

The Prospect of Inter-Data-Center Optical Networks

Xiaoxue Zhao, Vijay Vusirikala, Bikash Koley, Valey Kamalov, and Tad Hofmeister, Google Inc.

ABSTRACT

Mega data centers and their interconnection networks have drawn great attention in recent years because of the rapid public adoption of cloud-based services. The unprecedented amount of data that needs to be communicated between data centers imposes new requirements and challenges to inter-data-center optical networks. In this article, we discuss the traffic growth trends and capacity demands of Google's inter-data-center network, and how they drive the network architectures and technologies to scale capacities and operational ease on existing fiber plants. We extensively review recent research findings and emerging technologies, such as digital coherent detection and the flex-grid dense wavelength-division multiplexed channel plan, and propose practical implementations, such as C+L-band transmission, packet and optical layer integration, and a software-defined networking enabled network architecture for both capacity and operational scaling. In addition, we point out a few critical areas that require more attention and research to improve efficiency and flexibility of an inter-data-center optical network: optical regeneration, data rate mismatch between Ethernet and optical transport, and real-time optical performance monitoring.

INTRODUCTION

Google's mission is to organize the world's information, and make it universally accessible and useful. To achieve this goal, Google processes more than 3 billion search queries every day, of which 15 percent are new. Google has found over 30 trillion unique URLs on the web, and over 230 million web domains. To ensure the results returned to users' queries are as current as possible, Google has to crawl over 20 billion websites every day to refresh its index. All these computationally intensive tasks are done in warehouse-scale computers (WSCs), which are commonly known as mega data centers. Google offers services in 55 countries across the world in 146 languages, thereby driving the need for globally distributed computation resources and a global network (Fig. 1).

In addition to global reach, service availability is another important consideration. To ensure

that user experience is maintained to the extent possible during unplanned failures or planned maintenance events, many of Google services' backend designs maintain redundancy by keeping copies in multiple data centers. This combination of global reach, large scale, and inherent redundancy sets the fundamental requirements for Google's inter-data-center optical network.

Capacity scaling on existing fiber plants in the next 5–10 years is one of the main issues to address. Deployment of new fibers along long-haul and ultra-long-haul routes is time-consuming and capital-cost-intensive. Therefore, it is important to maximize the capacity of deployed fiber plants by utilizing various emerging techniques. Today, coherent 100-Gb/s polarization-multiplexed quadrature phase shift keying (PM-QPSK) technologies with increased number of dense wavelength-division multiplexing (DWDM) channels (through channel bandwidth reduction and guard band removal) can likely provide 12 Tb/s per fiber pair. As Internet traffic continues to grow at 50–60 percent year after year [1], solutions for capacity scaling beyond 12 Tb/s are needed. However, fiber capacity scaling is eventually bounded by the nonlinear Shannon limit [1]. The conventional paths explored in the past (e.g., data rate increase per wavelength channel) cannot easily be exploited going forward as we are getting close to the limit.

In addition to the critical task of capacity scaling, designing, deploying, and operating a global optical network on a tens-of-terabits scale has its own challenges. Network flexibility, agility, and automation are necessary features to ensure holistic network scaling.

The Holy Grail for network operators is to have a closed-loop network control and management system that includes monitoring, provisioning, commissioning, and configuration of the network all performed in an automated fashion. As shown in Fig. 2, this automated network control and management is required for new capacity adds as well as online optimization activities such as optical layer routing and spectrum allocation based on real-time monitored telemetry. We show later in this article that a logically centralized network operating system (OS) and a consolidated packet and optical data layer are desired to enable this software-defined networking (SDN) paradigm.

