

Broadening the scope of optical circuit networks

Malathi Veeraraghavan
University of Virginia
Charlottesville, VA 22904
Email: mv5g@virginia.edu

Inder Monga
ESnet, LBNL
Berkeley, CA 94720
Email: imonga@lbl.gov

Abstract—Advances in optical communications and switching technologies are enabling energy-efficient, flexible, higher-utilization network operations. To take full advantage of these capabilities, the scope of optical circuit networks can be increased in both the vertical and horizontal directions. In the vertical direction, some of the existing Internet applications, transport-layer protocols, and application-programming interfaces need to be redesigned and new ones invented to leverage the high-bandwidth, low-latency capabilities of optical circuit networks. In the horizontal direction, inter-domain control and management-protocols are required to create a global-scale interconnection of optical circuit-switched networks.

Keywords—Optical circuit networks; dynamic circuit services; data centers; applications

I. INTRODUCTION

The objective of this work is to provide an overall vision for broadening the scope of optical circuit networks. Our motivation for this study is driven bottom-up by the capabilities of various optical circuit networking technologies. The current use of optical circuit networks, which is to provide leased-line service, is rather limited in scope. The purpose of this paper is to examine how new dynamic circuit services can be leveraged on these networks in potentially new use cases involving datacenters and campus networks.

Section II provides background on data-, control-, and management-plane aspects of optical circuit networks. New dynamic circuit services and potential opportunities for using these services are described in Section III. Advances made in the research-and-education networking (REN) community are summarized in Section IV, and missing pieces needed to expand the scope of optical circuit networks are discussed in Section V. The paper is summarized in Section VI.

II. BACKGROUND

Background information is provided on data-, control-, and management-plane aspects of optical circuit networks, and an example of a deployed optical circuit network is presented.

A. Data plane

Optical network elements consist of amplifiers, dispersion compensation modules, transponders, multiplexers/demultiplexers including reconfigurable optical add/drop multiplexers (ROADMs), and optical crossconnects (OXC). Optical platforms available in the marketplace implement one or more of these elements. These platforms are deployed in

metro-area and wide-area core networks, and are typically used to provide point-to-point circuits between IP routers.

ROADMs and OXCs are Wavelength Division Multiplexed (WDM) circuit switches. The optical circuits are referred to as lightpaths, and are used to carry high-rate signals, e.g., 10 Gbps and higher. Electronic switches are used today to create sub-wavelength circuits. Two types of switches are in use today: Time-Division Multiplexed (TDM) and connection-oriented packet-multiplexed (also referred to virtual-circuit) switches. TDM switches based on the Optical Transport Network (OTN) hierarchy [1] are referred to as Digital cross-connects (DXCs). Virtual circuit (VC) technologies include MultiProtocol Label Switching (MPLS) Transport Profile [2], and Ethernet Virtual Local Area Network (VLAN) and Carrier Ethernet [3]. Newer technologies that support sub-wavelength circuits include spectrum-sliced elastic optical path network (SLICE) [4], and Digital SubCarrier Multiplexing (DSCM) [5].

Since Ethernet is the dominant technology used in Network Interface Cards (NICs), in campus and datacenter switches, and in provider IP routers, mapping techniques are required to transport Ethernet frames over circuits/VCs. Standards have been defined to carry Ethernet over MPLS VCs [6] or over TDM/WDM circuits [7]. Edge devices can be configured to map all Ethernet frames arriving on a port to a particular MPLS VC or TDM/WDM circuit, or tagged Ethernet frames with a particular VLAN ID can be directed to an MPLS VC or TDM/WDM circuit. These mapping techniques enable the creation of end-to-end circuits where the ends could be computers, Ethernet switches, or IP routers.

B. Control and management planes

Circuit/VC networks require a setup phase for circuit/VC provisioning prior to data transmission. Circuit setup/release can be handled using control-plane or management-plane solutions. The Path Computation Element (PCE) Communication Protocol (PCEP) [9] has been developed for determining the path for a circuit. The computed path is passed down to the ingress switch to initiate circuit provisioning using Resource reSerVation Protocol with Traffic Engineering (RSVP-TE) [10], which is part of the Generalized MPLS (GMPLS) standards.

