Evolving Traffic Grooming in Multi-Layer Flexible-Grid Optical Networks With Software-Defined Elasticity

Shuqiang Zhang, Massimo Tornatore, Gangxiang Shen, *Senior Member, IEEE*, Jie Zhang, and Biswanath Mukherjee

I. Introduction

TRAFFIC grooming was introduced for both single-wavelength and wavelength-division-multiplexing (WDM) rings [2], where multiple low-rate SONET/SDH circuit connections (such as OC-1, OC-3, etc.) are "groomed" onto a high-rate wavelength channel, e.g., OC-48 (2.5 Gb/s), OC-192 (10 Gb/s), or OC-768 (40 Gb/s). The objective is to minimize the network cost in terms of line terminating equipment (LTE), e.g., add/drop multiplexers. Then, the techniques for traffic grooming were extended to mesh topologies, which are widely deployed in today's backbone networks [3]. In an optical mesh network, the objective of grooming is usually either to maximize network throughput or to minimize network cost in terms of optical transponders, electronic processing, or power consumption.

The first generation of traffic grooming is realized by Layer-1 time-division-multiplexing circuit switching [e.g., SONET/SDH, now evolving to optical transport network (OTN)]. As the network is becoming more packet dominated (e.g., Layer-2 Ethernet, Layer-3 IP), the concept of grooming is evolving to aggregating multiple low-rate traffic flows (either packet or circuit flows) onto a high-rate wavelength channel. So, the function of grooming can also be realized by packet switching (e.g., IP/MPLS or Ethernet). Note that there is an important difference, often overlooked, in models used to optimally design optical networks for grooming packet and circuit traffic flows: packet flows can be split into fine flow granularity (as fine as a packet), while circuit flows can only be split into coarser-granularity flows (e.g., OC-1, OC-3, OC-12, etc.) by virtual concatenation (VCAT) if allowed. Please see optimization models for grooming circuit and packet traffic in [3] and [4], respectively. In this study, we use the two terms "connection" and "traffic flow" interchangeably.

In multi-layer networking, the concept of traffic grooming can be generalized, e.g., traffic flows requiring certain amount of bandwidth from any layer (e.g., TCP flows, or application-layer content flows) can be considered as being "groomed" onto any bandwidth container (e.g., MPLS or IP tunnels, wavelengths, superchannels, etc). In this study, we mainly consider network architectures and traffic grooming from Layer 0 (e.g., optical layer) to Layer 3 (e.g., IP layer).

A. *Essence of Current Grooming Paradigm*

There are three basic problems in communication networks [5]: 1) Traffic engineering, which is about putting the traffic where the bandwidth is; 2) Network engineering, which is about putting the bandwidth where the traffic is; and 3) Network planing, which is about putting the bandwidth where the traffic is forecasted to be. Traffic grooming can be involved in any of the three problems, depending on the phases of network operations.

Current grooming approaches based on Layer-1/Layer2 (e.g., SONET/SDH), Layer-2 (e.g., Ethernet), or Layer-3 (e.g., IP)

Manuscript received January 31, 2014; revised March 22, 2014; accepted March 31, 2014. Date of publication April 14, 2014; date of current version July 25, 2014. A short summarized version of this paper [1] was presented at the the European Conference and Exhibition on Optical Communication (ECOC), London, U.K., in Sep. 2013.

S. Zhang and B. Mukherjee are with the University of California Davis, CA 95616 USA (e-mail: sqzhang@ucdavis.edu; bmukherjee@ucdavis.edu).

M. Tornatore is with the Politecnico di Milano, Italy and University of California Davis, Davis, CA 95616 USA (e-mail: tornator@elet.polimi.it).

G. Shen is with the the Soochow University, Suzhou 215006, China (e-mail: shengx@suda.edu.cn).

J. Zhang is with Beijing University of Posts and Telecommunications, Beijing 100876, China (e-mail: lgr24@bupt.edu.cn).

Color versions of one or more of the figures in this paper are available online.

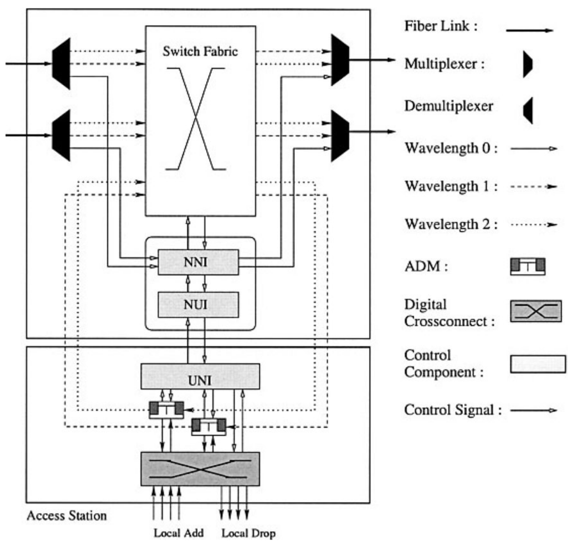

Fig. 1. Traffic grooming based on Layer-1 switching: SONET/SDH over WDM node architecture [5].

