Tzanakaki, A., Wosinska, L., Schares, L., Liu, H., & Simeonidou, D. (2018). Special issue on optical data center networks. *IEEE/OSA Journal of Optical Communications and Networking*, 10(7), ODC1-ODC3. <https://doi.org/10.1364/JOCN.10.00ODC1>

Peer reviewed version

Link to published version (if available):
[10.1364/JOCN.10.00ODC1](https://doi.org/10.1364/JOCN.10.00ODC1)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via OSA at <https://www.osapublishing.org/jocn/abstract.cfm?uri=jocn-10-7-ODC1>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Special Issue on Optical Data Center Networks

**Anna Tzanakaki^{1,5}, Lena Wosinska², Laurent Schares³, Hong Liu⁴,
Dimitra Simeonidou⁵**

*¹ National and Kapodistrian University of Athens, Greece, ² KTH, Sweden, ³ IBM, USA,
⁴ Google, USA, ⁵ University of Bristol, UK*

Abstract — To support the current trends in Cloud, Internet of Things, 5G etc, massive amounts of data need to be processed and stored with very different performance requirements in terms of capacity, granularity, delay etc. These requirements drive the need for both mega-size Data Centers (DCs) and distributed micro DCs interconnected with high speed communications links. This Special Issue focuses on the key role of optical networking in meeting these requirements through high-capacity, energy-efficient, and flexible connectivity and presents recent advancements in DC architectures, network solutions, and optical technologies needed to interconnect computing and storage resources within and between DCs.

Big data, cloud services, artificial intelligence, Internet of Things and 5G services must store and process massive amounts of data through advanced, interconnected computing resources and storage devices. The very diverse requirements of these services in terms of processing power, storage capacity, communication bandwidth, delay, mobility, etc., drive the need for a combination of a large number of distributed micro Data Centers (DCs) as well as both centralized and remote mega-size DCs scaling up to hundreds of thousands of server and storage modules. In view of this, high-speed communication across DCs is key. Additionally, for mega-DCs, scaling internal computing and storage resources and interconnection capacity is a major challenge. Traditionally, computing, storage and communication resources have been considered separately; however, there are significant measurable benefits when a unified allocation approach is considered and these heterogeneous resources are shared.

To address these challenges and take advantage of potential benefits, there is a clear need for novel architectural approaches, network solutions, and technologies to interconnect the distributed computing and storage resources, both within a DC and between DCs, referred to as intra- and inter-DC networking, respectively. In this context, it is important to ensure that sustainability and efficiency in the utilization of resources are maximized, thus reducing the associated capital and operational expenditures, as well as the overall energy consumption and CO₂ footprint.

Optical networking plays a key role in meeting these requirements by offering high-capacity, energy-efficient, and flexible connectivity. Moreover, to address the challenge of managing and operating these infrastructures efficiently, a high degree of agility and adaptability in the functions that the network can perform is required. These goals can be achieved by adopting open software and hardware platforms such as software-defined networking (SDN) and network function virtualization (NFV), as well as OpenCompute approaches. Exploiting open platforms could lead to new business and operational models such as multi-tenancy through slicing and virtualization.

This special issue aims at providing an overview of the current technical challenges in the area of DC networking and capturing the latest technology advancements that make optical networking a key technology enabler. It comprises 2 invited and 13 contributed papers addressing some of the key challenges in the context of inter- and intra- DC networking.

The first paper of the special issue is an invited paper entitled “Application-Triggered Automatic Distributed Cloud/Network Resource Coordination by Optically Networked Inter/Intra Data-Center” authored by Naoaki Yamanaka et al. This paper focuses on a distributed cloud architecture, where resources are dynamically combined through optical interconnections and focuses on edge/center cloud and edge/edge integration using virtual machine migration. The proposed architecture supports horizontal and vertical integration through inter-DC coordination, while intra-DC coordination is enabled by application triggering. This approach supports a variety of applications requiring very low latency and high availability with increased energy efficiency.

