1990 respectively; and his B.E. degree in Computer Science and Engineering in 1988 from Anna University, Chennai (Madras), India. While at SUNY Buffalo, he was a Presidential Fellow from 1988 to 1991.

His research interests include wireless networks, optical wavelength division multiplexed networks, and performance evaluation. He holds three patents in wireless networks and has published several research articles including more than twenty-five journal publications. He has published an edited book titled "Wireless Sensor Networks" in 2004 and an edited book titled "Optical WDM networks" in 2000. He is serving as a Guest Co-Editor for a special issue of ACM MONET on "Wireless Sensor Networks" in 2004 and, in the past, served as Guest Co-Editor for a special issue of ACM MONET on "Wireless Sensor Networks" (2003) and an issue of IEEE Journal on Selected Areas in Communications on optical WDM networks (2000). He is co-recipient of the Best Paper Award at the IEEE International Conference on Networks 2000 held in Singapore. His work has been supported by several sources including AFOSR, NSF, Cisco, Intel and Laboratory for Telecommunication Sciences. He is a member of the Editorial Board for ACM Wireless Networks Journal, IEEE Transactions on Mobile Computing, and KICS Journal of Computer Networks.

He is serving as Steering Committee Co-Chair for the First International Conference on Broadband Networks 2004 (www.broadnets.org); and as Technical Program Co-Chair for the First IEEE Conference on Sensor and Ad Hoc Communications and Networks (SECON) to be held in Santa Clara, CA in 2004. He has served as General Co-Chair for SPIE Opticomm 2003 (Dallas, TX) and for ACM Intl. Workshop on Wireless Sensor Networks and Applications (WSNA) 2003 held on conjunction with ACM MobiCom 2003 at San Diego, CA. He served as Technical Program Co-Chair of OptiComm conference at Boston, MA in July 2002. He is a Senior Member of IEEE and a member of ACM.

Suresh Subramaniam received the Ph.D. degree in electrical engineering from the University of Washington, Seattle, in 1997. He is an Associate Professor in the Department of Electrical and Computer Engineering at the George Washington University, Washington, DC. He is interested in a variety of aspects of optical and wireless networks including performance analysis, algorithms, and design. His research has been supported by DARPA, DISA, NSA, and NSF.

Dr. Subramaniam is a co-editor of the book "Optical WDM Networks: Principles and Practice" published by Kluwer Academic Publishers in 2000. He has been on the program committees of several conferences including IEEE Infocom, IEEE ICC, and IEEE Globecom, and is TPC Co-Chair for the 2004 Broadband Optical Networking Symposium, part of the First Conference on Broadband Networks (www.broadnets.org). He serves on the editorial boards of Journal of Communications and Networks and IEEE Communications Surveys and Tutorials. He is a co-recipient of the Best Paper Award at the 1997 SPIE Conference on All-Optical Communication Systems.

I

NETWORK ARCHITECTURES

Chapter 1

ENABLING ARCHITECTURES FOR NEXT GENERATION OPTICAL NETWORKS

Linda Cline, Christian Maciocco and Manav Mishra
Intel Labs, Hillsboro OR 97124
Email: linda.s.cline@intel.com, christian.maciocco@intel.com

Abstract As the demand grows for higher network access speeds, technologies such as optical fiber have begun to overtake traditional copper wire for data transport in short haul networks as well as long haul networks. Optical networking plays a growing role in next generation networks with new capabilities such as LCAS (Link Capacity Adjustment Scheme) and Virtual Concatenation (VC), and services such as dynamic provisioning and traffic grooming. While these emerging capabilities hold the promise of an intelligent optical network, there are still obstacles. Protocols and standards to support these capabilities are still evolving. In addition, in order to realize the new benefits, carriers and providers must invest in new optical equipment, as well as upgrades to existing equipment. In the current economic environment, a choice which leverages lower cost equipment with software which can provide advanced functionality is significantly more attractive than expensive alternatives. In addition, upgradeable software- based components provide future cost savings as well as flexibility in supporting new and changing protocols and standards. In this paper, we discuss each of these issues in detail and present a solution for optical services and applications, including Optical Burst Switching, using a network processor based platform to overcome the obstacles facing next generation optical networks.

Keywords: Optical Networking, SONET/SDH, Network Processors, GMPLS, UNI, Link Capacity Adjustment Scheme, Traffic Grooming, Optical Burst Switching.