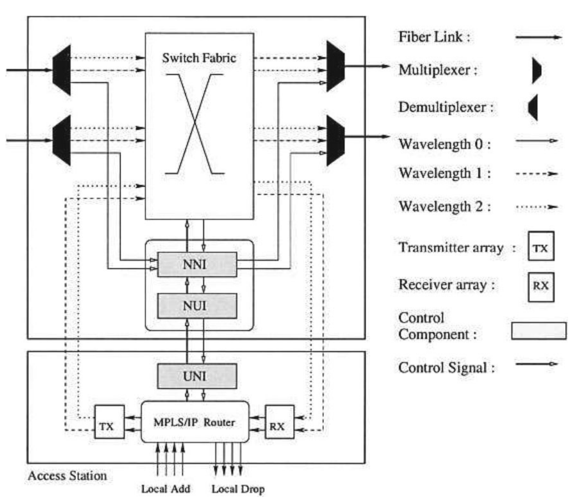

Fig. 2. Traffic grooming based on Layer-3 switching: IP over WDM node architecture [5].

switching are all realized by electronics. In general, grooming involves multi-layer node architectures, where an electrical layer (Layer 1, 2, or 3) and an optical layer (Layer 0) collaborate to support grooming. Figs. 1 and 2 show the node architectures for traffic grooming based on Layer-1 and Layer-3 switching, respectively. The traffic-grooming problem can be divided into two sub-problems:

1) electrical-layer traffic routing over a virtual topology formed by a set of virtual links (lightpaths). This sub-problem is similar to the traffic-engineering problem in IP networks and can be modeled as a multi-commodity flow problem, if packet switches are adopted for grooming.
2) optical-layer routing and wavelength assignment or routing and spectrum assignment (RSA) to establish the virtual links, which is a well-known NP-hard problem.

Fig. 3 shows an example illustrating how the two sub-problems interact with each other.

Traffic grooming can increase the utilization of optical-layer resources such as transponders and wavelengths, thereby

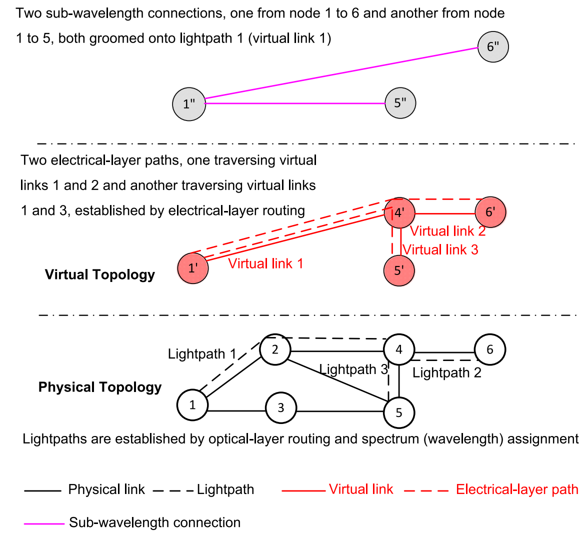

Fig. 3. An example view of multi-layer traffic grooming.

reducing network cost to satisfy a certain amount of traffic. There is a tradeoff between transponder cost (as well as related higher-layer electronic-processing cost, such as line cards of packet switches) and spectrum utilization: in an opaque network, where optical-electrical-optical (OEO) conversion and grooming are applied to traffic flows or connections at every intermediate network node, optical spectrum can be most-efficiently utilized. But transponder cost could be high since excessive OEO conversions and electronic processing are used to achieve the high spectrum utilization. On the other hand, if the network is optimized for transponder cost, spectrum utilization has to be sacrificed. So, network operators should design and operate their networks according to their specific situations. For example, in the network-planning phase, network operators can adjust the optimization objectives of traffic-grooming models to minimize either transponders or spectrum usage. And in the dynamic operation phase, traffic-grooming polices can be adjusted according to the real-time traffic and resource utilization.

Also, from the perspective of quality of service (QoS), traffic grooming could increase end-to-end latency for individual traffic flows because of indirect routing and intermediate electronic processing, which will be discussed in Section III.

The rest of this study is organized as follows. In Section II, the new scope of traffic grooming is introduced due to the disruptive innovation in the optical layer: flexible grid and elastic rates (ERs), and the role of traffic grooming in flexible-grid/ER networks is discussed. In Section III, we show the impact of sliceable optical layer on network architectures and traffic grooming. Section IV discusses how software-defined networking (SDN) would enable multi-layer convergence and application-centric networks. Section V concludes the study.

II. Grooming's Role in Flexible-Grid/ER Networks

A. Emergence of Flexible Grid and ER

Flexible grid: Mixed-line-rate (MLR) [6] signals are expected to co-exist on the same fiber plants of the network. The current

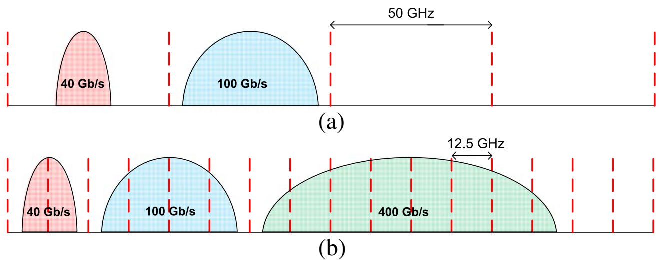

Fig. 4. Fixed grid versus flexible grid.

fixed-grid partitioning [see Fig. 4(a)] will not be able to accommodate the targeted rates in the coming years (e.g., 400 Gb/s and 1 Tb/s), because of insufficient inter-channel spacing. New and flexible grids are needed to accommodate future rate heterogeneity. It is expected that the standard ITU fixed grid will be replaced by a new recommendation [7] introducing the finer frequency slots of 12.5 GHz, where an optical channel's spectrum can span several frequency slots [see Fig. 4(b)].

ERs: Researchers are studying innovative transponders capable of dynamically tuning data rates or optical reach according to the network state, typically by adjusting modulation formats, number of subcarriers, or baud rates. These transponders are called elastic or software-defined, e.g., they allow flexible allocation of just enough optical bandwidth resources based on the current traffic demand, leading to unprecedented agility in optical networking. ER transponders can adapt their data rate to meet changing network conditions, whether planned (e.g., due to upgrades) or unexpected (e.g., due to failures). Also, an optimized tradeoff between spectrum efficiency and optical reach can be achieved on demand, e.g., the data rate of a transponder can be tuned up for shorter paths with lower physical impairments by using more advanced modulation formats.

