

# Preface

The ever-growing society of information has its backbone on a complex network of optical links among nodes where data are stored and processed. These nodes are mostly constituted by data centers. Here a large amount of traffic is handled which is unequally shared between the inter-data center traffic and out-of-the-data center traffic. This huge amount of data is exchanged by using optical communication technologies, thanks to their unique characteristics of high bandwidth, low power consumption, transparency to signal protocol, and enormous bit rate. Advanced optical network technologies are used in data transmission and are reaching the interior of the data centers to meet the increasing demands of capacity, flexibility, low latency, connectivity, and energy efficiency.

In fact, optical point-to-point interconnects support intra-data center networking of compute nodes, electrical packet switching fabrics, and storage equipments. The evolution is toward the use of optical switching to handle optically the data flow. It is difficult to predict when optical switching will be adopted in data centers, which architectures will be used, and which device technologies will be developed to implement these architectures. It is therefore important to make the point now of the technology to help move this transition. This is the scope of the present book.

Recently, many network architectures have been proposed, and many experiments have been realized to demonstrate intra-data center networking based on optical circuit switching, since it improves the performances of a packet communication network working in the electrical domain. Other networking experiments have been based instead on optical packet and burst switching, in which short data packets or longer data bursts are switched directly in the optical domain to further improve the networking resource utilization and energy efficiency. This reduces the electro-optical conversion to a minimum and improves the flexibility by using sub-wavelength bandwidth granularity and statistical multiplexing. In this framework, software-defined networking (SDN) and network function virtualization (NFV) are widely considered key enablers for flexible, agile, and reconfigurable optical data centers since they will provide the coordinated control of network resources and the capability to allocate dynamically the network capacity.

These demonstrations have been made possible by the development of photonic integrated circuits tailored to high-speed optical interconnects and low-cost integrated switches. Integrated approaches allow lowering the cost, footprint, and power consumption with respect to traditional discrete component-based counterparts. The various optical switching architectures make use of a specific optical platform for both transmission and switching, and some of them are based on transmission and switching of gray optical signals, while others exploit the advantages represented by wavelength switching.

This book introduces the reader to the optical switching technology for its application to data centers. In addition, it takes a picture of the status of the technology evolution and of the research in the area of optical networking in a data center. It is clear to the editors that there is still work to do in both the system architecture (toward a scalable architecture) and the device technology (toward high-performance and large-scale integration optical devices) before the introduction of optical switching in commercial networking equipments for data centers. However, the recent progress in the field make us confident that this technology is going to make a big impact on how the future data centers will be run.

The book is organized in four parts: the first part is focused on the system aspects of optical switching in intra-data center networking, the second part is dedicated to describing the recently demonstrated optical switching networks, the third part deals with the latest technologies developed to enable optical switching, and, finally, the fourth part of the book outlines the future prospects and trends.

In Chap. 1, the challenges in current and future data center architectures in terms of scalability, performances, and power consumption are discussed, and the need to develop new hardware platforms based on a tight integration of photonic ICs with electronic ICs and optoelectronic printed circuit boards is underlined. In this chapter also, a hybrid switch architecture based on small electrical switches interconnected by a wavelength router is presented, and the benefit of software-defined networking (SDN) for switch re-configurability and efficient bandwidth utilization is explained.

Chapter 2 reviews the optical circuit switching networks which have been recently proposed with the following main motivations: (a) improvement of the data center networking performances in terms of latency and power consumption by off-loading long-lived bulky data flow from the electrical switching domain to the optical switching networks, (b) provision of a flexible capacity to the intra-data center networking in order to increase the resource utilization, and (c) build of a high-capacity, future-proof networking infrastructure which is transparent to bit rate and protocol.

While an optical circuit switching layer has to operate in conjunction with a more dynamic electrical packet switching layer, optical packet/burst switching systems improve the bandwidth efficiency with sub-wavelength granularity and have the right dynamicity to handle effectively bursty traffic, eventually replacing completely the electrical packet switching layer. Chapter 3 presents and discusses optical packet/burst switching architectures, defines the challenges, and briefly introduces the enabling technologies.

Chapter 4 begins the second part of the book, which is dedicated to the system demonstrations. The chapter describes the implementation and performances of the OSA system architecture. OSA is an optical circuit switched network, which provides a highly flexible optical communication infrastructure between top-of-the-rack (ToR) switches. A first aggregation layer of wavelength-selective switches and a higher level of optical space switches constitute it. This network is able to adapt both the interconnect topology and the capacity to the changing traffic demand, and it supports on-demand connectivity avoiding or greatly reducing oversubscription.

The Hi-Ring architecture is described in Chap. 5. It is based on a multidimensional all-optical switching network interconnecting top-of-the-rack switches. The multidimensional switch comprises a lower layer of space switches, a medium layer of wavelength-selective switches, and a top layer of time-slot switches. While slower space and wavelength switches handle highly aggregated data flows, fast switches are used for time-slot switching of bursty traffic with sub-wavelength granularity. The use of multiple switching allows to implement an optimized network infrastructure with fewer nodes and links among servers with benefits in terms of power consumption, cost, and latency.

