NEW PARADIGMS IN
OPTICAL COMMUNICATIONS AND NETWORKS

Extending Software Defined Network Principles to Include Optical Transport

Steven Gringeri, Nabil Bitar, and Tiejun J. Xia – Verizon Laboratories

ABSTRACT

Software defined networks are based on the principle of a centralized control plane separated from the network forwarding or switching plane that it controls. The switching plane can be heterogeneous, composed of network elements from multiple vendors, and it can provide distinct services with different characteristics, configurations, and control at the packet and/or optical layers. Abstracting the control plane from the network elements allows network-platform-specific characteristics and differences that do not affect services to be hidden. In addition, software defined networking (SDN) is based on the principle that applications can request needed resources from the network via interfaces to the control plane. Through these interfaces, applications can dynamically request network resources or network information that may span disparate technologies. For instance, the application layer can dynamically request and obtain network resources at the packet flow, circuit, or even optical level based on application layer requirements. Current SDN implementations focus on Ethernet switching primarily for data center resource optimization. This article reviews the benefits and challenges of extending SDN concepts to various transport network architectures that include optical wavelength and fiber switches, circuit switches, and sub-wavelength optical burst switches. Control plane implementations for optical networks are more complex since they must account for physical constraints including optical signal reachability, bandwidth availability and granularity, light path routing, and light path reconfiguration speed. The long-term goal is to apply SDN concepts across multilayer multivendor networks in order to support a unified control structure.

INTRODUCTION

As communication networks have grown in size and complexity, streamlining the architecture and implementation to reduce costs, simplify management, improve service provisioning time, and improve resource utilization has become increasingly important. Ideally, an application or service could be completely decoupled from the underlying network infrastructure, but this is not always realistic. Parameters that affect application performance, such as bit rate, packet loss, and latency, are closely tied to the underlying network. To meet application performance objectives, it becomes necessary for the application or its proxy to ensure that the underlying network is aware of the application requirements and provides the necessary services. A network control plane component is often used to find and configure the required resources in the network, as well as map the application traffic to these resources. There are many approaches for applications to request and receive a service from a network. In all cases, the application interacts with some sort of a network control plane. The control plane can be separate from or embedded with the network forwarding or switching plane. To support transport networks with multiple administrative and technology segments, communication and interoperability between control planes is required. The mediation between control plane segments and layer boundaries (e.g., between packet and optical layers) is often a manual process that software defined networking (SDN) [1] could automate.

SDN is a paradigm that provides for separation between the control and the forwarding or switching planes, referred to simply as switching plane in the rest of this article. In its strictest form, SDN is about making the decision on how a flow (or a connection) needs to be set up across the network and configuring the network accordingly. That decision can be optimized based on network resource state, policies, and application demand. Once a decision is made, the control plane communicates the necessary actions to the switching plane via a network control protocol. OpenFlow [2, 3] is one such protocol that was initially defined for packet networks and services, but may be extended to the optical layer. More importantly, the SDN control paradigm should not be confused with OpenFlow or any other protocol used for configuring the network elements. In this article, network configuration and programmability are used interchangeably.

There have been several notions in the industry on the motives of SDN and the definition of SDN. Initially motivated by academic research

0163-6804/13/$25.00 © 2013 IEEE

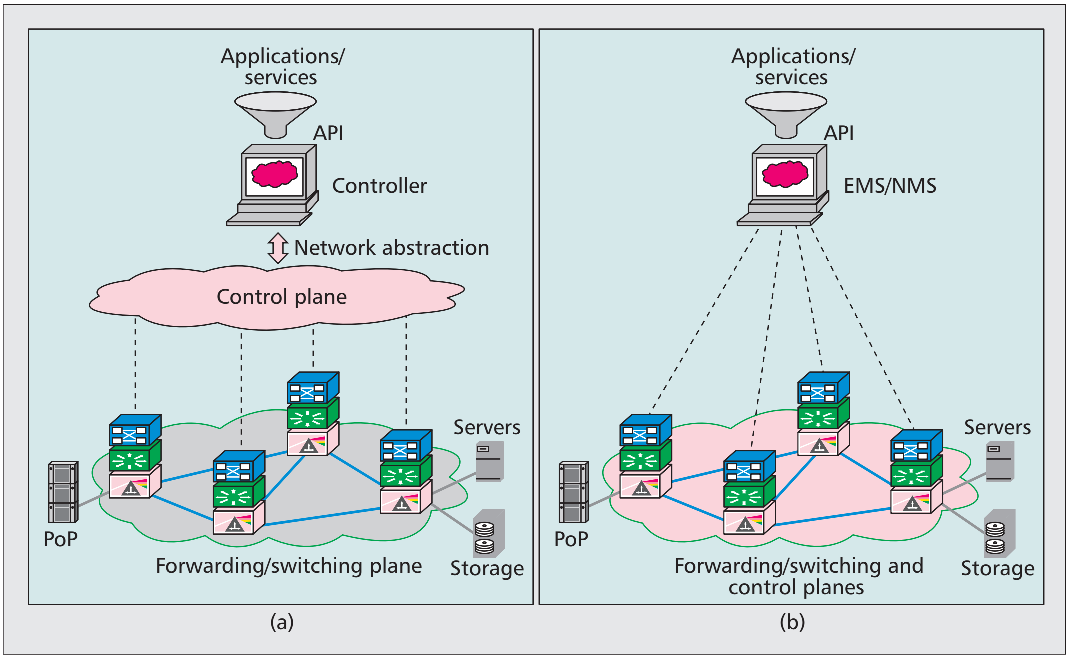

Figure 1. *Network architectures: a) SDN enabled with an API between the application and the SDN controller; b) GMPLS control plane with integrated control and switching on network elements and an API between the application and EMS/NMS.*