The second invited paper is entitled “High-Port and Low-Latency Optical Switches for Disaggregated Data Centers: the Hipolao switch architecture” authored by N. Terzenidis et al. The paper focuses on the network requirements in disaggregated systems and reviews state-of-the-art high-radix optical switch architectures and currently available disaggregated DC systems. It also presents a novel optical packet switch design combining an N-port Broadcast-and-Select (BS) approach and N×N Arrayed Waveguide Grating Router (AWGR)-based forwarding schemes. The proposed switch adopts a modified Spanke architecture and provides sub-microsecond latency and high throughput utilizing distributed control and optical feed-forward buffering. Feasibility of a 256-port layout with error-free operation was experimentally demonstrated at 10 Gb/s, while simulations were carried to further evaluate the performance of the proposed solution.

The paper “Experimental Demonstration of a 64-port Wavelength Routing Thin-Clos System for Datacenter Switching Architectures” authored by R. Proietti et al. proposes a Thin-Clos architecture for intra-datacenter connectivity, based on arrayed waveguide grating routers (AWGRs). More specifically, the paper presents results from the design, the experimental demonstration and evaluation of the proposed architecture that scales to 64 ports. The system that was demonstrated consumes 10 W and supports error-free performance at 10 Gbit/s per wavelength, utilizing 32 wavelengths in total, with a power penalty less than 3 dB in the worse-case crosstalk scenario.

The paper entitled “CBOSS: Bringing Traffic Engineering Inside Data Center Networks”, authored by N. Benzaoui et al. proposes an intra-data center network architecture adopting cloud burst optical-slot switching (CBOSS). This architecture leverages optical technologies and software defined network control to support traffic engineering. Sub-wavelength (time slot) switching granularity is demonstrated, supporting on-demand network slicing with specific QoS guarantees. A new integrated silicon photonics device that supports 1x2 Wavelength Selective Switch functionality and fast tunable laser technology is used.

The paper entitled “Inter-Domain Optimization and Orchestration for Optical Datacenter Networks” authored by G. Landi et al. present a hierarchical Software Defined Networking orchestration platform in order to support end-to-end inter-domain orchestration. The proposed platform treats the intra- and inter DC networks as different domains with domain specific orchestrators/controllers that can be used to achieve inter-domain orchestration. The paper showcases end-to end path establishment with dedicated capacity both in an emulated and a realistic test-bed.

The paper entitled “Reimagining Datacenter Topologies with Integrated Silicon Photonics” authored by Cyriel Minkenberg et al. presents the design and implementation of a switch ASIC architecture with in-package optical transceivers. A highly-integrated Packet Switching Optical Module prototype with twelve 100G Ethernet ports, integrated optics and L3 routing capability was demonstrated. The

proposed solution promises improved energy efficiency and reduced cost per bandwidth compared to conventional switch ASICs with pluggable optical transceivers.

The paper entitled “A Fully SDN Enabled All-Optical Architecture for Data Centre Virtualisation with Time and Space Multiplexing” is authored by K. Kondepu et al. This paper focuses on developing a Virtual Data Center architecture that allows dynamic slicing of network and compute resources in order to support multi-tenant operation over the physical network. The Data Centre Virtualization architecture adopts an SDN-controlled all-optical data plane comprising both Optical Circuit Switching and sub-wavelength optical switching through the Time Shared Optical Network solution. Experimental results show dynamic bandwidth provisioning suitable for TDM connections with different performance characteristics as well as low network delay and complexity.

The paper entitled “Dynamic Packet/Circuit Switch Placement for Optically Disaggregated Data Centres”, authored by Adaranijo Peters et al. This paper presents a reconfigurable hybrid disaggregated Data Center (dRedBox) architecture exploiting the concept of server resource disaggregation with software, electronic and optical technologies. The proposed approach aims at providing flexibility and connectivity through network function service programmability. In addition, algorithms and strategies for selecting and deploying electronic packet switching and optical circuit switching function services to implement Virtual Machine network requests were studied and the benefits of the proposed approach were quantified through simulations.

The paper entitled “TCP Congestion Control in Datacenter Optical Packet Network on Hybrid Switches” is authored by Artur Minakhmetov et al. and focuses on the potential of Optical Packet Switching technology in data center networks when shared electronic buffers are used and TCP congestion control algorithms are introduced to control the transport of optical packets. The Stop-And-Wait and the modified Additive Increase Multiple Decrease families of congestion control algorithms are studied. The presented analysis shows that these algorithms employed in optical packet switches with shared electronic buffers can provide improved network throughput compared to bufferless all-optical switches, matching the performance of all-electronic switches in Data Centers and Local Area Networks.

