

A Survey of Reconfigurable Optical Networks

Matthew Nance Hall^{a,*}, Klaus-Tycho Foerster^b, Stefan Schmid^b,
Ramakrishnan Durairajan^a

^a*Computer and Information Science,
University of Oregon, 1585 E 13th Ave, Eugene, OR 97403, USA*

^b*Faculty of Computer Science,
University of Vienna, Währinger Str. 29, 1090 Vienna, Austria*

Abstract

Reconfigurable optical networks have emerged as a promising technology to efficiently serve the fast-growing traffic produced by the digital society. This paper provides a survey of the field. We first review enabling optical hardware technologies in general and then consider technologies that are specific to data center networks and wide-area networks in more detail. We further provide an overview of the cost models used in the literature as well as the algorithmic problems introduced by these technologies, their first solutions, and discuss systems and implementation aspects. We conclude with a discussion of open challenges.

*Corresponding author

Contents

1	Introduction	2
1.1	Novelty and Contribution	3
1.2	Scope	4
1.3	Organization	5
2	Network Architectures	6
2.1	IP-over-Optical Transport Network	6
2.2	Data Center Architecture	8
2.3	Software Defined Networking	8
2.4	Elastic Optical Networks	9
2.5	Summary	9
3	Enabling Hardware Technologies	10
3.1	Wavelength Selective Switching	10
3.2	ROADMs	12
3.3	Bandwidth-variable Transponders	13
3.4	Silicon Photonics	15
3.5	Summary	16
4	Optically Reconfigurable Data Centers	16
4.1	DCN-specific Technologies	17
4.2	Cost Modeling	19
4.3	Algorithms	20
4.4	Systems Implementations	23
4.5	Summary	24
5	Reconfigurable Optical Metro and Wide-area Networks	24
5.1	Metro/WAN-Specific Challenges and Solutions	26
5.2	Cost Modeling	28
5.3	Algorithms	30
5.4	Systems Implementations	33
5.5	Summary	35
6	Open Challenges in Reconfigurable Optical Networks	36
7	Conclusion and Future Work	37

1. Introduction

The popularity of data-centric applications related to business, science, social networking or entertainment, as well as the rise of machine learning and artificial intelligence, leads to an explosive growth of communication traffic, especially inside data centers but also in the wide-area networks (WANs) connecting the data centers to users and each other. Optical communication is one

of the most promising technologies to cope with the resulting increase in communication requirements, due to the high bandwidth and energy efficiency it provides. Our society’s communication infrastructure is hence likely to become more and more optical soon. Accordingly, we currently witness many efforts to further advance optical interconnects and lightwave telecommunication in data centers and WANs.

Reconfigurable optical technologies are a particularly interesting innovation in this context. Reconfigurable optical networks enable adaptation: either of the topology itself (in case of data centers) or of the network capacity (in case of WANs). Such adaptations may be exploited by next-generation systems to improve performance and efficiency, *e.g.*, by making the network demand aware. For example, recent technologies based on free-space optics or optical circuit switches (OCSs) support very fast topology adaptations in data centers. Meanwhile, technologies based on reconfigurable optical add-drop multiplexers (ROADMs) can add or drop wavelengths carrying data channels from a transport fiber without the need to convert the signals to electronic signals and back. In both cases, the entire bandwidth assignment’s planning need not be carried out during the initial deployment of a network.

However, with reconfigurable optical networks being a relatively a new technology, the community is still discussing their conceptual fundamentals, benefits, and limitations.

1.1. Novelty and Contribution

To the best of our knowledge, this is the first and up-to-date survey on emerging reconfigurable optical networks, considering both data center and WAN settings. While there are surveys with broad overviews of reconfigurable optical technologies in these settings (*e.g.*, software-defined optical networks [1], routing and spectrum allocation [2], wavelength switching hardware architecture [3]), the functioning of such systems (*e.g.*, reconfigurable metropolitan networks [4] and data center networks [5]) requires a full-stack perspective on optical networks. We address this requirement, in this paper, by presenting an *end-to-end perspective* on reconfigurable optical networks by (a) emphasizing the interdependence of optical technologies with algorithms and systems and (b) identifying the open challenges and future work at the intersection of optics, theory, algorithms, and systems communities.

This survey is timely, as interest in dynamic optical layer networking technologies is gaining attention from the networked systems community. Upon reviewing the last five years of publications from five optical and systems network journals, we ran a clustering analysis to see how much overlap there has been between the two fields. Table 1 shows the journals, and their raw publication counts since 2015. The clustering analysis was conducted using CitNetExplorer software [6], and used a clustering resolution of 0.75. Connections between papers are first-order citations, and the publication records are from Web of

Science¹.

Figure 1 shows the results of the clustering analysis in the five largest clusters. Based on the largest cluster, 1, it appears there is a strong relationship between the Journal of Lightwave Technology (JLT) and the Journal of Optical Communications and Networking (JOCN) publications. This is no surprise, as the two journals have a strong emphasis on optical networking. Cluster 2 is mainly composed of IEEE Transactions on Communications (TCOM) papers, with a few IEEE/ACM Transactions on Networking (TON) papers. TCOM papers are generally concerned with physical layer networking, including radio signals and copper mediums in addition to optical transmission. Cluster 3 shows a strong relationship between TON and IEEE Transactions on Network and Service Management (TNSM) papers. These journals have a related interest in networked systems, applications, and management. The three predominantly physical-layer journals (JLT, TCOM, and JOCN) appear together in cluster 4. Cluster 5 shows a relation between physical-layer and systems journals, TCOM, and TON. Clusters 6 and beyond are mostly singleton clusters, comprising predominantly one journal and are therefore omitted from the figure. Our analysis underscores a division between optical and higher layer publishing venues.

The missing piece in this picture appears to be any significant overlap between the networked systems journals and optical journals. We hope this survey will illuminate common areas of interest between these two communities in order to bring more attention to cross-layer networking research, as cross-domain knowledge and expertise will be increasingly important for bringing greater flexibility and control to the management of optical networks.

Our survey is tutorial in nature and focuses on concepts rather than exhaustive related work, concentrating on selected articles. Hence, our paper targets students, researchers, experts, and decision-makers in the networking industry who would like to obtain an overview of the critical concepts and state-of-the-art results in reconfigurable optical networks. We start with an overview of the enabling optical hardware technologies. We explore where data center and WAN systems have integrated this hardware. We review cost models, discuss the novel algorithmic challenges and solutions in the literature, and elaborate on systems and implementation aspects. We also identify the major open issues which require further exploration and research to design the next generation reconfigurable optical networks.

1.2. Scope

There already exist some excellent surveys on optical networks which at least partially cover reconfigurable aspects, both in the context of data centers [7, 8, 9] and (to a lesser extent) in the context of WANs [10]. We extend these surveys while providing an up-to-date overview of the literature. Our paper aims to provide an understanding of the underlying fundamental issues and concepts and identify commonalities and differences in different reconfigurable optical

¹www.webofknowledge.com

Journal	Papers
Journal of Lightwave Technology (JLT)	3,988
IEEE Transactions on Communications (TCOM)	2,703
IEEE-ACM Transactions on Networking (TON)	1,226
Journal of Optical Communications (JOCN)	817
IEEE Transactions on Network and Service Management (TNSM)	533

Table 1: Papers published since 2015 in various networking journals.




Figure 1: Paper clusters among five networking and optical systems journals.

networks (spanning both data center and WANs). To this end, we proceed from practical technological constraints to theoretical models and solutions back to implementations. Within this domain, we restrict our coverage to enterprise and core networks. For those interested in last-mile passive optical networks, mobile front-haul, and multiple-access networks we defer to the following related surveys [11, 12, 13].

1.3. Organization

The remainder of this paper is organized as follows. Section 2 defines the network architecture model for optical networks and its connection to the packet-switched network model. Section 3 introduces the key hardware technologies that underpin reconfigurable optical networks. Section 4 showcases research on reconfigurable optical data center networks (DCNs) by highlighting DCN-specific hardware capabilities, cost modeling, algorithms, and systems implementations. Section 5 looks directly at reconfigurable WANs with an emphasis on WAN-specific challenges in addition to cost modeling, algorithms, and systems implementations. Section 7 concludes with the overarching open challenges for the field of reconfigurable optical networks spanning hardware, data centers, and WANs.




Figure 2: Metro, regional, and long-haul networks are connected by the IP-over-OTN standard.

2. Network Architectures

In this section, we briefly discuss two network architecture models that can leverage reconfigurable optics, IP-over-OTN networks and hybrid electric-optical data center networks. Our focus in this survey is to highlight and categorize reconfigurable optical networks in enterprise networks, and therefore leave last-mile optical networks, such as passive-optical networks and fiber-to-the-home networks beyond the scope of our discussion.

We also briefly outline principles leveraged in different contexts by reconfigurable optical networks, software defined networking and elastic optical networking. This discussion introduces key aspects for network designers to consider when building a reconfigurable optical network. This discussion reinforces our illustration of how full-stack perspective aids in the network design process. In sections 4 and 5, we look at specific implementations of reconfigurable optical networks in more detail.

2.1. IP-over-Optical Transport Network

IP-over-Optical Transport Networks (IP-over-OTN), defined in ITU-T G.709, is the standardized protocol that links metro, regional, and long-haul networks, as illustrated in Figure 2. Thus, we discuss IP-over-OTN when referring to the network’s IP and the optical layers. In IP-over-OTN, hosts (*e.g.*, data centers, points-of-presence or PoPs, servers, *etc.*) connect to routers, and these routers are connected through the optical transport network (OTN). A node in the optical layer is an Optical Cross-Connect (OXC). An OXC transmits data on modulated light through the optical fiber. The modulated light is called a lambda, wavelength, or circuit. The OXC can also act as a *relay* for other OXC nodes to transparently route wavelengths. When acting as a relay for remote hosts, an OXC provides optical switching capabilities, thus giving the network flexibility in choosing where to send transmitting lambdas over the OXC node.




Figure 3: IP-over-OTN network architecture model, showing the connection between IP and optical layers.

Figure 3 illustrates the connectivity at different layers of the IP-over-OTN model. The physical network connects points-of-presence (PoPs) with optical fiber spans. OXC nodes connects these PoPs with optical paths or circuits. The physical routes of the paths are abstracted away, and shown in color for reference. In the IP topology, the colors of light are also abstracted away, and we see a mesh IP network connected by routers and switches. Hosts connect to nodes at this layer, and their traffic travels down the optical paths in the physical network to reach its destination.

IP-over-OTN networks are not new. However, they are built at a great cost. Historically network planners have engineered them to accommodate the worst-case expected demand by (1) over-provisioning of dense wavelength division multiplexing (DWDM) optical channels and (2) laying redundant fiber spans as a fail-safe for unexpected traffic surges. These surges could come from user behavior changes or failures elsewhere in the network that forces traffic onto a given path. Only recently have reconfigurable optical systems begun to gain attention in the data center and wide-area network settings. For more information about early IP-over-OTN, we defer to Bannister *et al.* [14] and references therein, where the authors present work on optimizing WDM networks for node placement, fiber placement, and wavelength allocation.




Figure 4: Data center architecture proposed in c-Through [21]

2.2. Data Center Architecture

Historically, data centers relied on packet-switched networks to connect their servers; however, as scale and demand increased, the cost to build and manage these packet-switched networks became too large. As a result of this change, new reconfigurable network topologies gained more attention from researchers and large cloud providers. Many novel data center architectures with reconfigurable optical topologies have been proposed over the last decade. These architectures have in common that they reduce the static network provisioning requirements, thereby reducing the network’s cost by presenting a means for bandwidth between hosts to change periodically. Figure 4 shows one such example of a hybrid electrical-optical data center architecture. These architectures reduce cost and complexity via scheduling methods, which change bandwidth on optical paths in the data center. Various approaches have been demonstrated. Notable architectures employ fixed, and deterministic scheduling approaches [5, 15] or demand-aware changes that prioritize establishing optical paths between servers with mutual connectivity requests [16, 17]. Switching fabrics are also diverse for data center optical systems. These include fabrics based on nanosecond tunable lasers [18], digital micromirror devices (DMD) [19], and liquid crystal on silicon (LCOS) wavelength selective switches (WSS) [20].

2.3. Software Defined Networking

Modern data center, metro, and wide-area networks have been substantially influenced by developments in Software Defined Networking (SDN) [22], and this trend has also been making its way to optical networks [1]. The SDN paradigm decouples the control and data plane in network hardware, giving operators greater control and flexibility for controlling traffic within their network. Without this decoupling, it is more difficult to make lock-step changes to network functions, such as routing. SDN offers a logically centralized point of control for implementing policies across the network, thus enabling better network utilization for bandwidth, latency, security policies, *etc.* These concepts can also map further down the network stack to manage optical infrastructure, thereby




Figure 5: Example of fixed grid and flex grid spectrum allocation.

1) improving optical layer performance with technology, which we describe in Section 3, and 2) allowing management algorithms to adapt the optical paths in a demand-aware fashion, which we describe in Section 4 for data center networks and in Section 5 for metro and wide-area networks.

Notwithstanding, providing a standardized stable and reliable programmable optical physical layer control plane for SDNs is still an ongoing effort, as recently outlined by the *TURBO* project [23]. One important step in this direction is the development of virtual testbeds to evaluate the cross-layer operation of SDN control planes [24].

2.4. Elastic Optical Networks

A span of optical fiber enables transmission of data over a *spectrum* or set of wavelengths. These wavelengths can be allocated in a fixed or flexible (flex) grid. Networks that allow flex grid allocations are also called Elastic Optical Networks (EONs). For example, according to the ITU-T G.694.1 fixed grid standard, frequencies must be 12.5, 25, 50, or 100 GHz apart [1]. However, in elastic optical networks (EONs), also known as flex-grid networks, the frequency of a channel can be any multiple of 6.25 GHz away from the central frequency (193.1 THz) and have a width that is a multiple of 12.5 GHz. Figure 5 illustrates the difference between a flex-grid and fixed-grid allocation.

Flex grid networks can greatly improve the spectral efficiency of IP-over-OTN, allowing the network to pack data channels more densely within a span of optical fiber. However, they can also lead to unique challenges, particularly fragmentation. Fragmentation occurs when spectrum allocated on a fiber has gaps in it that are too narrow to be filled. Novel approaches to managing EONs with fragmentation-aware algorithms are covered in depth by Chatterjee *et al.* [25].

2.5. Summary

Our survey relates the latest developments in reconfigurable networks for data centers and WANs. The IP-over-OTN model is a useful framework for

reasoning about and managing optical metropolitan, regional, and wide-area networks. Similarly, we are seeing data center architectures become more reconfigurable and demand-aware with optical circuit switching. SDN is poised to bring substantial changes to the operation of optical networks in both domains by offering a centralized point for management and control for more network infrastructure, from routing of packets to routing of optical paths. Moreover, EONs are also enabling better spectral efficiency.

3. Enabling Hardware Technologies

In this section, we discuss hardware technologies that enable reconfigurable optical networks. In our end-to-end discussion on reconfigurable optical networks, the hardware is the foundational layer from which systems are built. Understanding these devices and their capabilities is crucial for designing and building real-world reconfigurable optical networking systems. We show examples of different optical technologies, including optical switches and transponders, and examples of systems that use them. We also highlight recent advances in silicon photonics, and the implications this may have for reconfigurable optical networks in the near future. Finally, we discuss open challenges in reconfigurable optical networks that might be solved with next-generation hardware.

3.1. Wavelength Selective Switching

In contrast to packet-switched networks, optically circuit-switched systems operate at a more coarse granularity. The transmission of information over a circuit requires an end-to-end path for the communicating parties. Although packet switching has generally prevailed in today’s Internet, recent research has revitalized the prospect of circuit switching for data centers and wide-area networks by illuminating areas in which flexible bandwidth benefits outweigh the start-up cost of circuit building.

Technological advancements for optical hardware, primarily driven by physics and electrical engineering research, have been instrumental in making circuit-switched networks a viable model for data center networks. Among these technologies are low-cost/low-loss hardware architectures. Here we give a brief overview of technological advancements in this domain that have had the most significant impact on networked systems.