1.1 Introduction

New capabilities and services for optical networks combined with optical fiber pushing toward the edge require continued investment in equipment and upgrades to support these new functions. This equipment needs to be flexible to support the networks of today as well as the capabilities for tomorrow. An architecture that is flexible enough to support this type of investment for the future is one that leverages software to augment less complex, and thus less expensive, hardware. Optical network nodes need to support changing network protocols and increased complexity in functionality. Use of a mass produced, inexpensive network processor that is optimized for network processing functions and completely programmable in software, provides an appropriate platform for these nodes. By implementing the complexity in software, there is increased adaptability to protocol upgrades for continued cost savings.

In this chapter, we discuss the problems and requirements of an intelligent optical network, and provide a solution describing the use of a software framework implemented on a network processor based optical platform.

In Section 1.2, we discuss several of the emerging optical services which are required by next generation optical networks, as well as some of the issues surrounding them. In Section 1.3, we provide an overview of network processors. Section 1.4 discusses the various software building blocks which can be used to implement the next generation optical services. In Section 1.5, we present a solution for Optical Burst Switching, which is a next generation optical application. Finally, Section 1.6 summarizes the choice of a network processor platform as an enabler for the continuously evolving optical networking technology.

1.2 Next-generation Optical Services

Next-generation optical services will support more customers and provide greater bandwidth in access networks. This capability requires new supporting services to be provided by the underlying networks. These services include automated optical provisioning, sophisticated traffic grooming, and services that ease management of networks with ever increasing complexity. These services are described in more detail in the subsequent sections.

1.2.1 Optical Provisioning

In current networks, setting up an optical connection to send SONET/SDH [11,12] frames from one location to another is a manual process. Typically, a Network Management System (NMS) is configured by one or more humans to add each new connection. It is not unusual for the turnaround time for a new connection to take up to six weeks to configure after the initial request has been submitted. Once a human has begun directly configuring the NMS software, the completion of the task may still take several minutes or hours. Provisioning that takes months or minutes may be acceptable, if not desirable, for setting up long haul connections which may be in place for long periods of time. However, as optical networking moves to the metro area network (MAN), this delay in provisioning connections becomes less acceptable. Access connections for the MAN have a finer granularity in bandwidth requirements and are more transient than long haul connections. Quantities of service connection or service modification requests will increase rapidly, which can swamp a provisioning system which is accomplished manually. Dynamic, automated provisioning is vital if service providers are going to meet the rigorous turnaround time and scalability requirements of MANs. Dynamic provisioning can also improve operational expenditures by reducing the need for human control, improving time to revenue for new services.

Support for dynamic provisioning is beginning to emerge, although today this is typically implemented using proprietary means. Such proprietary schemes make end to end automated provisioning not possible except where certain carriers control the complete paths. Efforts are underway in standards groups to define protocols for dynamic provisioning, which may solve the end to end problem eventually. Currently, these standards are moving targets, which magnifies the need for programmable network nodes which can easily be updated as new versions are defined or protocols modified. We talk about just a few of these protocols for illustration.

One aspect of automation in provisioning involves the configuration of end to end connections. In the past this has been primarily accomplished through manual means, but there are currently efforts underway to define standard signaling protocols such as the GMPLS (Generalized Multi-Protocol Label Switching) suite of protocols [1][2][3], to automate some of this process. One such standards effort is UNI (User-Network Interface) [4], defined at the Optical Internetworking Forum (OIF). In brief, UNI provides an interface by which a client may request services (i.e. establishment of connections) of an optical network. By supporting dynamic connection requests, end to end provisioning can be accomplished.

LCAS [8] is another area where efforts are being made in automation of provisioning. LCAS is a recent SONET based protocol that allows a particular connection to be resized (to adjust the capacity or bandwidth). It utilizes Virtual Concatenation (VC) [9], a method for providing SONET/SDH virtual connections in a variety of sizes, that supports flexibility as well as better bandwidth utilization. Combined, these two mechanisms can support dynamic changes to connections and their capacities, which allows new virtual connections to be easily integrated into the SONET/SDH multiplex, or existing connections to be given more or less bandwidth. Smaller granularities of bandwidth can be supported and increased dynamically, making SONET/SDH a viable alternative to Ethernet for metro carriers. Addition of bandwidth on demand will allow service providers to be much more responsive to transient customer bandwidth needs, enabling better utilization of empty fiber along with addition of premium services for short term bandwidth bursts.

Once connection provisioning can be automated, additional services can be developed that utilize this automation, such as intelligent protection and restoration schemes that do not rely on expensive hardware redundancy, and may provide better restoration by creating fall back routes which avoid points of failure. Network Management Systems (NMS) can take advantage of these services for more resilient and fine grained manageability of the optical network.