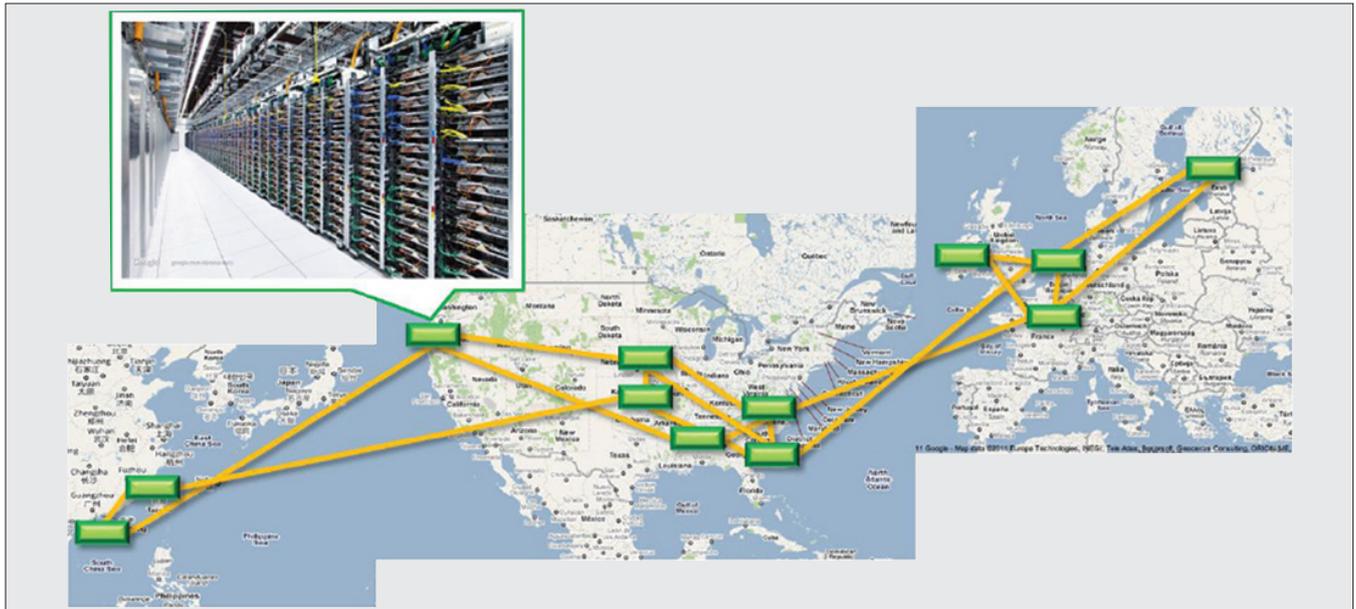


Figure 1. Google's global inter-data-center network. Each box represents a mega data center. Inset: close-up of a Google data center.

In the remainder of this article, we discuss in detail these areas for inter-data-center optical networks: capacity scaling, operational flexibility, and SDN. We review the latest technologies and research findings that can push the capacity-distance envelope on existing fiber plants. We explain how some of these technologies can provide operational ease to scale. Finally, we talk about key optical components and technologies that are necessary to evolve the network to SDN.

CAPACITY SCALING

100 Gb/s per wavelength DWDM systems with digital coherent detection have emerged as the new industry standard and are likely to be the main workhorse for the next several years due to their $10\times$ per fiber pair capacity improvement compared to 10G-based systems while maintaining an optical reach that is comparable to 10G systems. Novel fibers that have been reported in many research experiments, such as large-core, ultra-low-loss, hollow-core, or few-mode fibers, are very attractive for further capacity scaling. However, performing a global network-wide fiber infrastructure build requires a substantial capital investment and is unlikely to take place within the next decade. Therefore, capacity scaling on existing fiber plants remains of great interest to service providers and network operators, and deserves continued attention in both industry and academia as it can solve very practical challenges.

Capacity delivered through a single fiber pair using DWDM technology is determined by the product of the channel count and the per-channel data rate. It is also referred to as DWDM system capacity. Channel count is, in turn, determined by the fiber spectrum in use and the bandwidth allocated to each channel. Therefore, pragmatically, there are three major dimensions to look at for increasing capacity on a fiber pair:

- Number of channels
- Amplification-enabled fiber bandwidth
- Per-channel data rate

It should be noted that these dimensions are not independent. The number of channels is determined by the ratio of fiber bandwidth to the channel spacing (assuming that no guard band is allocated). Channel spacing is inversely proportional to the channel baud rate. And the channel baud rate together with the modulation format determines the channel data rate. Assuming a minimum optical reach, there is a trade-off between channel data rate and channel count given fixed fiber bandwidth, which makes system capacity scaling challenging. We review each of these three dimensions in detail in the rest of this section as they are closely related to DWDM system implementations.