Kachris *et al.* [26] have an in-depth look at optical switching architectures in data centers from 2012. In their survey, they primarily look at competing *data center* architectures and switch models. In this section, we choose to focus instead on those architectures’ physical manifestations (*i.e.*, the base components that make them up). Furthermore, exciting new developments have occurred since then, which we highlight in this section.

Polymer waveguides are a low-cost architecture for optical circuit switches. These have been fabricated and studied in depth over the last 20 years, including work by Taboada *et al.* [27] in 1999, Yeniay *et al.* [28] in 2004, and Felipe *et al.* [29] in 2018. Early implementations such as Taboada *et al.* [27] showed fabrication techniques for simple polymer waveguide taps. Multiple waveguide taps

can be combined to form an Array Waveguide Grating (AWG), and the signals traversing the AWGs can then be blocked or unblocked to create an optical circuit switch. A major inhibitor of the polymer waveguide architecture was signal-loss, which was as high as 0.2 dB/cm until Yeniay *et al.* [28] discovered an improvement on the state-of-the-art with ultra low-loss waveguides in 2004. Their waveguides, made with fluorocarbons, have $4\times$ less loss (0.05 dB/cm) than the next best waveguides at the time, made from hydrocarbons. Felipe *et al.* [29] demonstrate the effectiveness of a polymer waveguide-based switching architecture for reconfiguring groups of optical flows of up to 1 Tbps, proving that that AWG is a viable and competitive switching architecture for data centers. More recently, in 2020, AWGs were demonstrated to work in conjunction with sub-nanosecond tunable transmitters to create flat topologies, significantly reducing power consumption for data center networks due to the passive—no power required—nature AWGs [30]. Switching speeds below 820 ps have been demonstrated using a 1×60 AWG and tunable laser [18]. AWGs with as many as 512 ports have been demonstrated [31].

Microelectromechanical Systems (MEMS), introduced by Toshiyoshi *et al.* [32] in 1996, offered a lower-loss and more flexible alternative to polymer waveguide systems of the day. MEMS devices are made up of small mirrors, which can be triggered between states (*i.e.*, *on* and *off*). Therefore, in a MEMS system light is *reflected* rather than *guided* (as in the polymer waveguide systems). This distinction between reflection and guiding implies generally slower switching speeds for MEMS based systems, as the mirror must be physically turned to steer light out of the desired switch-port. Despite this limitation, MEMS systems evolved to be competitive with polymer waveguides in modern systems. Advances in MEMS technology have yielded wavelength selective switches (WSS) scalable to 32 ports with switching speeds under 0.5 ms [33]. Data center solutions leveraging MEMS based switches include Helios [34].

Liquid Crystal on Silicon (LCOS) was demonstrated as another viable optical switching architecture by Baxter *et al.* [35] in 2006. An LCOS switch is depicted in figure 6. Multiplexed optical signals enter the system from a fiber array. These signals are directed to a conventional diffraction grating where the different colors of light are spatially separated from each signal. These colors are then projected onto a unique position in the LCOS switching element. This element is divided into pixels or cells, and charged with an electrical current. The voltage applied to any cell in the switching element determines which output fiber a given channel will leave through. From there, the signal travels back through the system and into a different fiber in the array.

Switches based on this technology have a response time of 10 – 100 μ s [36]. Recent work by Yang *et al.* [37] demonstrates the construction of a 12×12 and 1×144 port WSS based on a 1×12 LCOS architecture. Chen *et al.* [38] developed an improved LCOS architecture with which they demonstrated a 16×16 optical switch. LCOS switches are commercially available and are recognized as a key enabler for reconfigurable optical networks [20].

Summary. Table 2 summarizes optical switch performance metrics. Each architecture comes with advantages under distinct circumstances. Highly scal-




Figure 6: Liquid crystal on silicon wavelength selective switch.

	Port Scalability	Switching Speed	
AWG	512×512	< 820 ps	Highly scalable with unsurpassed demonstrations for short-reach applications with tunable lasers.
MEMS	32×32	< 0.5 ms	Higher scalability and lower insertion loss, less crosstalk.
LCOS	16×16	$10 - 100$ μ s	Lower scalability and optical performance, but more modular design than MEMS.

Table 2: Summary of systems implementations of reconfigurable wide area networks

able data center architectures have been developed with sub-nanosecond tunable lasers and AWGs [5, 18, 30]. MEMS have generally better scalability, lower insertion loss, and less crosstalk over LCOS systems [39] but also demand higher precision manufacturing to ensure that all $N \times M$ mirrors configurations are accurately aligned. LCOS elements can also be packed more compactly into a modular unit due to the absence of moving parts that are present in MEMS.

3.2. ROADMs

Reconfigurable add-drop multiplexers, or ROADMs, are an integral component of IP-over-OTN networks. These devices have evolved over the years to provide greater functionality and flexibility to optical transport network operators. We briefly describe the evolution of ROADM architectures. Figure 7 shows a broadcast and select ROADM architecture. Please refer to [3] for more information about ROADM architectures.

Colorless (C). Early ROADMs were effectively programmable wavelength *splitter-and-blockers*, or *broadcast-and-select* devices. A wavelength splitter-and-blocker can be placed before an IP-layer switch. If the switch is intended to *add/drop* a wavelength (*i.e.*, transceive data on it), then the blocker prohibits light on the upstream path and enables light on the path to the switch. These splitter-and-blocker systems are better known as Colorless, or C-ROADMs, as




Figure 7: Broadcast and Select colorful ROADM. The add/drop node, R1, has ports for two optical channels. These channels are directed at the ROADM. The ROADM uses a splitter to *broadcast* the channels onto two outbound ports, where a wavelength blocker *selects* the appropriate channel for the next router.

the *splitter-and-blocker* architecture is independent of any specific frequency of light. To receive the maximum benefit from C-ROADMs, operators should deploy their networks with tunable transceivers as they allow more flexibility for the end hosts when connecting to remote hosts.

Colorless, Directionless (CD). The CD-ROADMs extend the architecture of C-ROADMs by pairing multiple C-ROADMs together in the same unit to allow for a wave to travel in one of many directions. One shortfall of this architecture is that the drop ports from each direction are fixed, and therefore if all of the drop ports are used from one direction, the remaining points from other directions cannot be used. Due to the limitation of drop ports in different directions, the CD architecture is not *contentionless*.

Colorless, Directionless, Contentionless (CDC). The CDC-ROADM solves the contention problem by providing a shared add/drop port for each direction of the ROADM. This allows contentionless reconfiguration of the ROADM as any drop-signal is routed to a common port regardless of the direction from which the wave begins/terminates.

Colorless, Directionless, Contentionless w. Flexible Grid (CDC-F). Flexi-grid, or elastic optical networks, are networks carrying optical channels with non-uniform grid alignment. This contrasts with a fixed-grid network, where different wavelengths are spaced with a fixed distance (*e.g.*, 50 GHz spacing). Wideband spacing allows signals to travel farther before becoming incoherent due to chromatic dispersion. Thus, CDC-Flex or CDC-F ROADMs enable the reconfiguration of wavelengths with heterogeneous grid alignments. These are most useful for wide area networks, with combinations of sub-sea and terrestrial circuits.

3.3. Bandwidth-variable Transponders

Before we discuss bandwidth-variable transponders, we must first take a moment to illuminate a common concept to all physical communications systems, not only optical fiber. This concept is modulation formats. Modulation formats determine the number of binary bits that a signal carries in one *symbol*. Two parties, a sender and receiver, agree on a symbol rate (baud), which determines a clock-speed to which the receiver is tuned when it interprets a symbol from the

sender. The simplest modulation format is on-off keying (OOK), which transmits one bit per symbol. In OOK, the symbol is sent via a high or low power level, as shown in figure 8A. A higher-order modulation technique is Quadrature Phase Shift Keying (QPSK), in which the symbol is a sinusoidal wave whose phase-offset relates the symbol. In this QPSK, there are four phase shifts agreed upon by the communicating parties, and therefore the system achieves two bits per symbol, or two baud, seen in figure 8B. A constellation diagram for QPSK is shown in figure 8B. As modulations become more complex, it is more useful to visualize them in the phase plane shown by their constellation diagram. Higher-order modulation formats are of the type, N -Quadrature Amplitude Modulation (QAM) techniques (Figure 8D), and these permit $\log_2(N)$ bits per symbol². In QAM, the symbol is denoted by phase and amplitude changes. Figure 8D shows an example of a constellation diagram for 16-QAM modulation, which offers 4 bits per symbol, or twice the baud of QPSK.

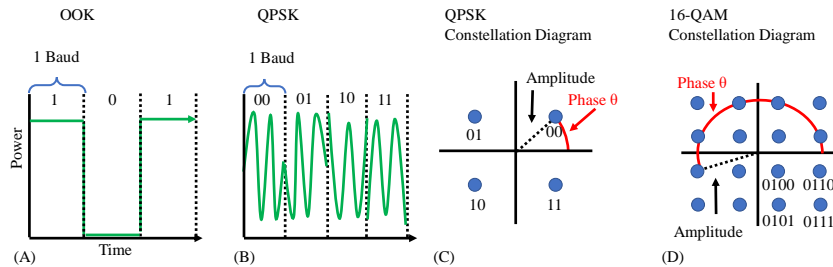


Figure 8: Modulation examples of on-off keying, quadrature phase shift keying (QPSK), quadrature amplitude modulation (QAM), and constellation diagrams for QPSK and 16-QAM.

Fiber optic communications are subject to noise. The noise level is Signal to Noise Ratio (SNR), and this metric determines the highest possible modulation format. In turn, the modulation format yields a potential capacity (Gbps) for an optical channel. For example, in [40], the authors claim that SNR of just 6 dB is sufficient to carry a 100 Gbps signal, while a circuit with an SNR of 13 dB can transmit 200 Gbps.

Bandwidth Variable Transponders (BVTs) [41] have recently proven to have significant applications for wide-area networks. These devices are programmable, allowing for the operator to choose from two or more different modulation formats, baud rates, and the number of subcarriers when operating an optical circuit. For example, the same transponder may be used for high-capacity/short-reach transmission (16-QAM or greater) or lower-capacity/longer-reach transmission (*e.g.*, QPSK). Higher modulation formats offer higher data rates. They are also more sensitive to the optical SNR, which decreases in a step-wise manner with distance, as illustrated in Figure 9. We note that BVTs enable network operators to meet the ever-growing demand in backbone traffic by increasing

²where N is generally a power of 2

optical circuits' spectral efficiency.

Low spectrum utilization, or waste, can be an issue for BVT circuits. For example, a BVT configured for a low-modulation circuit such as QPSK instead of 16-QAM has a potential for untapped bandwidth. Sambo *et al.* [42] introduced an improvement to the BVT architecture, known as Sliceable-BVT (S-BVT), which addresses this issue. They describe an architecture that allows a transponder to propagate numerous BVT channels simultaneously. Channels in the S-BVT architecture are sliceable in that they can adapt to offer higher or lower modulation in any number of the given subchannels.




Figure 9: Conceptualization of the trade off between modulation/data rate and distance/noise with BVT. Noise, which can be measured with bit error rate, Q factor, or SNR, increases with the distance covered by an optical circuit. As more noise is accumulated over greater distance, the highest-order modulation that the circuit can support, and thereby the data rate on that circuit, falls in a piece-wise manner.

3.4. Silicon Photonics

Various materials (*e.g.*, GaAs, Si, SiGe) can be used to make photonics hardware required for data transmission. These devices include photodetectors, modulators, amplifiers, waveguides, and others. Silicon (Si) is the preferred material for these devices due to its low cost. However, there are challenges to manufacturing these silicon devices, such as optical power loss and free carrier absorption. Other materials, notably GaAs, have better properties for propagating light; however, GaAs is more costly to manufacture. Despite these challenges, research into efficient and quality transmission using silicon-based photonic devices has boomed in the last decade. Early advances were made towards silicon photonics (SiP) in the 80s, particularly for waveguides, which are the basis for circuit switches and multiplexers. Today, SiP is an integral part of almost all optical hardware, including lasers, modulators, and amplifiers.

A significant challenge for power-efficient SiP transceivers is coupling loss between the laser source and passive waveguide on Si integrated circuit waveguides, which can be as high as 2.3 dB, or 25% power loss [43]. Recent work by Billah *et al.* [44] explores the integration of indium phosphide (InP) lasers on chips, demonstrating a coupling with only 0.4 dB of loss, or roughly 10%. InP

appears to be a promising compound for other SiP technology too, as evident by demonstrations of InP in-line amplification for WSS [45]. Costs are falling for optical hardware as more efficient and scalable manufacturing techniques are enabled by SiP [46], thus allowing network operators to deploy newer technology into their systems at a more advanced pace as the devices' quality and guarantees have continued to improve. For more information on silicon photonics, see the survey by Thomson *et al.* [47].

3.5. Summary

Hardware for reconfigurable optical networks is improving at rapid scales, where researchers are developing more scalable optical switches with faster response times year after year. These WSS architectures are quickly being integrated with ROADMs to offer CDC-F flexibility for networks. Meanwhile, improvements to transponder technology are also paving the way for reconfigurable optics at network endpoints. In particular, S-BVTs offer dramatic CAPEX savings as one transponder can deliver multiple modulated signals in parallel. These improvements are accelerated by silicon photonics, bringing CMOS manufacturing to optical hardware and greatly reducing the cost to deploy optical switches and upgraded transponders in networks.

4. Optically Reconfigurable Data Centers

In this section, we illuminate efforts to improve DCNs with reconfigurable optics. Related surveys on this subject include Foerster *et al.* [7] and Lu *et al.* [9]. We divide the state of reconfigurable optical DCNs into technology, cost modeling, and algorithms. In technology, we supplement the discussion from § 3 with hardware capabilities that currently exist only for DCNs. Such features include free-space optics and sub-second switching. Next, we highlight cost modeling research, whose goal is to derive formal estimates or guarantees on the benefit of reconfigurable optical networks over static topologies for DCNs. Finally, we survey the relevant algorithms for managing and optimizing reconfigurable optical networks in the data center. Many of these algorithms focus on the interdependencies between optical path set-up and routing and optimize them across layers. Notwithstanding, there is also work that optimizes the physical layer simultaneously as well, respectively focuses on the interplay between software defined networking (SDN) and the physical layer, as illustrated in Figure 10. We discuss these examples in more detail and also survey further related work across the next subsections.

A key challenge for data centers is to optimize the utilization of the data center network (DCN). In a DCN, many different services are running and competing for shared bandwidth. Communication patterns between top-of-rack (ToR) switches vary with the underlying applications that are running (*e.g.*, map-reduce, video stream processing, physics simulations, *etc.*). Thus, as future applications and user's needs change, it is challenging to predict where bandwidth will be needed.

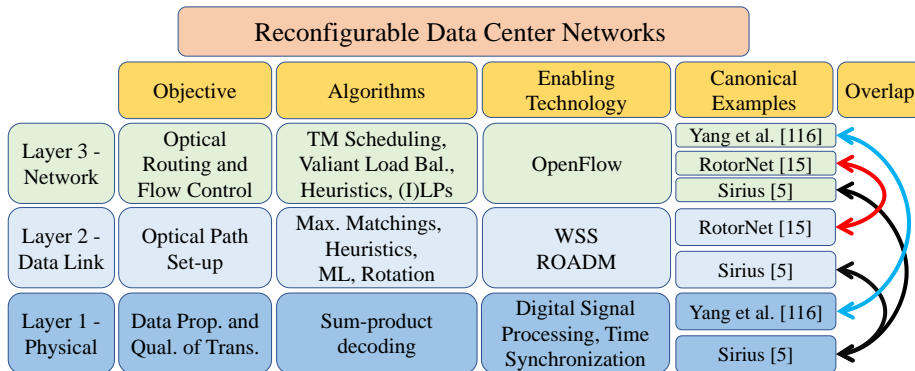


Figure 10: Solving the challenges involved in reconfigurable optics for data center layers networks requires bridging the gap between different technologies and goals for different layers of the network protocol stack.

Static and reconfigurable network solutions have been posed by research and industry to address this challenge. There is an assumption that the connectivity graph of the network cannot change in static network solutions. These solutions also assume fixed capacity (or bandwidth) on links. In reconfigurable network solutions, by contrast, these assumptions regarding connectivity and bandwidth are relaxed. Servers and switches (collectively referred to as nodes) may connect some subset of the other nodes in the network, and the nodes to which they are adjacent may change over time. Further, the bandwidth of a connection may also change over time.