The paper entitled “Location-Aware Energy Efficient Virtual Network Embedding in Software Defined Optical Data Center Networks” authored by Y. Zong et al. This paper focuses on location-aware virtual network embedding in the context of software-defined optical data center networks and proposes a mixed integer linear programming model with the objective to minimize the total power consumption by reducing the active data centers and power consuming network components. In addition, two algorithms are proposed based on the global topology node ranking method to improve performance in terms of parameters such as power consumption and acceptance ratio.

The paper entitled “Joint Jobs Scheduling and Lightpath Provisioning in Fog Computing Micro Datacenter Networks” authored by Zhen Liu et al. investigates a joint optimization algorithm of multiple jobs scheduling and lightpath provisioning with the aim to minimize the average completion time in elastic optical networks. Simulation results show that the proposed algorithm outperforms the scheduling-only and routing-only algorithms in terms of average completion time and frequency slots.

The paper entitled “Design and Verification of Large-Scale Optical Circuit Switch Using ULCF AWGs for Datacenter Application” authored by H. Nagai et al. proposes an optical-circuit-switch architecture that uses uniform-loss and cyclic-frequency (ULCF) arrayed-waveguide gratings (AWGs) as wavelength routers. The performance

of the proposed approach is experimentally evaluated through the development of part of a $1,536 \times 1,536$ optical circuit switch and a 96-wavelength 100-Gbps signal transmission in the full C-band. This experimental set-up employed a pair of 12×48 ULCF AWGs fabricated using PLC technology. The overall switch throughput reaches 153.6 Tbps.

The paper entitled “Silicon Photonics based 100Gbit/s, PAM4, DWDM Datacenter Interconnects” authored by Radhakrishnan Nagarajan et al. demonstrates a switch-pluggable, 4.5-W, 100-Gbit/s, silicon photonics based, PAM4, QSFP-28 module. This module is able to transport Ethernet data directly over dense WDM for connections between switches at datacenters located up to 120 km apart. This approach eliminates the need for a separate optical transport layer. The module used is based on direct detect modulation relying on combination of Si photonics for the highly integrated optical components, and highspeed Si CMOS for signal processing. These are important for the implementation of low-cost, low-power, switch pluggable optical modules suitable for massive interconnections between regional datacenters.

The paper entitled “HiFOS: A Novel Scalable and Low Latency Hybrid Data Center Network Architecture Based on Flow Controlled Fast Optical Switches” authored by Fulong Yan et al. investigates the performance of a hybrid Fast Optical Switch and Top-of-Rack Switch Data Center Network architecture. This eliminates the need of optical buffers for fast intra-cluster interconnection, while inter-cluster connectivity is obtained exploiting direct TORs interconnections. The performance of the architecture is investigated in terms of latency, packet loss, and throughput. Numerical results show that the proposed architecture interconnects up to tens of thousands of servers with low end-to-end latency, a low packet loss and improved power consumption and cost.

The paper entitled “Applying Multi-Controller Collaboration in Fine-Grained All-Optical Intra-Data Center Networks” authored by Yinqiu Jia et al. proposes a multi-controller control scheme for fine-grained all-optical datacenter networks. The concept of optical time slice switching is adopted and partitioning rules for network and controllers’ responsibility, routing algorithm and singling process are discussed. The proposed control scheme can reduce the computation and processing requirements of each controller and improve network performance. Emulation results show that the multi-controller scheme can reduce the connection setup delays, the request and the data failed rate.

We hope that the articles of this special issue will serve as a useful resource for researchers who would like to get up-to-date views on the latest research efforts. We would like to thank the authors and reviewers, whose dedicated efforts maintain the high technical standard of this journal. We also thank the JOCN Editor-in-Chief, Jane Simmons, for the continuous and valuable support she has provided throughout the preparation of this special issue, and the JOCN editorial staff, especially Keith Jackson, who have produced a high-quality print volume under the tight schedule required for a special issue of this kind.