Current DWDM systems typically use fixed 50 GHz spacing and support 80, 88, or 96 channels in the C-band (ranging from 1530–1565 nm) defined by the International Telecommunication Union (ITU). As shown in Fig. 3a, this channel plan was designed originally to accommodate a guard band allocated between adjacent channels to minimize intersymbol interference (ISI) and nonlinear crosstalk, and also to allow optical filtering of a single channel for direct detection systems. However, the use of advanced modulation formats such as QPSK together with polarization multiplexing (PM) and coherent detection obviates the need for a fixed 50 GHz grid. With the use of high-speed digital-to-analog converters (DACs), pulse shapes in the time domain can be synthesized digitally to provide sharp spectral shapes in the frequency domain and minimize ISI between successive symbols. Commonly used spectral shapes are square, raised-cosine, and square-root raised-cosine [2]. Spectral shaping makes the channel bandwidth occupation close to the Nyquist limit for ISI-free transmission, which coincides with the baud rate. In the case of 100 Gb/s PM-QPSK, the typical baud rate is around 32 Gbaud/s (25 Gbaud/s data rate plus ~ 25 percent error correction coding overhead). Therefore, the same 150 GHz fiber bandwidth used to fit three individual wavelength channels, as shown in Fig. 3a, can now fit four spectrally shaped channels, as shown in

With coherent systems and Raman amplification, the optimal signal launch power is much lower (~ 3-4 dB less compared to OOK). Hence, we believe that the SRS effect would not be a major limitation for in-service upgrade to L-band on an existing terrestrial C-band coherent system.

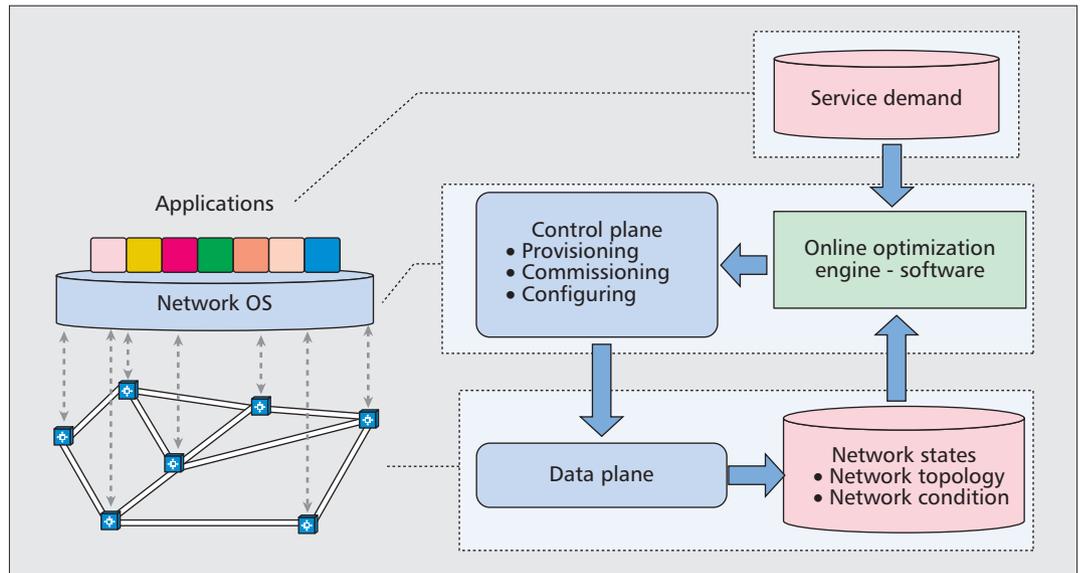


Figure 2. The concept of a software-defined packet-optical inter-data-center network. Each node in the network is a packet-optical integrated device. Software functions and their closed-loop interactions are shown on the right.

Fig. 3b. Channels closely packed together can be viewed as a super-DWDM channel and routed together at a reconfigurable optical add/drop multiplexing (ROADM) node provided the wavelength selective switches (WSSs) in the ROADM are flex-grid (defined by ITU-T SG15/G.694.1) enabled. Then the size of the super-DWDM channel is determined by the WSS setting, and is reconfigurable at a granularity of 12.5 GHz with current technologies and finer granularities in the future. The wider the super-DWDM channel, the higher the overall spectral efficiency, as the guard band is needed only between super-channels, not between individual channels. This can bring 20 to 30 percent more capacity to the same fiber spectrum in use with the same channel data rate.