1.2.2 Traffic Grooming

Another service which has great importance in the next generation optical network, especially for access networks, is traffic grooming. Traffic grooming refers to efficient multiplexing at the ingress of a network. Typically, it is used to group lower-rate traffic flows into higher-rate traffic flows in such a manner that add/drop operations are minimized. Grooming is a composite solution employing various traffic pattern, engineering, topology and routing schemes. Grooming can be employed at MAN gateways to exhaustively utilize bandwidth in an intelligent manner. There are three main components of traffic grooming for next generation optical networks: admission control, traffic management, and LCAS/VC.

Admission control ensures that the customers adhere to their Traffic Conditioning Agreements as specified by their SLAs (Service Level Agreements). This helps to support Authentication, Authorization and Accounting (AAA) of the customers. It also supports policing of the customer traffic flows and enforcement of domain policies. If a customer's flow exceeds the SLA, then a back pressure message (i.e. Ethernet PAUSE flow control message) can be sent to the customer to initiate a slow-down in the rate of traffic.

Once traffic has been authenticated and authorized, traffic management deals with queuing and scheduling of the incoming traffic flows onto the various egress queues available. The scheduler usually doubles as a shaper as well and thereby ensures that the traffic is pumped onto the network based on a profile characteristic to the network.

Use of the LCAS/VC feature of next-generation SONET networks allows the service provider to over-provision bandwidth on existing channels, which ensures rapid provisioning of services to customers. This feature also enables the service provider to add new customers to its clientele without making forklift or cumbersome upgrades to the network infrastructure.

Figure 1.1 illustrates a deployment scenario for traffic grooming at a metro gateway where numerous gigabit Ethernet lines are aggregated and provisioned over an outgoing PoS (Packet over SONET) or EoS (Ethernet over SONET) line for transport across the core of the network. Unlike traffic engineering, which is end-to-end, traffic grooming is done primarily at the ingress of the

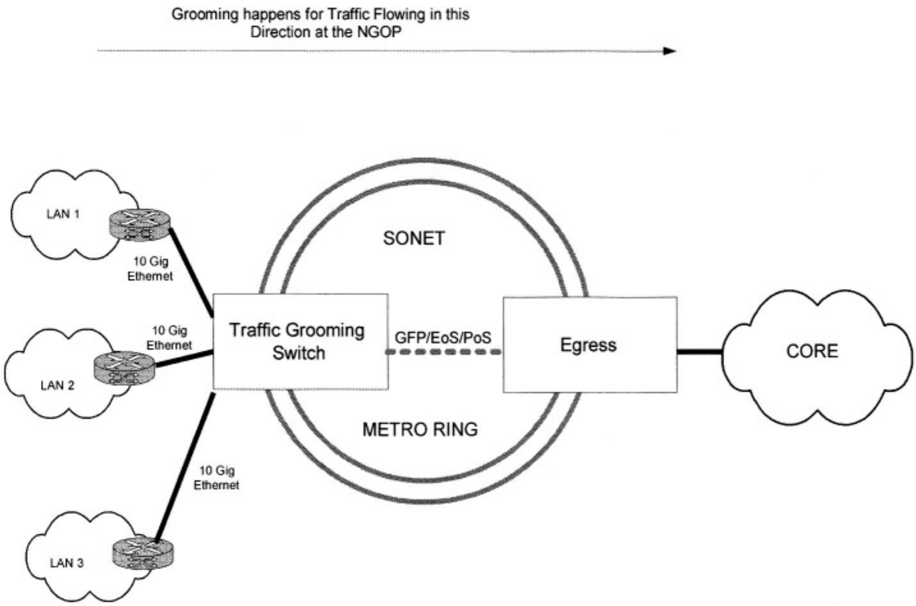

Figure 1.1. Traffic grooming switch.

MAN, as this is where a major aggregation of trunk lines and gigabit Ethernet lines happens.

1.2.3 Automated Device Control

Network management in optical networks has traditionally been implemented as a centralized control. As complexity in optical devices and networks increases, and the number of managed devices grows, it becomes an increasingly difficult management problem to centralize all functions. Network elements can off-load some of the NMS tasks, if they are capable of handling additional processing. This may include better statistics gathering and alarm/event correlation, support at the device level for some levels of automated provisioning, support for some policy administration at the device level, and higher level, easy to use interfaces for device configuration to ease the work of administrators.

