Evolving to an SDN-Enabled ISP Backbone: Key Technologies and Applications

Martin Birk, Gagan Choudhury, Bruce Cortez, Alvin Goddard, Narayan Padi, Aswatnarayan Raghuram, Kathy Tse, Simon Tse, Andrew Wallace, and Kang Xi

The authors focus on mission-critical large scale backbone networks and discuss the key technologies and applications for ISPs' SDN adoption. Taking an evolutionary viewpoint, they discuss strategies and approaches to integrate SDN capabilities with traditional technologies to achieve high performance traffic engineering, fast service provisioning, and rapid failure restoration required for existing network operations.

Abstract

Software-defined networking provides a promising approach to build and operate high-performance networks with high efficiency and affordable cost. While SDN is of great importance to ISPs, it is infeasible to replace the entire existing infrastructure with a clean slate of SDN design. Most early SDN practices concentrate on small-scale greenfield designs. Here, we focus on mission-critical large-scale backbone networks, and discuss the key technologies and applications for ISPs' SDN adoption. Taking an evolutionary viewpoint, we discuss strategies and approaches to integrate SDN capabilities with traditional technologies to achieve the high-performance traffic engineering, fast service provisioning, and rapid failure restoration required for existing network operations.

Introduction

The rapid growth of network services has brought great challenges to large Internet service providers (ISPs). Sophisticated traffic engineering (TE) is essential to deal with dynamic traffic changes under various network conditions and make efficient use of network resources. Fast and automated multi-layer resource provisioning is increasingly important to meet agile service demands. With extensive growth in both capacity and footprint, ISP networks also face unprecedented complexities to achieve fast but resource-efficient failure restoration. Although traditional technologies can solve some of the problems, they are insufficient or inefficient to tackle all of the challenges, especially in an agile and cost-effective manner.

Software-defined networking (SDN) is based on separating the control plane from the data plane [1, 2]. In a simplified example, an SDN network could include a set of dumb switches or common off-the-shelf (COTS) components controlled by a centralized controller. Built with low-cost commodities chips, such switches support high-speed packet forwarding but have limited local control capability in the sense that not all classic routing protocols are fully implemented. The centralized controller takes a global view of the network to make intelligent decisions on routing, TE, service provisioning, failure restoration, and so on. It then enforces the decisions by configuring and directing the switches accordingly. SDN creates opportunities to reduce capital expenses by using low-cost commodity switches and improving resource utilization. It can also save operational expenses by simplifying, automating, and optimizing network operations [3].

Early SDN practice in data center networks and small-scale wide area networks (WAN) have shown good performance, low complexity, and high agility [4–6]. While these pioneering works well justify the benefits of SDN, such clean-slate designs cannot be directly applied to ISP networks where unique requirements and properties create nontrivial challenges. Built and upgraded over many years, ISP networks typically include a variety of equipments supporting heterogeneous protocols. Replacing the entire infrastructure with a clean slate design will inevitably disrupt certain protocols and services, let alone incur tremendous cost. Therefore, an ISP must take an evolutionary approach that keeps backward compatibility, fully utilizes the existing protocols and equipment, and progressively reaps SDN benefits. Also, many early SDN practices are used in an environment where the network has the privilege to schedule or throttle certain types of traffic when necessary. For example, many bulky data transfers in a data center or a private WAN are not delay- or bandwidth-sensitive and thus can be scheduled to avoid peak hours. Unfortunately, ISPs do not have similar privilege in that they must ensure service level agreements (SLAs) for heterogeneous customer applications. Delay-sensitive traffic must be routed along the minimum delay paths without congestion. Predetermined bandwidth must be guaranteed under both normal and failure conditions. Disruptions caused by failures must be restored within a pre-determined time interval. Such requirements and constraints not only add significant complexity to SDN adoption but also reduce the potential benefits. For example, nearly 100 percent bandwidth utilization could be achieved on many links in a private SDN network by properly scheduling and throttling upper layer applications and allowing long disruptions. In contrast, it is impossible to reach such high utilization in large ISP backbones because redundant capacity has to be

The authors are with AT&T Labs, except for Kang Xi, who was at AT&T when this work was being done.