Optical fiber provides an enormous amount of bandwidth as a transmission media. For high-performance communication, the low-loss window around 1550 nm has been the main focus. C-band is particularly popular due to the availability of low-cost erbium-doped fiber amplifiers (EDFAs). L-band, ranging from 1565–1625 nm, has a similar low loss property but is not as widely utilized because it requires either L-band EDFAs or Raman amplification. Also, chromatic dispersion (CD) in L-band is larger than it is in C-band, which is less desired in traditional dispersion-managed systems.

With the channel speed going up to 100 Gb/s, increasing link optical signal-to-noise ratio (OSNR) becomes the key to achieving the desired reach. One obvious approach is to deploy low-noise amplifiers, and it has been demonstrated that distributed Raman amplifiers can have much lower effective noise figures than comparable EDFAs because they leverage the transmission fiber itself as the gain media where the nonlinear Raman effect can take place as the signal power propagates down the line. Therefore, the signal power would not yet fully drop to its end-of-span value before getting any amplification, as opposed to the case of an EDFA, which in turn minimizes the impact of amplified

spontaneous emission on the degradation of OSNR [2]. It has also been demonstrated that 100 nm (C+L-band) flat-gain Raman amplifiers can be achieved [3]. Therefore, Raman amplifiers can both improve OSNR performance and allow use of both C and L bands. The latest 100 Gb/s digital coherent systems compensate CD in bulk at the receiver in the digital signal processing (DSP) chip, which removes the need for complex link dispersion management. Dispersion-unmanaged links are much more desired for coherent systems as large CD helps mitigate nonlinear effects through transmission by allowing strong DWDM channel walk-off. Hence, the larger CD value in L-band compared to C-band reduces nonlinear impairments. Previous implementations of C+L-band transmission systems had limited maximum reach compared to C-band-only systems because of the penalty from spectrum tilt due to severe stimulated Raman scattering (SRS) in the transmission fiber as the number of channels increased [4]. However, this scenario was based on 10 Gb/s on/off keying (OOK) systems where signal launch power was much higher. With coherent systems and Raman amplification, the optimal signal launch power is much lower (~ 3-4 dB less compared to OOK). Hence, we believe that the SRS effect would not be a major limitation for in-service upgrade to L-band on an existing terrestrial C-band coherent system.

Given that C+L-band transmission using coherent optics provides a promising approach for capacity scaling, it is important for the entire C+L-band ecosystem to commercially mature to be robust and cost-competitive. With the additional available fiber bandwidth in L-band, the fiber capacity can be more than doubled.

Spectral efficiency (SE), which characterizes the number of bits per second per hertz, is an insightful way to evaluate the efficiency of fiber spectrum utilization for digital communication. One direct approach to increase SE is to apply advanced modulation formats such as higher-order quadrature amplitude modulation (QAM). Figure 4 shows a

few well-known examples of QAM constellations. Each filled circle represents a symbol state. Different constellations carry different numbers of bits per symbol. If all symbols in the constellation are equally utilized, the number of bits carried per symbol is determined by $\log_2(M)$, where M is the number of symbols in the constellation. In the example constellations shown in Fig. 4, binary phase shift keying (BPSK) carries 1 b/symbol, quaternary phase shift keying (QPSK) or 4-QAM carries 2 b/symbol, 8-QAM carries 3 b/symbol, and 16-QAM carries 4 b/symbol. Assuming everything else remains the same, SE (and capacity) increases linearly with the number of bits per symbol. However, it also increases the required OSNR threshold at the receiver as indicated by the reduced Euclidean distance between symbols on the constellation from BPSK to higher-order QAM (less tolerance to noise). This means the increase of fiber capacity is at the cost of reducing transmission reach. It has been demonstrated that a $2\times$ increase of symbols in the constellation will result in about a 3 dB increase in the receiver OSNR threshold [2], which will translate to reach penalty. However, if the transmitter is implemented with a DAC, different modulation formats can be selected through system software configuration. Therefore, a single transponder can cover different link distances in the network by operating in different modes.