One of the scalability problems with a large optical network is the sheer volume of statistics and events that must be analyzed and processed at the NMS. A single hardware failure can escalate into a large number of alarms which need to be handled with great efficiency to isolate the failure and select a solution or a workaround. A link failure can cause these alarm notifications to be generated from all affected network elements. As the size of the network grows and the number of elements increases, this can swamp a centralized management system. The network element can handle some correlation of multiple alarms and events if it can accommodate the analysis processing. Gathering of fine grained statistics and coherent summarization can be supported at the network element level and propagated in summary form to the NMS. This can relieve some of the obvious scalability problems.

While end to end path provisioning may be better served by a central NMS, some levels of provisioning as well as policy administration can be supported at the network element level. The NMS could delegate select policies for administration directly at the optical device level. This might include SLA information for traffic flows, or certain admission control policies. Local decisions about LCAS initiation and processing based on information provided by the NMS could support lower levels of automated provisioning directly at the device level.

Obviously, as the capabilities of optical network devices become more sophisticated, there is additional complexity in programming these devices. Remote administration, compatibility with existing as well as emerging management interface standards, and high level, easy to use functions with fine grained control are among the requirements of the interfaces supported by the optical network element.

All of these new capabilities and services that will be present in the next generation optical platform impose new demands on the network and the equipment used to support it. In addition, standards for the new protocols and interfaces to support these new services are still in development and subject to industry acceptance. Some of these standards are described in the next section.

1.2.4 Standards for Tomorrow's Optical Networks

We briefly describe some of the standards efforts that are geared toward improving automation of provisioning and adding intelligent capabilities to the optical network. The primary efforts in this space are GMPLS and UNI.

GMPLS Overview. GMPLS [1-3] is not a single protocol, but a collection of protocols being defined by the IETF, offering a consolidated control plane and extending topology awareness and bandwidth management across all network layers, thus enabling new, more efficient and cost effective core network architectures. The potential of GMPLS is that it makes possible the evolution to peer-based networks where all network elements have information about all other elements. The GMPLS suite of protocols is applicable to all types of traffic and provides mechanisms for data forwarding, signaling and routing on a variety of data plane interfaces. GMPLS enhances MPLS to additionally support Packet Switched Capable interfaces (PSC), Time Division Multiplexing Capable interfaces (TDMC), Lambda Switch Capable interfaces (LSC) and Fiber Switch Capable interfaces (FSC).

GMPLS extends IP technology to control and manage lower layers. In order to establish a connection, the GMPLS control plane is using a routing protocol, e.g. OSPF or IS-IS, to maintain route information, and also a signaling protocol, e.g. RSVP-TE or CR-LDP, to provide the messaging functionality. GMPLS also manages TE (Traffic Engineering) links, where combining multiple data links for routing purposes forms a single TE link, through the use of LMP (Link Management Protocol) running between neighboring nodes.

The intelligent optical network uses GMPLS to dynamically establish, provision, maintain / tear down, protect and restore, groom and shape traffic to make efficient use of SONET/SDH, WDM, or OTN networks. These features allow operators/service providers to offer new services over these networks such as bandwidth-on-demand, efficient traffic grooming, etc. GMPLS supports the establishment of labels with traffic engineering attributes between end-points. A GMPLS domain consists of two or more Label Edge Routers (LER) connected by Label Switched Routers (LSR). Label Switched Paths (LSPs) are established between pair of nodes along the path to transfer packets across the domain. An LSP consists of an ingress LER, one or more LSRs or OCX (Optical Cross-Connect) and an egress LER. For all packets/frames received at the ingress, the LER determines which packets should be mapped to which particular LSP based on packet classifications, i.e. destination address, source address, protocol port, etc. Multiple LSPs can be established between any ingress and egress LER pair. Figure 1.2 shows a GMPLS domain consisting of two LERs and two LSRs.