If the frequency of circuit setup/release actions is low, centralized management-plane solutions are sufficient. Most switch vendors offer network management systems to execute

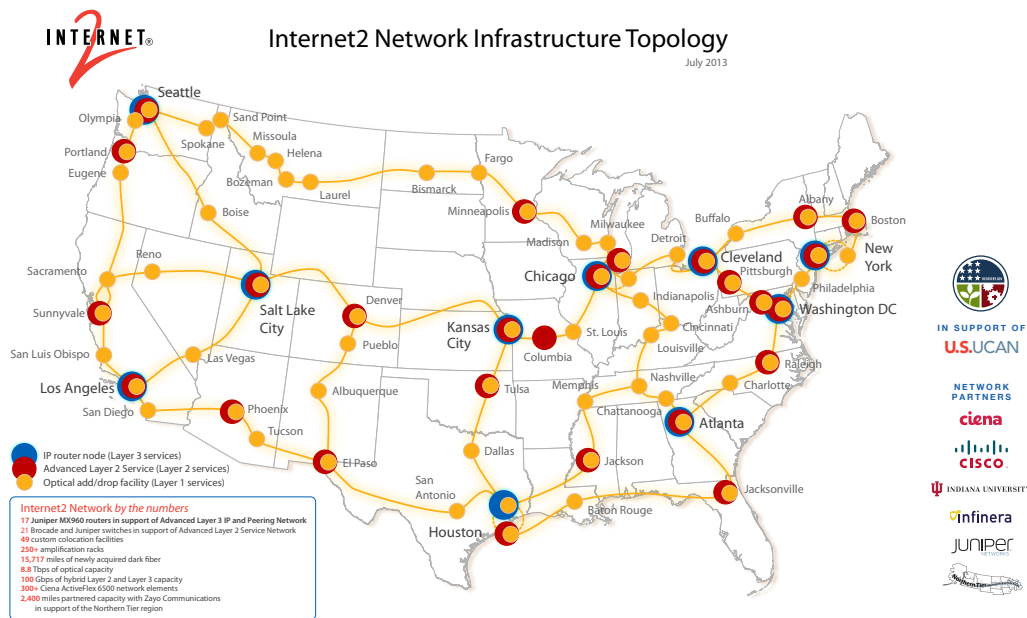


Figure 1: Internet2’s optical circuit (Layer 1), Ethernet-switched (Layer 2), and IP-routed (Layer 3) networks [8]

circuit provisioning actions, which falls under the “configuration” category of the umbrella term Fault, Configuration, Accounting, Performance and Security (FCAPS) used to characterize the management plane. Typically, circuit provisioning for leased-line service is executed using these centralized network management systems.

Control-plane protocols to support advance reservation of circuits/VCS are under development in the Open Grid Forum, referred to as Network Services Interfaces (NSI) connection services [11]. An earlier protocol, Inter-Domain Controller Protocol (IDCP) [12], has been implemented and is in use in RENs. Centralized servers (one per domain¹) handle all advance-reservation circuit requests, and use IDCP or NSIv2 for inter-domain server-to-server communications. Just before the scheduled start time of a circuit, each centralized server initiates intra-domain MPLS VC provisioning, which occurs in a hop-by-hop manner using the built-in RSVP-TE software in the MPLS switches.

More recently, Software Defined Networking (SDN) and OpenFlow² are being leveraged to develop centralized solutions for circuit/VC provisioning.

In summary, the RSVP-TE signaling protocol is part of the control-plane software implemented in switches to enable distributed circuit provisioning. All three types of systems, i.e., network management systems, advance-reservation circuit schedulers, and PCE engines, are typically implemented in servers external to the switches. Similarly, SDN controllers, run on external servers, communicate via the OpenFlow pro-

ocol, which is implemented as part of switch software.