Flexible grid and ER are evolutionary technologies for optical networks, and they can coexist. Table I summarizes the new scope of optical networking where traffic grooming can be applied. Most previous works on traffic grooming focused on fixed-grid (WDM) optical networks with single type of fixed-rate (FR) transponders, e.g., single-line rate (SLR). Fixed-grid networks with multiple types of FR transponders (e.g., 10, 40, and 100 Gb/s) are often referred to as MLR networks [6]. The MLR concept can be extended to flexible-grid networks. The grooming model for MLR networks (both fixed grid and flexible grid) is similar to that in SLR networks (see [6] for the design of an MLR network with traffic grooming). Studies on ER in fixed-grid networks have also been performed to minimize power consumption by re-adapting network capacity to varying traffic intensity [8].

The value of ER lies more in its dynamic-adaptation capability. It is expected that the potential of ER can be fully exploited by a software-defined elastic optical network with a separated control plane. The most agile form of future optical networks would combine both flexible grid and ER (e.g., software-defined Flexible-grid/ER networks). To realize this, bandwidth-variable reconfigurable optical add/drop multiplexers (BV-ROADMs) are needed to dynamically switch heterogeneous optical channels. Flexible grid and ER can benefit each other: elasticity of ER would be limited without flexible grid since the spectrum of an optical channel is strictly limited by the fixed grid; and, without ER, the spectrum-efficiency benefits of flexible grids cannot be fully exploited, e.g., under dynamic traffic, two optical channels cannot efficiently share spectra without ER [9].

B. Re-examining Traffic Grooming's Role

What is the role of traffic grooming if we have the most agile optical layer with flexible grid and ER? It may be argued that traffic grooming can be eliminated in future Flexible-grid/ER networks, as they can provide fine-granular optical channels. However, optical channels must be separated by guard bands to avoid interference when switched by BV-ROADMs. So, provisioning each low-speed connection with a separate optical channel may lead to high spectrum wastage by guard bands. Also, low-speed optical channels usually have lower spectrum efficiency (bit/s/Hz) than high-capacity ones. So, from the spectrum utilization point of view, provisioning connections without traffic grooming is inefficient. If high-capacity transponders are adopted to increase spectrum efficiency of optical channels, dedicated traffic-grooming methods, such as the spectrum-reservation scheme proposed in [10], should be adopted to avoid low utilization of high-capacity ER transponders. Traffic grooming was shown to significantly increase spectrum utilization by saving guard bands in both static [11] and dynamic [12] traffic models. Traffic grooming can also significantly increase ER-transponder utilization, if proper spectrum-(re)assignment schemes are used (e.g., spectrum reservation [10] for grooming or defragmentation [13]). It was also shown in [10] that the optimization tradeoff between transponder cost and spectrum utilization still stands in Flexible-grid/ER networks. But without grooming, both transponder and spectrum utilization would suffer in Flexible-grid/ER networks. So, traffic grooming plays a similar role in Flexible-grid/ER networks as in the traditional fixed-grid networks under the current grooming paradigm. But note that, in case of distance-adaptive transponders, shorter reach implies higher spectrum efficiency, which, in turn, means that a larger number of transponders might be preferable with respect to network scenarios with non-adaptive transponders [14].

C. Grooming in the Spectral Domain

So, if traffic grooming is beneficial with respect to both transponder and spectrum utilization, can we implement its functionality in other ways instead of just using electronic circuit/packet switching (both of which groom traffic in the time domain)?

To increase the capacity of single-carrier transponders beyond 100 Gb/s, either symbol rates need to be increased or more bits should be coded on each symbol, or both. However, the symbol rates supported by current development of integrated electronics cannot satisfy the demand for bandwidth growth, and introducing more advanced modulation formats to increase the bits per symbol would result in a reduced maximum optical reach.

TABLE I
SCOPE OF TRAFFIC GROOMING

	Fixed grid (WDM)	*Flexible grid*
Fixed-rate (FR) transponder	Single-line rate (SLR)/Mixed-line rate (MLR)	MLR
Elastic-rate (ER) transponder	Adaptive data rates and optical reach	Software-defined Flexible-grid/ER optical networks

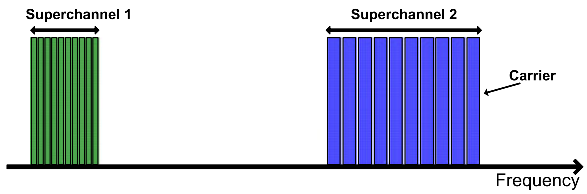

Fig. 5. Superchannels.

Alternatively, superchannel is a promising candidate for moving the line-side speed beyond 100 Gb/s. As shown in Fig. 5, a superchannel is a composite line-side signal that is a combination of multiple optical carriers which are treated by network nodes (e.g., BV-ROADMs) as a single switching unit. The optical carriers within a superchannel can be tightly grouped together in the spectral domain without guard bands, which makes superchannels highly spectrum-efficient. Superchannels can be implemented by different multi-carrier technologies, such as Nyquist-WDM and orthogonal frequency-division multiplexing (OFDM) [15].

From the perspective of traffic grooming, multiple clients' connection requests, each of which may occupy a certain number of carriers, can be spectrally groomed onto a superchannel without the involvement of traditional electrical-layer (e.g., Layer 1, 2, 3) switching. It is obvious that such kind of spectral grooming can achieve higher spectrum efficiency than the case without such spectral grooming, e.g., client signals are transmitted using multiple guard-band-separated low-speed optical channels.

However, only client requests with the same source and destination can be spectrally groomed together, which is not sufficient to replace the traditional electrical-layer (e.g., Layer 1, 2, 3) grooming. From the multi-layer network architecture perspective, the optimization model of traffic grooming is unchanged: a superchannel can be viewed just as a "lightpath" or a "virtual link" in the traditional multi-layer network, although some grooming tasks (e.g., grooming connections that share the same source and destination) can be offloaded to the optical layer, e.g., transponders or muxponders.