In Chap. 6, the LIONS optical network switch is presented, and its experimental demonstrations are discussed. LIONS is a very low-latency, high-bandwidth, energy-efficient switch that interconnects many servers and is implemented in two versions: passive architecture and active architecture. Both types of systems are based on array waveguide router (AWGR) devices. LIONS exploits the AWGR property of de-multiplexing into different output ports a comb of wavelengths received at each input port and multiplexing in a cyclical manner, at each output port, the wavelengths coming from different input ports. There is no need to use fast optical switching fabric, and the wavelength switching is performed by fast tunable laser diodes. Active LIONS is an all-optical packet switch, while in passive LIONS, packet switching is performed in the electrical domain at the network edge with the AWGR performing wavelength routing.

The torus photonic data center is presented in Chap. 7. The top-of-the-rack switches are connected to a network of hybrid optoelectronic routers (HOPRs) interconnected with a torus topology and controlled by a centralized network controller. Such network architecture is characterized by flexible scalability since it can be expanded by simply adding nodes in a plug-and-play manner. In this way, robust redundancy of the links due to the many alternative routes can be made available. Moreover, it does not require high-radix optical switches. The torus network supports optical packet switching (OPS), optical circuit switching (OCS), and the novel virtual optical circuit switching (VOCS).

LIGHTNESS is another switching network for communication among ToRs that combines OPS and OCS in an interchangeable manner with OPS switching short-lived data flows and OCS handling long-lived data flows, and it is controlled by an SDN-enabled control plane. This network is dealt with in Chap. 8.

In Chap. 9, two network architectures are presented. The first is a hybrid optical/electrical packet switching (OPS/EPS) network in which the data packets are separated in small data packet to be handled in the electrical domain and large data

packet to be handled in the optical domain. Short packets are forwarded by using conventional protocol, while long packets are processed in an aggregation node by converting each of them into a photonic frame (adding label, guard gap and scrambling.) before sending them to the optical packet switch. The second network is the pure photonic packet switching network that is a synchronous (time-slotted) OPS, handling all types of packets, and is based on a photonic frame wrapper and on the separation of the control path and the data path.

The last recently demonstrated intra-data center network is the optical pyramid data center (OPMDC), which is discussed in Chap. 10. It is a recursive network, based on a pyramid construct, interconnecting ToR switches. OPMDC comprises three tiers of wavelength-selective optical switching nodes; the first is a reconfigurable optical add/drop multiplexer (ROADM) directly connected to the ToR switches, and the upper tiers are wavelength cross-connects (WXC). This network enables extensive wavelength reuse and efficient allocation of wavelength channels, managed by a centralized SDN controller, in order to support packet-based and circuit-based data transfer with low latency.

The third part of the book, dedicated to the enabling technologies, starts with Chap. 11 that reviews the commercially available optical switch technologies. Microelectromechanical system (MEMS), piezoelectric, liquid crystal,  $\text{LiNbO}_3$ , semiconductor optical amplifier (SOA), and photonic lightwave circuit (PLC)-based switches are presented and discussed. A table is included for comparing the key parameters.

Chapter 12 explains the physical effects and mechanisms for optical switching in silicon and presents the different types of switching cells used in large-scale integration silicon photonic switch matrices. The most used silicon photonic matrix architectures are presented and discussed, and three types of matrices are considered: those with switching speed in the range of microseconds, those with switching speed in the range of nanoseconds, and the wavelength-selective switch matrices. The recently demonstrated matrices are here reviewed and compared.

The other key enabling technology for the introduction of optical switching in data centers is the optical transceiver technology. High-speed, low-cost, short-reach optical interconnects must be deployed with efficient modulation formats and photonic integration. Two chapters are focused on this aspect. Chapter 13 presents the trend in high-speed interconnects reviewing the multidimensional modulation formats that allow increasing the transmission rate with respect to on-off key modulation (OOK) without the need of using costly coherent detection systems. The evolution of the transceiver architecture toward a high-dimensional format from 1D to 4D is discussed, and the digital signal processing functions enabling these types of modulations and their direct detection are briefly described.

Chapter 14 reviews the techniques, capabilities, and future potential of InP monolithic integrated technology for the implementation of optical transceivers and optical switches for data centers.

Finally, the fourth part of the book presents, in Chap. 15, an overview of the recent and future trends in technologies and architectures for high-performance optically switched interconnects. The different aspects are discussed: on-chip, on-board, and

rack-to-rack optical interconnects and optical switching. Recent research is addressed on the development of new technologies for increasing capacity and performance of optical networks while providing high flexibility and high energy efficiency to support future cloud applications.

We are grateful to our past and present colleagues, students, and friends at Ericsson and at the Nanoscience Laboratory of the Department of Physics of the University of Trento, for maintaining an environment of scientific excellence and friendship over the years. We owe special thanks to the authors of the various chapters for their excellent work. In addition to thanking the authors, we would like to thank Brinda Megasyamalan and Mary E. James for the help, assistance, and patience.

Pisa, Italy  
Trento, Italy  
April 2017

Francesco Testa  
Lorenzo Pavesi

Optical Switching in Next Generation Data Centers

Testa, F.; Pavesi, L. (Eds.)

2018, XIV, 336 p. 179 illus., 139 illus. in color.,

Hardcover

ISBN: 978-3-319-61051-1