Under the assumption of a static physical topology, different network architectures and best practices have been established. Some of these architectures include Clos, fat-tree, and torus topologies. Best practices include (over-)provisioning all links such that the expected utilization is a small fraction of the total bandwidth for all connections. These solutions can incur high cabling costs and are inefficient.

Reconfigurable network solutions circumvent the limitations of the static network solutions by reducing cabling costs or reducing the need to over-provision links. The flexibility of light primarily empowers these reconfigurable solutions. Some of these flexibilities include the steering of light (*e.g.*, with MEMs or polymer waveguides) and the high capacity of fiber-optics as a medium (*e.g.*, dense wavelength division multiplexing, or DWDM, enables transmitting $\mathcal{O}(Tb/s)$ on a single fiber).

4.1. DCN-specific Technologies

Innovations in reconfigurable optical networks are enabled by hardware’s evolution, as discussed in Section § 2. There is a subset of innovations that are well-suited for data centers only. These are *free-space optics* and *sub-second switching*. Although we have separated these below, there may be overlaps between free-space optics and sub-second switching systems as well.




Figure 11: Free-space optics switching architecture for data centers [50]

Free-space Optics. In free-space optics systems, light propagates through the air from one transceiver to another. Free-space optics enables operators to reduce their network’s complexity (a function of cabling cost). These closed environments and their highly variable nature of intra-data center traffic make such solutions appealing, we refer to the overview by Hamza *et al.* [48] for further application scenarios. Recent works such as Firefly [49] have demonstrated that free space optics are capable of reducing latency for time-sensitive applications by routing high-volume/low-priority traffic over the wireless optical network while persistently serving low-volume/high-priority traffic on a packet-switched network. High fan-out (1-to-thousands) for free-space optics is enabled with DMDs, or Dense Micro-mirror Devices, as shown by Project-ToR [50]. The DMDs are placed near Top-of-Rack (ToR) switches and pair with disco-balls, fixed to the ceiling above the racks. The DMD is programmed to target a specific mirror on the disco-ball, guiding the light to another ToR in the data center. Figure 11 illustrates the main properties of the free space optics deployment proposed in [50]. The deployment and operation of a free-space optics data center are fraught with unique challenges, *e.g.*, geometrical placement as investigated in 2D in OWCell [51] and in 3D in Diamond [52], but also particularly for keeping the air clear between transceivers and DMDs. Any particulate matter that the light comes into contact with can severely degrade performance and cause link failures should they persist. This phenomenon is known as atmospheric attenuation [53]. Another aspect is misalignment due to, *e.g.*, vibrations, requiring active alignment systems [54] respectively a tradeoff between beamwidth and received power density, depending on the distance covered [49]. In summary, even though free-space optics is an attractive alternative for many scenarios [48], and can be seen as “*fiber without the fiber*” [54], these technologies “*are not used in commercial data centers yet*” [55], and hence the main challenge is working towards their practical deployment. We refer to two recent specialized surveys for more details [55, 56].

Sub-second Switching. In data centers, distances are short between hosts, and therefore they do not lose their strength to such a degree that mid-line devices such as amplifiers are necessary. Therefore, applications can benefit from

all of the agility of optical layer devices without accounting for physical-layer impairments, which can slow down reconfiguration times in wide-area networks. Research has shown that micro-second switching of application traffic is possible in data center environments [57, 58, 59]. The ability to conduct circuit switching at microsecond timescales has illuminated further intrigue, particularly for transport protocols running on top of these networks. In c-Through [21], the authors observed that throughput for TCP applications dropped when their traffic migrated to the optical network. They showed how to mitigate this by increasing the queue size for optical circuit switches and adjusting the host behaviors. Mukerjee *et al.* [60] augmented their solution by expanding TCP for reconfigurable data center networks. Another method to deal with rapid reconfiguration times at a micro-second level is using traffic matrix scheduling, as we will further elaborate in Section § 4.3.

However already *e.g.*, Alistarh *et al.* [61] showcased the possibility of switching in the order of nano-seconds in a thousand port 25 Gbps+ optical switch design. Notwithstanding, a challenging question is how to make use of such fast reconfiguration times, when accounting for computation and routing update delays. Mellette *et al.* follow an intriguing design choice with their rotor switches [62], by creating demand-oblivious connections that change in the order of micro-seconds, in turn pre-configuring the routing in RotorNet [15] and Opera [63]. Project Sirius expands such ideas to the sub-nano-second level [18, 30], resulting in a demand-oblivious design that can perform end-to-end reconfigurations in less than 4 nano-seconds at 50 Gbps [5]. We further discuss these strategies in Section § 4.3.

Summary. Unlike in the WAN, data center technologies allow extremely fast switching times and high fan-out across the whole network, the latter in particular in the case of free-space optics. Hence especially the algorithmic design ideas allow substantially more flexibility and often differ fundamentally, as we will see in Section § 4.3.

4.2. Cost Modeling

Momentum has been building for data centers to move to optically switched and electrical/optical hybrid networks. However, there is a general reluctance to walk away from the old paradigm of a packet-switched-only network (PSO) due to the additional complexity of optical circuit switching (*e.g.*, the control plane management of optical circuits with shifting demand, and the variety of optical switching architectures available). Further, without a quantitative measure of *value-added* by optical switching over PSO, DCN operators are understandably reluctant to spend capital on an unvetted system. A discussion on the cost differences between optically and electrically switched data center networks can be found in the work of Kassing *et al.* [64], with an analysis for non-wired topologies in the works of Shin *et al.* [65] and Terzi and Korpeoglu [55].

To address the concerns surrounding complexity and value while raising awareness for the necessity of optically switched interconnects, researchers have constructed cost models to demonstrate the benefit of optical switching and hybrid architectures. Wang *et al.* [66] developed one such model. They conducted

intra-DC traffic measurements, which consisted of mixed workloads (*e.g.*, Map-Reduce, MPI, and scientific applications). They then played the traces back in simulation, assuming that three optical circuits could be created and reconfigured between racks every 30 seconds. Their data center with seven racks showed that rack-to-rack traffic could be reduced by 50% with circuit switching.

The following sections present more cost modeling work in the context of algorithmic simulations and systems implementations.

4.3. Algorithms

The capability of optical circuit switching for data center networks comes with the need to define new algorithms for optimizing utilization, bandwidth, fairness, latency, or any other metric of interest. Research has presented many different approaches for optimizing the metric relevant to the network operator in static networks. Traffic Engineering (TE) generally refers to the determination of paths for flows through the network, and the proportion of bandwidth levied for any particular flow. If the data center has a static network topology (*e.g.*, fat-tree), then TE is simple enough that switches can conclude how to route flows. However, introducing reconfigurable paths complicates the process of TE significantly: network elements (*e.g.*, switches) must now also determine with whom and when to establish optical paths, and when to change them.

Overview. The current algorithmic ideas to establish such optical paths can be classified into roughly five different areas, which we will discuss next. Due to the inherent hardware constraints (forming circuits), all of them rely on 1) matchings, where on its own the main idea is to maximize matching’s weight, *e.g.*, representing throughput, latency, etc. However maximum matchings can be slow to compute, and hence there has been interest in 2) demand-oblivious approaches, cycling through different network designs, 3) traffic matrix scheduling, to batch-compute a whole set of matchings ahead of time, and also leveraging the speed-up of 4) machine learning algorithms. Lastly, another way of quickly reacting to demand changes is by borrowing ideas from 5) self-adjusting data structures, in particular adapting the aspect of purely local circuit changes.

Matchings can be computed quickly [67] and often provide a good approximation, especially in settings where the goal is to maximize single-hop throughput along with reconfigurable links. Matching algorithms hence frequently form the basis of reconfigurable optical networks, *e.g.*, Helios [68], c-Through [21] and [69] rely on maximum matching algorithms. If there exist multiple reconfigurable links (say b many), it can be useful to directly work with a generalization of matching called b -matching [70]: b -matchings are for example used in Proteus [71] and its extension OSA [72], as well as in BMA [73] which relies on an online b -matching algorithm; BMA also establishes a connection to online (link) caching problems. In some scenarios, for example, when minimizing the average weighted path length under segregated routing, maximum b -matching algorithms even provide optimal results [74, 75]. This however is not always true, *e.g.*, when considering non-segregated routing policies [74, 75], which require heuristics [49, §5.1], [76].

Oblivious Approaches. Matchings also play a role in reconfigurable networks which do not account for the traffic they serve, *i.e.*, in *demand-oblivious* networks. The prime example here is RotorNet [15] which relies on a small set of matchings through which the network cycles endlessly: since these reconfigurations are “dumb”, they are fast (compared to demand-aware networks) and provide frequent and periodic direct connections between nodes, which can significantly reduce infrastructure cost (also known as “bandwidth tax”) compared to multihop routing, see also Teh *et al.* [77]. In case of uniform (delay-tolerant) traffic, such single-hop forwarding can saturate the network’s bisection bandwidth [15]; for skewed traffic matrices, it can be useful to employ Valiant load balancing [78] to avoid underutilized direct connections, an idea recently also leveraged in Sirius [5] via Chang *et al.* [79]. Opera [63] extends RotorNet by maintaining expander graphs in its periodic reconfigurations. Even though the reconfiguration scheduling of Opera is deterministic and oblivious, the pre-computation of the topology layouts in their current form is still randomized. Expander graphs (and their variants, such as random graphs [80]) are generally considered very powerful in data center contexts. An example of a demand-aware expander topology was proposed in Tale of Two Topologies [81], where the topology locally converts between Clos and random graphs.

Traffic Matrix Scheduling. Another general algorithmic approach is known as *traffic matrix scheduling*: the algorithmic optimizations are performed based on a snapshot of the demand, *i.e.*, based on a traffic matrix. For example, Mordia [82] is based on an algorithm that reconfigures the network multiple times for a single (traffic demand) snapshot. To this end, the traffic demand matrix is scaled into a bandwidth allocation matrix, which represents the fraction of bandwidth every possible matching edge should be allocated in an ideal schedule. Next, the allocation matrix is decomposed into a schedule, employing a computationally efficient [83] Birkhoff-von-Neumann decomposition, resulting in $O(n^2)$ reconfigurations and durations. This technique also applies to scheduling in hybrid data center networks which combine optical components with electrical ones, see *e.g.*, the heuristic used by Solstice [84]. Eclipse [85] uses traffic matrix scheduling to achieve a $(1 - 1/e^{(1-\varepsilon)})$ -approximation for throughput in the hybrid switch architecture with reconfiguration delay, but only for direct routing along with single-hop reconfigurable connections. Recently Gupta *et al.* [86] expanded similar approximation guarantees to multi-hop reconfigurable connections, for an objective function closely related to throughput.

While Eclipse is an offline algorithm, Schwartz *et al.* [87] presented online greedy algorithms for this problem, achieving a provable competitive ratio over time; both algorithms allow to account for reconfiguration costs. Another example of traffic matrix scheduling is DANs [88, 89, 90, 91] (short for demand-aware networks, which are optimized toward a given snapshot of the demand). DANs rely on concepts of demand-optimized data structures (such as biased binary search trees) and coding (such as Huffman coding) and typically aim to minimize the expected path length [88, 89, 90, 91], or congestion [89]. In general, the problem features intriguing connections to the scheduling literature, *e.g.*, the work by Anand *et al.* [92], and more recently, Dinitz *et al.* [93] and Kulkarni *et*

al. [94]; the latter two works however are not based on matchings or bipartite graphs. In Dinitz *et al.* [93], the demands are the edges of a general graph, and a vertex cover can be communicated in each round. Each node can only send a certain number of packets in one round. The approach by Kulkarni *et al.* [94] considers a model where communication requests arrive online over time and uses an analysis based on LP relaxation and dual fitting.

Self-Adjusting Datastructures. A potential drawback of traffic matrix scheduling algorithms is that without countermeasures, the optimal topology may change significantly from one traffic matrix snapshot to the next, even though the matrix is similar. There is a series of algorithms for reconfigurable networks that account for reconfiguration costs, by making a connection to self-adjusting data structures (such as splay trees) and coding (such as dynamic Huffman coding) [90, 95, 96, 97, 98, 99, 100, 101]. These networks react *quickly and locally* to two new communication requests, aiming to strike an optimal trade-off between the benefits of reconfigurations (*e.g.*, shorter routes) and their costs (*e.g.*, reconfiguration latency, energy, packet reorderings, *etc.*).

To be more specific, the idea of the self-adjusting data structure based algorithms is to organize the communication partners (*i.e.*, the destinations) of a given communication source in either a static binary search or Huffman tree (if the demand is known), or in a dynamic tree (if the demand is not known or if the distribution changes over time). The tree optimized for a single source is sometimes called the *ego-tree*, and the approach relies on combining these ego-trees of the different sources into a network while keeping the resulting node degree constant and preserving distances (*i.e.*, low distortion). The demand-aware topology resulting from taking the union these ego-trees may also be complemented with a demand-oblivious topology, *e.g.*, to serve low-latency flows or control traffic; see the ReNet architecture for an example [99].

Machine Learning. Another natural approach to devise algorithms for reconfigurable optical networks is to use machine learning. To just give two examples, xWeaver [17] and DeepConf [102] use neural networks to provide traffic-driven topology adaptation. Another approach is taken by Kalmbach *et al.* [103], who aim to strike a balance between topology optimization and “keeping flexibilities”, leveraging self-driving networks. Finally, Truong-Huu *et al.* [104] proposed an algorithm that uses a probabilistic, Markov-chain based model to rank ToR nodes in data centers as candidates for light-path creation.

Accounting for Additional Aspects. Last but not least, several algorithms account for additional and practical aspects. In the context of shared mediums (*e.g.*, non-beamformed wireless broadcast, fiber³ (rings)), contention and interference of signals can be avoided by using different channels and wavelengths. The algorithmic challenge is then to find (optimal) edge-colorings on multi-graphs, an NP-hard problem for which fast heuristics exist [106]. However, on specialized topologies, optimal solutions can be found in polynomial

³In the context of data center proposals, shared fiber is the more popular medium, *e.g.*, in [72, 105, 82].

time, *e.g.*, in WaveCube [107]. Shared mediums also have the benefit that it is easier to distribute data in a one-to-many setting [108]. For example, on fiber rings, all nodes on the ring can intercept the signal [105, §3.1]. One-to-many paradigms⁴ such as multicast can also be implemented in other technologies, using *e.g.*, optical splitters for optical circuit switches or half-reflection mirrors for free-space optics [111, 112, 113, 114, 115].

4.4. Systems Implementations

There have been many demonstrations of systems for reconfigurable optics in data centers. Many of the papers that we discuss in Section 4.3 are fully operational systems. Another notable research development that does not fit into algorithms is the work by Mukerjee *et al.* [60]. They describe amendments to the TCP protocol to increase the efficiency of reconfigurable data center networks. These amendments include dynamic buffer re-sizing for switches and sharing explicit network feedback with hosts. Moreover, Yang *et al.* [116] showcase an interesting cross-layer aspect where the physical layer itself is controlled by SDN, in the sense that they allow for transceiver tuning in real-time. Their main contributions relate to new SDN control modules and interfaces, being orthogonal to (scheduling) algorithms. Much of the other work on reconfigurable DCNs are summarized in Table 3.

We see two main conceptual differences in current reconfigurable data center network designs, namely concerning 1) the demand-aware or -oblivious circuit control plane and the 2) all-optical or hybrid fabric. Sirius [5], Opera [63], and RotorNet [15] all propose a demand-oblivious optical layer, in essence rotating through a set of topologies, letting the higher layers take advantage of the changing optical connections. To this end, there is no computational delay, but on the other hand, specifically skewed demands can suffer from performance degradation. Demand-aware control planes can adapt to any demands but need careful tuning to avoid scaling and prediction issues, which then again can be inferior to demand-oblivious network designs, depending on the scenario. Notwithstanding, the three listed demand-oblivious designs currently rely on specialized and experimental hardware. Regarding the choice of fabric, hybrid designs are highly beneficial for small and short-lived flows, and hence a combination of packet and circuit switching, such as in RotorNet [15] or Eclipse [16], can combine the best of both worlds. Notwithstanding, provisioning for both types of networks leads to overheads in cost and terms of cross-fabric efficiency, and thus are not a silver bullet solution. An intriguing design in this context is Opera [63], as it always provisions a small diameter network with optical links, emulating classic DCN properties inside their circuit choices. However, as mentioned above, this design choice comes with the price of demand-obliviousness, and it would be interesting to see how other all-optical demand-aware systems, such as *e.g.*, OSA [72], can implement such properties as well

⁴Conceptually similar challenges arise for coflows [109, 110].