Reducing space and power consumption in terrestrial long-haul links is critical as those resources can be limited in certain geographic locations. In addition, it is part of the overall effort to improve data center infrastructure energy efficiency. Therefore, techniques to increase optical reach under any particular modulation format are extremely valuable to reduce the need for frequent optical-electrical-optical (OEO) regeneration. Advanced forward error correction (FEC) and various nonlinear compensation techniques are widely researched to help gain a few additional dB of OSNR budget. High-performance soft-decision (SD) FECs such as turbo product code (TPC), low-density parity check (LDPC) convolutional code, and LDPC code concatenated with hard-decision (HD) FEC all show net coding gain (NCG) > 10 dB. K. Sugihara *et al.* recently demonstrated NCG of 12 dB using a spatially coupled LDPC code concatenated with BCH HD-FEC [5]. The same team also proposed that closer interplay of FEC and nonlinear equalization can be very useful to increase OSNR performance (hence reach), especially at high bit rate [6]. Similar techniques have been demonstrated experimentally for high-SE submarine transmission with a record high capacity-reach product of 203 Pb/s-km [7]. Digital back propagation (DBP), a technique for reversing the nonlinear impairments incurred during optical transmission in the digital domain using DSP, was proposed and extensively studied to perform compensation for various nonlinear impairments in optical fibers [8]. However, the number of DBP stages and gate count required for such computation prevent it from being implemented in the DSP chip within a reasonable power and space footprint. There have been proposals to reduce the complexity of DBP. For example, a perturbation-based algorithm can treat a multi-span transmission link as a single stage, and the complexity

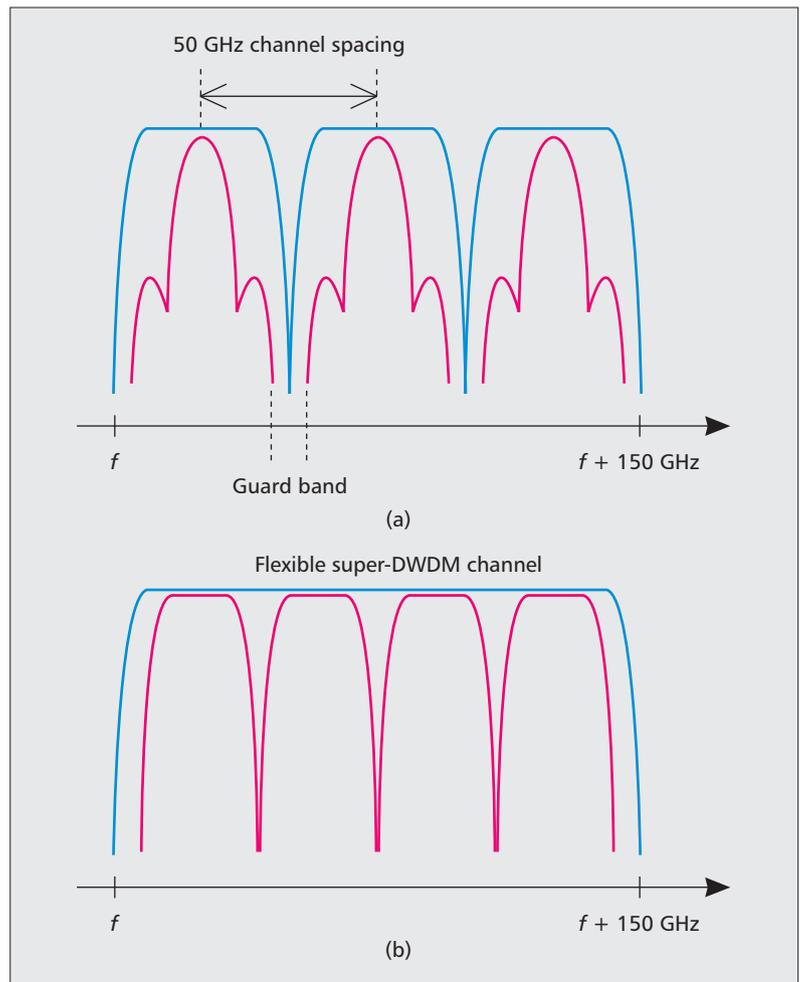


Figure 3. Spectral shaping and flex-grid concept for capacity scaling: a) traditional 50 GHz spaced DWDM channel plan; b) spectrally shaped wavelength channels can be packed tighter within a flexible super-DWDM channel (defined by flex-grid WSS) to allow more capacity within the same fiber bandwidth.

per stage can be further reduced by algorithms such as perturbation back propagation or perturbation pre-distortion [9]. It was shown that the computation can be drastically simplified without degradation of performance. Another interesting approach is correlated back propagation, which correlates the signal powers of neighboring symbols in the time domain to determine the amount of phase shift for self-phase modulation (SPM) compensation together with a much smaller number of DBP stages [10]. This method can achieve the same performance with more than 70 percent reduction in complexity. Another observation worth paying attention to is the interaction between CD and nonlinear effects. As CD alters the intensity profile of the optical symbols in the time domain, it will result in optical pulses with very high peak power, which generates strong nonlinear effects. It was observed that CD pre-compensation can help improve link performance by reducing nonlinear effects [11]. Therefore, it is worth investigating the optimal ratio of pre- and post-compensation of CD to improve the performance of coherent systems further.