An OFDM-based optical grooming scheme was proposed in [16], aiming to eliminate OEO conversion and electronic processing at intermediate nodes. This scheme uses the same optical-grooming principle as superchannels at the transmitter side, but relaxes the constraint that only the connections with the same source and destination can be groomed. By utilizing optical nodes with a broadcast-and-select structure, a superchannel can be optically split into multiple parts in the spectral domain, each of which is a group of carriers and can be transmitted to a separate destination. This "optical source grooming" scheme can save transmitters (but not receivers) compared to the normal point-to-point superchannel schemes. But compared to the current grooming paradigm based on electronic circuit/packet switching that can groom connections from/to any source/destination, optical source grooming's benefits with respect to both spectrum and transponder utilization is limited.

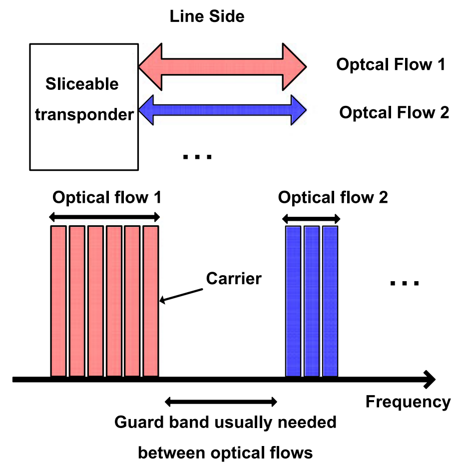

Fig. 6. Sliceable transponder.

III. POTENTIAL IMPACT OF SLICEABLE OPTICAL LAYER ON FUTURE GROOMING PARADIGM

Sliceable (or multi-flow, virtualized) optical transponders [15], [17] were recently proposed to offload IP traffic to the optical layer and reduce the number of IP router ports and transponders. Fig. 6 shows the logical function of a sliceable transponder, which can be realized using multiple sub-transceivers exploiting multi-carrier techniques (e.g., OFDM or Nyquist-WDM), similar to superchannel transponders. However, unlike point-to-point superchannels, a single sliceable optical transmitter/receiver (composed by an array of sub-transmitters/sub-receivers) can transmit/receive multiple 'optical flows' to/from multiple nodes without any intermediate electronic switching. Each optical flow can be switched as a single independent unit by BV-ROADMs and formed by an arbitrary number of carriers, as long as the total number of carriers of all optical flows transmitted/received by a single transponder does not exceed its capacity. Note that the optical carriers within an optical flow can be tightly grouped together in the spectral domain without guard bands. This is what makes sliceable transponder more spectral-efficient than just using multiple traditional low-speed transponders. However, different optical flows usually need guard bands between each other in order to be switched by BV-ROADMs, unless they have the same source and destination pair, in which

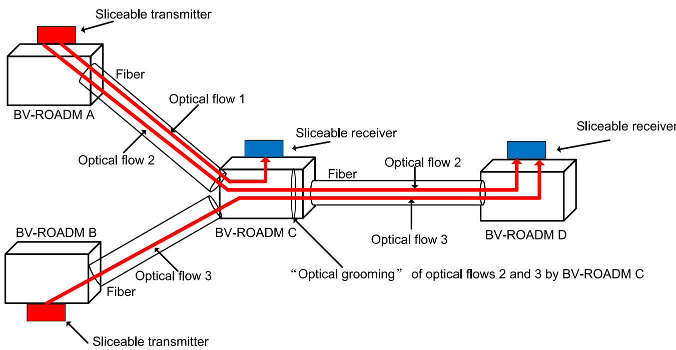

Fig. 7. Illustration of "optical grooming".

case they can be spectrally "groomed" into a larger optical flow or a superchannel.

For example, in Fig. 7, a single sliceable transmitter at BV-ROADM A can transmit optical flows 1 and 2 to the two sliceable receivers at BV-ROADMs C and D, respectively, and a single sliceable receiver at BV-ROADM D can receive optical flows 2 and 3 from BV-ROADMs A and B simultaneously. This capability of sliceable transponders, together with BV-ROADMs that can switch optical flows with arbitrary spectrum widths, provides the similar function as electronic grooming: e.g., multiple traffic (optical) flows transmitted from different sliceable transponders (possibly at different locations) can be optically "groomed" by an intermediate BV-ROADM (including the one at the destination) and thereafter switched as a single unit in the network and received by a single sliceable transponder. Note that *the "optical grooming" here is achieved by switching optical flows to the same output port of BV-ROADM, and there are still guard bands between those optical flows from different sources (e.g., optical flows 2 and 3 in Fig. 7)*. So, from spectrum utilization point of view, it cannot achieve the same benefit as electronic grooming. However, the advantage of using sliceable transponders is that they can reduce the need for intermediate electronic processing, thus potentially reducing end-to-end latency and, at the same time, reduce the total number of transponders and electronic ports, thus potentially reducing the total network cost. So, it can achieve the similar function as electronic grooming in terms of increasing transponder utilization.

A. Network Architectures With Sliceable Optical Layer

To verify the benefits of sliceable optical layer, we consider three different network architectures. First is the traditional packet-over-optical network architecture (Traditional), shown in Fig. 8(a). Packet switches (IP/MPLS, or Ethernet switches) are used to aggregate clients' traffic and groom transit traffic from other nodes, if necessary. Traditional point-to-point optical transponders are used in this architecture. The optical channel established in this case is often called a "lightpath".