0163-6804/16/$25.00 © 2016 IEEE

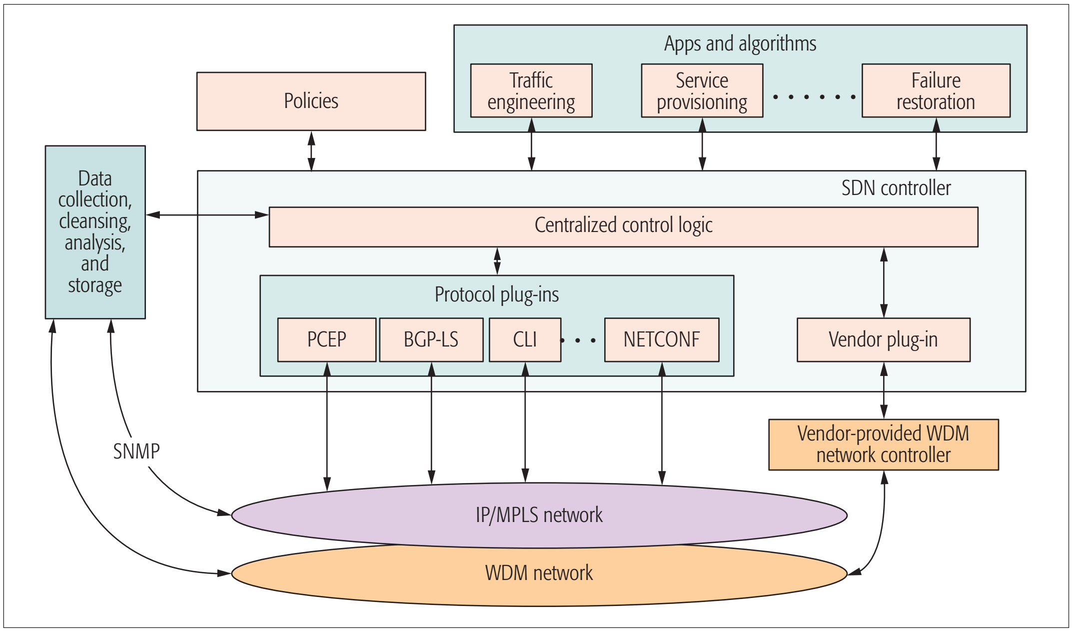

Figure 1. Schematic SDN network architecture.

The IP/MPLS routers support traditional distributed routing. They also provide open interfaces enabling centralized SDN control. MPLS still plays a critical role to ISPs in many operations, including routing, traffic engineering, priority control, service provisioning, and failure restoration.

deployed to tolerate failures and traffic changes.

AT&T embraces SDN and network functions virtualization (NFV) with a transformative initiative called Domain 2.0 [7, 8]. From an ISP's perspective, we present an implementable architecture to evolve current large-scale backbone networks with SDN technologies. We concentrate our discussions on how to leverage the architecture to enhance TE, service provisioning, and failure restoration.

Architecture and Key Technologies

Figure 1 illustrates the overall architecture of an SDN-enabled backbone network. The data plane is an IP/multiprotocol label switching (MPLS) over wavelength-division multiplexing (WDM) network. The optical layer uses reconfigurable optical add-drop multiplexers (ROADMs) that provide open interfaces. Such interfaces enable centralized control of the optical network to dynamically and remotely establish/release wavelength circuits, thus adjusting the layer 3 logical topology and capacity. Extending SDN control from layer 3 to the optical layer is important for automating and optimizing network provisioning and operation tasks, such as capacity augmentation and failure restoration. Many vendors support such control interfaces in their new products.

The IP/MPLS routers support traditional distributed routing. They also provide open interfaces enabling centralized SDN control. MPLS still plays a critical role to ISPs in many operations, including routing, TE, priority control, service provisioning, and failure restoration. Accordingly, the SDN control in this architecture works at the level of label-switched paths (LSPs). Some prior works perform fine-grained traffic control based on 5-tuple flows identified by source/destination IP addresses, protocol, and source/destination ports. While such an approach is beneficial to certain systems (e.g., data center networks), it is unnecessary and technically challenging in large-scale backbone networks.

Separate from the data plane, a global SDN controller communicates with the routers and ROADMS through various interfaces to perform monitoring and control tasks. For the optical layer, the key control capability is to establish/release wavelength circuits to adjust the layer 3 logical topology. The controller performs such operations through a vendor-provided control module, which then proceeds with lightpath computation and ROADM configuration. For the IP/MPLS layer, the key control function is to monitor traffic load and link state to perform globally optimized LSP operations, such as establishing/releasing LSPs, adjusting the paths/bandwidth of LSPs, and multi-path routing. The controller provides interfaces to support customized policies and various applications, such as centralized TE, dynamic service provisioning, and optimized failure restoration. The logically centralized controller could be implemented in a physically distributed way to achieve scalability and resilience. The specific controller design and deployment are beyond the scope of this article. Interested readers should refer to related materials, such as [4, 9, 10].

The interfaces between the control and data planes are extensible to support multiple protocols. While OpenFlow is frequently discussed in the SDN community [11], so far it is not the most critical protocol in our architecture. Figure 1 shows a few indispensable protocols to realize the applications covered in this article.