C. Example deployment

As an example of today’s use of optical circuit networks, we show a US-wide core optical circuit network and describe how it is used. Fig. 1 shows Internet2’s deployed optical circuit network, which consists of over 60 ROADMs (Layer-1)³, while in contrast there are only 9 IP routers and 26 Ethernet switches. Static point-to-point optical circuits are configured through this network of ROADMs to interconnect IP routers (Layer 3) and Ethernet switches (Layer 2). These inter-router/inter-switch circuits are intra-domain, but inter-domain circuits are also used, e.g., a Univ. of Memphis IP router connects to the Chicago Internet2 IP router via an optical circuit passing through six ROADMs from Memphis to Chicago (backhaul access link). ESnet⁴ runs an optical circuit network on the same fiber footprint as Internet2, and uses this network for creating static circuits between its IP routers.

A traceroute executed using Internet2’s routerproxy Web site⁵ from the Internet2 Seattle router to Univ. of California (UC) at Davis (in northern California) showed the path as traversing from Internet2’s Seattle router to Internet2’s Los Angeles router (in southern California), which has a direct peering with a CENIC (regional REN in California) router. From the CENIC Los Angeles router, the path travels back up north to reach a CENIC router in the Bay Area that is peered with UC Davis’ campus router. This example illustrates that even though path latencies are longer in this multi-layer

¹The term “domain” is used here to represent an autonomous system, i.e., a network under one organization’s ownership.

²<https://www.opennetworking.org/>

³Some of the 300+ Ciena network elements listed in Fig. 1 are regeneration points without add/drop capability.

⁴<http://www.es.net>

⁵<http://routerproxy.grnoc.iu.edu/internet2/>




Figure 2: Communication services possible with optical circuit networks

topology, fewer IP routers can be deployed when compared to ROADMs to limit costs.

III. OPPORTUNITIES

Given the type of large-scale optical-circuit network deployment described above, we consider the question of what opportunities, if any, exist for broadening the scope and value of these networks. First, we describe how two relatively new types of services can be supported on optical-circuit networks, and then consider the use cases enabled by these services.

A. Services

Optical circuit networks are used today primarily to support leased-line and Virtual Private Network (VPN) services. Fig. 2 shows two other communication services that could potentially be offered on a large-scale (inter-domain) basis on optical circuit networks: Scheduled Dynamic Circuit Service (SDCS) and Unscheduled Dynamic Circuit Service (UDCS).

First we consider how leased-line/VPN service differs from SDCS and UDCS, and then we describe differences between SDCS and UDCS [13]. With leased-line/VPN service, customers purchase a single contract for point-to-point circuits/VCs or a multipoint VPN with specified endpoints, rate, and duration, which is typically on the order of months to years. In contrast, SDCS and UDCS require each customer to purchase an access link on a long-term (e.g., two years) contract to connect its endpoint to the service provider’s switch, and then customers can request circuits/VCs with specified rates to any other SDCS/UDCS customer for short-term usage.

Now we consider how UDCS differs from SDCS. UDCS is the same as Plain Old Telephone Service (POTS) except that UDCS supports flexible circuit rates unlike POTS in which all circuits are 64 kbps. UDCS is effectively a queueing based service while SDCS is a reservation based service (e.g., people stand in queues to get on to city buses, while they make reservations for flights). In a queueing based service, there is no scheduler. Instead jobs/customers request immediate service upon arrival but may be required to wait in a queue if there are other jobs/customers ahead of them. In POTS however, there is no buffer to hold calls because a circuit requires channel allocations on multiple links of the end-to-end path, which means that if a channel is assigned to a call on one link while the call is waiting for channel allocations at other links the assigned channel is being wasted. Therefore, POTS is set up to be a call blocking system where calls are blocked immediately if no channels are available on any single link




Figure 3: Possible use cases for optical circuit networks

of the end-to-end path. Effectively POTS, and by extension UDCS, are bufferless queueing systems. In contrast, SDCS is a reservation system, i.e., a scheduler is present to accept and process advance reservation requests in which endpoints, circuit rate, call duration, and starting times are specified.