Second is the direct sliceable optical layer network architecture (DS) [see Fig. 8(b)], in which the packet layer is eliminated. The optical channel established between a pair of sliceable transponders is called an "optical flow". The function of aggregating client traffic is accomplished by the N × M sliceable

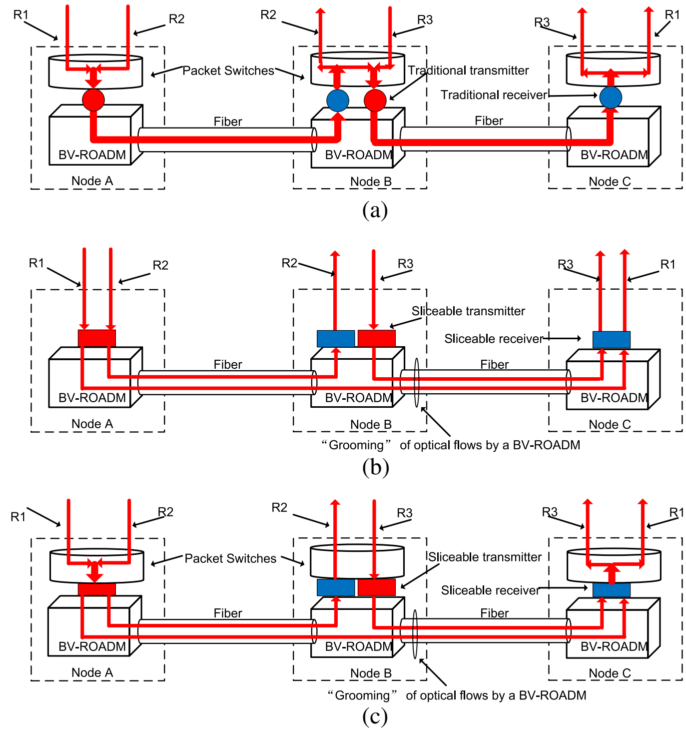

Fig. 8. Three network architectures.

optical transponder [15], which has N client ports and supports M optical flows on the line side. From the perspective of the client, it is similar to the traditional muxponder. The advantage of DS is that it eliminates the packet layer, so it can reduce the latency and network cost. Also, thanks to the sliceable transponder and BV-ROADM, multiple optical flows from different sources can be received by a single sliceable receiver, providing benefits similar to electronic traffic grooming. These benefits are provided without OEO conversions, thus eliminating any latency due to electronic processing.

However, the aggregation pattern of client traffic is fixed: a client's node must be directly connected to a transponder and any reconfiguration of the connection pattern (e.g., re-multiplexing a client traffic to another unused port of a transponder) would involve manual intervention. Also, each client's connection request is *unsplittable*: e.g., all its traffic must be sent by a single sliceable transmitter and received by a corresponding sliceable receiver due to the fixed physical connection. In contrast, with packet switches, a client's traffic can be mapped to one or multiple transmitter/receivers, since a client is not directly facing an optical transponder, but a packet-switch port.

Moreover, because the BV-ROADMs in the DS architecture can only "groom" optical flows, the grooming granularity of the DS architecture is coarse and is limited by the bandwidth granularity of the optical flows. A low-speed traffic flow may still request a bandwidth that is smaller than that of an optical flow. For example, if two clients, from two different sources and going to the same destination, both request a bandwidth that is smaller than half of the capacity of an optical flow, their traffic can *not* be groomed into a single optical flow in DS, which is possible if a packet switch is used for performing electronic grooming.

So, to add flexibility, the third architecture is similar to the second but with an additional packet layer, which is called packet-over-sliceable (PoS) optical network architecture [see Fig. 8(c)]. Now, the sliceable transponder has only one client port (e.g., 1 × M sliceable optical transponder [15]) that is connected to the packet switch (or integrated with packet switches). Electronic traffic grooming can be accomplished by the packet switches, if necessary, besides the "optical-grooming" function provided by the sliceable optical layer.

Fig. 8 shows an example of three clients' connection requests: 1 to 3. Each connection requests 40-Gb/s bandwidth, and both traditional point-to-point and sliceable optical transponders have 400-Gb/s capacity. In the traditional case, Requests 1 and 3 can be electronically groomed by the packet switch at node B and received by a single traditional receiver at node C. In the other two cases with the sliceable optical layer, the similar function of "electronic grooming" can be achieved by "optical grooming", e.g., Requests 1 and 3 can be directly "groomed" by the BV-ROADM at node B and received by a single sliceable receiver at node C. For Request 1, intermediate electronic processing is eliminated.

As mentioned before, despite the grooming-alike functions by the sliceable optical layer, it alone cannot provide the grooming granularity as fine as electronic packet switching. So, electronic packet switches should coexist with the sliceable optical layer [see Fig. 8(c)]. But, with the sliceable optical layer, the granularity gaps between optical channels (optical flows) and traffic flows become much smaller. The sliceable optical layer has the potential to affect future grooming paradigm by offloading significant amount of traffic load from upper electrical layers (originally due to the need of electronic grooming) to the optical layer, thereby reducing the total network cost. Also, from the perspective of QoS, end-to-end latency for individual traffic flows can be significantly reduced because of less electronic processing.

B. Illustrative Numerical Examples

We propose optimal designs of the three network architectures by formulating integer linear programs (ILP), which try to minimize the the number of transponders as the first priority and then minimize either the average normalized latency or spectrum usage as the second priority (see [18] for the details of the formulations). In the numerical example, the capacities of the traditional and sliceable optical transmiters/receivers are assumed to be the same: 400 Gb/s. The sliceable transponders can slice the capacity into as many as 10 optical flows (carriers), each with bandwidth granularity g of 40 Gb/s. We use the same spectrum-width parameters (shown in Table II) as in [15] for different data rates (e.g., multiples of 40 Gb/s) of optical flows. Note that the spectrum width in Table II for each data rate includes the guard bands. We consider a six-node network (see Fig. 9) and 14-node The National Science Foundation Network (NSFNET) (see Fig. 10) as test networks, in which the link weight represents the link's length in kilometers. In the six-node network, we consider that there is a directional connection request from each node to all the other five nodes. In NSFNET, the number of connection requests from each source is uniformly distributed in [1], [3], and their destinations are randomly chosen. In both examples, the bandwidth of each connection request is 40 Gb/s multiplied by an integer number (e.g., multiples of 40 Gb/s) and the integer number is uniformly distributed in [1, x], $1 \leq x \leq 10$. For instance, when $x = 3$, the requested bandwidth of each connection is uniformly distributed in [40, 120] Gb/s, with a step of 40 Gb/s.