- Border Gateway Protocol-Link State (BGP-LS) reports link state to the controller, such as up/down status, total capacity, and bandwidth available to reserve for each traffic class.
- Path Computation Element Communication Protocol (PCEP) supports interactions between the controller and routers to perform LSP operations. The routers use PCEP to report LSP status to the controller, such as path and bandwidth. The controller issues PCEP commands to adjust LSP paths and bandwidth. PCEP and BGP-LS are essential to support centralized TE and fail-

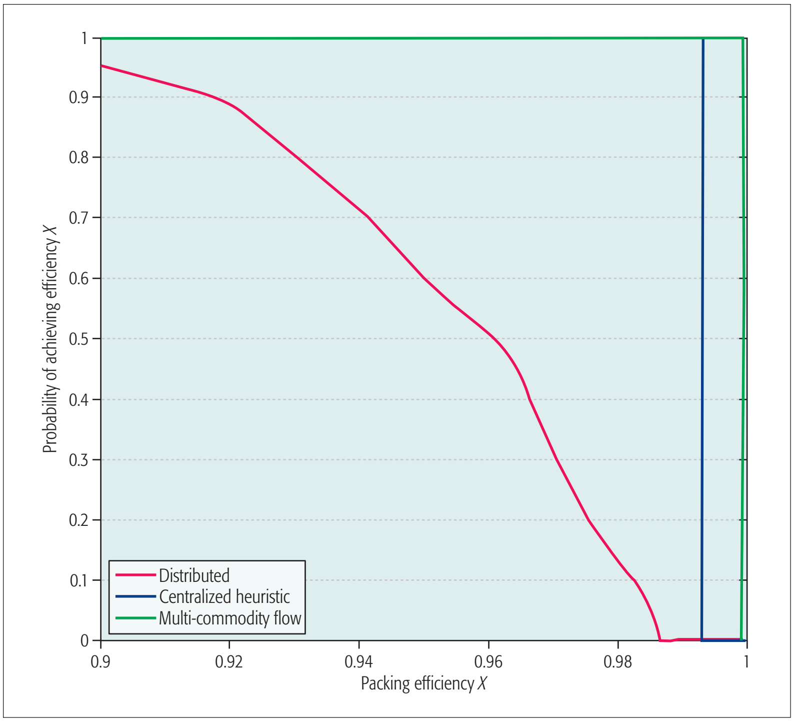

Figure 2. Distributed routing cannot guarantee good performance compared to centralized routing.

ure restoration.

•Network Configuration Protocol (NETCONF) and Command Line Interface (CLI) allow the controller to install and modify router configuration. NETCONF is generally preferred over CLI since it simplifies automation, programmatic access, and code maintenance. However, CLI is still necessary for operations not supported by NETCONF, especially vendor-specific ones.

•Simple Network Management Protocol (SNMP) reports traffic measurement on each link and LSP at minute-level frequency. The measurement data can complement BGP-LS and PCEP to achieve global optimal TE. The control plane also maintains a module for data collection, cleansing, analysis, and storage. Historical data can be retrieved to perform traffic analytics and forecast for various purposes.

Traffic Engineering

The objective of TE is to efficiently and optimally utilize network resources to route traffic and meet service requirements. In an MPLS network, the essential approach is to manipulate the routing and bandwidth allocation of each LSP to avoid congestion and meet the bandwidth and latency requirements. TE plays an important role in achieving SLAs, improving resource utilization, and reducing capital expense.

MPLS-TE using distributed constrained shortest path first (CSPF) is commonly adopted in ISP networks. In a nutshell, the head-end router of an LSP obtains the network topology and available bandwidth on each link, computes the shortest path that meets the LSP's bandwidth requirement, and then deploys the LSP. Such computation is performed by all the routers in a distributed way to adapt to traffic changes. In case of congestion, high-priority LSPs could preempt low-priority ones. Low-priority LSPs could get longer paths and/or reduced bandwidth. Most ISPs design their network topology, link capacity, and link weight settings to optimize the performance of CSPF. Therefore, MPLS-TE performs relatively well under common traffic scenarios.