UDCS can be operated at high utilization and low call blocking probability only if the number of shared channels per link is high. For example, if the number of shared channels per link is as low as 10 (e.g., only 10 wavelengths or sub-wavelength channels are multiplexed on a link), and the desired utilization is 80%, associated call blocking probability is 23%. One way to reduce call blocking probability when the number of channels is small is to deploy SDCS because the advance-reservation scheduler can then use the time dimension as a waiting bay for calls instead of blocking them [14].

B. Use cases

Fig. 3 classifies potential use cases. The most commonly deployed use cases are the ones in which IP routers are the endpoints of circuits/VCs as described in the Internet2 example. The notation “leased line service” is added for these use cases in Fig. 3 because even though there has been much interest in using dynamic circuits to interconnect IP routers [15], there has been limited deployment success potentially for the following reasons. Other than in response to failures, triggers for setting up a new circuit and for modifying or releasing an existing circuit have been difficult to determine. Also any modifications to the optical circuits underlying the IP-network topology will impact IP routing protocols leading to potential instabilities, and will also influence TCP congestion control mechanisms for ongoing flows.

On the other hand, if the circuit/VC endpoints are computers as shown in Fig. 3, leased-line service is not feasible because of scalability. There are far more computers than routers and therefore any use of circuits/VCs terminating on computers requires the use of SDCS and/or UDCS. The mapping technologies to carry Ethernet frames in circuits/VCs, as described in Section I, make it feasible to have circuits/VCs terminate on the Ethernet Network Interface Cards (NICs) of computers.

Data centers are of increasing importance in the Internet. In addition to cloud computing, a number of Internet applications rely on data-center servers. These include Web/Content Delivery Network (CDN), email, document sharing, and social networking services. Both intra- and inter-datacenter applications are of interest as potential users of optical circuit

networks as shown in Fig. 3. Hybrid packet-switched/optical-circuit networks such as Helios [16] have been proposed for data centers. Recent work [17] shows how a tight integration of Hadoop scheduling and circuit setup/release enables the use of optical circuit networks within datacenters.

Inter-datacenter communications can be intra- or inter-domain as shown in Fig. 3. A recent example of an intra-domain inter-datacenter use case was reported by Google [18]. By rate limiting applications running on the datacenter computers in combination with traffic engineering to direct flows to specific paths, close to 100% utilization was reported for the inter-datacenter network. Effectively, this example shows that SDCS can be used to create circuits/VCs between computers in distant datacenters for certain applications such as bulk-data transfer. Setting up end-to-end circuits between data-center servers owned and operated by the same provider is much simpler than using circuits between campus computers with multiple domains in between for administrative reasons.

Fig. 3 lists both intra- and inter-domain use cases when the computers are located in campuses. Applications being developed to use the Dynamic Network System (DYNES) equipment⁶ are primarily inter-domain in nature, but later we propose a deployment strategy that starts with intra-domain applications.

IV. ADVANCES MADE BY THE REN COMMUNITY

Several advances have been made in developing dynamic circuit services, both SDCS and UDCS, in the REN community with support of US agencies such as the National Science Foundation (NSF) and the Department of Energy (DOE). These advances are described below.

A. SCDS

Through the DOE-funded On-Demand Secure Circuits and Advance Reservation System (OSCARs) project [19], and the NSF-funded Dynamic Resource Allocation via GMPLS Optical Networks (DRAGON) project [20], control-plane software was implemented and deployed in ESnet, Internet2, and in several universities and regional RENs through the NSF DYNES project. Another control-plane software called Open Exchange Software Suite (OESS)⁷ has been deployed on Internet2 and DYNES. Using OSCARs or OESS software clients or a Web GUI, users/applications can make advance-reservation requests for circuits by specifying the endpoints, rate, duration, and start time. If the request is granted, the OSCARs/OESS controller will initiate circuit provisioning just before the scheduled start time by communicating with the switches. Currently VLAN and MPLS switches are supported with this software. OESS currently supports only intra-domain circuits, while OSCARs supports both intra- and inter-domain circuits. Both OSCARs and OESS controllers support the OpenFlow 1.0 protocol, and therefore can be used to configure a variety of switches. The IDCP protocol is also implemented in AutoBAHN [21], a GEANT2 circuit service. ESnet has

experimented with optical SDN [22] and multilayer provisioning⁸. Finally, several scheduling algorithms have been proposed for advance reservations [23].