TABLE II
SPECTRUM USAGE FOR OPTICAL FLOWS WITH DIFFERENT DATA RATES [15]

Bit rate (Gb/s)	40	80	120	160	200
Spectrum width (GHz)	25	50	50	75	75
Bit rate (Gb/s)	240	280	320	360	400
Spectrum width (GHz)	100	100	125	125	150

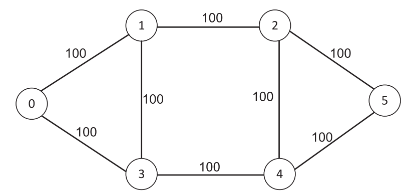

Fig. 9. Six-node network.

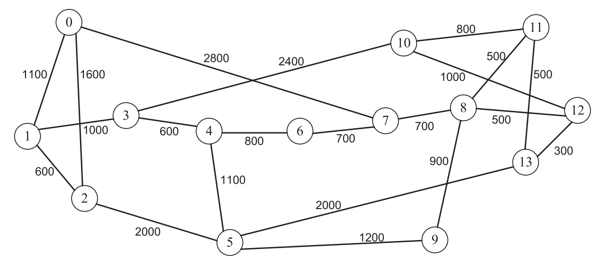

Fig. 10. NSFNET.

We have already shown some preliminary results of the transponder and latency performance for the 6-node network in [18]. In this study, we demonstrate additional results on NSFNET and also compare the network architectures' performance in terms of spectrum usage.

Fig. 11 shows the number of transmitters/receivers needed to provision all the requests in the NSFNET example. In the traditional architecture, the number of transmitters is always equal to that of receivers, so we use a single bar to indicate this number. It can be observed that, even with electronic traffic grooming, the traditional architecture consumes more transmitters and receivers than DS and PoS, thus potentially it has higher cost (considering that sliceable transponders are implemented using similar multi-carrier technologies as non-sliceable ones and they have comparable costs [15]).

Because the latencies of connections depend on the actual implementations of network equipment (transponders, switches,

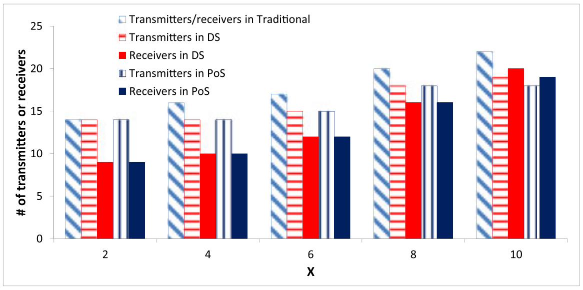

Fig. 11. Number of transmitters or receivers in NSFNET.

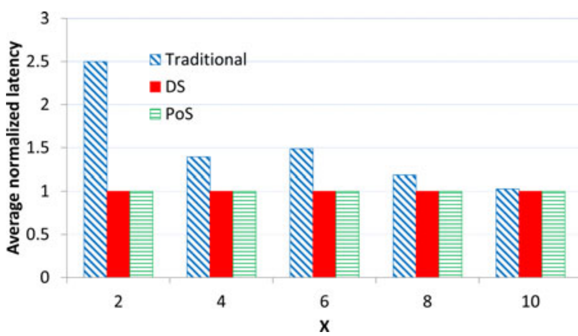

Fig. 12. Average normalized latency in NSFNET.

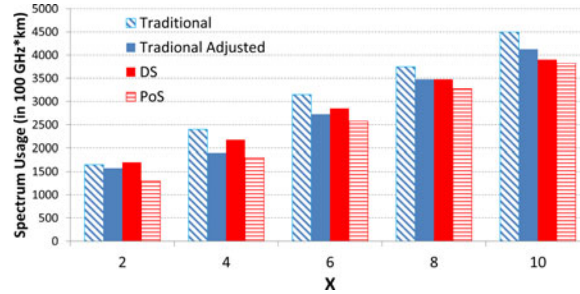

Fig. 13. Spectrum usage in six-node network.

etc.), we define a metric, called average (over all requests) normalized latency, denoted by D, to evaluate the overall latency performance affected by the network architectures. The normalized latency of a connection is defined as the actual latency of the connection normalized by the "lowest-possible" latency of the connection in the corresponding network architecture. The latency of a connection is "lowest possible" when it is established using an end-to-end optical channel (e.g., a lightpath or an optical flow) through the shortest-distance path with no intermediate electronic processing.

Fig. 12 shows the average normalized latency of the requested connections in the NSFNET example. *Note that the latency is achieved as a secondary objective when the transponder usage is guaranteed to be minimum for all the architectures.* In the traditional architecture, traffic grooming by intermediate packet switches is indispensable to reduce the number of transmitters/receivers, which would inevitably increase the end-to-end latency. In the DS and PoS architectures, the average normalized latency is always equal to one, since there is no intermediate electronic processing for each traffic request in the optimal design and we assume that transponders are scarcer resources than optical spectrum (in other words, we assume there is sufficient optical spectrum). In this example, network architectures with a sliceable optical layer can achieve up to 60% latency reduction, compared to the traditional architecture with traffic grooming.

Note that the traffic load used in our numerical study is relatively small (e.g., only one or two requests originate from each node), because we need to fairly compare the *optimal* designs of the network architectures using ILPs which are not scalable to large inputs. If the traffic load is higher, we expect that the advantages of DS and PoS would be more prominent in terms of both transponder usage and latency.

Fig. 13 shows the total spectrum usage of different network architectures. The 6-node network is adopted in this example because the ILPs, which minimize transponder usage as the first priority and the spectrum usage as the second, are not scalable to NSFNET on our computing platform. We define the spectrum usage of an optical channel (e.g., either a lightpath or an optical flow) to be the *product* of its occupied spectrum width (in GHz) and its transmission distance (in 100 kilometers). Note that the spectrum usage in Fig. 13 is achieved when the transponder usage is guaranteed to be minimum for each architecture. Otherwise, the minimum spectrum usage can be always achieved by opaque network architectures in which OEO conversions are applied to all connections at all nodes, while consuming much more transponders. The difference between the "Traditional" and "Traditional Adjusted" bars in Fig. 13 is that "Traditional" assumes that every lightpath of 400-Gb/s capacity would fully occupy 150 GHz, while "Traditional Adjusted" can "adjust" the spectrum widths of the lightpaths that are not fully utilized. For example, in "Traditional Adjusted", if only 40 Gb/s of a lightpath (with capacity of 400 Gb/s) is needed, the lightpath's spectrum width can be adjusted to 25 GHz.