Whether a network implements an SDN paradigm, a fully distributed control plane embedded in the network elements, or a hybrid will depend on the mix of services supported on the network and their requirements, considering the trade-offs between these paradigms.

and coupled with the evolution of OpenFlow, SDN was about creating an environment that allows for experimentation with different control plane intelligence not supported in commercial switching hardware platforms. For some, it was about reducing the switching platform cost or implementing the control plane on generic computer server hardware. Others saw the significance in service enablement. In this article we focus on service enablement, and specifically the ability to optimize network resource usage and automating certain functions that often require a manual or cumbersome sequence of interactions using current paradigms and procedures. This article does not discuss the impact on hardware cost or hardware commoditization since SDN by itself is unlikely to significantly affect either. Rather, SDN is a paradigm that offers flexibility in optimizing and implementing the control and switching planes, and can make using commodity hardware more feasible in certain cases. In addition, this article does not make the assertion that SDN replaces other existing distributed control planes embedded in network elements performing traffic switching. Rather, it asserts that the SDN paradigm is applicable in whole or in part to solving issues that exist with such distributed control paradigms, and the applicability depends on the services being addressed. Whether a network implements an SDN paradigm, a fully distributed control plane embedded in the network elements, or a hybrid will depend on the mix of services supported on the network and their requirements, considering the trade-offs between these paradigms.

Figure 1a shows the proposed SDN architecture where the control plane is implemented in a centralized controller that communicates with the switching hardware using a standardized interface that abstracts hardware-specific implementation details. For Ethernet, IP, and multiprotocol label switching (MPLS) networks, OpenFlow is one protocol that could provide a standardized interface to the hardware, but does not cover circuit or wavelength-based equipment. Applications can interact with the network through the controller using an application programming interface (API) to allow for better orchestration of network resources, optimize overall network performance, and provide for the application needs. The network orchestration function provides a logical to physical mapping of network services so that network resources can be virtually allocated for the application at various levels of granularity. The application layer can dynamically request and/or control network resources at wavelength, circuit, or flow granularity to seamlessly reconfigure resources based on higher-layer requirements.

The control structure for a traditional optical, circuit, or packet switch that uses a centralized network management system (NMS) can be modified by adding an application interface to support application-driven (re)configurability of network resources. The physical limitations of each architecture will determine overall flexibility. It should be noted that traditional optical networks have been managed using an element management system (EMS/NMS) that requires network operators to design circuits and consequently drive the configuration of network elements. Some implementations of optical networks have evolved to support a dynamic control plane distributed to network elements to expedite the circuit design process and improve provisioning times. Figure 1b shows an optical network that uses a generalized MPLS (GMPLS) dynamic control plane with integrated control and switching planes. It is still possible for the application layer to request resources on this network, but such interaction would use an API that represents the northbound interface of the EMS/NMS. Control is limited to the functions that the EMS/NMS can access since the control plane is still embedded in the network elements.

This article reviews how an SDN control paradigm can be applied to optical networks that support various transport and switching technologies. It starts with a review of application requirements and service characteristics that can take advantage of SDN control, including connectivity, bandwidth, and resiliency. Next, the article describes sample architectures that use optical and/or electrical switching to provide wavelength, circuit, and packet-based services to satisfy application requirements under SDN control. The specific aspects of each architecture that are programmable or controllable are described, and service characteristics are compared for the various architectures. Finally, the article compares distributed and centralized control, and shows an SDN application for a multilayer network.

Since both the packet and TDM approaches use switching in the electrical domain, the underlying optical infrastructure is hidden at the service switching layer, albeit some characteristics of the underlying optical paths are often factored into the selection of the TDM circuit or packet path.

APPLICATIONS AND SERVICE CHARACTERISTICS

Many applications have connectivity, bandwidth, quality of service (QoS), and resiliency requirements that the network must support for an application to perform properly. This section reviews two applications in transport networks that can make use of an SDN paradigm: connection path computation and setup, and transport virtual private networks (VPNs). To provide connectivity, the network establishes communication paths among the application's endpoints. The communication paths must be selected such that the bandwidth needs of the application and the associated QoS requirements for jitter, packet loss, and delay are met. Resiliency defines the survivability requirement of the application, in particular the impact of network failures on the application performance and the required recovery time.

Some applications require guarantees for certain performance objectives. As an example, an application may require a connection with guaranteed bandwidth, an upper bound on one-way delay, very low packet loss, and network protection. This connection may be a flow on a packet network such as a label switched path (LSP) on an MPLS network, a circuit on a time-division multiplexed (TDM) network such as Ethernet mapped into an optical transport network (OTN) or synchronous optical network (SONET) circuit, or an optical wavelength service. The application may also have time-variant requirements, such as high bandwidth demand during the day and lower bandwidth requirements at night, that will require the network to dynamically adapt the service it provides for the application. The overall goal is a joint optimization of application performance and network resource utilization to maximize network monetization and network efficiency, and minimize service costs. As an example, without dynamic allocation of resources, the worst-case usage of an application must be considered, and corresponding resources must be allocated to that application, leading to underutilized resources. On the other hand, dynamic allocation of resources could allow unused resources assigned to an application or a service to be reallocated to other applications or services. Finally, certain operational environments or applications may require continuous monitoring of connections for liveliness and performance in order to take action when service objectives are not met.