In long terrestrial links where regeneration is required, doing full reamplification, reshaping, retiming (3R) OEO regeneration imposes signif-

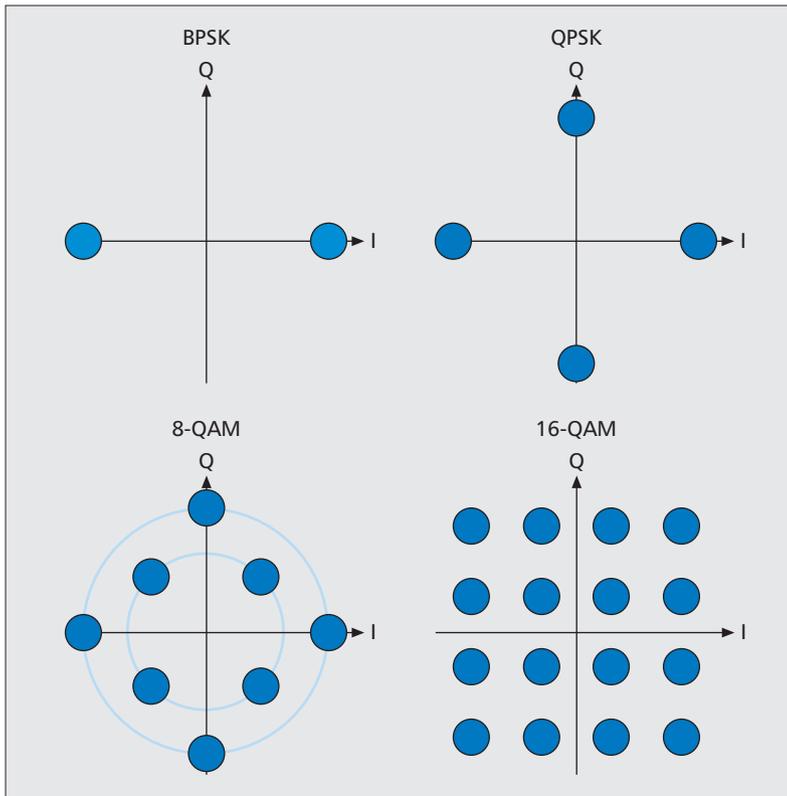


Figure 4. Examples of constellations using both amplitude and phase of the optical field.

icant challenges on the space and power required. In the pre-coherent era, there were many interesting all-optical regeneration ideas to address the space and power issues of OEO regeneration [12]. However, with coherent detection and phase modulation, many of these approaches are no longer applicable since they do not preserve signal phase information. We believe that there is great potential and need for all-optical regeneration techniques to address the space and power challenge faced by high-capacity optical networks. One recent demonstration uses a hybrid optical/digital scheme that combines the signal and its phase-conjugated copy via four-wave mixing (FWM) to cancel out nonlinear noise [13]. However, FWM generates a phase-conjugated idler at a different wavelength, which takes up usable spectrum in the fiber, thus reducing system capacity. Ideally, a viable optical regeneration method should have the following characteristics:

- It can regenerate all WDM channels at once.
- It does not impose unnecessary wavelength conversion.
- It is compatible with ROADM architectures.
- The implementation should be space- and power-efficient.
- It does not reduce the fiber capacity.

OPERATIONAL FLEXIBILITY

In addition to enabling capacity increases (through tighter spectrum packing and higher-order modulation formats), transmitter DAC implementation also significantly increases the flexibility of network operation. For example,

the super-channel architecture needs to be supported by WSS with a fine resolution of 12.5 GHz. With this feature, the channel plan is essentially flexible and opens up rich design options that can be selected based on the link characteristics without introducing any additional hardware variants. For example, super-channels can be created with a different number of wavelength channels depending on traffic demand between a source-destination pair, or conventional 50 GHz spaced channels can coexist with super-channels if the transponders are a mix of DAC and non-DAC implementations.