It can be observed from Fig. 13 that PoS can achieve the lowest spectrum usage because it can exploit both spectrum-efficient optical flows and electronic traffic grooming. "Traditional" architecture can utilize electronic traffic grooming but still consumes the highest spectrum usage because it wastes much spectrum when there is not enough traffic to fully utilize the capacity it provisions. But if we adjust the spectrum of "Traditional", we can observe that "Traditional Adjusted" can achieve lower spectrum usage than DS in most cases. This is because, without electronic traffic grooming, different optical flows in DS must be separated by guard bands, whose spectrum wastage outweighs the benefits of individual spectrum-efficient optical flows.

Note that the low spectrum usage of PoS in Fig. 13 is achieved when minimizing spectrum usage is the second priority of our ILP formulation (e.g., spectrum usage has a much smaller weight than transponder usage in the objective function), which is not the situation when the "lowest-possible" latency is achievable. In fact, the low spectrum usage of PoS is achieved by using electronic traffic grooming which would inevitably increase the latency. However, we can see that PoS is very flexible to adapt

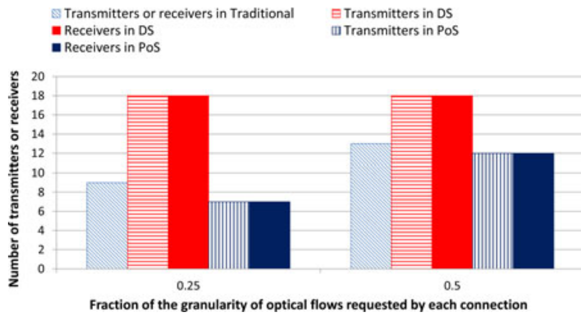

Fig. 14. The benefits of packet switches for fine-granularity connections.

to different needs of network operators, either provisioning the lowest latency or lowest spectrum usage, and at the same time the lowest usage of transponders is always guaranteed.

As shown in Fig. 11, the DS architecture achieves almost the same performance as PoS in terms of the number of transmitters/receivers. But, if clients request much smaller bandwidth than the granularity provided by the sliceable transponder, PoS could still benefit from electronic traffic grooming to reduce the number of transmitters/receivers compared to DS. Fig. 14 shows the benefits of packet switches for grooming fine-granularity connections. In this example, each connection requests only a fraction (e.g., 0.25 or 0.5) of the minimum granularity of optical flows. To make the total traffic demand large enough to affect the transponder usage, we adjust the minimum granularity of optical flows to 200 Gb/s while the capacity of each sliceable transponder is still 400 Gb/s and supports two optical flows. We also use the six-node network and the same connection pattern as the previous example. We can observe from Fig. 14 that DS is not able to groom fine-granularity traffic and it consumes much more transponders than PoS and Traditional. The finer the traffic granularity, the larger the gap is. PoS still performs the best because it can take advantages of both optical and electronic grooming.

IV. Application-Centric Multi-Layer Convergence Enabled by SDN

The ultimate goal of networking is to satisfy applications' QoS requirements, while keeping the cost as low as possible. New application trends call for an evolution in existing multi-layer networking practices. Traffic patterns induced by emerging applications are highly dynamic, and managing such traffic requires flexible bandwidth management and traffic engineering. The rise of cloud computing, in which users request on-demand access to software, storage, and infrastructure, also requires on-demand bandwidth. The current trend of "big data" applications, as well as data-center networks and content delivery networks (CDN), also need dynamic "big bandwidth". However, current multi-layer networks are inflexible and incapable of adapting to such new computing trends. The agility and flexibility offered by SDN can be an attractive solution to the problems faced by current networks when they try to adapt to future application trends.

The developments in the optical layer, such as flexible grid and ER, and SDN are highly complementary, because SDN-based control and centralized optimization is necessary to fully utilize the new flexibilities offered by the novel elastic optical layer. The key concept of SDN, namely separation of control and data planes, is not new to optical networking. Industry practitioners and researchers have been working on dynamic optical networking based on a separate control plane for decades, but with only limited success in deployment so far. This is because the optical layer is usually isolated from the applications and major clients of optical layer are packet layers. Due to the inflexible nature of upper-layer packet networks, there were not enough motivations to deploy dynamic optical networking. However, new application trends are driving packet networks to become more agile. So, this ongoing SDN evolution is also an enabler for dynamic optical networking, and SDN-based converged multi-layer networks could bring the optical layer closer to applications.

A. *SDN-Enabled Traffic Grooming*

Packet and optical networks should be jointly optimized and operated to achieve the application-centric convergence. While applications' requirements are diverse, bandwidth and latency are the two most fundamental and important QoS metrics. The resources of multi-layer networks can be jointly managed and controlled in the SDN paradigm, providing applications with virtualized bandwidth with guaranteed latency.

One important benefit of extending SDN to optical networks is joint dynamic resource provisioning and optimization of multiple layers (e.g., IP/MPLS, OTN, and optical layer). Based on SDN, dynamic traffic grooming [10] can become a reality. If an application needs on-demand connectivity with QoS guarantees (e.g., bandwidth, latency, etc.), SDN controllers or their connected path computation elements, based on the collected global information of multiple layers, can calculate a qualified path through multiple layers, and dynamically establish the connection by grooming the connection onto existing or immediately-established lower-layer tunnels (e.g., OTN tunnels, or optical paths). For example, ESnet, Brocade, and Infinera have collaborated to demonstrate multi-layer bandwidth provisioning [19], although SDN was extended to the OTN layer only and the optical layer was still static. So, SDN is an unique opportunity to optimize the way we plan and operate traffic grooming.