B. UDCS

In the NSF-funded Circuit-Switched High-speed End-to-End Transport architecture (CHEETAH) project [24], on a wide-area Ethernet-SONET network using off-the-shelf switches with built-in GMPLS signaling protocols, the use of UDCS was demonstrated by integrating an RSVP-TE client [20] with a Web caching application [25] running on top of Circuit TCP (CTCP) [26].

In another NSF-funded project, we demonstrated an implementation of RSVP-TE in FPGA hardware [27]. A deployment of switches with this hardware-accelerated signaling engine will reduce circuit setup delay to the round-trip propagation delay since call processing delays will be in μ s. Even if all-optical switch reconfiguration takes ms, it is worthwhile reducing call processing delays. For example, in the off-the-shelf switch used in CHEETAH, the time for processing an RSVP-TE Path message was 91 ms [24].

V. MISSING PIECES

This section identifies components, protocols, applications and strategies that need further development in order to advance the scope of optical circuit networks.

A. Network components and protocols

In the control plane, while protocols have been defined for inter-domain circuit provisioning as described above, there is no community-wide effort to define an inter-domain routing protocol equivalent to BGP. While there is much interest in an integrated control plane for IP and optical circuit services, the type of information sufficient for IP inter-domain routing, i.e., address reachability, may not be sufficient for SDCS, where path information may be required for future start times [28].

On the management plane, support is needed for the FCAPS functions. For example, there are no simple diagnostic tools such as traceroute to isolate the source of a problem on circuits/VCs. While OSCARs and OESS control-plane software offer users the option of reserving a backup circuit, circuits/VCs are viewed as being less robust when compared to IP-routed service.

With regards to API, since most applications on computers use TCP/IP sockets, even if the end-to-end path is a provisioned circuit/VC, TCP/UDP over IP datagrams are carried within Ethernet frames. While the use of sockets simplifies application programming, it adds a task of IP address configuration for the circuits/VCs at the end hosts.

B. Applications

A concerted effort is required to develop new applications suited for circuits and/or to modify existing applications to leverage circuit characteristics. These applications should be

⁶<https://spaces.internet2.edu/display/dynedoc/DYNES+switch>

⁷<http://globalnoc.iu.edu/sdn/oess.html>

⁸<http://www.sdncentral.com/event/brocade-infinera-esnet-sdn-demo/>

integrated with clients that communicate with the SDCS controllers or UDCS signaling clients as needed.

Bulk-data transfers: There have been many demonstrations of this application over circuits [18], [25], [26], including the use of Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) over wide-area Ethernet VLANs [29]. Since RoCE implements the InfiniBand transport- and network-layer protocols in the interface card hardware, CPU utilization to maintain multi-Gbps transfer rates is significantly lower when compared to TCP/IP.

Reliable multicast of data files: In the scientific community, there is a need for pushing data to multiple recipients. For example, the University Corporation of Atmospheric Research (UCAR) collects and disseminates real-time meteorology data to over 240 institutions. Arguing that problems with IP multicast (such as complexity of multicast routing protocols, malicious users without credentials joining multicast groups, and congestion control) could be overcome with VCs, we developed a reliable multicast transport protocol for virtual circuits called VCMTP [30]. Since users join and leave multicast groups, dynamic adjustments of VCs will be required. As optical-layer multicast becomes available, VCMTP can be adopted to run over multicast optical circuits.