Another application at the transport layer is network slicing whereby paths are set up with constraints on bandwidth, delay and jitter and dedicated to a client to form a transport VPN. Network slicing can happen at multiple layers in the network hierarchy. For instance, wavelengths may be set up and dedicated to a client. TDM circuits of the client could then be restricted to be set up over corresponding client-dedicated wavelengths.

SAMPLE OPTICAL NETWORK ARCHITECTURES

Software defined networking principles can be applied to many different optical transport network architectures, but the underlying network characteristics will determine the implementation's suitability to various applications. In this section, some key network technologies and architectures that are applicable to the SDN control paradigm are presented. All of the architectures considered here provide a switching function that allows the network to be reconfigured or reprogrammed based on application layer traffic demands and performance requirements. Switching can be supported in both the electrical and optical domains. Electrical switching can be at the frame granularity for packet services, or at the time slot level for TDM circuits with guaranteed bandwidths. TDM is generally defined by the OTN optical data unit (ODUx) standards (or synchronous optical network standards) with a rigid bandwidth hierarchy, and provides fixed and guaranteed bandwidth for client ports irrespective of whether traffic is present or not. The packet-based approach provides more flexibility and finer bandwidth granularity than with TDM. Individual packets, as opposed to fixed time slots, are switched, and statistical variations in bandwidth can be leveraged to increase network utilization. However, guaranteeing bandwidth and latency requires the appropriate packet queuing and scheduling mechanisms to be implemented on network elements. Since both the packet and TDM approaches use switching in the electrical domain, the underlying optical infrastructure is hidden at the service switching layer, albeit some characteristics of the underlying optical paths are often factored into the selection of the TDM circuit or packet path.

Optical architectures can be constructed to switch fibers or ports that may contain multiple channels over a wide spectrum (> 5 THz bandwidth) or to provide wavelength selective switching with fixed or flexible slices of spectrum. Fiber switches are more applicable in a data center or small metropolitan network with abundant