However, MPLS-TE has several limitations. First, it does not guarantee near-optimal performance because each router performs greedy local optimization. Thus, the network could suffer from corner traffic scenarios where MPLS-TE fails to resolve congestion. Under heavy load, it is not uncommon to see LSPs taking extra-long paths due to this limitation. Figure 2 compares the performance between distributed and centralized routing in an example core network, where packing efficiency shows the relative performance of a routing method compared to the optimal solution. Multi-commodity flow can arbitrarily split each traffic unit and route them optimally and thereby achieve a globally optimal solution. We define the efficiency of multi-commodity flow packing as 100 percent and compare other routing methods with respect to that. Distributed routing has a high probability to achieve 90 percent efficiency, but the probability drops quickly toward higher efficiency, which is caused by the aforementioned corner cases. A centralized heuristic performs much better in that it guarantees good packet efficiency up to 99 percent. Second, MPLS-TE does not provide an easy way to fully leverage multipath routing. When an LSP is heavily loaded, an effective way is to employ tunnel splitting by establishing multiple LSPs between the same source/destination pair, optimizing multipath routing, and then distributing traffic among these LSPs. If the traffic load decreases, tunnel merging is performed by concentrating the traffic back to a single LSP and removing the other LSPs. Traditional MPLS-TE does not have built-in splitting and merging capability.

SDN overcomes the MPLS-TE limitations. The controller can exploit a global view of network status and traffic load to coordinate routing and bandwidth allocation to achieve near-optimal TE. It can also easily create, reroute, and remove LSPs through NETCONF, CLI, and PCEP interfaces to realize dynamic multipath routing. Figure 3 illustrates a case where the congestion appearing as a deadlock to distributed routing is easily resolved by centralized control. Link M2–D is 50 Gb/s, and all the other links are 100 Gb/s. Figure 3a shows the initial state where neither of the two LSPs carry heavy traffic, and the network is congestion-free. Figure 3b shows that traffic surge on LSP2 causes congestion. Figure 3c shows that simple rerouting is sufficient to resolve the congestion. However, distributed routing cannot resolve the congestion due to lack of global coordination. For LSP2, its head-end router observes only 60 Gb/s available capacity on the alternative path S2–M1–D, which is insufficient to meet LSP2's bandwidth demand of 80 Gb/s. Thus, it will keep LSP2 on the current path. Similarly, the headend router for LSP1 sees only 10 Gb/s available on alternative path S1–M2–D and thus has no way to reroute LSP1. As a result, neither of the routers is able to perform

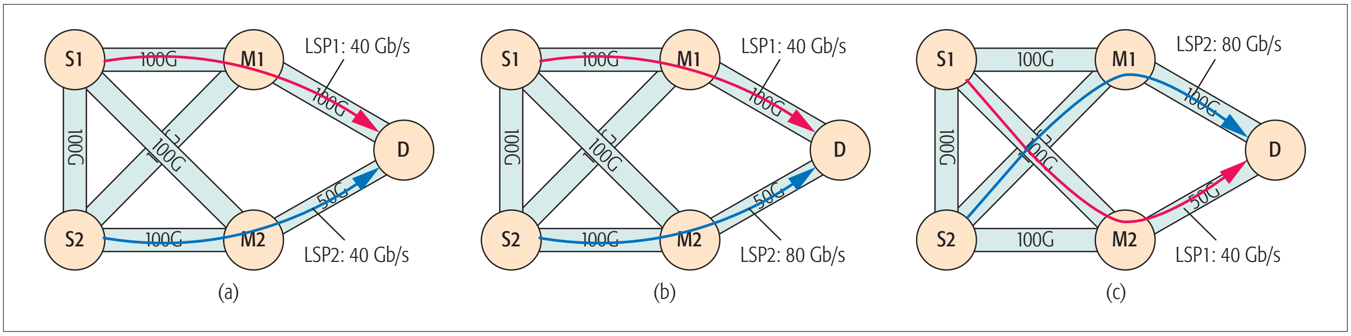

Figure 3. Resolving congestion using centralized SDN TE.

rerouting to resolve the congestion, creating a deadlock. With a global view, an SDN controller can easily find the solution in Fig. 3c and complete the transition by manipulating routing and bandwidth in a coordinated way. A similar problem is discussed in [12].

In the extreme case, the controller could administer the routing of all the LSPs and react to each traffic change by adjusting the LSPs to their optimal or near-optimal paths. However, this approach is neither simple nor necessary. First, centralized control of all the LSPs in a large nationwide backbone with multiple traffic classes is difficult to scale. Letting the controller manipulate all the LSPs in real time and react to all traffic changes would pose a great challenge on scalability. Second, optimal TE is not always necessary. Globally optimal TE incurs high computation complexity and may require frequent LSP manipulations. During off-peak periods or in the absence of failures, it is usually sufficient to keep the network congestion-free by just relying on sub-optimal distributed TE. Even if congestion develops during peak traffic periods or during failure scenarios, it is usually enough to use a fast centralized heuristic and avoid the very high computational complexity of globally optimal solutions (Fig. 2). Hence, a good solution for ISPs is to combine traditional distributed MPLS-TE and centralized SDN control in the following complementary way.