Parallel job scheduling within datacenters: Scientific applications such as Climate simulations require thousands of cores for acceptable execution times, and communications is increasingly becoming the bottleneck. Since job schedulers such as Portable Batch Scheduler are used in cluster computing systems, circuits/VCs can be set up between the computers assigned to a job. Circuits/VCs can reduce communication latencies/jitter and improve program performance.

Scientific computational workflow scheduling across datacenters: As an example, consider a large scale workflow system, PANDA⁹, used in the Large Hadron Collider (LHC) ATLAS project. The distributed compute resources and data associated with the computations are centrally managed. Jobs submitted to the central manager are allocated compute resources that are matched with dataset locations. If the number of jobs that require certain datasets exceeds a threshold, the PANDA system automatically initiates the creation of a copy of the required datasets to another location. This system manages 10s of thousands of jobs a day that include movement of hundreds of terabytes of data/day. For effective workflow scheduling, task completion times need to be predictable. For predictable transfer delays, optical circuit networks are more suitable than IP networks. Therefore, OSCARS/OESS clients could be integrated into PANDA software to leverage SDCS as part of workflow management.

Intra- and inter-datacenter VM migration: VMs are migrated live from one physical machine to another with small service interruptions, if any, in order to optimize one or more cost functions [31]. Since the amount of data moved can be large, optical circuits are well suited for this application [32].

Parallel and wide-area file systems: There is interest in

leaving the data in its current place of storage and using file systems while running computations on another server instead of moving large datasets. Filesystem protocols such as NFS are notoriously chatty requiring many back-and-forth messages. For both high-performance parallel file systems and wide-area file systems, rate-guaranteed circuits could yield better performance.

Interactive applications: Connectionless (IP) networking has an advantage over circuit networking for interactive applications such as VoIP, gaming, and Web browsing because unlike with circuits, bandwidth is not wasted during silence periods. However, for interactive TCP applications such as Web browsing, TCP's Slow Start results in multiple round-trip times (RTTs) being incurred for small transfers. On large-RTT paths, delays can be significant. If circuit setup delay can be reduced to round-trip propagation delay (e.g., with hardware signaling engines as described in Section IV-B), and the number of channels per link is large, then one can envision using circuits for interactive applications. Once a circuit is setup, a Web page with all its images bundled in could be streamed at the circuit rate instead of requiring multiple GET request-response messages for image downloads. A 100 KB page transmitted on a 10 Mbps circuit requires 80ms, which could be lower than multiple back-and-forth messages sent across a wide-area path with large propagation delays. Remote instrument control, tele-surgery, network-controlled robots, audio/video, and other delay-sensitive applications could be implemented to take advantage of circuits.

C. Deployment strategy

The current REN strategy for deploying SDCS on campus networks has been to leverage the already deployed campus packet switches/routers. However, the number of channels per link is limited if VLAN based VCs are used because of policing and packet scheduling limitations of most packet switches. It could be potentially easier to implement circuit multiplexers when compared to virtual circuits for supporting large numbers of channels per link. To enable the use of circuits by interactive applications, optical circuit switches with support for thousands of channels/link and hardware signaling should be considered. Furthermore, intra-datacenter and intra-campus applications should be developed first to allow for independent adoption by autonomous systems. Inter-domain usage will follow when the number of campus/datacenter deployments reaches a critical mass.

VI. SUMMARY

Optical circuit networks are used primarily for leased-line service to create static circuits between IP routers. To broaden the scope of optical circuit networks, Scheduled and Unscheduled Dynamic Circuit Services (SDCS and UDCS), are proposed for use cases in which circuit/virtual circuit (VC) endpoints are computers, not IP routers. Computers could be located in datacenters or on campuses. Intra- and inter-datacenter applications and intra- and inter-campus applications are identified for circuits/VCs. Management-plane

⁹<https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA>

solutions for fault management and performance monitoring are required, along with inter-domain routing protocols for a horizontal extension of optical circuit networks. For vertical extension, new applications, transport protocols and APIs are required to take advantage of the properties of circuits.