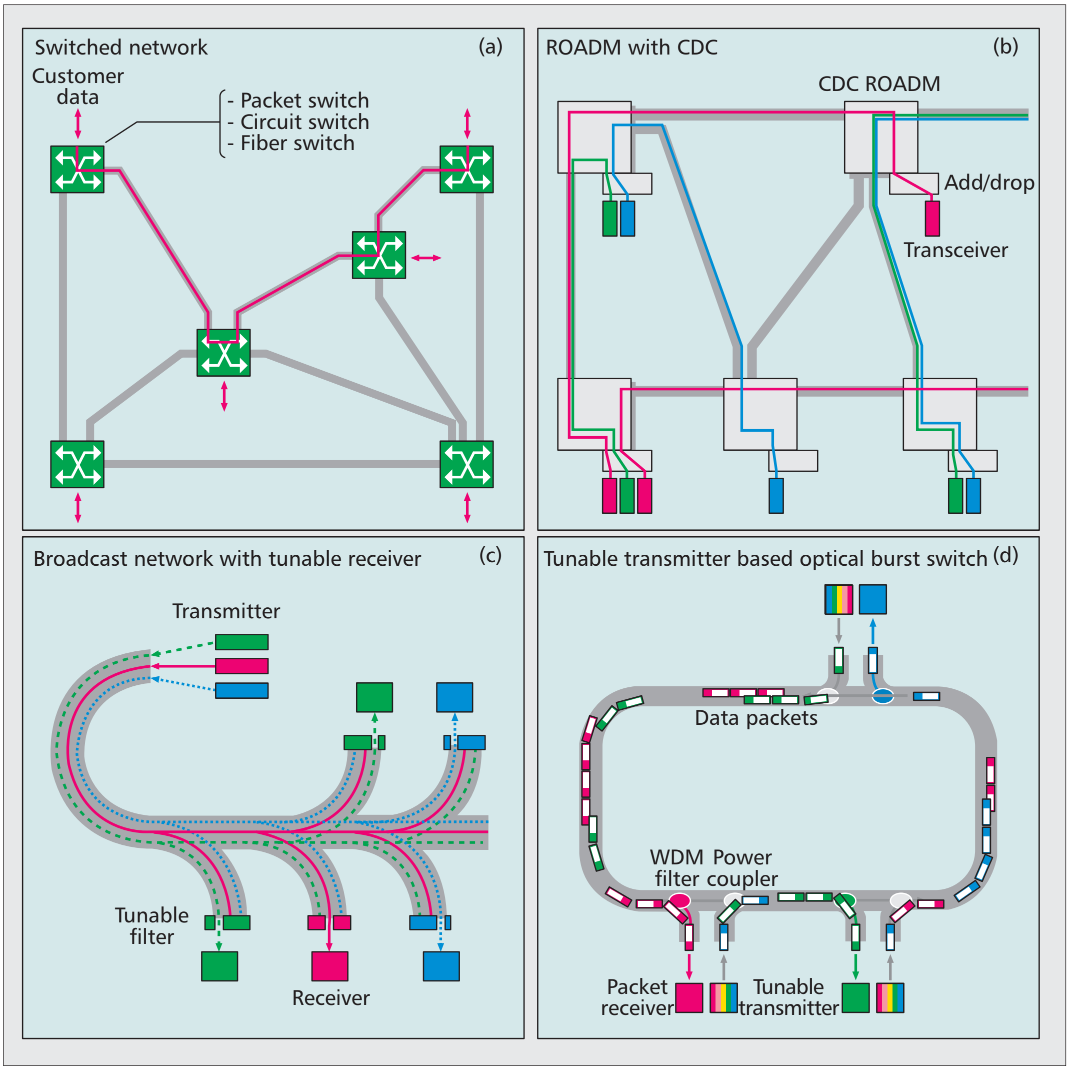

Figure 2. *a) Generic packet/TDM/fiber switching; b) CDC ROADM-based network; c) sub-wavelength switched network (broadcast network with tunable receiver); d) sub-wavelength switched network (tunable-transmitter-based optical burst switch).*

Without modifying the ROADM and amplifier hardware, provisioning times can be improved by optimizing the power control algorithm and establishing engineering guidelines for tolerances that allow faster adjustment and provisioning.

fiber resources since the switching granularity is so coarse (the whole fiber). Optical switching causes transients that can affect transmission performance and in addition not all paths through the network can be guaranteed based on optical impairments. Figure 2a shows a generic switched network where the switching granularity can be packet, time slot, or fiber.

Wavelength switching is typically implemented as a reconfigurable add/drop multiplexer (ROADM) that, as its name indicates, can add, drop, or express wavelengths through the node [4]. There are various add/drop structures possible, from the simplest fixed multiplexer to a structure that can add or drop wavelengths without restrictions on color, ingress/egress direction, or color reuse from different directions. This is referred to as a colorless, directionless, and contentionless (CDC) add/drop structure. Figure 2b shows a sample network where a wavelength service is added at an ingress node, switched through the network, and dropped at an egress node. This architecture allows automated setup of a wavelength between two endpoints, but the control structure is generally slow and requires minutes to complete the setup of the wavelength path. The circuit reach is limited by the optical performance of the transceiver and network impairments. When necessary, signal regeneration can be used to extend reach. Power levels are increased slowly to minimize amplifier transients. Without modifying the ROADM and amplifier hardware, provisioning times can be improved by optimizing the power control algorithm and establishing engineering guidelines for tolerances that allow faster adjustment and provisioning. This is simplest for metropolitan networks that use fewer amplifiers and do not use Raman amplification.

Traditionally, transceivers had a fixed performance based on the optical components that were selected. With the introduction of digital signal processing in the new generation of coherent transceivers, it is now possible to design flexible hardware that will support trade-offs between parameters such as optical reach, bit rate, and spectral efficiency under software control [5]. Figure 3 shows an illustration of such a device where software “knobs” could be configured to achieve the desired performance. Param-

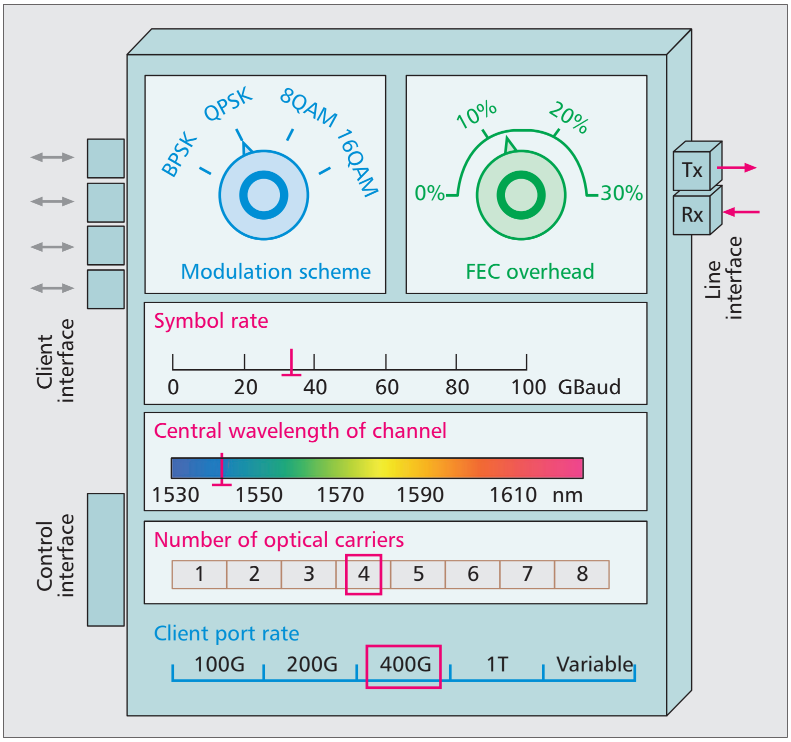

Figure 3. *Schematic representation of a tunable transceiver that illustrates various parameters that can be configured.*

eters such as the modulation scheme — quadrature phase shift keying (QPSK), binary PSK (BPSK), 16-quadrature amplitude modulation, QAM) — forward error correction (FEC) coding overhead (7 , 13 , 20 , 25 percent), symbol rate, and number of carriers can be varied under software control to achieve the desired performance. For example, selecting 16-QAM instead of QPSK doubles the bit rate, but with a significant loss in unregenerated reach. Rates greater than 100 Gb/s (such as 400 Gb/s or even 1 Tb/s) can be supported for a desired reach by using various combinations of modulation scheme, carriers, and symbol rates [6]. When many optical carriers are packed together and switched as a group, a super-channel is formed. Flexible spectral allocation will also be required since more than the typical 50 GHz of spectrum will be required to achieve reasonable reach for most super-channel applications. Both the symbol rate and FEC overhead knobs can be used to fine-tune the super-channel's bandwidth. There are several operating modes possible depending on which parameter is fixed. For instance, with a fixed amount of spectrum, bit rate could be varied to optimize reach. Alternatively, for a fixed bit rate, spectrum could be increased to meet a reach requirement. These parameters, which determine the optical performance, can be under the control of an SDN controller to meet application needs and maximize network resource utilization.