- Traditional distributed MPLS-TE continues to function and provides a baseline for further optimization. The SDN controller monitors the network status and optimizes TE only if necessary. Typically, the centralized TE algorithm is invoked either periodically or by congestion events. Instead of doing a complete routing makeover, the controller tries to adjust a small subset of LSPs to resolve congestion and reach near-optimal TE. Such adjustments include re-routing LSPs along alternative paths, limiting allocated bandwidth to certain LSPs, and splitting large LSPs to leverage multipath routing.
- The control over an LSP has dual modes. The default mode is distributed routing where the headend router controls an LSP to adjust its routing and bandwidth. Once the controller decides to perform adjustment, it sends PCEP commands to the head-end router to manipulate LSP routing and bandwidth allocation. The algorithm needs to be well designed to avoid flapping, where the path of an LSP changes frequently due to inconsistent routing decisions made by centralized and distributed control, respectively.
- When a head-end router receives an LSP manipulation command, it exploits the traditional Resource Reservation Protocol (RSVP) to notify the downstream routers to perform LSP operations. This simplifies the process because the controller does not need to deal with MPLS label assignment and configuration of the intermediate and tail-end routers.

Fast Service Provisioning

We address two types of service provisioning that can benefit from SDN. The first one is called bandwidth calendaring (BWC), which provisions on-demand LSPs according to requested bandwidth and schedules. The second one is capacity planning and bandwidth augmentation, where wavelength circuits are established or released to adjust the layer 3 topology and link capacity based on traffic analysis and forecast.

The traffic load in an ISP network is mostly uneven. High-priority traffic, such as that from enterprise services, peaks during working hours. Low-priority traffic, including lots of video streaming, peaks from evening to midnight. Overall, the network load demonstrates a statistical daily periodic pattern. This pattern creates opportunities for BWC. A typical BWC application is to schedule business bulk transfers, such as data replication, between data centers at the appropriate time to meet the bandwidth demand without causing congestion.

The operation at the IP/MPLS layer is straightforward through the NETCONF/CLI and PCEP interfaces: establish an LSP, complete data transfer, and remove the LSP. The critical challenge is to design an algorithm module on top of the controller to perform effective traffic analysis and service admission. The goal is to minimize the odds that the quality assurance of an admitted BWC service request cannot be fulfilled. For example, a non-elastic BWC request asks for an LSP with guaranteed bandwidth throughout the entire scheduled duration. An elastic request allows the allocated bandwidth to be adjusted, but the accumulated capacity through the LSP's life cycle is sufficient to complete the scheduled data transfer. For this purpose, the algorithm needs to perform thorough analysis that covers historical data analysis, traffic load modeling and forecast, network maintenance scheduling, and even failure recovery.

ISPs periodically perform capacity planning and augmentation to accommodate traffic growth. The common approach includes several steps. First, historical traffic analysis is performed

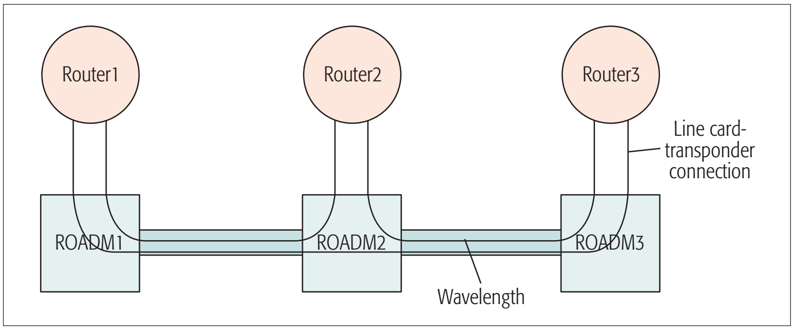

Figure 4. New ROADMs allow setting up wavelength circuits without color, direction, and contention constraints.

to forecast traffic growth. Second, based on the traffic forecast, an augmentation plan is made to add capacity to existing layer 3 links or create new layer 3 links. Finally, the augmentation plan is implemented by installing new equipment (e.g., router line cards, ROADM transponders, regenerators) and establishing new wavelength circuits.

The traditional approach has a long cycle time. The main reason is that many of the operations heavily rely on manual and non-integrated processes that are inherently lengthy. In particular, it is virtually impossible to substantially speed up the process of labor-intensive wavelength provisioning with traditional ROADMs where many constraints such as wavelength color, direction, and contention complicate installation and configuration. The traditional approach incurs high capital expense and high operational expense.