On the other hand, traffic grooming represents a powerful tool to exploit the full potential of SDN. Bandwidth and latency requirements of new services are highly dynamic due to new workloads (e.g., cloud/mobile computing, social media, and big-data analytics). Rapidly reprovisioning the bandwidth and latency parameters can be achieved by acting on the traffic-grooming policies that can be effectively adapted to the network state and the service requirements through SDN. Without extending SDN to the optical layer and traffic grooming, bandwidth and latency guarantees are difficult to achieve by higher electrical layers alone, even if SDN is applied to them.

So, all the optical-networking specific problems, such as RSA, traffic grooming, and spectrum management and defragmentation [20] are not independent of higher packet-layer problems. They should be jointly considered under the umbrella of application-centric multi-layer resource control and management enabled by SDN.

V. Conclusion

The evolution of traffic grooming was discussed with emphasis on its role in Flexible-grid/ER networks. It was found that the optimization tradeoff between transponder cost and spectrum utilization still stands in Flexible-grid/ER networks, and traffic grooming plays a similar role in Flexible-grid/ER networks as in traditional fixed-grid networks, in terms of both increasing transponder and spectrum utilization. The optical-layer technologies facilitating grooming in the spectral domain, such as superchannels and "optical source grooming", were found beneficial but not sufficient to have a major impact on the current grooming paradigm. Sliceable optical layer was identified to have a significant impact on the future grooming paradigm by offloading a significant part of electronic-grooming function to the optical layer. Through numerical examples, it was verified that the PoS network architecture can always achieve the lowest usage of transponders and is adaptive to different needs of network operators, either provisioning the lowest average latency for connections or lowest spectrum usage. Extending SDN to the optical layer can enable dynamic traffic grooming and the automated joint control and management of the multi-layer resources with the objective of satisfying the applications' needs and reducing the total network cost.

References

[1] S. Zhang, M. Tornatore, G. Shen, and B. Mukherjee, "Evolution of traffic grooming from SDH/SONET to flexible grid," in *Proc. Eur. Conf. Exhibit. Opt. Commun.*, 2013, pp. 1–3.

[2] O. Gerstel, P. Lin, and G. Sasaki, "Wavelength assignment in a WDM ring to minimize cost of embedded SONET rings," in *Proc. IEEE Conf. Comput. Commun.*, 1998, pp. 94–101.

[3] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, Jan. 2002.

[4] R. Dutta and G. N. Rouskas, "Traffic grooming in WDM networks: Past and future," *IEEE Netw.*, vol. 16, no. 6, pp. 46–56, Nov./Dec. 2002.

[5] B. Mukherjee, *Optical WDM Networks*. New York, NY, USA: Springer, 2006.

[6] A. Nag, M. Tornatore, and B. Mukherjee, "Optical network design with mixed line rates and multiple modulation formats," *IEEE/OSA J. Lightw. Technol.*, vol. 28, no. 4, pp. 466–475, Feb. 15, 2010.

[7] (2012). G.694.1: Spectral grids for WDM applications: DWDM frequency grid [Online]. Available: http://www.itu.int/rec/T-REC-G.694.1/

[8] A. Morea, O. Rival, N. Brochier, and E. Le Rouzic, "Datarate adaptation for night-time energy savings in core networks," *IEEE/OSA J. Lightw. Technol.*, vol. 31, no. 5, pp. 779–785, Mar. 1, 2013.

[9] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Time-varying spectrum allocation policies and blocking analysis in flexible optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 13–25, Jan. 2013.

[10] S. Zhang, C. Martel, and B. Mukherjee, "Dynamic traffic grooming in elastic optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 4–12, Jan. 2013.

[11] Y. Zhang, X. Zheng, Q. Li, N. Hua, Y. Li, and H. Zhang, "Traffic grooming in spectrum-elastic optical path networks," in *Proc. Opt. Fiber Commun. Conf.*, 2011, pp. 1–3.

[12] S. Zhang and B. Mukherjee, "Energy-efficient dynamic provisioning for spectrum elastic optical networks," in *Proc. IEEE Int. Conf. Commun.*, 2012, pp. 3031–3035.

[13] R. Wang and B. Mukherjee, "Provisioning in elastic optical networks with non-disruptive defragmentation," *IEEE/OSA J. Lightw. Technol.*, vol. 31, no. 15, pp. 2491–2500, Aug. 1, 2013.

[14] C. Rottondi, M. Tornatore, A. Pattavina, and G. Gavioli, "Routing, modulation level, and spectrum assignment in optical metro ring networks using elastic transceivers," *J. Opt. Commun. Netw.*, vol. 5, no. 4, pp. 305–315, 2013.

[15] M. Jinno, H. Takara, Y. Sone, K. Yonenaga, and A. Hirano, "Multiflow optical transponder for efficient multilayer optical networking," *IEEE Commun. Mag.*, vol. 50, no. 5, pp. 56–65, May 2012.

[16] G. Zhang, M. De Leenheer, and B. Mukherjee, "Optical traffic grooming in OFDM-based elastic optical networks [invited]," *J. Opt. Commun. Netw.*, vol. 4, no. 11, pp. B17–B25, 2012.

[17] O. Gerstel, M. Jinno, A. Lord, and S. Yoo, "Elastic optical networking: A new dawn for the optical layer," *IEEE Commun. Mag.*, vol. 50, no. 2, pp. 12–20, Feb. 2012.

[18] S. Zhang, R. Wang, U. Mandal, M. Habib, and B. Mukherjee, "Connecting the clouds with low-latency, low-cost virtual private lines enabled by sliceable optical networks," presented at the IEEE Global Telecommun. Conf., Atlanta, GA, USA, 2013.

[19] SDN multi-layer provisioning and optimization demonstration. (2013). [Online]. Available: http://blog.infinera.com/2013/10/14/sdn-multi-layer-provisioning-and-optimization-demonstration/

[20] R. Wang and B. Mukherjee, "Spectrum management in heterogeneous bandwidth optical networks," *Opt. Switching Netw.*, vol. 11, pp. 83–91, 2014.