Wavelength networks can be made more dynamic by adding fast wavelength tunability to either the transmitter or receiver to implement sub-wavelength switching. All-optical sub-wavelength switching allows sharing of wavelengths in the time domain so that the same transmitter or receiver (depending on the implementation) can communicate with multiple nodes. Transmitter or receiver tunability allows the network to be reconfigured based on traffic demands if a wavelength is assigned based on the traffic destination or source. The advantage of these architectures is that transceiver resources are not fixed between two endpoints but can be dynamically reused based on traffic demands. Figure 2c shows a sub-wavelength switched network that uses optical broadcast in a tree topology to broadcast to each destination and a fast tunable receiver to select the desired wavelength. Note that a ring topology could be used instead of the broadcast tree. The selection can be made using a tunable filter or by tuning the local oscillator in a coherent receiver. Overall performance is determined by the filter or local oscillator tuning time plus the time for the receiver to lock the incoming signal. The narrow line width external cavity lasers typically used in coherent receivers can take tens of seconds to tune. Faster-tuning narrow line width lasers are just becoming available, but tuning speed and line width can be traded off based on the desired performance.

Figure 2d shows a sub-wavelength switched network based on a tunable transmitter burst switching approach [7], where destinations have fixed colors, and the transmitter tunes to the color based on the desired destination. This is generally practical for ring topologies of several hundred kilometers in diameter, and a scheduling algorithm or carrier sensing hardware is required to prevent multiple transmitters from simultaneously transmitting to the same destination (on the same wavelength). The design tradeoff is to maximize link utilization while keeping the probability of collision very low. These architectures can support fast connection setup between any of the endpoints, but the number of endpoints and the network topology are constrained by the transceiver's un-regenerated optical reach between all combinations of endpoints and potentially the scheduling algorithm. A GMPLS control plane may be used to manage these networks, but reaction to dynamic changes in bandwidth over short time intervals between endpoints is likely to impose scale and performance challenges. The current approach models the network as a distributed Ethernet switch with an optical fabric. It hides the underlying optical infrastructure from the application since the connection between any two endpoints will be triggered by demand measured in the data plane at an ingress node, and the connection is set up based on a distributed scheduling algorithm across that optical fabric which is distributed over the nodes.

Service Characteristics Comparison for the Sample Architectures

The service characteristics that can be provided to the application layer will vary based on the underlying network technology, as shown in Table 1. Based on the application requirements, the network can be constructed with the appro-

Service characteristics	Network architecture							
Function	**Switching**			**Transport (optional λ routing)**			**Transport and switching**	
Type	Ethernet or MPLS switch	OTN switch	Fiber or Port switch	λ transport (fixed)	λ transport & switching (ROADM with CDC	l Transport & switching with programmable transceiver	Sub λ switching with tunable receiver	Sub λ switching with tunable transmitter
Switching domain	Electrical	Electrical	Physical port	None	λ	λ	λ	λ
Service type	Packet connection	Circuit	Fiber or port	λ	λ	λ	Sub λ	Sub λ
Switching granularity	Frame or Packet	ODU	Fiber	NA	λ	λ (flexible bit rate)	Sub λ	Sub λ
Minimum practical rate	Mb/s	1.25 Gb/s	Bit rate independent	10 Gb/s	40–100 Gb/s	40–100 Gb/s	Based on burst size	Based on burst size
Network physical topology	Mesh	Mesh	Mesh	Linear or Ring	Mesh	Mesh	Ring or broad-cast tree	Ring
Distance sensitivity	No	No	Yes	Yes	Yes	Yes (flexible reach)	No (within reach limit)	No (within reach limit)
Flow, circuit, or λ reconfiguration time	ms	ms	ms to s	Hours	10s	10s	ms to s	ms
Number of nodes	Thousands	Thousands	Hundreds	Tens	Hundreds	Thousands	Tens	Tens

Table 1. *Comparison of service characteristics for various network architectures.*

priate combination of electrical switching, optical switching, and transport functions. Network architectures that use electrical switching do not restrict reachability based on optical impairments when the underlying optical network is pre-provisioned. These architectures can support networks over large geographic areas with many endpoints, but a static (non-dynamically reconfigurable) optical layer will limit link bandwidth flexibility and overall network efficiency. As the underlying wavelength transport networks become dynamic through the use of ROADM technology, the link capacity can be dynamically optimized. Wavelength paths can be set up, torn down, or rerouted based on the application level requirements. For this application, SDN can provide multiple-layer control to optimize between the electrical and optical domains. By managing a flexible transceiver, an SDN controller can optimize capacity and efficiency by selecting transmission parameters based on application requirements. As an example, wavelength capacity could be maximized based on the reach requirement.