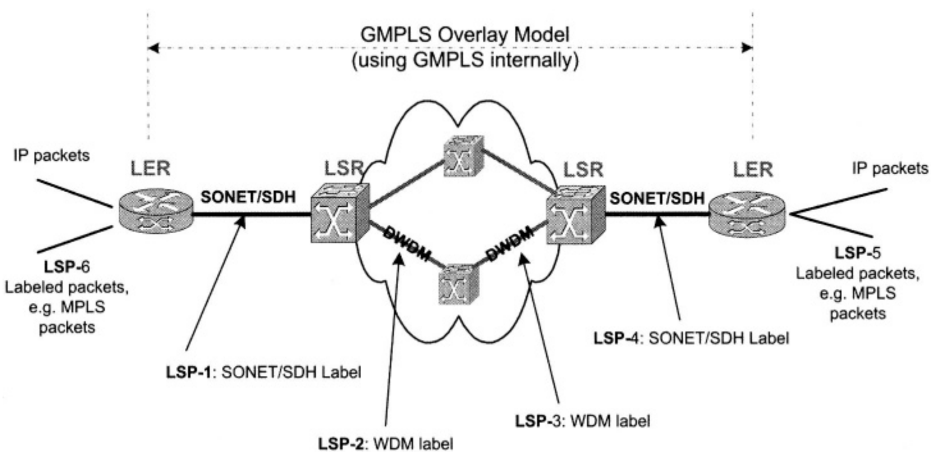

Figure 1.2. GMPLS Overlay Model (UNI).

UNI overview. New technologies in optical networking such as Dense Wave Division Multiplexing (DWDM) have evolved to provide a cost effective means to increase the bandwidth capacity of optical fiber networks as well as create a new optical network layer that can provide intelligent transport services to allow IP routers, MPLS label switching routers, ATM switches, etc., to interconnect using SONET/SDH or other future interfaces. To support provisioning of end to end connections across multi-vendor networks, a standard method of signaling to create these connections is required. The Optical Internetworking Forum (OIF) has been working toward defining such standards. The standard interface that provides a service control interface between the transport network and the client is called the User- Network Interface (UNI) [4]. The initial specification developed is Optical UNI 1.0.

UNI 1.0 allows a client from the user network to dynamically initiate and establish an optical connection with a remote node using GMPLS signaling. A neighbor discovery mechanism is defined, which allows a client (termed UNI-C) and the network node (termed UNI-N) which supports the client's request, to discover each other. In addition, a service discovery mechanism is included, which allows the client to discover the services offered by an optical network.

UNI 2.0 specification is underway and addresses features such as security, bandwidth modification, extension to physical layers such as Ethernet, and the ability to establish multiple connections with a single request.

While the work with UNI holds promise for paving the way to fully automated provisioning and advanced optical network services and management, it is still very much a work in progress. OIF and IETF are separate standards organizations, and thus may not be completely in agreement or consensus. IETF has not settled on a single signaling protocol for GMPLS, but OIF has selected CR-LDP as the signaling protocol for UNI. IETF supports a peer to peer model for GMPLS, whereas OIF, heavily guided by telecom companies, defines an overlay model. It can be difficult for equipment vendors to decide how to develop their implementations if they wish to become early adopters. The standards are currently moving targets, which means that any implementation of the protocols in these standards needs to accommodate frequent modifications and updates. The most flexible and cost-effective solution for this is to provide as much of the implementation in software as possible.

1.2.5 Optical Services: Advantages and Issues

In summary, emerging optical services provide the following advantages and features:

- Automated provisioning for finer grained and more efficient network management

- Traffic grooming for better traffic and customer management

- Network management capabilities to support scalable and more resilient and flexible networks
- Standards to allow inter-vendor interoperation

However, the following obstacles remain:

- Standards specifications are still unstable
- Implementions require increased complexity
- Upgrades to legacy optical equipment may be unavoidable

Next, we talk about network processors and the solution they provide for overcoming these obstacles.

1.3 Programmable Silicon and Network Processors

Historically, manufacturers have employed fixed-function ASICs (application specific integrated circuits) to perform packet processing in network devices at line rates. However, ASICs can be expensive to revise if adding or changing protocol functionality, and can be difficult to program. The complexity of required packet processing is directly related to the number and sophistication of the supported protocols. For IP networks, protocols such as IPv6 and MPLS are being added to the required lineup of supported protocols in network devices. For optical networks, the list grows with the addition of protocols such as LCAS, VC, GMPLS, CR-LDP, etc.

Instead of utilizing fixed-function ASICs to support growing and evolving packet processing functions, a network processor is a fast, but more flexible and programmable device that could constitute a better choice. Network processors are capable of processing packets at line rates, but provide users with the programmability of a generic processor. Use of commercial off-the-shelf network processors can eliminate development time for custom ASICs, and better support adaptation to changing customer requirements or evolving standards by allowing software updates [6].

A typical lineup of packet processing functions handled by a network processor could include header classification, deep packet inspection and analysis, packet processing, policing, statistics, and traffic management. The network processor could work in conjunction with a control plane processor to handle control and exception packets that are detected. Routing tables can be stored in SRAM, TCAM, or DRAM. Packets can be stored internally or in an external DRAM. Standard interfaces such as SPI-3, SPI- 4, CSIX [5], etc. can be used to connect network processors to framer or switch fabric devices, or to co-processors such as those used to perform encryption, etc.