4.5. Summary

There is a wide range of data center specific technology and algorithmic ideas that enable efficient circuit switching in data center networks, with newer developments focusing on leveraging the benefits of faster circuit reconfigurations. In contrast, there has also been some recent work [117] that discusses the idea of robust topology engineering, *e.g.*, adapting the circuits only every few minutes or even days [118]. Notwithstanding, scaling current system designs can be problematic, in particular, due to the speed of the control plane and fan-out restrictions. Whereas one solution for the latter is free-space optics, those still face significant practical deployment issues in data center contexts. On the other hand, demand-oblivious system designs inherently overcome such control plane delays, but cannot adapt well to skewed demands. In their current form, they are not available as off-the-shelf hardware. Designing scalable demand-aware reconfigurable data centers is hence one of the main next challenges.

5. Reconfigurable Optical Metro and Wide-area Networks

In this section, we survey recent research in reconfigurable optics in metropolitan (metro) and wide-area networks (WAN). Reconfigurable optics refers to dynamism in the physical-layer technology that enables high-speed and high-throughput WAN communications, fiber optics. We divide reconfigurable optical innovations into two sub-categories, rate-adaptive transceivers, and dynamic optical paths. Rate adaptive transceivers, or bandwidth-variable transceivers (introduced in section 3.3) are optical transceivers that can change their modulation format to adapt to physical layer impairments such as span-loss and noise. Dynamic optical paths refer to the ability to *steer light*, thus allowing the edges of the network graph to change (*e.g.*, to avoid a link that has failed).

Many groups have studied the programmability and autonomy of optical networks. Gringeri *et al.* [119] wrote a concise and illuminating introduction to the topic. In it, the authors propose extending Software Defined Network (SDN) principles to optical transport networks. They highlight challenges, such as reconfiguration latency in long-haul networks, and provide a trade-off characterization of distributed vs. centralized control for an optical SDN system. They claim that a tiered hierarchy of control for a multi-regional network (*e.g.*, segregated optical and network control loops) will offer the best quality solution. Further, they argue that centralized control should work best to optimize competing demands across the network, but that the controller's latency will be too slow to react to network events, *e.g.*, link outages quickly. Therefore, the network devices should keep some functionality in their control plane to respond to link failures in a decentralized manner, *e.g.*, reallocating the lost wavelengths by negotiating an alternative path between the endpoints.

The question of centralized vs. distributed network control is just one example of the many interesting questions that arise when considering reconfigurable optical networks for metro and wide areas. This space is unique because many of the solutions here require understanding and sharing of information across

	Fabric	Demand-Aware	Novelty
Helios [68]	Hybrid	✓	First hybrid system using WDM for busy low-latency traffic
c-Through [21]	Hybrid	✓	Enlarged buffers for optical ports increases utilization
ProjecToR [50]	Hybrid/FSO	✓	Introduces DMDs for free-space switching thus enabling a fan-out potential to thousands of nodes
Proteus [71]	All-optical	✓	Design of an all-optical and reconfigurable DCN.
OSA [72]	All-optical	✓	Demonstrates greater reconfiguration flexibility and bisection bandwidth than hybrid architectures
RotorNet [15]	Hybrid	×	An all-optical demand-oblivious DCN architecture for simplified network management
Opera [63]	All-optical	×	Extends Rotornet to include expander graphs rotations
Flat-tree [81]	Hybrid	✓	A hybrid of random graphs and Clos topologies brings reconfigurable optics closer to existing DCNs.
Solstice [84]	Hybrid	✓	Exploits sparse traffic patterns in DCNs to achieve fast scheduling of reconfigurable networks.
Eclipse [16]	Hybrid	✓	Outperforms Solstice by applying submodular optimization theory to hybrid network scheduling.
xWeaver [17]	Hybrid	✓	Trains neural networks to construct performant topologies based on training data from historic traffic traces.
DeepConf [102]	Hybrid	✓	Presents a generic model for constructing learning systems of dynamic optical networks
WaveCube [107]	Hybrid	✓	A modular network architecture for supporting diverse traffic patterns.
Sirius [5]	All-optical	×	Achieves nanosecond-granularity reconfiguration for thousands of nodes

Table 3: Summary of systems implementations of reconfigurable data center networks

layers of the network stack. For example, figure 12 illustrates interdependence between the objectives for communication across different layers of the stack; these features include algorithms, enabling technologies. We highlight several canonical examples of systems that exist in those domains and across different layers. In this section, we will explore these examples more deeply along with other related efforts.

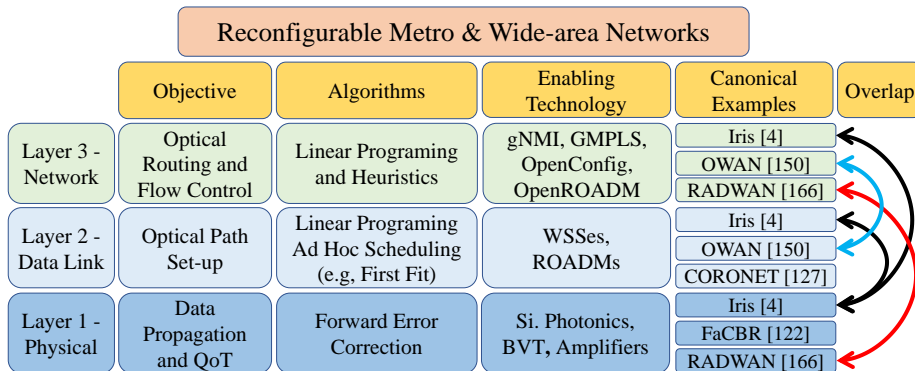


Figure 12: To deploy and operate reconfigurable optical networks in metro and wide-area networks require expertise spanning the bottom three layers of the network stack, including algorithms and enabling technology. We highlight several canonical examples of systems that exist in this space and explore other related works along with these systems more deeply in this section.

5.1. Metro/WAN-Specific Challenges and Solutions

There are many reasons for the prevalence of optical fiber as the de-facto leader for long-distance communications. First, it has incredible reach compared to copper—optical signals can propagate 80 to 100 km before being amplified. Second, it has an incredibly high bandwidth compared to the radio spectrum. Third, optical fiber itself has proved to be a robust medium over decades, as improvements to the transponders at the ends of the fiber have enabled operators to gain better value out of the same fiber year after year.

To design a WAN, the network architect must solve several difficult challenges, such as estimating the demand on the network now and into the future, optimal placement of routers and quantity of ports on those routers within the network, and optimal placement of amplifiers in the network.

Many design challenges solve more easily in a static WAN, where optical channels are initialized once and maintained for the network’s life. For example, amplifiers carrying the channel must have their gain set in such a way that the signal is transmitted while maximizing the signal-to-noise ratio (SNR). This calculation can take minutes or hours depending on the network’s characteristics (*e.g.*, the number of indeterminate hosts and the number of distinct channels on shared amplifiers).

Dynamic optical networks must rapidly address these challenges (in sub-second time frames) to achieve the highest possible utilization, posing a significant challenge. For example, it requires multiple orders of magnitude increases in the provisioning time for optical circuits beyond what is typically offered by hardware vendors. Therefore, several research efforts have explored ways to automate WAN network elements’ configuration concerning physical layer impairments in a robust and time-efficient manner.

Chromatic Dispersion. DWDM makes efficient use of optical fiber by putting as many distinct optical channels, each identified by a frequency (or

lambda λ) onto the shared fiber. Each of these lambdas travels at a different speed relative to the speed of light. Therefore, two bits of information transmitted simultaneously via two different lambdas will arrive at the destination at two different times. Further, chromatic dispersion is also responsible for pulse-broadening, which reduces channel spacing between WDM channels and can cause FEC errors. Therefore, DWDM systems must handle this physical impairment.

Amplified Spontaneous Emission (ASE) Noise. A significant limitation of circuit switching is the latency of establishing the circuit due to ASE noise constraints [120]. Although SDN principles can apply to ROADMs and WSSs (to automate the control plan of these devices), physical layer properties, such as Noise Figure (NF) and Gain Flatness (GF) complicate the picture. When adding or removing optical channels to or from a long-haul span of fiber, traversing multiple amplifiers, the amplifiers on that path must adjust their gain settings to accommodate the new set of channels. To this end, researchers have worked to address the challenge of dynamically configuring amplifiers. Oliveira *et al.* [121] demonstrated how to control gain on EDFAs using GMPLS. They evaluated their solution on heterogeneous optical connections (10, 100, 200, and 400 Gbps) and modulations (OOK, QPSK, and 16-QAM). They used attenuators to disturb connections and allow their GMPLS control loop to adjust the amplifier's gains. They show that their control loop helps amplifiers to adjust while transmitting bits with BER below the FEC threshold for up to 6 dB of added attenuation.

Moura *et al.* [122] present a machine learning approach for configuring amplifier gain on optical circuits. Their approach uses case-based reasoning (CBR) as a foundation. The intuition behind CBR is that the gain setting for a set of circuits will be similar if similar circuits are present on a shared fiber. They present a genetic algorithm for configuring amplifiers based on their case-based reasoning assumption. They show that their methodology is suitable for configuring multiple amplifiers on a span with multiple optical channels. In a follow-up study, they present FAcCBR [123], an optimization of their genetic algorithm, which yields gain recommendations more quickly by limiting the number of data-points recorded by their algorithm.

Synchronization. Managing a WAN requires coordinating services (*e.g.*, end-to-end connections) among diverse sets of hardware appliances (transponders, amplifiers, routers), logically and consistently. The Internet Engineering Task Force (IETF) has defined protocols and standards for configuring WAN networks. As the needs and capabilities of networks have evolved, so have the protocols. Over the years, new protocols have been defined to bring more control and automation to the network operator's domain. These protocols are Simple Network Management Protocol (SNMP) [124] and Network Configuration Protocol (NETCONF) [125]. Additionally, network operators and hardware vendors have been working to define a set of generalized data models and configuration practices for automating WAN networks under the name OpenConfig [126]. Although OpenConfig is not currently standardized with the IETF, it is deployed and has demonstrated its value in several unique settings.

In addition to the standardized and proposed protocols for general-purpose WAN (re)configuration, there has been a push by various independent research groups to design and test protocols specifically for reserving and allocating optical channels in WAN networks.

One protocol was developed in conjunction with the CORONET [127] program, whose body of research has led to several other developments in reconfigurable optical WANs. The proposal, by *Skoog et al.* [128], describes a three-way handshake (3WHS) for reserving and establishing optical paths in single and multi-domain networks. In the 3WHS, messages are exchanged over an optical supervisory channel (OSC)—an out-of-band connection between devices isolated from user traffic. The transaction is initiated by one Optical Cross-Connect (OXC^A) and directed at a remote OXC, OXC^Z . At each hop along the way, the intermediate nodes append the available channels to the message. Then, OXC^Z chooses a channel via the first-fit strategy [129] and sends a message to OXC^A describing the chosen channel. Finally, OXC^A activates the chosen channel and begins sending data over it to OXC^Z . This protocol is claimed to meet the CORONET project standard for a setup time of 50 ms + RTT between nodes. Bit arrays are used to communicate the various potential channels between nodes and are processed in hardware. The blocking probability is 10^{-3} if there is one channel reserved between any two OXC elements so long as there are at least 28 total channels possible between OXCs [128].

5.2. Cost Modeling

Fiber infrastructure for wide-area networks is incredibly costly. Provisioning of fiber in the ground requires legal permitting processes through various governing bodies. As the length of the span grows beyond metropolitan areas, to connect cities or continents, the number of governing bodies with whom to acquire the legal rights to lay the fiber grows [130]. Then, keeping the fiber lit also incurs high cost; power requirements are a vital consideration for wide-area network provisioning [131]. Therefore, reliable cost models are necessary for deploying and managing wide-area networks. In this section, we look at cost modeling efforts particularly suited for reconfigurable optical networks.

An early study on the cost comparison of IP/WDM vs IP/OTN networks (in particular: European backbone networks) was conducted by Tsirilakis *et al.* in [132]. The IP/WDM network consists of core routers connected directly over point-to-point WDM links in their study. In contrast, the IP/OTN network connects the core routers through a reconfigurable optical backbone consisting of electro-optical cross-connects (OXCs) interconnected in a mesh WDM network.

Capacity planning is a core responsibility of a network operator in which they assess the needs of a backbone network based on the projected growth of network usage. Gerstel *et al.* [133] relates the capacity planning process in detail, which includes finding links that require more transponders and finding shared-risk-link-groups that need to be broken-up, among other things. They note that in this process, the IP and Optical network topologies are historically optimized separately. They propose an improvement to the process via multi-layer optimization, considering the connection between IP and optical layers.

They save 40 to 60% of the required transponders in the network with this multi-layer approach. The networks they looked at were Deutsche Telekom [134] and Telefonica Spain core networks. These authors' work provides a strong motivation for jointly optimizing IP and Optical network layers and sharing of information between the two.

Papanikolaou *et. al.* [135] propose a cost model for joint multi-layer planning for optical networks. Their paper presents three network planning solutions; dual-plane network design, failure-driven network design, and integrated multi-layer survivable network design. They show that dual-plane and failure-driven designs over-provision the IP layer, leaving resources on the table that are only used if link failures occur. They show that integrated multi-layer survivable network design enables a significant reduction in CapEx and that the cost savings increases beyond dual plane and failure driven designs.

Cost models for evaluating C-ROADM vs. CDC-ROADM network architectures are described by Kozdrowski *et al.* [136]. They show that for three regional optical networks (Germany, Poland, USA), CDC-ROADM based networks can offer 2 to 3 \times more aggregate capacity over C-ROADM based networks. They evaluate their model with uniform traffic matrices (TMs) and apply various scalar multipliers to the TM. Their model accounts for many optical hardware related constraints, including the number of available wavelengths and cost factors associated with manual-(re)configuration of C-ROADM elements. However, their model doesn't include an optical-reach constraint. They limit solver computation time to 20 hours and present the best feasible solution determined in that amount of time.

Service velocity refers to the speed with which operators may grow their network as demand for capacity grows. Woodward *et al.* [137] tackles the problem of increasing service velocity for WANs. In this context, they assume a network of colorless non-directional ROADMS (CN-ROADMs)⁵, in which any incoming wavelength can be routed on any outgoing fiber. They claim that one of the largest impedances for network growth in these networks is the availability of *re-generators*. To solve this problem, they present three algorithms for determining regenerators' placement in a network as service demand grows. The algorithms are: locally aware, neighbor aware, and globally aware. Each algorithm essentially considers a broader scope of the network, which a node uses to determine if an additional regenerator is needed at the site at a particular time. They show, via Monte Carlo simulations, varying optical reach and traffic matrices. The broadest scope algorithm performs the best and allocates enough regenerators at the relevant sites without over-provisioning. This work shows that service velocity is improved with demand forecasting, enabling infrastructure to be placed to meet those projected demands.

Programmable and elastic optical networks can also work together with Network Function Virtualization (NFV) to offer lower-cost service-chaining to

⁵CN-ROADMs are also called CD-ROADMs in other papers. These both refer to the same ROADM architecture.

users. Optimal strategies have been demonstrated, with heuristic algorithms, to quickly find near-optimal solutions for users and service brokers by Chen *et al.* [138]. In their work, they take a game-theoretic approach to modeling the competition among service brokers—who complete offering the lowest cost optical routs and service chains, and between users—who compete to find the lowest cost and highest utility service chains among the brokers. They demonstrate both parties’ strategies, which converge on low-latency service chain solutions with low blocking probability for optical paths.