Sub-wavelength switching architectures (burst switches) integrate switching and optical transport as a single layer but are restricted to smaller geographic areas. Optical impairments are hidden, and fast reconfiguration (almost the speed of the electrical switch) can be supported. The reconfiguration performance is bounded by the burst switching scheduling processes to minimize packet loss. Sub-wavelength switching architectures can adapt to application level requirements since wavelength paths are not fixed between any two nodes but rather dynamically configured based on the traffic demands.

COMPARISON OF CENTRALIZED AND DISTRIBUTED CONTROL

An objective of SDN is to guarantee application performance and resource needs while maximizing network efficiency, as previously discussed. SDN control mechanisms and controllable network attributes depend on the underlying network technology. Packet services can be connectionless or connection-oriented. A connectionless packet service such as IP routing can use a connection-oriented service, but rarely vice versa. Connection-oriented technologies such as GMPLS can provide transport services to applications or connectionless networks. An LSP between two endpoints may have constraints on protection, delay, and bandwidth to satisfy the needs of services running over it. In this case, it becomes prudent to select a path that satisfies those constraints, establish the path on the network, and monitor it. Those actions could be performed by the network elements in the path in a distributed control plane environment, or via a centralized SDN controller. The preference for one approach or another could be operational in nature or more fundamentally dependent on technical and cost trade-offs.

In a distributed control plane, one endpoint of the connection (or circuit) is told what con-

Topology information can be stored in the form of a database or learned from the network using routing protocols and updated based on network events. Delay information can be measured or computed and can be part of the database or topology information.

nection it needs to set up, and the attributes of the connection in terms of bandwidth, protection and some quality of service objectives. If there are multiple independent endpoints in the network setting up connections at the same time, it is possible that the network may not be utilized in the best possible way. Some links may be more loaded than others. Link utilization may also be fragmented such that no single link may be able to satisfy a new demand even though the network is underutilized. There are also cases where the service cannot be satisfied by a network element because it or the network lacks the necessary capabilities to learn or disseminate relevant topology information. For instance, a distributed control plane has difficulty in maintaining diversity between dynamically computed paths across network elements since the computation of each path is an independent process. A centralized controller could help address these problems if it has the right capabilities. For this application, an SDN controller must be aware of the network topology, link attributes such as bandwidth and delay, and dependencies between links such as shared risk link groups. Topology information can be stored in the form of a database or learned from the network using routing protocols and updated based on network events (e.g., link failures or link additions). Delay information can be measured or computed and can be part of the database or topology information.

An SDN controller, using topology information, can compute a path on a network that satisfies connection, circuit, or wavelength constraints. In addition, it can better globally optimize the placement of connections (or circuits) on the network as the computation is done centrally for all as opposed to being distributed where every connection (or circuit) is set up independently. Thus, links could be more efficiently utilized, leading to network cost reductions. Specifically, an SDN controller can move connections to new paths while satisfying their constraints in order to satisfy the constraints of new connection(s) that would have been blocked otherwise. It can also compute diverse paths irrespective of whether they share the same endpoints. That is, the SDN controller approach has benefits over a control plane that is solely distributed to network elements when global network resource optimization is required or there are interdependent connection constraints across nodes. Centralized control can overcome the shortcomings of a distributed control plane with its global knowledge of application requirements and network resources. Finally, an SDN controller may compensate for deficiencies with the distributed control plane to satisfy the application requirements, such as computing a path across Internet Gateway Protocol (IGP) areas, or computing a path that satisfies a latency constraint. It can also configure the performance monitoring capabilities of these paths, and collect appropriate measurements on delay and packet loss. If the performance objectives are violated due to degraded links, congestion, or increased delay, it may recompute a path that avoids the problematic segments. The power of centralized SDN control also manifests itself in coordinating the application network needs across disparate technologies, each with its own controllable attributes, to realize a network service. Such a situation arises in multitenant data centers where tenant networks need to be established or modified dynamically as nodes on the network need to be selected and connectivity established. It also arises in multilayer networks composed of packet, TDM, and optical layers.

An SDN controller generally provides two APIs, one for applications to request resources from the network and the other to connect directly to the network elements or to an EMS that controls the network. While centralized controllers provide for global network optimization and coordination among connections, they may be slower to react to network changes (e.g., link failure) than a distributed control plane. This can stem from the time to propagate the change from the network to the SDN controller and from the response time at the SDN controller, which is potentially processing events or requests pertaining to multiple network elements, unlike a control plane processor dedicated per network element.

It should be noted that scale, responsiveness to network events, and reliability are likely to drive the deployment of multiple geographically distributed CPUs that together form a unified intelligent SDN control cluster. Coordination within the cluster becomes important to maintain decision centralization and reap the benefits of centralized control. Similarly, control across operational or authoritative domain boundaries controlled via individual SDN controllers will necessitate communication among the controllers to establish services that cross these domains. It is likely that these SDN controllers will belong to different vendors; thus, inter-SDN controller interoperability is also required. These are challenges that do not arise in a constrained and homogeneous environment such as data centers operated by one cloud service provider, but are likely to arise in larger or interconnected networks from different service providers and/or vendors. Similar problems that arise today in interoperability and standards protocol definition will need to be addressed to enable SDN on a larger scale, although complexity may be reduced by using simple APIs across SDN controllers or mature protocols.