An SDN-enabled optical network helps to reduce the cycle time for service provisioning. Although manual installation is still necessary, it is less complex with colorless, directionless, and contentionless ROADMs with open interfaces. In particular, the analysis and wavelength provisioning can be automated by using the optical network control interfaces. Figure 4 shows an example where a logical ring topology is created among the three routers. With traditional ROADMs, the two line card-transponder connections between Router2 and ROADM2 are distinct in that they are in opposite directions and cannot be swapped. With new ROADMs that are colorless, directionless, and contentionless, the same two connections can be flexibly configured to connect in either direction.

By reducing cycle time and automating wavelength provisioning, SDN makes capacity management less reliant on the accuracy of traffic forecast. In fact, it has become increasingly complex for ISPs to accurately forecast traffic growth. A product release or an update from a major content provider could cause considerable traffic growth in just a few weeks. For example, shortly after a major social media company introduced a product promoting video streaming, the traffic growth accelerated considerably and became harder to forecast. Traditionally, ISPs have to add extra capacity in each augment plan to compensate for forecast inaccuracy. While SDN cannot eliminate the problems of forecast inaccuracy, it can greatly reduce cycle time (e.g., from a few months to a few weeks). With short cycle time for capacity management, it is easier to tolerate traffic forecast inaccuracy and thus avoid adding too much capacity. SDN can also leverage spare ports to dynamically provision capacity to deal with unexpected traffic surges. To prepare for failures, it is common practice to pre-deploy spare line cards, transponders, regenerators, and wavelengths. Upon an unexpected traffic surge, the SDN controller could be more efficient than current approaches in using such spare resources to quickly provision additional bandwidth for temporary usage. Unlike scheduled capacity augmentation, this process can be fully automated because all the hardware is deployed in advance.

Failure Restoration

In large-scale networks, it is virtually impossible to avoid failures caused by various factors, such as hardware faults, software bugs, power outages, misconfigurations, and human errors. Failure restoration is thus a critical consideration in the design, deployment, and operation of ISP networks to ensure reliable service provisioning. The first major goal of restoration is to minimize the duration of service disruption caused by a failure. For example, traditional synchronous optical network (SONET) automatic protection switching (APS) restores from a fiber cut within 50 ms. The second major goal is to minimize the impact of failure on the loss of bandwidth and meet the SLAs of various services. For example, a high-priority service may require minimum bandwidth guarantee even during failover. While restoration could sacrifice the latency and bandwidth of low-priority services, it is highly preferable to minimize the loss of bandwidth for such services.

In traditional IP/MPLS over WDM networks, a common approach to achieve restoration is MPLS fast reroute (FRR) followed by distributed CSPF re-convergence. FRR leverages local protection to achieve fast restoration comparable to SONET APS. To use FRR, protection LSPs are established for potential link or node failures. When a failure occurs, the adjacent router immediately directs traffic to the protection LSP to resume packet forwarding. The limitation of FRR is that it does not perform TE in case of failure. Solely relying on FRR to fully meet an SLA would incur high cost in that considerable redundant capacity has to be allocated to the protection LSPs. To compensate FRR, CSPF re-calculates better routing once the post-failure topology update is advertised throughout the network. CSPF is not as fast as FRR, and the re-convergence could take up to tens of seconds to complete. In many cases, setting up protection wavelengths at the optical layer is not adopted for failure restoration because such redundancy introduces high cost. Dynamically setting up wavelength circuits after a failure is traditionally infeasible due to considerable manual operations.

Failure restoration can benefit from SDN in several ways.

- SDN automates the planning and setting up of FRR LSPs. The protection LSPs need to be carefully planned according to the network traffic load. If an LSP traverses over a heavily loaded link, it could cause severe congestion and hurt SLA during failure restoration. Traditionally,

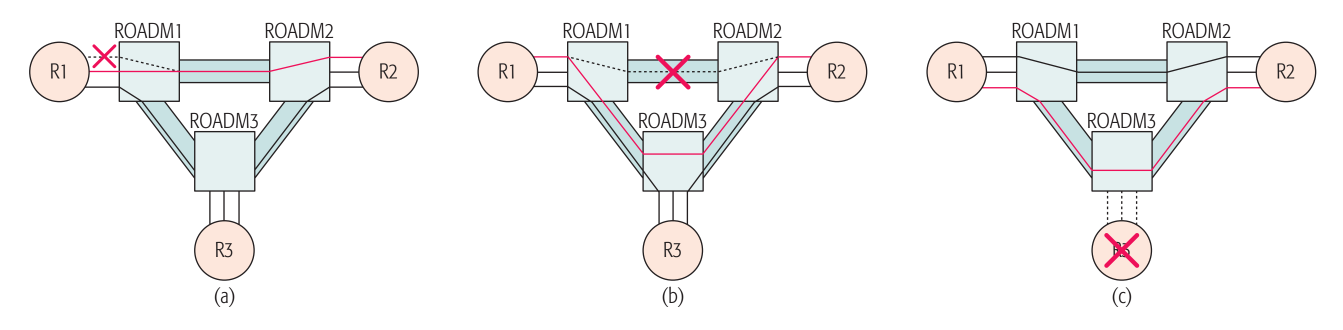

Figure 5. Optical layer restoration (dotted lines are broken lightpaths, and red lines are restored lightpaths).