VII. ACKNOWLEDGMENT

The University of Virginia work was supported by the U.S. DOE grant DE-SC0007341, and NSF grants CNS-1116081, OCI-1127340, ACI-1340910, and CNS-1405171. The ESnet portion of the work was supported by the Director, Office of Science of the U.S. DOE under Contract no. DE-AC02-05CH11231.

REFERENCES

- [1] "ITU-T Recommendation G.872: Architecture of optical transport networks," International Telecommunication Union, Tech. Rep., Nov. 2001. [Online]. Available: <http://www.itu.int/rec/T-REC-G.872/en>
- [2] M. Bocci, S. Bryant, D. Frost, L. Levrau, and L. Berger, "A Framework for MPLS in Transport Networks," RFC 5921 (Informational), Internet Engineering Task Force, Jul. 2010, updated by RFC 6215. [Online]. Available: <http://www.ietf.org/rfc/rfc5921.txt>
- [3] "IEEE 802.1Q Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks," Tech. Rep., Aug 2011. [Online]. Available: <http://www.ieee802.org/1/pages/802.1Q.html>
- [4] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Mat-suoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies," *Communications Magazine, IEEE*, vol. 47, no. 11, pp. 66–73, 2009.
- [5] R. Hui, W. Huang, Y. Zhang, M. Hameed, M. Razo, M. Tacca, and A. Fumagalli, "Digital subcarrier optical networks and cross-connects," *Journal of High Speed Networks*, vol. 19, no. 1, pp. 55 – 69, 2013.
- [6] L. Martini, E. Rosen, N. El-Aawar, and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks," RFC 4448 (Proposed Standard), Internet Engineering Task Force, Apr. 2006, updated by RFC 5462. [Online]. Available: <http://www.ietf.org/rfc/rfc4448.txt>
- [7] "ITU-T G.709 Interfaces for the Optical Transport Network (OTN)," International Telecommunication Union, Tech. Rep., Mar. 2003. [Online]. Available: <http://www.itu.int/rec/T-REC-G.809/en>
- [8] Internet2. [Online]. Available: <http://www.internet2.edu/>
- [9] J. Vasseur and J. L. Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)," RFC 5440 (Proposed Standard), Internet Engineering Task Force, Mar. 2009. [Online]. Available: <http://www.ietf.org/rfc/rfc5440.txt>
- [10] D. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209 (Proposed Standard), Internet Engineering Task Force, Dec. 2001, updated by RFCs 3936, 4420, 4874, 5151, 5420, 5711. [Online]. Available: <http://www.ietf.org/rfc/rfc3209.txt>
- [11] Open Grid Forum, "Network Services Interface (NSI) v2 Connection Services," Tech. Rep. [Online]. Available: http://redmine.ogf.org/dmsf_files/12969
- [12] T. Lehman, C. Guok, A. Lake, R. Krzywiania, and M. Balkcerkiewicz, "Inter-domain Controller (IDC) Protocol Specification," Open Grid Forum (OGF) GFD-I.170, Oct 2010.
- [13] M. Veeraraghavan, M. Karol, and G. Clapp, "Optical dynamic circuit services," *Communications Magazine, IEEE*, vol. 48, no. 11, pp. 109 –117, november 2010.
- [14] X. Zhu and M. Veeraraghavan, "Analysis and Design of Book-ahead Bandwidth-Sharing Mechanisms," *IEEE Transactions on Communications*, vol. 56, no. 12, pp. 2156–2165, Dec. 2008.
- [15] M. Chamania and A. Jukan, "A comparative analysis of the effects of dynamic optical circuit provisioning on IP routing," *Networking, IEEE/ACM Transactions on*, vol. PP, no. 99, pp. 1–1, 2013.
- [16] N. Farrington, G. Porter, S. Radhakrishnan, H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat, "Helios: a hybrid electrical/optical switch architecture for modular data centers," in *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4. ACM, 2010, pp. 339–350.
- [17] G. Wang, T. E. Ng, and A. Shaikh, "Programming your network at run-time for big data applications," in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 103–108. [Online]. Available: <http://doi.acm.org/10.1145/2342441.2342462>