While centralized SDN control can address the coordination of application requirements with the network, large networks will still need to rely on a distributed embedded control plane to provide fast reaction to network events. This embedded control plane is more scalable, and leverages mature and reliable technologies that were developed to replace early centralized control approaches. Connectionless IP networks are a good example of truly distributed control. The centralized SDN control approach brings us full circle, but does not resolve all of the underlying limitations of centralized control. Thus, to implement large scalable networks that are coupled to the application layer, it becomes essential that SDN-based centralized control and distributed embedded control coexist on the same network. The combination of centralized and distributed control requires careful design since functions are distributed based on processing requirements, and some data synchronization is required between the SDN controller and the embedded control processors.

SDN APPLICATIONS FOR MULTILAYER NETWORKS

Modern optical architectures provide a number of controllable attributes that can be soft-tuned based on application needs and network state. An SDN controller can provide for such control. For instance, an application requiring optical bandwidth between any two points may trigger an SDN controller to compute and select such a path with appropriate setting of the various optical parameters, depending on the underlying technology as discussed in the previous sections. In addition to computing and possibly provisioning the optical path, the controller could also manage a software programmable transceiver setting parameters such as modulation scheme, FEC overhead, symbol rate, and number of carriers based on the required performance and optimum resource utilization. Such a controller may need knowledge of vendor-specific attributes of the underlying technology, which is difficult to model for all possible vendors. Thus, an SDN control model must allow at minimum a two-tier hierarchy of SDN controllers whereby one coordinates among specific technology controllers that hide some of the underlying technology characteristics, keeping the interface between the first and second tiers at the service level as if the first tier were an application. The overall application in this case may be a demand placed by another network and automatically orchestrated via a multilayer SDN controller. It may also be a portal for a wavelength service. With demands changing at small timescales, automatic control will be needed, and that control must take into account the engineering policies and configuration times for the network.

In order to illustrate the role SDN can play in multilayer networks, consider the following scenario depicted in Fig. 4. Three routers, R1, R2, and R3, are connected to transport nodes T1, T2, and T3, respectively. Assume that the traffic demand between R1 and R2 is 10 Gb/s, and so is the traffic demand between R1 and R3 and R2 and R3. Furthermore, corresponding transport circuits are allocated between T1, T2, and T3 to carry the 10 Gb/s links. If there is an additional 5 Gb/s of bandwidth required between R1 and R3, the SDN controller is asked to enable 5 Gb/s worth of capacity between R1 and R3 on a path with a particular delay bound. The SDN controller first attempts to set up an ODU2 circuit between T1 and T3, but realizes that there is no capacity available that would satisfy the delay bound. The path through T4 and T5 has more hops, but has less delay than the path through T2 and can meet the delay constraint. The SDN controller enables a wavelength path between T1 and T3 via T4 and T5, and computes an ODU2 path between T1 and T3 such that the total latency requirement is met. The SDN controller computes an LSP path between R1 and R3, and configures an LSP with 5 Gb/s of bandwidth between these two routers. This sequence of operations takes place without requiring any level of protocol interaction between the router nodes and the transport nodes, or requiring such protocols to exist and/or be enabled in the network.

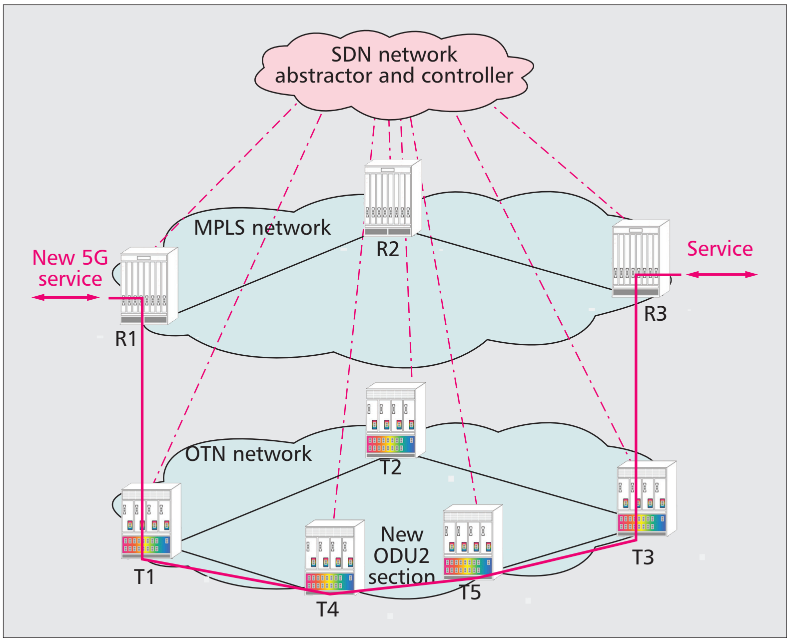

Figure 4. *Services provided by the SDN abstractor and controller.*

CONCLUSION

SDN-based control supports application-driven control of the traffic path and associated network services while providing for global optimization of network resource utilization. It provides a framework for applications to request and receive services from the network via a centralized SDN controller. That controller can use multiple processors distributed across multiple locations for improved scalability and reliability while maintaining global network optimization. A centralized control plane may impose performance and scaling limitations compared to distributed control networks. Performance can be improved by adding processing power and precomputing protection paths, but getting failure notifications and effecting dynamic reroute will be slower than with a distributed control plane. An SDN controller can also provide for coordinating services across multiple network layers (e.g., packet, TDM, and optical) and among non-interoperable technologies (e.g., proprietary embedded control plane or incompatible features) in the same layer, capitalizing on network abstractions at the various layers and hardware-independent control.