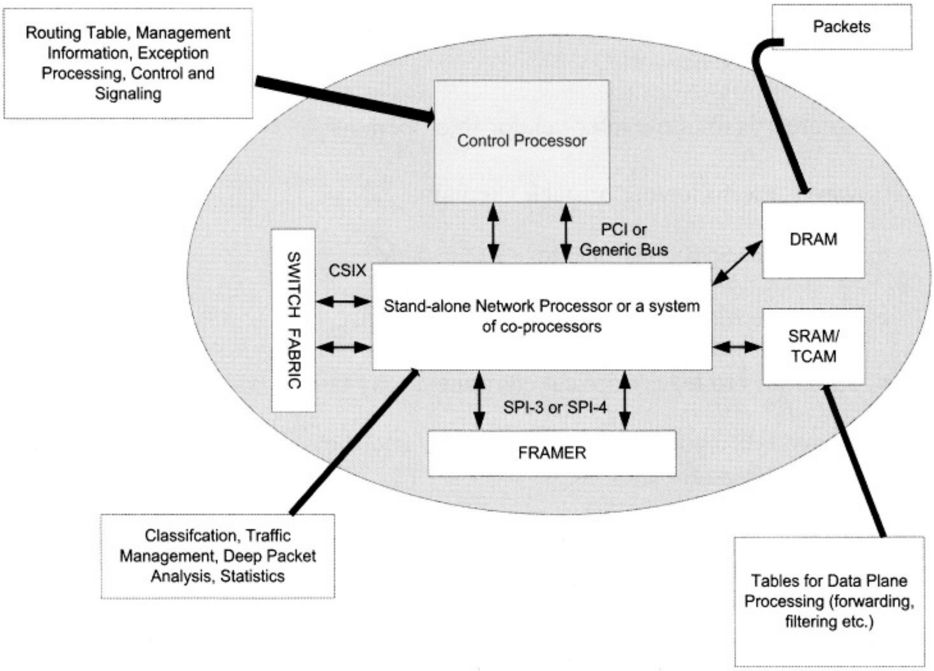

Figure 1.3. A generic network processor.

PCI or generic bus interfaces can be used to connect the network processor to a control plane processor. Figure 1.3 shows the architecture of a generic network processor and the functionality of the various components.

1.3.1 Network processor programming and software framework

With a software based platform, it becomes possible to create modular and reusable components that can form the basis for more sophisticated processing functions. These components become software building blocks which are aggregated in different ways to suit different and more sophisticated applications. This type of software development strategy can be followed and implemented in software development kits that are made available with network processors.

Such development kits emphasize network elements which support separate but interoperating control and data planes. This allows independent development of control and data path software. Typically, a host network processor, host operating system, higher-layer software, and client APIs run on a control plane processor. Software which runs on the control plane includes protocol processing for routing and signaling, exception handling and control protocols. Data plane software includes packet processing functions which need to handle data at line rates, also known as fast-path code. The fast-path code runs on a network processor. Many network nprocessor vendors supply reference fast-path code as part of their development kits.

Another model that is an integral part of the development strategy is the use of pipelined processing to handle packet data along the fast-path. An ingress software module receives packets from a hardware interface, and passes these along to one or more modules for classification, filtering, policing and shaping before passing to an egress module for transmission on hardware once again. As network processors gain in capability and speed, this type of fast-path handling could grow from simple functions such as filtering and forwarding, to deep packet analysis and more. If these processing modules represent different aspects of protocol handling, they can be replaced, or mixed and matched with future modules for updated or layered protocol handling. Vendors can take modules with standard protocol behaviors and insert their own special processing functions within a pipeline for value added functionality, without having to implement all of the standard behaviors themselves.

The software building block and pipelined processing model has the flexibility and potential to support a suite of network processing and protocol functions that can grow and adapt with the needs of future networks. We next discuss some specific examples of building blocks that can help solve some of the problems with development for next generation optical networks.

1.4 Software Building Blocks for the Intelligent Optical Network

The control and data plane separation and the pipelined functional building blocks models as described above provide an architectural platform implemented in software which utilizes the flexibility of a network processor and supports development for the services and applications of the next generation optical networks. Designed well, these functional blocks are modular and reusable, and can be used to build various combinations to suit a wide variety of networking functions. When protocols are updated to reflect changes in evolving standards, or to include value added vendor processing, only the affected blocks need to be updated. Common blocks can be provided with equipment in the form of libraries or reference code, and can be utilized as is, or modified for differences or additions in supported features. In this section, examples will be presented which illustrate in detail the way in which several next generation optical services can be implemented.