Modeling *opportunity cost* of optically switched paths is explored by Zhang *et al.* [139]. In their work, they present an algorithm for quickly evaluating the opportunity cost of a wavelength-switched path. Given a request and a set of future requests, the opportunity cost for accommodating the initial request is the number of future requests blocked as a result of the accommodation. Thus, the network operator’s goal is to minimize opportunity cost by permitting connections that interfere with the fewest future requests.

5.3. Algorithms

Jointly optimizing both the optical and the network layer in wide-area networks leads to new opportunities to improve performance and efficiency, while introducing new algorithmic challenges. In contrast to the previously discussed data center networks, it is impossible to create new topological connections in a wide-area network (without deploying more fiber. Free space optics solutions don’t apply here). Instead, reconfigurability is possible by adjusting and shifting bandwidth capacities along the fiber edges, possibly over multiple hops. Hence, we need a different set of algorithmic ideas that optimize standard metrics such as throughput, completion time, blocking probability, and resilience. In this section, we discuss recent papers that tackle these issues, starting with some earlier ones. Moreover, there is the need for some central control to apply the routing, policy, lightpath *etc.* changes, for which we refer to recent surveys [140, 141].

Routing aspects are explored intensively in this context. Algorithmic approaches to managing reconfigurable optical topologies have been studied for a decade, but are recently gaining new attention. In early work by Kodialam *et al.* [142] explores IP and optical wavelength routing for a series of connection requests. Their algorithm determines whether a request should be routed over the existing IP topology, or if a new optical path should be provisioned for it. Interesting work by Brzezinski and Modiano [143] who leverage matching algorithms and Birkhoff–von Neumann matrix decompositions and evaluate multi-versus single-hop routing⁶ in WDM networks under stochastic traffic. However, the authors mostly consider relatively small networks, *e.g.*, with three to six nodes. For larger networks, shortest lightpath routing is a popular choice [145]. Another fundamental aspect frequently considered in the literature regards resilience [146, 147, 148]. For example, Xu *et al.* [146] investigate resilience in the context of shared risk link groups (SLRGs) and propose a method on how

⁶See also the idea of lightpath splitting in Elastic Optical Networks [144].

to provision the circuits in a WAN. To this end, they construct Integer Linear Programs to obtain maximally SLRG-diverse routes, which they then augment with post-processing for DWDM system selection and network design issues. We now introduce further selected algorithmic works, starting with the topic of bulk transfers [149].

In *OWAN* [150], Jin *et al.* optimize bulk transfers in a cross-layer approach, which leverages both the optical and the network layer. Their main objective is to improve completion time; while an integer linear program formulation would be too slow, the authors rely on a simulated annealing approach. A local search shifts wavelength allocations, allowing heuristic improvements to be computed at a sub-second scale. The scheduling of the bulk transfer then follows the standard shortest job first approaches. When updating the network state, if desired, *OWAN* can extend prior consistent network update solutions [151] by introducing circuit nodes in the corresponding dependency graphs. *OWAN* also considers deadline constrained traffic, implementing the earliest deadline first policy. Follow-up work extended *OWAN* in two directions, via theoretic scheduling results and for improvements on deadline-constrained transfers.

In *DaRTree* [152], Luo *et al.* develop an appropriate relaxation of the cross-layer optimization problem for bulk transfers under deadlines. Their approach relies on a non-greedy allocation in an online setting, which allows future transfers to be scheduled efficiently without needing to reallocate currently utilized wavelengths. To enhance multicast transfers (*e.g.*, for replication), they develop load-adaptive Steiner Tree heuristics.

Jia *et al.* [153] design various online scheduling algorithms and prove their competitiveness in the setting of *OWAN* [150]. The authors consider the minimum makespan and sum completion time, analyzing and extending greedy cross-layer scheduling algorithms, achieving small competitive ratios. Dinitz and Moseley [154] extend the work of Jia *et al.* by considering a different objective, the sum of flow times in an online setting. They show that resource augmentation is necessary for acceptable competitive bounds in this setting, leading to nearly (offline) optimal competitive ratios. While their algorithms are easy to implement (*e.g.*, relying on ordering by release time or by job density), the analysis is complicated and relies on linear program relaxations. Moreover, their algorithm also allows for constant approximations in the weighted completion time setting, without augmentations.

Another (algorithmic) challenge is the integration of cross-layer algorithms into current traffic engineering systems. Such TEs are tried and tested, and hence service providers are reluctant to adapt their designs. To this end, Singh *et al.* [155] propose an abstraction on how dynamic link capacities (*e.g.*, via bandwidth variable transceivers) can be inserted into classic TEs. Even though the TE is oblivious to the optical layer, an augmentation of the IP layer with fake links enables cross-layer optimization via the TE. A proposal [156] for a new TE for such dynamic link capacities is discussed in the next Section 5.4. Singh *et al.* [155] also discuss consistent update methods [157] for dynamic link capacities, which Tseng [158] formalizes into a rate adaption planning problem, providing intractability results and an LP-based heuristic.

OptFlow [159] proposes a cross-layer abstraction for programmable topologies as well, but focuses on shifting wavelengths between neighboring fibers. Here, the abstraction concept is extended by not only creating fake links but also augmenting the traffic matrix with additional flows. As both links and flows are part of the input for TEs, *OptFlow* enables the compilation of optical components into the IP layer for various traffic engineering objectives and constraints. Concerning consistent updates, classic flow-based techniques [157] carry over, enabling consistent cross-layer network updates too.

Optimizing reconfigurable optical networks for circuit provisioning and per-flow rate allocation is a complex and challenging endeavor; the static routing and wavelength allocation problem is NP-complete [160]. Recent work by Guo *et al.* [161] explores the potential for an artificial intelligence (AI) implementation of a network controller using deep-learning. They describe a network control agent based on deep-learning which determines where and when to activate and deactivate a limited set of circuits given a snapshot of demand between hosts in the network. They also explore inherent drawbacks and precautions to consider settings in which such an agent is deployed. Their study offers insights for the potential benefit of an AI-assisted optical network controller, and novel challenges to consider for their given model.

Algorithms that optimize optical network topology for higher-layer applications, such as virtual network functions (VNF) have recently gained attention. In particular, VNF network embedding (VNF-NE) has been studied by various groups [162, 163]. VNFs are an abstraction of resources in networks that have traditionally been deployed as hardware devices (*e.g.*, intrusion detection systems, firewalls, load-balances, *etc.*). Now, instead of monolithic hardware appliances many of these devices are deployed as software on commodity servers, giving more flexibility to add and remove them at will and yielding cost-savings for network operators. Network embedding is a physical layer abstraction for creating end-to-end paths for network applications or network function virtualization (NFV) service chains. Paths have requirements for both bandwidth and CPU resources along the service chain. Wang *et al.* [162] proves this problem to be NP-complete for elastic optical networks. Soto *et al.* [163] provides an integer linear program (ILP) to solve the VNF-NE problem. The ILP solution is intractable for large networks. Thus they provide a heuristic that uses a ranking-system for optical paths. Their heuristic ranks optical paths by considering a set of end-to-end connection requests. Paths with higher rank satisfy a more significant proportion of the demand for bandwidth and CPU among all of the requests.

Optical layer routing with traffic and application constraints is a difficult problem. The running theme has been that linear programming solutions can find provably optimal solutions [164], but take too long to converge for most use cases. However, network traffic is not entirely random and therefore has an underlying structure that may be exploited by offline linear program solvers, as shown by Kokkinos *et al.* [165]. They use a two-stage approach for routing optical paths in an online manner. Their technique finds periodic patterns over an epoch (*e.g.*, daily, weekly, or monthly) and solves the demand characterized

	BVT	Network Design	Amps.	Algorithms
CORONET [127]	×	×	×	ROLEX protocol
OWAN [150]	×	×	×	Simulated Annealing
FACcBR [122]	×	×	✓	Case Based Reasoning
RADWAN [166]	✓	×	×	Linear Program
DDN [167]	×	✓	✓	Time-slotted packet scheduling
Iris [4]	×	✓	✓	Shortest path for any failure scenario

Table 4: Summary of systems implementations of reconfigurable wide area networks

within the epoch with an offline linear program. Then, their online heuristic makes changes to the topology to accommodate random changes in demand within the epoch.

5.4. Systems Implementations

The integration of reconfigurable optics with WAN systems has been impracticable due to its cost and a lack of convergence on cross-layer APIs for managing the WAN optical layer with popular SDN controllers. However, some exciting work has demonstrated the promise for reconfigurable optics in closed settings. Notably, RADWAN [166] and CORONET [127] for bandwidth-variable WAN systems and systems with dynamic optical paths, respectively. In this section, we explore reconfigurable optical WAN systems more deeply in these two contexts. Table 4 summarizes these systems.

Bandwidth Variable Transceivers. A team of researchers at Microsoft evaluates bandwidth variable transponders’ applicability for increased throughput in Azure’s backbone in North America [40]. They find that throughput for the WAN can increase if they replace the fixed-rate transponders in their backbone network with three-way sliceable transponders. They also show that for higher-order slices, bandwidth gran increases at diminishing returns.

Traffic Engineering with rate-adaptive transceivers was recently proposed by Singh *et al.* [156]. The authors are motivated by a data-set of Microsoft’s WAN backbone Signal-to-Noise ratio from all transceivers in the North-American backbone, over two and a half years. They note that over 60% of links in the network could operate at $0.75\times$ higher capacity and that 25% of observed outages due to SNR drops could be mitigated by reducing the modulation of the affected transceivers. They evaluate the reconfigurability of Bandwidth-Variable Transponders, showing that reconfiguration time for the transceivers could be reduced from minutes to milliseconds by *not* turning the transceivers off. Then, they propose a TE objective function via linear-programming, to minimize churn, or impact due to SNR fluctuations, in a WAN. Finally, they evaluate their TE controller on a testbed WAN and show that they improve network throughput by 40% over a competitive software-defined networking controller, SWAN [168].

Dynamic Optical Paths. In the early aughts, researchers explored the benefit of dynamic optical paths for networks in the context of *grid-computing*. Early efforts by Figueira *et al.* [169] addressed how a system might manage dynamic optical paths in networks. In this work, the authors propose a web-based interface for submitting optical reconfiguration requests and a controller for optimizing the requests' fulfillment. They evaluate their system on OMNInet [170], a metropolitan area network with 10 Gbps interconnects between 4 nodes and Wavelength Selective Switches between them. They claim that they can construct optical circuits between the OMNInet nodes in 48 seconds. Further, they show that amortized setup time and transfer is faster than packet-switching for files 2.5 Gb or larger (assuming 1 Gbps or greater optical interconnect and 300 Mbps packet switching throughput). They go on to evaluate file transfer speeds using the optical interconnect and show that they can archive average transfer speeds of 680 Gbps. Iovanna *et al.* [171] address practical aspects of managing multilayer packet-optical systems. They present a set of useful abstractions for operating reconfigurable optical paths in traffic engineering using an existing management protocol, GMPLS.

Stability is an important feature of any network. An interesting question about reconfigurable optical networked systems arises regarding the stability of optically switched paths. That is if the topology can continuously change to accommodate random requests, what service guarantees can the network make? Can the fluctuation of the optical layer be detrimental to IP layer services? Chamania *et al.* [172] explore this issue in detail, providing an optimal solution to keep quality of service guarantees for IP traffic while also improving performance beyond static optical layer systems.

Blocking probability is a crucial metric for assessing the flexibility of an optical network. It is the probability that a request for an end-to-end lightpath in the network cannot be provisioned. Turkcü *et al.* [173] provides analytical probability models to predict the blocking probability in ROADM based networks with tunable transceivers and validate their models with simulation considering two types of ROADM architecture in their analysis, namely *share-per-node* and *share-per-link*. In share-per-link, each end of a link has a fixed number of transponders that can use it. In share-per-node, a node has a fixed set of transponders that may use any incident links. The authors show that a low tunable range (4 to 8 channels, out of 32 possible) is sufficient for reducing blocking probability in two topologies, NSF Net (14 Nodes), and a ring topology with 14 nodes. As the tunable range moves beyond 8 and up to 32, there is little to no benefit for split-per-node and share-per-link architectures. As the load on the network increases, blocking probability increases, as well as the gap between blocking probability of split-per-node and split-per-link decreases.

Bandwidth-on-demand (BoD) is an exciting application of reconfigurable networks. Von Lehmen *et al.* [127] describe their experience in deploying BoD services on CORONET, DARPA's WAN backbone. They implement protocols for add/dropping wavelengths in their WAN with a novel 3-way-handshake protocol. They demonstrate how their system can utilize SWAN [168] Traffic Engineering Controller as one such application that benefits from the BoD

service.

More recently, there has been a resurgence of academic work highlighting the potential benefit of dynamic optical paths in the WAN. One such system, called OWAN (Optical Wide-Area Network) [174], proposes how to use dynamic optical paths to improve the delivery time for bulk transfers between data centers. They build a testbed network with home-built ROADMs and implement a TE controller to orchestrate bulk transfers between hosts in a mesh optical network of nine nodes. They compare their results with other state-of-the-art TE systems, emphasizing that OWAN delivers more transfers *on time* than any other competing methods.

Dynamic optical paths increase the complexity of networks and capacity planning tasks because any optical fiber may need to accommodate diverse and variable channels. However, this complexity is rewarded with robustness or tolerance to fiber link outages. Gossels *et al.* [175] propose dynamic optical paths to make long-haul networks more robust and resilient to node and link failures by presenting algorithms for allocating bandwidth on optical paths dynamically in a mesh network. Their objective is to protect networks from any single node or link failure event. To this end, they present an optimization framework for network planners, which determines where to deploy transponders to minimize costs while running a network over dynamic optical paths.

Another effort in reducing the complexity of dynamic optical path WAN systems was presented by Dukic *et al.* [4]. Their system, *Iris*, exploits a unique property of regional connectivity, *i.e.*, the vast abundance of optical fiber in dense metropolitan areas [176]. They find that the complexity of managing dynamic optical paths is greatly reduced when switching at the fiber-strand level versus the (sub-fiber) wavelength level. To this end, they detail their design trade-off space for inter-data center connectivity across metropolitan areas. They deploy their system in a hardware testbed to emulate connectivity between three data centers, verifying that optical switching can be done in 50 to 70 ms over three amplifiers. They obviate amplifier reconfiguration delays by conducting fiber-level switching rather than wavelength-level. Thus, the amplifiers on a fiber path are configured once for the channel that traverses it. When a circuit changes its path, away from one data center and towards another, it uses a series of amplifiers that have been pre-configured to accommodate the loss of that given circuit.

Inter-data center network connectivity over a regional optical backbone was also investigated by Benzaoui *et al.* [167]. Their system, *Deterministic Dynamic Network (DDN)*, imposes strict constraints for application layer latency and jitter. They show that they can reconfigure optical links in under 2 ms, and guarantee consistent latency and jitter through their time-slotted scheduling approach.

5.5. Summary

Reconfigurable optics for metro and wide-area networks have gained substantial attention in the last decade. This push requires cross-domain collaboration as demand aware changes at the optical layer are influenced by physical-

layer impairments (signal-loss, chromatic dispersion, noise, *etc.*), in addition to higher-layer performance metrics (latency, demand, congestion, *etc.*). There are various novel works that have addressed several fundamental questions in reconfigurable optical networks. Cost-modeling efforts predict network performance with various classes of reconfigurable hardware. Algorithmic work suggests efficient methods for efficiently managing network layer and optical layer elements in the face of shifting traffic demands. Researchers have proposed and prototyped several systems for reconfigurable optical networks in recent years, but much of this work is still in the design and proof-of-concept phase. All in all, there are still many open challenges ahead to widely deploy and efficiently utilize reconfigurable optics in production networks, as we discuss next.

6. Open Challenges in Reconfigurable Optical Networks

Hardware technologies. The development of hardware for reconfigurable optical networking is a burgeoning field in engineering and research. While CDC-F ROADMs exist today, they are costly to produce, and their capabilities are found lacking. In particular, the benefit of integrating CDC-F ROADMs with optical transport networks is limited by cascading fiber impairments, signal loss at WSS modules, and wavelength and fiber collision [177]. We expect silicon photonics to bring down the cost of transport hardware, thereby increasing access to such devices and lowering entry barriers for research and development.