FRR planning and configuration are not fully automated and cannot be dynamically updated to adapt to traffic changes. With SDN, the global controller can periodically re-optimize FRR LSP settings based on updated traffic load and near-term forecast. If any changes are needed, the controller can easily make such changes through the open interfaces.

•SDN can perform post-failure TE to meet SLA and optimize network performance. In an aggressive way, the SDN controller can replace CSPF re-convergence to perform post-failure TE. Similar to the situation in failure-free scenarios, SDN outperforms CSPF in that it always achieves near-optimal TE, which is particularly important for post-failure routing when the network is likely burdened with heavy load. In a conservative way, if the SDN controller cannot guarantee fast reaction in failure scenarios, it can wait until CSPF re-convergence completes and then re-optimize the results of distributed routing. It is hard for the SDN controller to achieve millisecond-level restoration by taking a post-failure reactive approach. FRR, with proactive configuration and local switching over, is still indispensable to rapid restoration.

•SDN automates speedy optical layer restoration without incurring extra cost. With access to the optical layer control plane, the SDN controller can dynamically provision wavelength for failure restoration. The key is to ensure that the algorithm only uses the available line cards, transponders, wavelengths, and regenerators to establish restoration wavelength circuits. Without human intervention, the restoration speed mainly depends on the physical layer configuration, such as tuning lasers, which typically takes a couple of minutes. The available resources used in this case include both standby components and components newly released from the failure. It is particularly beneficial to leverage resources released from a failure because no extra resource redundancy is required in advance. Figure 5 shows three cases to leverage optical layer restoration. In Fig. 5a, a line card failure on Router1 disconnects Router1 and Router2. The standby connection between Router1 and ROADM1 is used to reconnect the two routers. In Fig. 5b, a failed fiber disconnects Router1 and Router2. On both ends the line card-transponder connections become idle. The controller can reuse such resources to make a new wavelength along the alternative path, where an idle wavelength is available. In Fig. 5c, Router3 fails completely. The SDN controller configures ROADM3 to reuse the released wavelengths and transponders to add new capacity between Router1 and Router2, which helps to deliver traffic previously relayed by Router3.

Conclusions

SDN is vital to ISPs to improve network performance, reduce complexity, and cut cost. A critical issue is that the adoption of SDN in large-scale ISP networks cannot take a clean slate design approach. Instead, it is preferred to leverage both SDN and traditional network equipment and protocols. On one hand, this strategy takes advantage of the time-tested traditional systems to maintain stable services and fully utilize the existing equipment. On the other hand, it provides flexibility to exploit SDN to achieve optimization, automation, and new service provisioning that would not be possible otherwise. From an ISP's perspective, we discuss the architecture, protocols, and methods to optimize TE, speed up service provisioning, and improve failure restoration. With a proper evolutionary strategy, large-scale backbone networks can use SDN to transition to agile, efficient, and highly automated infrastructure. The impact of such transition will not be limited to the network itself. Instead, such a network will inspire and enable future service design and creation that dynamically integrate multiple types of resources including network connectivity, network functions, and cloud resources.

References

[1] ONF, "Software-Defined Networking: the New Norm for Networks," white paper, Apr. 2012.

[2] M. Casado *et al.*, "Ethane: Taking Control of the Enterprise," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 37, no. 4, 2007, pp. 1–12.

[3] N. Feamster, J. Rexford, and E. Zegura, "The Road to SDN: An Intellectual History of Programmable Networks," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 44, no. 2, 2014, pp. 87–98.

[4] S. Jain *et al.*, "B4: Experience with a Globally-Deployed Software Defined WAN," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 43, no. 4, 2013, pp. 3–14.

[5] A. Singh *et al.*, "Jupiter Rising: A Decade of clos Topologies and Centralized Control in Google's Datacenter Network," *Proc. 2015 ACM Conf. Special Interest Group on Data Commun.*, 2015, pp. 183–97.

[6] M. Kobayashi *et al.*, "Maturing of OpenFlow and Software-Defined Networking through Deployments," *Computer Networks,* vol. 61, no. 0, 2014, pp. 151–75.

[7] AT&T, "AT&T Domain 2.0 Vision White Paper," white paper, 2013.