1.4.1 Traffic Grooming Data Plane Blocks

An example architecture of a traffic grooming switch is shown in Figure 1.4 below. The figure shows control plane based software modules which are re-

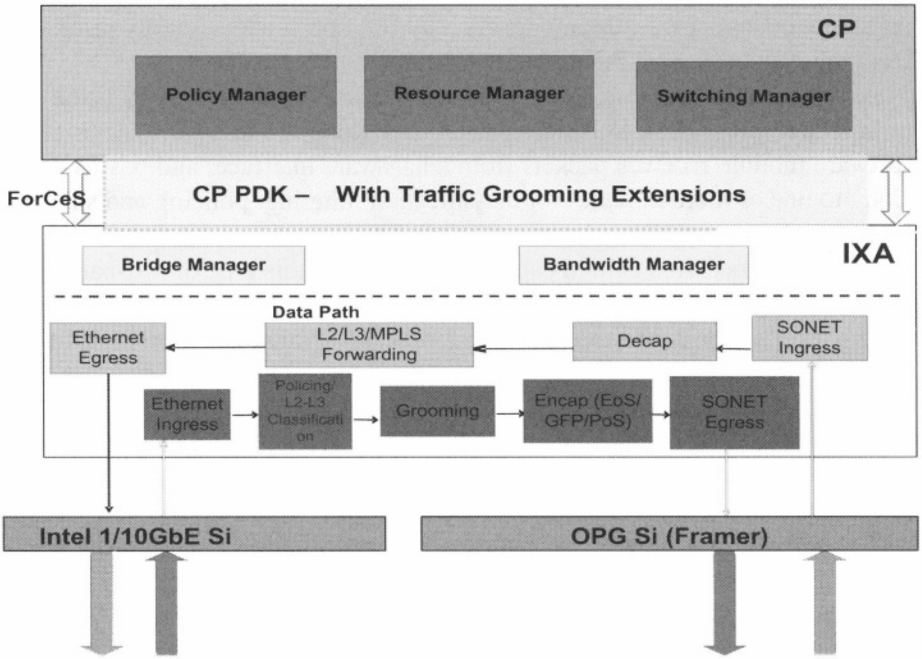

Figure 1.4. Traffic grooming architecture.

sponsible for policies, signaling, and management. These communicate with the data plane blocks through the ForCES [7] (Forwarding and Control Element Separation) protocol defined by IETF. The traffic grooming data plane blocks are responsible for authenticating, authorizing, accounting, provisioning of QoS of the incoming packet data, based on information downloaded from the control and management planes. The various data plane components for traffic grooming are shown in Figure 1.5. A brief overview on these blocks is provided below.

Classifier. A classifier is a functional data path element which consists of filters that select matching and non matching packets. Based on this selection, packets are forwarded along the appropriate data path within the gateway. Therefore, a classifier splits a single incoming traffic stream into multiple outgoing streams. This traffic grooming solution could employ an IEEE 802.1 p/q based classifier, or a Layer 3 classifier or a MPLS classifier. The IEEE 802.1 p/q classifier is employed for metro gateways with incoming Gigabit Ethernet or 10-Gigabit Ethernet links, as the VLAN ID in the header maps to the client's identity, and the priority bits map to the class of service for the flow.

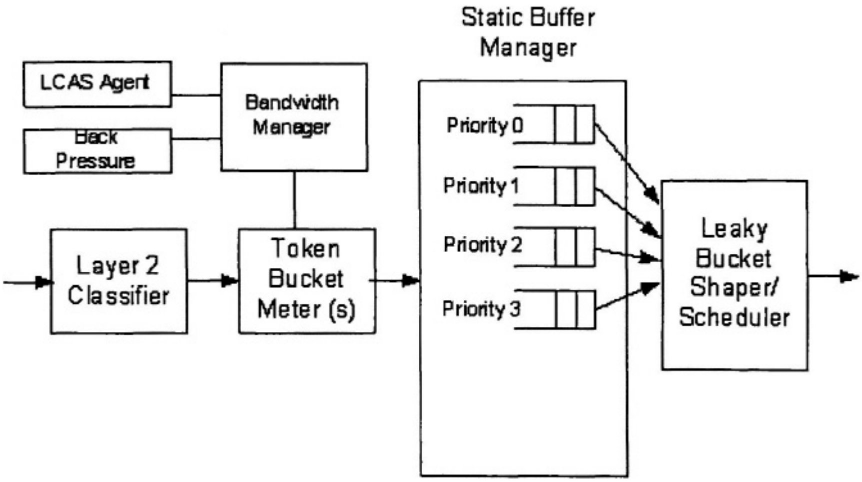

Figure 1.5. Traffic grooming data plane blocks.