- [18] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hözlze, S. Stuart, and A. Vahdat, "B4: Experience with a globally-deployed software defined wan," in *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, ser. SIGCOMM '13. ACM, 2013, pp. 3–14.
- [19] On-Demand Secure Circuits and Advance Reservation System (OSCARS). [Online]. Available: <http://www.es.net/OSCARS/docs/index.html>
- [20] X. Yang, C. Tracy, J. Sobieski, and T. Lehman, "GMPLS-based dynamic provisioning and traffic engineering of high-capacity ethernet circuits in hybrid optical/packet networks," in *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, 2006, pp. 1–5.
- [21] AutoBAHN: Automated Bandwidth Allocation across Heterogeneous Networks. [Online]. Available: <http://geant3.archive.geant.net/service/autobahn/pages/home.aspx>
- [22] A. Sadasivarao, S. Syed, P. Pan, C. Liou, A. Lake, C. Guok, and I. Monga, "Open transport switch: A software defined networking architecture for transport networks," in *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 115–120. [Online]. Available: <http://doi.acm.org/10.1145/2491185.2491192>
- [23] N. Charbonneau and V. M. Vokkarane, "A survey of advance reservation routing and wavelength assignment in wavelength-routed WDM networks," *Communications Surveys Tutorials, IEEE*, vol. 14, no. 4, pp. 1037–1064, 2012.
- [24] X. Zhu, X. Zheng, and M. Veeraraghavan, "Experiences in implementing an experimental wide-area GMPLS network," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 3, pp. 82–92, 2007.
- [25] X. Fang and M. Veeraraghavan, "Internetworking circuit and connection-less networks," in *Advanced Communication Technology, 2009. ICACT 2009. 11th International Conference on*, vol. 01, 2009, pp. 63–68.
- [26] A. P. Mudambi, X. Zheng, and M. Veeraraghavan, "A transport protocol for dedicated end-to-end circuits," in *Communications, 2006. ICC '06. IEEE International Conference on*, vol. 1, 2006, pp. 18–23.
- [27] H. Wang, M. Veeraraghavan, R. Karri, and T. Li, "Design of a high-performance RSVP-TE hardware signaling accelerator," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 8, pp. 1588–1595, 2005.
- [28] M. Veeraraghavan and D. Starobinski, "A routing architecture for scheduled dynamic circuit services," in *Proceedings of the Re-Architecting the Internet Workshop*, ser. ReARCH '10. New York, NY, USA: ACM, 2010, pp. 6:1–6:6. [Online]. Available: <http://doi.acm.org/10.1145/1921233.1921241>
- [29] E. Kissel, M. Swany, B. Tierney, and E. Pouyoul, "Efficient wide area data transfer protocols for 100 Gbps networks and beyond," in *Proceedings of the Third International Workshop on Network-Aware Data Management*, ser. NDM '13. New York, NY, USA: ACM, 2013, pp. 3:1–3:10. [Online]. Available: <http://doi.acm.org/10.1145/2534695.2534699>
- [30] J. Li and M. Veeraraghavan, "A reliable message multicast transport protocol for virtual circuits," in *International Conference on Communications, Mobility, and Computing (CMC)*, May 21–23 2012.
- [31] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *Proceedings of the 2Nd Conference on Symposium on Networked Systems Design & Implementation - Volume 2*, ser. NSDI'05. Berkeley, CA, USA: USENIX Association, 2005, pp. 273–286. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1251203.1251223>
- [32] F. Travostino, P. Daspit, L. Gommans, C. Jog, C. de Laat, J. Mambretti, I. Monga, B. van Oudenaarde, S. Raghunath, and P. Y. Wang, "Seamless live migration of virtual machines over the MAN/WAN," *Future Generation Computer Systems*, vol. 22, no. 8, pp. 901 – 907, 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167739X06000483>