Extensions and additional standardization are required to allow the SDN controllers to manage TDM circuit and wavelength-based architectures reviewed in this article. With these extensions, LSPs, OTN circuits, and wavelength paths could be abstracted and managed by the SDN controller. As the optical layer has become more flexible and dynamic with the introduction of ROADMs with CDC add/drop capabilities and

Mechanisms that support a combination of centralized and distributed control over a multiple layer network that contains both electrical and optical switching will be required to support and optimize network efficiency for evolving applications.

software configurable transceivers, manual operations can be automated. Sub-wavelength optical switching architectures will provide additional bandwidth granularity and dynamic routing capabilities that can be used to further optimize network utilization and adapt to changing application service requirements over a short timescale. Using a flexible optical layer, SDN can evolve from its initial packet applications to support more generalized switching and transport applications. However, for some network technologies and services, a more distributed routing and signaling control plane is needed to provide for faster reaction to data plane events, and to scale to support more endpoints and larger geographic area networks. Both control paradigms have their trade-offs, and both are likely to need to coexist in the same network. Mechanisms that support a combination of centralized and distributed control over a multilayer network that contains both electrical and optical switching will be required to support and optimize network efficiency for evolving applications.

REFERENCES

[1] G. Goth, "Software-Defined Networking Could Shake Up More than Packets," *IEEE Internet Computing*, vol. 15, no. 4, 2011, p. 6.

[2] N. McKeown *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *ACM SIGCOMM Computer Commun. Rev.*, vol. 38, no. 2, Apr. 2008, pp. 69–74.

[3] ONF, "Software-Defined Networking: The New Norm for Networks," Mar. 13, 2012, https://www.opennetworking.org/images/stories/downloads/white-papers/wp-sdn-newnorm.pdf.

[4] S. Gringeri *et al.*, "Flexible Architectures for Optical Transport Nodes and Networks," *IEEE Commun. Mag.*, vol. 48, no. 7, 2010, p. 40.

[5] K. Roberts and C. Laperle, "Flexible Transceivers," *ECOC '12 Proc.*, We.3.A.3.

[6] S. Gringeri *et al.*, "Technical Considerations for Supporting Data Rates Beyond 100 Gb/s," *IEEE Commun. Mag.*, vol. 50, no. 2, 2012, p. S21.

[7] P. Bayvel and M. Dueser, "Optical Bust Switching: Research and Application," *Proc. OFC*, 2004.

BIOGRAPHIES

STEVEN GRINGERI (steven.gringeri@verizon.com) received his B.S. and M.S. degrees in electrical engineering from Worcester Polytechnic Institute in 1986 and 1989, respectively. He joined General Electric in 1986, where he developed signal processing algorithms and systems for machine health monitoring. In 1990 he joined GTE Laboratories, now Verizon Laboratories, where he is currently a principal member of technical staff. His research and development activities have varied from video compression to the design and specification of metropolitan and long-haul networks. He is currently responsible for the architecture, specification, and deployment of next-generation optical networks.

NABIL BITAR holds B.S., M.S. and Ph.D. degrees in electrical engineering from Boston University, Massachusetts. He is currently a principal member of technical staff at Verizon. He leads the end-to-end IP-MPLS network and services architecture, including next generation broadband IP services (subscriber management, IP data and video, IPv6), CDN, VPN (IP and Ethernet), packet aggregation and transport, services routing, and DC/cloud network services virtualization and interconnection. He also led or co-led architecture for carrier Ethernet services, converged packet/OTN/TDM transport, IMS and VoIP services routing, wireless-wireline convergence, and consumer data center evolution. Prior to joining Verizon in 2004, he worked at Ascend/Lucent as a system architect for ATM and IP-MPLS services (forwarding, traffic management, signaling, and routing). Prior to Ascend/Lucent he worked at GTE Laboratories for five years on wireless AIN, ADSL architecture, IP IntServ and DiffServ, MPLS, VoIP, and traffic management. He is a regular contributor to the IETF and Broadband Forum, and a frequent speaker and panelist at various fora. He has (co)authored several technical papers, IETF RFCs and drafts, and holds several patents on traffic management, routing, signaling, and policy control. He is Co-Chair of the L2VPN working group at IETF and serves on the IETF Routing Area Directorate.

TIEJUN J. XIA is an expert on photonic technologies and optical fiber communications in research, development, and technology innovations. He is a Distinguished Member of Technical Staff with Verizon. He is also an adjunct professor at Miami University, Ohio. Recently he served as director for network technology development at Chorum Technologies. Prior to Chorum, he worked with MCI. Before joining MCI, he was a research faculty member in the Department of Electrical Engineering at the University of Michigan. He was an adjunct professor at the University of Texas at Dallas from 2005 through 2009. He is cofounder and current president of the Advanced Fiber Connectivity & Switching Forum. He has served on Technical Program Committees of OFC/NFOEC, APOC, ACP, and OECC. He serves on the Executive Board of the IEEE Dallas Section. He holds his Ph.D. degree in physics from the Center for Research and Education in Optics and Lasers (CREOL) at the University of Central Florida, his M. S. degree from Zhejiang University, China, and his B.S. degree from the University of Science and Technology of China. He has published more than 100 technical papers, given numerous invited talks, and holds more than 60 granted or pending U.S. patents. He is a Fellow of OSA. In 2011, he was featured as "Verizon Innovator" on YouTube.