Data center networks. Our understanding of algorithms and topologies in reconfigurable networks is still early, but first insights into efficient designs are being published. One front where much more research is required concerns the modeling (and dealing with) reconfiguration costs. Indeed, existing works differ significantly in their assumptions, even for the same technology, making it challenging to compare algorithms. Related to this is also the question of how reconfigurations affect other layers in the networking stack, and how to design (distributed) controllers. In terms of algorithms, even though a majority of problems are intractable to solve optimally, due to integral connection constraints, the question of approximation guarantees is mostly open. For example, consider designing a data center with minimum average weighted path length. A logarithmic approximation is easy to achieve by simply minimizing the diameter of a (constant-degree) static topology. However, computing an optimal solution is NP-hard. So, can we obtain polynomial approximation algorithms with constant performance trade-offs? Similarly, do good (fixed) parameter characterizations enable efficient run times, and what can we expect from *e.g.*, linear time and distributed algorithms? Moreover, beyond general settings, how do specific (oblivious) network designs enable better algorithms, and how does their design interplay with topologies of the same equipment cost?

Next, going beyond scheduling, how can the framework of online algorithms be leveraged in this context? Ideally, we want a reconfigurable link to exist *before* the traffic appears. How can we balance this from a worst-case perspective? In this context, traffic prediction techniques might reduce the possible solution

space massively, but we will still need extremely rapid reaction times to new traffic information.

Another open challenge is the efficient interplay between reconfigurable and non-reconfigurable network parts. Theory for specific reconfigurable topologies (*e.g.*, traffic matrix scheduling for a single optical switch) has seen much progress. However, more general settings, particularly non-segregated routing onto both network parts, are still an open issue, beyond an abstract view of the combination with a single packet switch.

Metro and Wide-area Networks. Metro and wide-area optical networks are rich with open challenges. The works presented in this section highlight significant developments that have been made towards reconfigurable WAN systems and illuminate great benefits for such systems. However, programmability, cross-layer information sharing, and physical properties of light still must be solved. On the programmability front, efforts such as OpenConfig [126], OpenROADM [178], and ONOS [179] are working to provide white-box system stacks for optical layer equipment. If these are widely adopted and standardized, this will open the door for agile and efficient use of wide-area networks for a variety of applications (*e.g.*, new tools to combat DDoS [180]). Other challenges include wrangling with the physical constraints of efficient and rapidly reconfigurable WANs, for example, coordination of power adjustments across amplifiers for long-haul circuits.

7. Conclusion and Future Work

Reconfigurable optical networks are a young technology, and much of their potential and limitations are not well-understood today. In this paper, we have specifically considered data center and wide-area networks. Still, many other networks may benefit from similar technologies, and even in our context, the tradeoffs between costs and benefits (*e.g.*, in terms of resilience, performance, efficiency) are not well understood. In particular, these tradeoffs also depend on the specific technology, *e.g.*, on the reconfiguration time, as well as the traffic pattern; for example, demand-aware reconfigurable networks may only be useful if the traffic pattern exhibits temporal and spatial structure [181]. We currently specifically lack models for reconfiguration costs, and these costs, in turn, depend on the control plane, which is another open research challenge. It is not clear whether decentralized control planes are always superior to centralized ones, or whether hybrid designs are required. It is also not clear how to optimally design such control planes. From an algorithmic point of view, reconfigurable optical networks present a mostly uncharted complexity landscape. Whereas classic networking problems can largely rely on decades of optimization and graph theory, reconfiguration adds new and different twists to networking problems.

We hence hope that our survey can help to put the new concepts, technologies and challenges of reconfigurable optical networks into perspective and hence help researchers to bootstrap and contribute to this emerging field.

Funding

European Research Council (ERC) (864228); National Science Foundation (NSF) (1850297); UO Faculty Research Award; Ripple Faculty Fellowship.

References

References

- [1] A. S. Thyagaturu, A. Mercian, M. P. McGarry, M. Reisslein, W. Kellerer, Software defined optical networks (sdons): A comprehensive survey, *IEEE Communications Surveys & Tutorials* (2016).
- [2] B. C. Chatterjee, N. Sarma, E. Oki, Routing and spectrum allocation in elastic optical networks: A tutorial, *IEEE Communications Surveys Tutorials* 17 (2015) 1776–1800.
- [3] D. M. Marom, P. D. Colbourne, A. D’errico, N. K. Fontaine, Y. Ikuma, R. Proietti, L. Zong, J. M. Rivas-Moscoso, I. Tomkos, Survey of photonic switching architectures and technologies in support of spatially and spectrally flexible optical networking, *IEEE/OSA Journal of Optical Communications and Networking* 9 (2017) 1–26.
- [4] V. Dukic, G. Khanna, C. Gkantsidis, T. Karagiannis, F. Parmigiani, A. Singla, M. Filer, J. L. Cox, A. Ptasznik, N. Harland, W. Saunders, C. Belady, Beyond the mega-data center: networking multi-data center regions, in: *SIGCOMM*, ACM, 2020, pp. 765–781.
- [5] H. Ballani, P. Costa, R. Behrendt, D. Cletheroe, I. Haller, K. Jozwik, F. Karinou, S. Lange, K. Shi, B. Thomsen, H. Williams, Sirius: A flat datacenter network with nanosecond optical switching, in: *SIGCOMM*, ACM, 2020, pp. 782–797.
- [6] N. J. Van Eck, L. Waltman, Citation-based clustering of publications using citnetexplorer and vosviewer, *Scientometrics* 111 (2017) 1053–1070.
- [7] K.-T. Foerster, S. Schmid, Survey of reconfigurable data center networks: Enablers, algorithms, complexity, *SIGACT News* 50 (2019) 62–79. doi:10.1145/3351452.3351464.
- [8] A. Celik, B. Shihada, M.-S. Alouini, Optical wireless data center networks: potentials, limitations, and prospects, in: *Broadband Access Communication Technologies XIII*, volume 10945, International Society for Optics and Photonics, 2019, p. 109450I.
- [9] Y. Lu, H. Gu, Flexible and scalable optical interconnects for data centers: Trends and challenges, *IEEE Communications Magazine* 57 (2019) 27–33. doi:10.1109/MCOM.001.1900326.

- [10] A. Willner, *Optical fiber telecommunications*, volume 11, Academic Press, 2019.
- [11] R. A. Butt, M. W. Ashraf, M. Faheem, S. M. Idrus, A survey of dynamic bandwidth assignment schemes for tdm-based passive optical network, *Journal of Optical Communications* 41 (2020) 279–293.
- [12] I. A. Alimi, A. L. Teixeira, P. P. Monteiro, Toward an efficient c-ran optical fronthaul for the future networks: A tutorial on technologies, requirements, challenges, and solutions, *IEEE Communications Surveys Tutorials* 20 (2018) 708–769. doi:10.1109/COMST.2017.2773462.
- [13] M. P. McGarry, M. Reisslein, M. Maier, Ethernet passive optical network architectures and dynamic bandwidth allocation algorithms, *IEEE Communications Surveys Tutorials* 10 (2008) 46–60. doi:10.1109/COMST.2008.4625804.
- [14] J. A. Bannister, L. Fratta, M. Gerla, Topological design of the wavelength-division optical network, in: *IEEE INFOCOM'90*, IEEE Computer Society, 1990, pp. 1005–1006.
- [15] W. M. Mellette, R. McGuinness, A. Roy, A. Forencich, G. Papen, A. C. Snoeren, G. Porter, RotorNet: A Scalable, Low-complexity, Optical Datacenter Network, in: *SIGCOMM*, 2017.
- [16] S. B. Venkatakrisnan, M. Alizadeh, P. Viswanath, Costly circuits, sub-modular schedules and approximate carathéodory theorems, in: *SIGMETRICS*, ACM, 2016, pp. 75–88.
- [17] M. Wang, Y. Cui, S. Xiao, X. Wang, D. Yang, K. Chen, J. Zhu, Neural network meets DCN: traffic-driven topology adaptation with deep learning, *POMACS* 2 (2018) 26:1–26:25.
- [18] S. Lange, A. Raja, K. Shi, M. Karpov, R. Behrendt, D. Cletheroe, I. Haller, F. Karinou, X. Fu, J. Liu, A. . Lukashchuk, B. Thomsen, K. Jozwik, P. Costa, T. J. Kippenberg, H. Ballani, Sub-nanosecond optical switching using chip-based soliton microcombs, in: *Optical Fiber Communication Conference (OFC'20)*, The Optical Society (OSA), 2020.
- [19] M. Ghobadi, R. Mahajan, A. Phanishayee, P.-A. Blanche, H. Rastegarfar, M. Glick, D. Kilper, Design of Mirror Assembly for an Agile Reconfigurable Data Center Interconnect, Technical Report MSR-TR-2016-33, 2016.
- [20] A. Shakeri, M. Garrich, A. Bravalheri, D. Careglio, J. Solé-Pareta, A. Fumagalli, Traffic allocation strategies in wss-based dynamic optical networks, *Journal of Optical Communications and Networking* 9 (2017) B112–B123.

- [21] G. Wang, D. G. Andersen, M. Kaminsky, K. Papagiannaki, T. E. Ng, M. Kozuch, M. Ryan, *c-through: Part-time optics in data centers*, in: *Proceedings of the ACM SIGCOMM 2010 conference*, 2010, pp. 327–338.
- [22] W. Xia, Y. Wen, C. H. Foh, D. Niyato, H. Xie, *A survey on software-defined networking*, *IEEE Communications Surveys & Tutorials* 17 (2014) 27–51.
- [23] D. Kilper, K. Bergman, *TURBO: Terabits/s using reconfigurable bandwidth optics (final report)*, 2020. URL: <https://www.osti.gov/biblio/1618041>.
- [24] B. Lantz, A. A. Díaz-Montiel, J. Yu, C. Rios, M. Ruffini, D. Kilper, *Demonstration of software-defined packet-optical network emulation with mininet-optical and onos*, in: *Optical Fiber Communication Conference (OFC) 2020*, Optical Society of America, 2020, p. M3Z.9.
- [25] B. C. Chatterjee, S. Ba, E. Oki, *Fragmentation problems and management approaches in elastic optical networks: a survey*, *IEEE Communications Surveys & Tutorials* 20 (2017) 183–210.
- [26] C. Kachris, I. Tomkos, *A survey on optical interconnects for data centers*, *IEEE Communications Surveys & Tutorials* 14 (2012) 1021–1036.
- [27] J. M. Taboada, J. J. Maki, S. Tang, L. Sun, D. An, X. Lu, R. T. Chen, *Thermo-optically tuned cascaded polymer waveguide taps*, *Applied physics letters* 75 (1999) 163–165.
- [28] A. Yeniay, R. Gao, K. Takayama, R. Gao, A. F. Garito, *Ultra-low-loss polymer waveguides*, *Journal of lightwave technology* 22 (2004) 154–158.
- [29] D. De Felipe, M. Kleinert, C. Zawadzki, A. Polatynski, G. Irmscher, W. Brinker, M. Moehrl, H.-G. Bach, N. Keil, M. Schell, *Recent developments in polymer-based photonic components for disruptive capacity upgrade in data centers*, *Journal of Lightwave Technology* 35 (2016) 683–689.
- [30] K. A. Clark, D. Cletheroe, T. Gerard, I. Haller, K. Jozwik, K. Shi, B. Thomsen, H. Williams, G. Zervas, H. Ballani, et al., *Synchronous subnanosecond clock and data recovery for optically switched data centres using clock phase caching*, *Nature Electronics* 3 (2020) 426–433.
- [31] S. Cheung, T. Su, K. Okamoto, S. J. B. Yoo, *Ultra-compact silicon photonic 512×512 25 ghz arrayed waveguide grating router*, *IEEE Journal of Selected Topics in Quantum Electronics* 20 (2014) 310–316. doi:10.1109/JSTQE.2013.2295879.
- [32] H. Toshiyoshi, H. Fujita, *Electrostatic micro torsion mirrors for an optical switch matrix*, *Journal of Microelectromechanical systems* 5 (1996) 231–237.

- [33] J. Tsai, M. C. Wu, A high port-count wavelength-selective switch using a large scan-angle, high fill-factor, two-axis mems scanner array, *IEEE Photonics Technology Letters* 18 (2006) 1439–1441. doi:10.1109/LPT.2006.877235.
- [34] N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, A. Vahdat, Helios: a hybrid electrical/optical switch architecture for modular data centers, in: *Proceedings of the ACM SIGCOMM 2010 conference*, 2010, pp. 339–350.
- [35] G. Baxter, S. Frisken, D. Abakoumov, H. Zhou, I. Clarke, A. Bartos, S. Poole, Highly programmable wavelength selective switch based on liquid crystal on silicon switching elements, in: *2006 Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference*, IEEE, 2006, pp. 3–pp.
- [36] M. Wang, L. Zong, L. Mao, A. Marquez, Y. Ye, H. Zhao, F. J. Vaquero Caballero, Lcos slm study and its application in wavelength selective switch, in: *Photonics*, volume 4, Multidisciplinary Digital Publishing Institute, 2017, p. 22.
- [37] H. Yang, B. Robertson, P. Wilkinson, D. Chu, Small phase pattern 2d beam steering and a single lcos design of 40 1 × 12 stacked wavelength selective switches, *Opt. Express* 24 (2016) 12240–12253. URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-24-11-12240>. doi:10.1364/OE.24.012240.
- [38] H. Chen, N. K. Fontaine, R. Ryf, D. T. Neilson, Lcos-based photonic crossconnect, in: *Optical Fiber Communication Conference (OFC) 2019*, Optical Society of America, 2019, p. Th1E.6. URL: <http://www.osapublishing.org/abstract.cfm?URI=OFC-2019-Th1E.6>. doi:10.1364/OFC.2019.Th1E.6.
- [39] Z. Zhang, Z. You, D. Chu, Fundamentals of phase-only liquid crystal on silicon (lcos) devices, *Light: Science & Applications* 3 (2014) e213.
- [40] M. Filer, J. Gaudette, M. Ghobadi, R. Mahajan, T. Issenhuth, B. Klinkers, J. Cox, Elastic optical networking in the microsoft cloud, *IEEE/OSA Journal of Optical Communications and Networking* 8 (2016) A45–A54.
- [41] M. Jinno, B. Kozicki, H. Takara, A. Watanabe, Y. Sone, T. Tanaka, A. Hirano, Distance-adaptive spectrum resource allocation in spectrum-sliced elastic optical path network [topics in optical communications], *IEEE Communications Magazine* 48 (2010) 138–145. doi:10.1109/MCOM.2010.5534599.
- [42] N. Sambo, P. Castoldi, A. D’Errico, E. Riccardi, A. Pagano, M. S. Moreolo, J. M. Fabrega, D. Rafique, A. Napoli, S. Frigerio, et al., Next generation sliceable bandwidth variable transponders, *IEEE Communications Magazine* 53 (2015) 163–171.