Meter. A meter is a data path functional element that monitors the temporal characteristics of a flow and applies different actions to packets based on the configured temporal profile for that flow. A token bucket meter is used for the traffic grooming solution as it measures the conformance of a flow profile against the traffic profile specified by the SLA.

Bandwidth Manager. The bandwidth manager is a control block responsible for: Resource Management Admission Control through Ethernet backpressure based throttling Provisioning of Bandwidth using LCAS, when oversubscription occurs and excess bandwidth is available.

Buffer Manager. A buffer manager is a queuing element which modulates the transmission of packets belonging to the different traffic streams and determines their ordering, possibly storing them temporarily or discarding them. Packets are usually stored either because there is a resource constraint (e.g., available bandwidth) which prevents immediate forwarding, or because the queuing block is being used to alter the temporal properties of a traffic stream (i.e., shaping). A simple, static buffer manager serves the purpose of traffic grooming.

Scheduler. A scheduler is an element which gates the departure of each packet that arrives at one of its inputs, based on a service discipline. It has one or more inputs and exactly one output. Each input has an upstream element to which it is connected, and a set of parameters that affects the scheduling of packets received at that input. The service discipline (also known as a scheduling algorithm) is an algorithm which might take any of the parameters such as relative priority associated with each of the scheduler's inputs (or) the absolute token bucket parameters for maximum (or) the minimum rates associated with each of the scheduler's inputs (or) the packet length or 802.1p QoS bits of the packet (or) the absolute time and/or local state as its input.

A leaky bucket based scheduler is employed as it implicitly supports traffic shaping on each outgoing queue according to a leaky bucket profile associated with that queue.

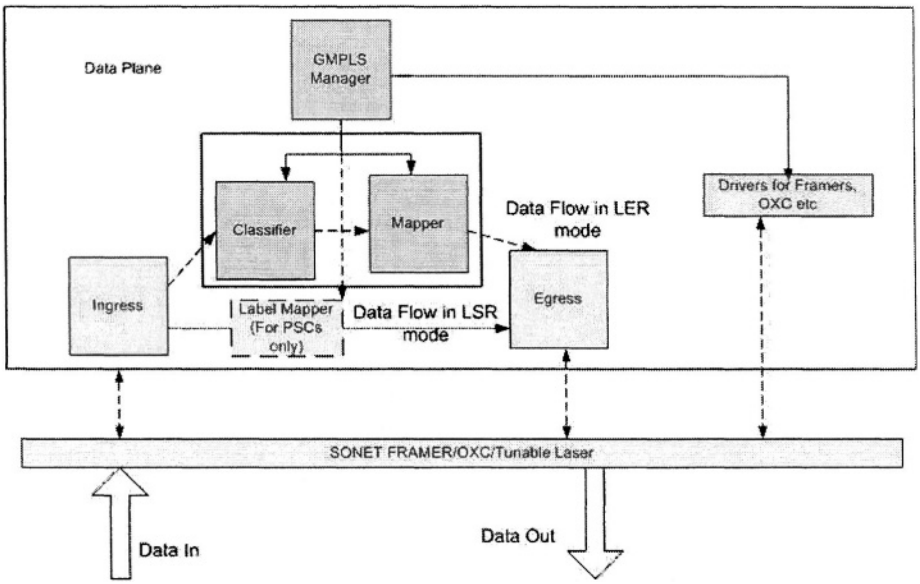

Figure 1.6. GMPLS/UNI data plane blocks.

1.4.2 UNI/GMPLS Example

An example architecture for a network element which supports UNI / GMPLS is described in this section. At present, UNI / GMPLS is a developing technology with continuously evolving specifications, and illustrates an application which would particularly benefit from a flexible software architecture which supports frequent protocol updates. GMPLS networks support three distinct functionalities, namely, ingress LER processing, LSR switching and egress LER processing. The various GMPLS data plane blocks used to manifest the GMPLS switching functionalities is shown in Figure 1.6. A brief discussion on the functionality of these data plane blocks is provided below.