[8] D. John, "Continuing Our Software Transformation with Managed Internet on Demand," Sept. 2015; http://about.att.com/innovationblog/09302015internet, accessed June 14, 2016.

[9] https://www.opendaylight.org, accessed on June 14, 2016.

[10] P. Berde *et al.*, "ONOS: Towards an Open, Distributed SDN OS," *Proc. 3rd Wksp. Hot Topics in Software Defined Networking*, 2014, pp. 1–6.

[11] https://www.opennetworking.org/sdn-resources/openflow, accessed June

14, 2016.
[12] C.-Y. Hong *et al.*, "Achieving High Utilization with Software-Driven WAN." *ACM SIGCOMM Comp. Commun. Rev.*, vol. 43, no. 4, 2013, pp. 15–26.

Biographies

Martin Birk [SM] (birk@att.com) received his Master's and doctorate degrees from Germany's University of Ulm in 1994 and 1999, respectively. Since 1999, he has been with AT&T Labs in New Jersey, working on high-speed optical transmission at data rates of 40 Gb/s, 100 Gb/s, and above.

Gagan Choudhury [F] (gchoudhury@att.com) is a lead inventive scientist at AT&T Labs, Middletown, New Jersey. He received a Ph.D. in electrical engineering from the State University of New York at Stony Brook in 1982. His research interests are in the optimization, analysis, and design of software defined networks and mobility networks. He became an AT&T Fellow in 2009 for "outstanding contributions to performance analysis and robust design and their application to improving the performance, reliability and scalability of AT&T's networks."

Bruce Cortez (bc2752@att.com) is a director member of technical staff at AT&T Labs, Middletown New Jersey. Before joining AT&T, he received a B.S. in physics from Caltech and a Ph.D. in physics from Harvard. He has worked on many projects related to restoration of service on the AT&T network, using intelligent optical switches to get sub-second restoration. He is currently leading a group that is responsible for transitioning AT&T Core Network technologies to SDN control.

Al Goddard (agoddard@att.com) is a principal member of technical staff at AT&T Labs, New Jersey. He received his B.S. and M.S. degrees, both in computer science, from the New Jersey Institute of Technology in 1979 and 1999, respectively. In 2000 he joined AT&T Labs and has pursued work in the area of automation. Currently, he is in the Domain 2.0 Architecture and Design department working on software defined networking.

Narayan Padi (np2698@att.com) is a principal member of technical staff at AT&T Labs. He received his M.C..A. in computer science in 1997 and his B.Sc in mathematics in 1993, both from Osmania University, Hyderabad, India. In 2009 he joined AT&T Labs with rich analytic experience. He has contributed in several software defined networking (SDN) modules including L3 optimization. His current contribution/focus is in prototyping multi-layer integration of layer 0 and layer 3 in the SDN realm.

Aswatnarayan Raghuram (araghuram@att.com) is a lead member of technical staff at AT&T. He received his B.E. from University Visvesvaraya College of Engineering, Bangalore, India, his M.E. from the Indian Institute of Science, Bangalore, and his Ph.D. in electrical and computer engineering from Drexel University. He joined AT&T Bell laboratories in 1988, and has worked on X.25, Frame Relay, ATM, IP, and MPLS networks. His research includes network architecture, network design/optimization, traffic engineering, and routing.

Kathleen Hartley Tse (kt1371@att.com) received a Ph.D. degree in engineering from Brown University and a B.Sc. from Cornell University. She received AT&T's Medal of Science and Technology for her work on AT&T's ULH backbone. She has worked on a number of groundbreaking transformations in AT&T's network. Her current role at AT&T includes developing strategy for next generation photonic technologies including metro and long-haul technologies, and she leads a small team of engineers across the country.

Simon Tse (stse@att.com) is a director of inventive science at AT&T, New Jersey. He received a B.S. in engineering from Brown University, an M.S. and a Ph.D. from Harvard University in applied sciences, and an M.B.A. from the Wharton School of the University of Pennsylvania. He began his career in 1985 with AT&T Bell Laboratories. He currently manages a group of technical professionals in network topology designs, network traffic management, and software defined network controllers for multi-layer network resource optimization.

Andrew Wallace (afwallace@att.com) is a lead member of technical staff at AT&T, Middletown, New Jersey. Before joining AT&T in 1997, he received his Ph.D. in physics from Yale University.

Kang Xi was formerly a principal inventive scientist at AT&T Labs. He was assistant/associate professor in the Department of Electrical and Computer Engineering at New York University from 2005 to 2014. He received his B.S., M.S., and Ph.D. from Tsinghua University in 1998, 2000, and 2003, respectively, all in electrical engineering. He received the ICCCN 2010 best paper award. His research includes SDN, NFV, data center and cloud computing, and network resilience.