- [43] N. Hatori, T. Shimizu, M. Okano, M. Ishizaka, T. Yamamoto, Y. Urino, M. Mori, T. Nakamura, Y. Arakawa, A hybrid integrated light source on a silicon platform using a trident spot-size converter, *Journal of Lightwave Technology* 32 (2014) 1329–1336. doi:10.1109/JLT.2014.2304305.
- [44] M. R. Billah, M. Blaicher, T. Hoose, P.-I. Dietrich, P. Marin-Palomo, N. Lindenmann, A. Nestic, A. Hofmann, U. Troppenz, M. Moehrle, S. Randel, W. Freude, C. Koos, Hybrid integration of silicon photonics circuits and inp lasers by photonic wire bonding, *Optica* 5 (2018) 876–883. URL: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-5-7-876>. doi:10.1364/OPTICA.5.000876.
- [45] T. Matsumoto, T. Kurahashi, R. Konoike, K. Tanizawa, K. Suzuki, A. Uetake, K. Takabayashi, K. Ikeda, H. Kawashima, S. Akiyama, S. Sekiguchi, In-line optical amplification for silicon photonics platform by flip-chip bonded inp-soas, in: *2018 Optical Fiber Communications Conference and Exposition (OFC)*, 2018, pp. 1–3.
- [46] K. Williams, X. Liu, M. Matters-Kammerer, A. Meighan, M. Spiegelberg, J. van der Tol, M. Trajkovic, M. Wale, W. Yao, X. Zhang, Indium phosphide photonic circuits on silicon electronics, in: *Optical Fiber Communication Conference (OFC) 2020*, Optical Society of America, 2020, p. M3A.1. URL: <http://www.osapublishing.org/abstract.cfm?URI=OFC-2020-M3A.1>. doi:10.1364/OFC.2020.M3A.1.
- [47] D. Thomson, A. Zilkie, J. E. Bowers, T. Komljenovic, G. T. Reed, L. Vivien, D. Marris-Morini, E. Cassan, L. Viro, J.-M. Fédéli, et al., Roadmap on silicon photonics, *Journal of Optics* 18 (2016) 073003.
- [48] A. S. Hamza, J. S. Deogun, D. R. Alexander, Classification framework for free space optical communication links and systems, *IEEE Commun. Surv. Tutorials* 21 (2019) 1346–1382.
- [49] N. H. Azimi, Z. A. Qazi, H. Gupta, V. Sekar, S. R. Das, J. P. Longtin, H. Shah, A. Tanwer, Firefly: a reconfigurable wireless data center fabric using free-space optics, in: *SIGCOMM*, ACM, 2014, pp. 319–330.
- [50] M. Ghobadi, R. Mahajan, A. Phanishayee, N. R. Devanur, J. Kulkarni, G. Ranade, P. Blanche, H. Rastegarfar, M. Glick, D. C. Kilper, Projector: Agile reconfigurable data center interconnect, in: *SIGCOMM*, ACM, 2016.
- [51] A. S. Hamza, S. Yadav, S. Ketan, J. S. Deogun, D. R. Alexander, Owcell: Optical wireless cellular data center network architecture, in: *ICC*, IEEE, 2017, pp. 1–6.
- [52] Y. Cui, S. Xiao, X. Wang, Z. Yang, S. Yan, C. Zhu, X. Li, N. Ge, Diamond: Nesting the data center network with wireless rings in 3-d space, *IEEE/ACM Trans. Netw.* 26 (2018) 145–160.

- [53] S. Bloom, E. Korevaar, J. Schuster, H. Willebrand, Understanding the performance of free-space optics, *Journal of optical Networking* 2 (2003) 178–200.
- [54] A. Trichili, M. A. Cox, B. S. Ooi, M.-S. Alouini, Roadmap to free space optics, *J. Opt. Soc. Am. B* 37 (2020) A184–A201.
- [55] C. Terzi, I. Korpeoglu, 60 ghz wireless data center networks: A survey, *Computer Networks* 185 (2021) 107730.
- [56] A. S. Hamza, J. S. Deogun, D. R. Alexander, Wireless communication in data centers: A survey, *IEEE Commun. Surv. Tutorials* 18 (2016) 1572–1595.
- [57] N. Farrington, G. Porter, Y. Fainman, G. Papen, A. Vahdat, Hunting mice with microsecond circuit switches, in: *HotNets*, ACM, 2012, pp. 115–120.
- [58] G. Porter, R. Strong, N. Farrington, A. Forencich, P. Chen-Sun, T. Rosing, Y. Fainman, G. Papen, A. Vahdat, Integrating microsecond circuit switching into the data center, in: *SIGCOMM’13*, 2013, pp. 447–458.
- [59] N. Farrington, A. Forencich, G. Porter, P. . Sun, J. E. Ford, Y. Fainman, G. C. Papen, A. Vahdat, A multiport microsecond optical circuit switch for data center networking, *IEEE Photonics Technology Letters* 25 (2013) 1589–1592. doi:10.1109/LPT.2013.2270462.
- [60] M. K. Mukerjee, C. Canel, W. Wang, D. Kim, S. Seshan, A. C. Snoeren, Adapting tcp for reconfigurable datacenter networks, in: *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, 2020, pp. 651–666.
- [61] D. Alistarh, H. Ballani, P. Costa, A. Funnell, J. Benjamin, P. M. Watts, B. Thomsen, A high-radix, low-latency optical switch for data centers, *Computer Communication Review* 45 (2015) 367–368.
- [62] W. M. Mellette, G. M. Schuster, G. Porter, G. Papen, J. E. Ford, A scalable, partially configurable optical switch for data center networks, *Journal of Lightwave Technology* 35 (2017) 136–144. doi:10.1109/JLT.2016.2636025.
- [63] W. M. Mellette, R. Das, Y. Guo, R. McGuinness, A. C. Snoeren, G. Porter, Expanding across time to deliver bandwidth efficiency and low latency, in: *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, 2020, pp. 1–18.
- [64] S. Kassing, A. Valadarsk, G. Shahaf, M. Schapira, A. Singla, Beyond fat-trees without antennae, mirrors, and disco-balls, in: *SIGCOMM*, 2017.

- [65] J. Shin, E. G. Sirer, H. Weatherspoon, D. Kirovski, On the feasibility of completely wireless datacenters, *IEEE/ACM Trans. Netw.* 21 (2013) 1666–1679.
- [66] G. Wang, D. G. Andersen, M. Kaminsky, M. Kozuch, T. S. E. Ng, K. Pagiannaki, M. Glick, L. B. Mummert, Your data center is a router: The case for reconfigurable optical circuit switched paths, in: *HotNets*, ACM SIGCOMM, 2009.
- [67] J. Edmonds, Paths, trees and flowers, *Canad. J. Math* 17 (1965) 449–467.
- [68] N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, A. Vahdat, Helios: a hybrid electrical/optical switch architecture for modular data centers, in: *SIGCOMM*, ACM, 2010, pp. 339–350.
- [69] W. Dai, K.-T. Foerster, D. Fuchssteiner, S. Schmid, Load-optimization in reconfigurable networks: Algorithms and complexity of flow routing, in: *Proc. 38th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (PERFORMANCE)*, 2020.
- [70] M. Müller-Hannemann, A. Schwartz, Implementing weighted b-matching algorithms: Insights from a computational study, *ACM Journal of Experimental Algorithmics* 5 (2000) 8.
- [71] A. Singla, A. Singh, K. Ramachandran, L. Xu, Y. Zhang, Proteus: a topology malleable data center network, in: *HotNets*, ACM, 2010.
- [72] K. Chen, A. Singla, A. Singh, K. Ramachandran, L. Xu, Y. Zhang, X. Wen, Y. Chen, OSA: an optical switching architecture for data center networks with unprecedented flexibility, *IEEE/ACM Trans. Netw.* 22 (2014) 498–511.
- [73] M. Bienkowski, D. Fuchssteiner, J. Marcinkowski, S. Schmid, Online dynamic b-matching with applications to reconfigurable datacenter networks, in: *Proc. 38th International Symposium on Computer Performance, Modeling, Measurements and Evaluation (PERFORMANCE)*, 2020.
- [74] K.-T. Foerster, M. Ghobadi, S. Schmid, Characterizing the algorithmic complexity of reconfigurable data center architectures, in: *ANCS*, IEEE/ACM, 2018.
- [75] K.-T. Foerster, M. Pacut, S. Schmid, On the complexity of non-segregated routing in reconfigurable data center architectures, *ACM SIGCOMM Computer Communication Review (CCR)* (2019).
- [76] T. Fenz, K.-T. Foerster, S. Schmid, A. Villedieu, Efficient non-segregated routing for reconfigurable demand-aware networks, in: *18th IFIP Networking Conference (IFIP Networking)*, 2019. URL: <http://eprints.cs.univie.ac.at/6007/>.

- [77] M. Y. Teh, Y.-H. Hung, G. Michelogiannakis, S. Yan, M. Glick, J. Shalf, K. Bergman, Tago: Rethinking routing design in high performance reconfigurable networks, in: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '20, IEEE Press, 2020.
- [78] L. G. Valiant, A scheme for fast parallel communication, *SIAM J. Comput.* 11 (1982) 350–361.
- [79] C. Chang, D. Lee, Y. Jou, Load balanced birkhoff-von neumann switches, part I: one-stage buffering, *Comput. Commun.* 25 (2002) 611–622.
- [80] A. Singla, C. Hong, L. Popa, P. B. Godfrey, Jellyfish: Networking data centers randomly, in: NSDI, 2012. URL: <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/singla>.
- [81] Y. Xia, X. S. Sun, S. Dzinamarira, D. Wu, X. S. Huang, T. S. Eugene Ng, A tale of two topologies: Exploring convertible data center network architectures with flat-tree, in: SIGCOMM, ACM, 2017.
- [82] G. Porter, R. D. Strong, N. Farrington, A. Forencich, P. Sun, T. Rosing, Y. Fainman, G. Papen, A. Vahdat, Integrating microsecond circuit switching into the data center, in: SIGCOMM, ACM, 2013, pp. 447–458.
- [83] A. Goel, M. Kapralov, S. Khanna, Perfect matchings in $o(n \log n)$ time in regular bipartite graphs, *SIAM J. Comput.* 42 (2013) 1392–1404.
- [84] H. Liu, M. K. Mukerjee, C. Li, N. Feltman, G. Papen, S. Savage, S. Seshan, G. M. Voelker, D. G. Andersen, M. Kaminsky, G. Porter, A. C. Snoeren, Scheduling techniques for hybrid circuit/packet networks, in: CoNEXT, ACM, 2015, pp. 41:1–41:13.
- [85] S. B. Venkatakrisnan, M. Alizadeh, P. Viswanath, Costly circuits, submodular schedules and approximate carathéodory theorems, *Queueing Syst.* 88 (2018) 311–347.
- [86] H. Gupta, M. Curran, C. Zhan, Near-optimal multihop scheduling in general circuit-switched networks, in: CoNEXT, ACM, 2020, pp. 31–45.
- [87] R. Schwartz, M. Singh, S. Yazdanbod, Online and offline greedy algorithms for routing with switching costs, arXiv preprint arXiv:1905.02800 (2019).
- [88] C. Avin, K. Mondal, S. Schmid, Demand-aware network designs of bounded degree, in: DISC, 2017.
- [89] C. Avin, K. Mondal, S. Schmid, Demand-aware network design with minimal congestion and route lengths, in: Proc. IEEE INFOCOM, 2019.

- [90] C. Avin, S. Schmid, Toward demand-aware networking: A theory for self-adjusting networks, in: ACM SIGCOMM Computer Communication Review (CCR), 2018.
- [91] C. Avin, A. Hercules, A. Loukas, S. Schmid, rdan: Toward robust demand-aware network designs, in: Information Processing Letters (IPL), 2018.
- [92] S. Anand, N. Garg, A. Kumar, Resource augmentation for weighted flow-time explained by dual fitting, in: Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, SIAM, 2012, pp. 1228–1241.
- [93] M. Dinitz, B. Moseley, Scheduling for weighted flow and completion times in reconfigurable networks, arXiv preprint arXiv:2001.07784 (2020).
- [94] J. Kulkarni, S. Schmid, P. Schmidt, Scheduling opportunistic links in two-tiered reconfigurable datacenters, arXiv preprint arXiv:2010.07920 (2020).
- [95] S. Schmid, C. Avin, C. Scheideler, M. Borokhovich, B. Haeupler, Z. Lotker, Splaynet: Towards locally self-adjusting networks, IEEE/ACM Trans. Netw. 24 (2016) 1421–1433.
- [96] B. Peres, O. Goussevskaia, S. Schmid, C. Avin, Concurrent self-adjusting distributed tree networks, in: Proc. International Symposium on Distributed Computing (DISC), 2017.
- [97] C. Avin, B. Haeupler, Z. Lotker, C. Scheideler, S. Schmid, Locally self-adjusting tree networks, in: Proc. 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2013.
- [98] B. Peres, O. A. de Oliveira Souza, O. Goussevskaia, C. Avin, S. Schmid, Distributed self-adjusting tree networks, in: INFOCOM, IEEE, 2019.
- [99] C. Avin, S. Schmid, Renets: Statically-optimal demand-aware networks, in: Proc. SIAM Symposium on Algorithmic Principles of Computer Systems (APOCS), 2021.
- [100] C. Avin, K. Mondal, S. Schmid, Dynamically optimal self-adjusting single-source tree networks, in: Proc. Latin American Theoretical Informatics Symposium (LATIN), 2020.
- [101] C. Avin, I. Salem, S. Schmid, Working set theorems for routing in self-adjusting skip list networks, in: Proc. IEEE INFOCOM, 2020.
- [102] S. Salman, C. Streiffer, H. Chen, T. Benson, A. Kadav, Deepconf: Automating data center network topologies management with machine learning, in: Proceedings of the 2018 Workshop on Network Meets AI & ML, NetAI'18, ACM, New York, NY, USA, 2018, pp. 8–14.

- [103] P. Kalmbach, J. Zerwas, P. Babarczy, A. Blenk, W. Kellerer, S. Schmid, Empowering self-driving networks, in: Proc. ACM SIGCOMM 2018 Workshop on Self-Driving Networks (SDN), 2018.
- [104] T. Truong-Huu, P. M. Mohan, M. Gurusamy, Virtual network embedding in ring optical data centers using markov chain probability model, *IEEE Transactions on Network and Service Management* 16 (2019) 1724–1738.
- [105] L. Chen, K. Chen, Z. Zhu, M. Yu, G. Porter, C. Qiao, S. Zhong, Enabling wide-spread communications on optical fabric with megaswitch, in: 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17), USENIX Association, Boston, MA, 2017, pp. 577–593. URL: <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/chen>.
- [106] J. Misra, D. Gries, A constructive proof of vizing’s theorem, *Inf. Process. Lett.* 41 (1992) 131–133.
- [107] K. Chen, X. Wen, X. Ma, Y. Chen, Y. Xia, C. Hu, Q. Dong, Y. Liu, Toward a scalable, fault-tolerant, high-performance optical data center architecture, *IEEE/ACM Transactions on Networking* 25 (2017) 2281–2294.
- [108] H. Wang, Y. Xia, K. Bergman, T. S. E. Ng, S. Sahu, K. Sripanidkulchai, Rethinking the physical layer of data center networks of the next decade: using optics to enable efficient *-cast connectivity, *Computer Communication Review* 43 (2013) 52–58.
- [109] X. S. Huang, X. S. Sun, T. S. E. Ng, Sunflow: Efficient optical circuit scheduling for coflows, in: CoNEXT, ACM, 2016, pp. 297–311.
- [110] H. Wang, X. Yu, H. Xu, J. Fan, C. Qiao, L. Huang, Integrating coflow and circuit scheduling for optical networks, *IEEE Transactions on Parallel and Distributed Systems* (2019). doi:10.1109/TPDS.2018.2889251.
- [111] J. Bao, D. Dong, B. Zhao, Z. Luo, C. Wu, Z. Gong, Flycast: Free-space optics accelerating multicast communications in physical layer, *Computer Communication Review* 45 (2015) 97–98.
- [112] X. S. Sun, T. S. E. Ng, When creek meets river: Exploiting high-bandwidth circuit switch in scheduling multicast data, in: ICNP, IEEE Computer Society, 2017, pp. 1–6.
- [113] X. S. Sun, Y. Xia, S. Dzinamarira, X. S. Huang, D. Wu, T. S. E. Ng, Republic: Data multicast meets hybrid rack-level interconnections in data center, in: ICNP, IEEE Computer Society, 2018, pp. 77–87.
- [114] Y. Xia, T. S. E. Ng, X. S. Sun, Blast: Accelerating high-performance data analytics applications by optical multicast, in: INFOCOM, IEEE, 2015, pp. 1930–1938.

- [115] L. Luo, K. Foerster, S. Schmid, H. Yu, Splitcast: Optimizing multicast flows in reconfigurable datacenter networks, in: INFOCOM, IEEE, 2020.
- [116] Mingwei Yang, H. Rastegarfar, I. B. Djordjevic, Physical-layer adaptive resource allocation in software-defined data center networks, *IEEE/OSA Journal of Optical Communications and Networking* 10 (2018) 1015–1026.
- [117] M. Y. Teh, S. Zhao, P. Cao, K. Bergman, COUDER: robust topology engineering for optical circuit switched data center networks, *CoRR abs/2010.00090* (2020).
- [118] M. Y. Teh, S. Zhao, K. Bergman, METTEOR: robust multi-traffic topology engineering for commercial data center networks, *CoRR abs/2002.00473* (2020).
- [119] S. Gringeri, N. Bitar, T. J. Xia, Extending software defined network principles to include optical transport, *IEEE Communications Magazine* 51 (2013) 32–40. doi:10.1109/MCOM.2013.6476863.
- [120] K. Christodoulopoulos, P. Kokkinos, E. M. Varvarigos, Indirect and direct multicost algorithms for online impairment-aware rwa, *IEEE/ACM Transactions on Networking* 19 (2011) 1759–1772.
- [121] J. Oliveira, A. Caballero, E. M. aes, U. Moura, R. Borkowski, G. Curiel, A. Hirata, L. Hecker, E. Porto, D. Zibar, J. M. ao, I. T. Monroy, J. Oliveira, Demonstration of edfa cognitive gain control via gmpls for mixed modulation formats in heterogeneous optical networks, in: *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2013*, Optical Society of America, 2013, p. OW1H.2. URL: <http://www.osapublishing.org/abstract.cfm?URI=OFC-2013-OW1H.2>. doi:10.1364/OFC.2013.OW1H.2.
- [122] U. Moura, M. Garrich, H. Carvalho, M. Svolenski, A. Andrade, A. C. Cesar, J. Oliveira, E. Conforti, Cognitive methodology for optical amplifier gain adjustment in dynamic dwdm networks, *Journal of Lightwave Technology* 34 (2016) 1971–1979.
- [123] U. Moura, M. Garrich, A. C. Cesar, J. Oliveira, E. Conforti, Execution time improvement for optical amplifier cognitive methodology in dynamic wdm networks, in: *2017 SBMO/IEEE MTT-S International Microwave and Optoelectronics Conference (IMOC)*, IEEE, 2017, pp. 1–5.
- [124] J. Case, M. Fedor, M. L. Schoffstall, J. Davin, Rfc1157: Simple network management protocol (snmp) (1990).
- [125] R. Enns, M. Bjorklund, J. Schoenwaelder, A. Bierman, Network configuration protocol (netconf) (2011).

- [126] A. Shaikh, T. Hofmeister, V. Dangui, V. Vusirikala, Vendor-neutral network representations for transport sdn, in: 2016 Optical Fiber Communications Conference and Exhibition (OFC), 2016, pp. 1–3.
- [127] A. V. Lehmen, R. Doverspike, G. Clapp, D. M. Freimuth, J. Gannett, K. Kim, H. Kobrinski, E. Mavrogiorgis, J. Pastor, M. Rauch, K. K. Ramakrishnan, R. Skoog, B. Wilson, S. L. Woodward, Coronet: Testbeds, cloud computing, and lessons learned, in: Optical Fiber Communication Conference, Optical Society of America, 2014, p. W4B.1. URL: <http://www.osapublishing.org/abstract.cfm?URI=OFC-2014-W4B.1>. doi:10.1364/OFC.2014.W4B.1.
- [128] R. A. Skoog, A. L. Neidhardt, A fast, robust signaling protocol for enabling highly dynamic optical networks, in: 2009 Conference on Optical Fiber Communication-incudes post deadline papers, IEEE, 2009, pp. 1–3.
- [129] H. Zang, J. P. Jue, B. Mukherjee, et al., A review of routing and wavelength assignment approaches for wavelength-routed optical wdm networks, *Optical networks magazine* 1 (2000) 47–60.
- [130] R. Durairajan, P. Barford, J. Sommers, W. Willinger, Intertubes: A study of the US long-haul fiber-optic infrastructure, in: SIGCOMM, ACM, 2015, pp. 565–578.
- [131] C.-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu, R. Alimi, K. N. B., C. Bhagat, S. Jain, J. Kaimal, S. Liang, K. Mendelev, S. Padgett, F. Rabe, S. Ray, M. Tewari, M. Tierney, M. Zahn, J. Zolla, J. Ong, A. Vahdat, B4 and after: Managing hierarchy, partitioning, and asymmetry for availability and scale in google’s software-defined wan, in: Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM ’18, Association for Computing Machinery, New York, NY, USA, 2018, p. 74–87.
- [132] I. Tsirilakis, C. Mas, I. Tomkos, Cost comparison of ip/wdm vs. ip/otn for european backbone networks, in: Proceedings of 2005 7th International Conference Transparent Optical Networks, 2005., volume 2, IEEE, 2005, pp. 46–49.
- [133] O. Gerstel, C. Filsfil, T. Telkamp, M. Gunkel, M. Horneffer, V. Lopez, A. Mayoral, Multi-layer capacity planning for ip-optical networks, *IEEE Communications Magazine* 52 (2014) 44–51.
- [134] M. Gunkel, A. Autenrieth, M. Neugirg, J.-P. Elbers, Advanced multilayer resilience scheme with optical restoration for ip-over-dwdm core networks, in: 2012 IV International Congress on Ultra Modern Telecommunications and Control Systems, IEEE, 2012, pp. 657–662.
- [135] P. Papanikolaou, K. Christodoulopoulos, E. Varvarigos, Joint multilayer planning of survivable elastic optical networks, in: Optical Fiber Communication Conference, Optical Society of America, 2016, pp. M2K–3.

- [136] S. Kozdrowski, M. Żotkiewicz, S. Sujecki, Optimization of optical networks based on cdc-roadm technology, *Applied Sciences* 9 (2019) 399.
- [137] S. L. Woodward, M. D. Feuer, I. Kim, P. Palacharla, X. Wang, D. Bihon, Service velocity: Rapid provisioning strategies in optical roadm networks, *Journal of Optical Communications and Networking* 4 (2012) 92–98.
- [138] X. Chen, Z. Zhu, R. Proietti, S. J. B. Yoo, On incentive-driven vnf service chaining in inter-datacenter elastic optical networks: A hierarchical game-theoretic mechanism, *IEEE Transactions on Network and Service Management* 16 (2019) 1–12.
- [139] X. J. Zhang, S. Kim, S. S. Lumetta, Opportunity cost analysis for dynamic wavelength routed mesh networks, *IEEE/ACM Transactions on Networking* 19 (2011) 747–759.
- [140] A. S. Thyagaturu, A. Mercian, M. P. McGarry, M. Reisslein, W. Kellerer, Software defined optical networks (sdons): A comprehensive survey, *IEEE Commun. Surv. Tutorials* 18 (2016) 2738–2786.
- [141] R. Casellas, R. Martínez, R. Vilalta, R. M. noz, Control, management, and orchestration of optical networks: Evolution, trends, and challenges, *J. Lightwave Technol.* 36 (2018) 1390–1402. URL: <http://jlt.osa.org/abstract.cfm?URI=jlt-36-7-1390>.
- [142] M. Kodialam, T. V. Lakshman, Integrated dynamic ip and wavelength routing in ip over wdm networks, in: *IEEE INFOCOM*, volume 1, 2001, pp. 358–366.
- [143] A. Brzezinski, E. Modiano, Dynamic reconfiguration and routing algorithms for ip-over-wdm networks with stochastic traffic, *Journal of Lightwave Technology* 23 (2005) 3188.
- [144] Z. Zhong, N. Hua, M. Tornatore, J. Li, Y. Li, X. Zheng, B. Mukherjee, Provisioning short-term traffic fluctuations in elastic optical networks, *IEEE/ACM Trans. Netw.* 27 (2019) 1460–1473.
- [145] S. K. Patri, A. Autenrieth, D. Rafique, J.-P. Elbers, C. M. Machuca, Hecson: Heuristic for configuration selection in optical network planning, in: *Optical Fiber Communication Conference (OFC) 2020*, Optical Society of America, 2020, p. Th2A.32. URL: <http://www.osapublishing.org/abstract.cfm?URI=OFC-2020-Th2A.32>. doi:10.1364/OFC.2020.Th2A.32.
- [146] D. Xu, G. Li, B. Ramamurthy, A. L. Chiu, D. Wang, R. D. Doverspike, On provisioning diverse circuits in heterogeneous multi-layer optical networks, *Comput. Commun.* 36 (2013) 689–697.

- [147] C. Mas, I. Tomkos, O. K. Tonguz, Failure location algorithm for transparent optical networks, *IEEE Journal on Selected Areas in Communications* 23 (2005) 1508–1519.
- [148] J. Chen, L. Wosinska, C. M. Machuca, M. Jaeger, Cost vs. reliability performance study of fiber access network architectures, *IEEE Communications Magazine* 48 (2010) 56–65.
- [149] L. Luo, H. Yu, K.-T. Foerster, M. Noormohammadpour, S. Schmid, Interdatacenter bulk transfers: Trends and challenges, *IEEE Network to appear* (2020).
- [150] X. Jin, Y. Li, D. Wei, S. Li, J. Gao, L. Xu, G. Li, W. Xu, J. Rexford, Optimizing bulk transfers with software-defined optical WAN, in: *SIGCOMM*, 2016. URL: <http://doi.acm.org/10.1145/2934872.2934904>. doi:10.1145/2934872.2934904.
- [151] X. Jin, H. H. Liu, R. Gandhi, S. Kandula, R. Mahajan, M. Zhang, J. Rexford, R. Wattenhofer, Dynamic scheduling of network updates, in: *SIGCOMM*, ACM, 2014, pp. 539–550.
- [152] L. Luo, K. Foerster, S. Schmid, H. Yu, Deadline-aware multicast transfers in software-defined optical wide-area networks, *IEEE Journal on Selected Areas in Communications* (2020).
- [153] S. Jia, X. Jin, G. Ghasemiefteh, J. Ding, J. Gao, Competitive analysis for online scheduling in software-defined optical wan, in: *INFOCOM*, 2017.
- [154] M. Dinitz, B. Moseley, Scheduling for weighted flow and completion times in reconfigurable networks, in: *INFOCOM*, 2020.
- [155] R. Singh, M. Ghobadi, K. Foerster, M. Filer, P. Gill, Run, walk, crawl: Towards dynamic link capacities, in: *HotNets*, ACM, 2017.
- [156] R. Singh, M. Ghobadi, K.-T. Foerster, M. Filer, P. Gill, Radwan: rate adaptive wide area network, in: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, ACM, 2018, pp. 547–560.
- [157] K. Foerster, S. Schmid, S. Vissicchio, Survey of consistent software-defined network updates, *IEEE Communications Surveys and Tutorials* 21 (2019) 1435–1461.
- [158] S. Tseng, Perseverance-aware traffic engineering in rate-adaptive networks with reconfiguration delay, in: *ICNP*, IEEE, 2019, pp. 1–10.
- [159] K.-T. Foerster, L. Luo, M. Ghobadi, Optflow: A flow-based abstraction for programmable topologies, in: *SOSR*, ACM, 2020.

- [160] I. Chlamtac, A. Ganz, G. Karmi, Lightpath communications: an approach to high bandwidth optical wan's, *IEEE Transactions on Communications* 40 (1992) 1171–1182. doi:10.1109/26.153361.
- [161] J. Guo, Z. Zhu, When deep learning meets inter-datacenter optical network management: Advantages and vulnerabilities, *J. Lightwave Technol.* 36 (2018) 4761–4773. URL: <http://jlt.osa.org/abstract.cfm?URI=jlt-36-20-4761>. doi:10.1364/JLT.36.004761.
- [162] Y. Wang, Z. McNulty, H. Nguyen, Network virtualization in spectrum sliced elastic optical path networks, *Journal of Lightwave Technology* 35 (2017) 1962–1970.
- [163] P. Soto, P. Maya, J. F. Botero, Resource allocation over eon-based infrastructures in a network virtualization environment, *IEEE Transactions on Network and Service Management* 16 (2019) 13–26.
- [164] A. E. Ozdaglar, D. P. Bertsekas, Routing and wavelength assignment in optical networks, *IEEE/ACM Transactions on Networking* 11 (2003) 259–272.
- [165] P. Kokkinos, P. Soumplis, E. A. Varvarigos, Pattern-driven resource allocation in optical networks, *IEEE Transactions on Network and Service Management* 16 (2019) 489–504.
- [166] R. Singh, M. Ghobadi, K.-T. Foerster, M. Filer, P. Gill, Radwan: Rate adaptive wide area network, in: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '18*, ACM, New York, NY, USA, 2018, pp. 547–560. URL: <http://doi.acm.org/10.1145/3230543.3230570>. doi:10.1145/3230543.3230570.
- [167] N. Benzaoui, M. S. Gonzalez, J. M. Estarán, H. Mardoyan, W. Lautenschlaeger, U. Gebhard, L. Dembeck, S. Bigo, Y. Pointurier, Deterministic dynamic networks (ddn), *Journal of Lightwave Technology* 37 (2019) 3465–3474.
- [168] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, R. Wattenhofer, Achieving high utilization with software-driven wan, *ACM SIGCOMM Computer Communication Review* 43 (2013) 15–26.
- [169] S. Figueira, S. Naiksatam, H. Cohen, D. Cutrell, P. Daspit, D. Gutierrez, D. B. Hoang, T. Lavian, J. Mambretti, S. Merrill, et al., Dwdm-ram: Enabling grid services with dynamic optical networks, in: *IEEE International Symposium on Cluster Computing and the Grid, 2004. CCGrid 2004.*, IEEE, 2004, pp. 707–714.
- [170] E. Bernier, M. Vukovic, D. Goodwill, P. Daspit, G. Wang, Omninet: A metropolitan 10 gb/s dwdm photonic switched network trial, in: *Optical Fiber Communication Conference*, Optical Society of America, 2004, p. WH4.

- [171] P. Iovanna, R. Sabella, M. Settembre, A traffic engineering system for multilayer networks based on the gmpls paradigm, *IEEE Network* 17 (2003) 28–37.
- [172] M. Chamania, M. Caria, A. Jukan, Achieving ip routing stability with optical bypass, *Optical Switching and Networking* 7 (2010) 173–184. URL: <http://www.sciencedirect.com/science/article/pii/S1573427710000305>. doi:<https://doi.org/10.1016/j.osn.2010.05.005>.
- [173] O. Turkcu, S. Subramaniam, Performance of optical networks with limited reconfigurability, *IEEE/ACM Trans. Netw.* 17 (2009) 2002–2013. URL: <http://dx.doi.org/10.1109/TNET.2009.2014158>. doi:10.1109/TNET.2009.2014158.
- [174] X. Jin, Y. Li, D. Wei, S. Li, J. Gao, L. Xu, G. Li, W. Xu, J. Rexford, Optimizing bulk transfers with software-defined optical wan, in: *Proceedings of the 2016 ACM SIGCOMM Conference*, ACM, 2016, pp. 87–100.
- [175] J. Gossels, G. Choudhury, J. Rexford, Robust network design for ip/optical backbones, *IEEE/OSA Journal of Optical Communications and Networking* 11 (2019) 478–490.
- [176] S. K. Mani, M. Nance Hall, R. Durairajan, P. Barford, Characteristics of metro fiber deployments in the us, in: *Proceedings of the Network Traffic Measurement and Analysis Conference*, 2020.
- [177] H. Li, H.-y. Zhang, L. Wang, Y.-b. Li, J.-s. Lai, R. Tang, W.-y. Zhao, B.-b. Wu, D. Wang, X. Zhao, et al., Field trial of network survivability based on otn and roadm hybrid networking, in: *Asia Communications and Photonics Conference, Optical Society of America*, 2017, pp. M3C–2.
- [178] S. Oda, M. Miyabe, S. Yoshida, T. Katagiri, Y. Aoki, J. C. Rasmussen, M. Birk, K. Tse, A learning living network for open roadm networks, in: *ECOC 2016; 42nd European Conference on Optical Communication, VDE*, 2016, pp. 1–3.
- [179] P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O’Connor, P. Radoslavov, W. Snow, et al., Onos: towards an open, distributed sdn os, in: *Proceedings of the third workshop on Hot topics in software defined networking*, 2014, pp. 1–6.
- [180] M. Nance Hall, G. Liu, R. Durairajan, V. Sekar, Fighting fire with light: Tackling extreme terabit ddos using programmable optics, in: *Proceedings of the Workshop on Secure Programmable Network Infrastructure, SPIN ’20*, Association for Computing Machinery, New York, NY, USA, 2020, p. 42–48. URL: <https://doi.org/10.1145/3405669.3405824>. doi:10.1145/3405669.3405824.
- [181] C. Avin, M. Ghobadi, C. Griner, S. Schmid, On the complexity of traffic traces and implications, *Proc. ACM Meas. Anal. Comput. Syst.* 4 (2020) 20:1–20:29.