In a meshed network architecture, ROADMs are deployed in a multi-degree junction node to allow signals to optically express through the node between different fiber degrees in addition to local add/drop. Today, ROADMs are composed of a multi-degree WSS and fixed-grid passive multiplexers/demultiplexers (mux/demux) for local add/drop. The design is directional (or directed), meaning each direction has its own dedicated add/drop mux/demux and transponders/muxponders (Fig. 5a). With a directional design, wavelength contention is avoided and network planning is simplified. However, this approach has a couple of limitations:

- It is cumbersome and error-prone operationally, especially for a junction node with a large number of degrees.
- Transponders are tied to specific fiber degrees and cannot be remotely reconfigured to transmit on different degrees if demands change.

The operational complexity with this architecture can be appreciated by noting that for currently deployed systems, a WSS supports 8 degrees, and each degree can support 88 channels, which results in 704 color-specific channels in a fully loaded junction node. This complexity could be exacerbated in the future when a WSS can support up to 20 degrees.

With coherent detection and super-channels, one can achieve colorless add/drop operation by using power splitters instead of a mux/demux, and utilizing the local oscillator on the coherent receiver to filter out an individual wavelength channel. Moving forward, having a directionless design where all degrees share the same pool of transponders will simplify network planning and operation tremendously (Fig. 5b). However, this can introduce wavelength contention when different degrees use the same wavelength color for add/drop. There are proposals to achieve colorless, directionless, and contentionless (CDC) wavelengths simultaneously, but cost and scalability prevent it from being widely adopted today. With expected technology progression for a more streamlined solution and lower costs, a CDC ROADM is expected to be a viable building block for an operationally scalable optical network in the future.

Variable bit rate per channel by changing the modulation format (enabled by transmitter DAC) is another attractive feature that allows the network capacity to be optimized per link without introducing additional hardware variants. However, with packet and optical layer segregation in current network architectures, it is difficult to truly take advantage of the variable bit rate feature because of the fixed client interface mapping

between the packet and the optical transport devices, as shown in Fig. 6a. For data center networks, the packet processing device (e.g., a router) is usually deployed with an Ethernet medium access control (MAC). IEEE Ethernet standards define fixed-rate MACs at 10, 40, and 100 Gb/s now, and 400 Gb/s in the near future. These rates do not match the optical transport variable line rate supported by different levels of modulation (e.g., 50, 100, 150, and 200 Gb/s). Also, the MAC can only operate at a single static rate without the support of in-service rate change as on the transport side. The incompatibility between Ethernet MAC and optical transport results in inefficient use of transport line side capacity and limits the ability to exploit the value of variable rate transmission. This mismatch can be addressed by tightly integrating the packet layer device with the optical transport device, as shown in Fig. 6b. This allows better interoperability or even consolidation of the Ethernet MAC and optical transport network (OTN) framing schemes.

SOFTWARE-DEFINED PACKET-OPTICAL NETWORK

As capacity increases, the optical transport cost is the dominant contributor to the overall global inter-data-center wide area network (WAN) cost. Therefore, improving efficiency and optimizing the network architecture become exceedingly critical. SDN in the WAN separates the control plane from the data plane, and allows intelligent application software running on the network to exercise greater control and enable network-wide optimization. Many of the aforementioned operational advantages (e.g., real-time change of data rate and DWDM channel planning based on traffic demands) are not practically feasible without SDN in the WAN. A differentiating aspect of SDN relative to the conventional networking paradigm is the decoupling of network hardware (transponders, ROADMs, etc.) from control software (path optimization, data rate selection, etc.). This separation would allow one to choose hardware based on necessary features and software based on protocol and control requirements. With the decoupling, a logically centralized control plane with intelligent software and a global view of the network can be realized, which leads to a more deterministic and fault-tolerant WAN with the potential for a high degree of automation. This SDN paradigm also enables a platform where network operators can use innovative approaches to network optimization and automation that were previously not possible with a traditional system where the control software was embedded in and coupled with a specific hardware implementation.

As shown earlier in Fig. 2, all WAN-layer packet-optical devices communicate with a logically centralized network OS. The nodes operate at the data plane layer (packet forwarding and optical transport layer), while the network OS implements all the control functions such as monitoring, data collection, optimization, and configuration. The network OS is aware of the topology, bandwidth availability, and network redundancy options, and by communicating with the applications the network OS provides a topology-aware

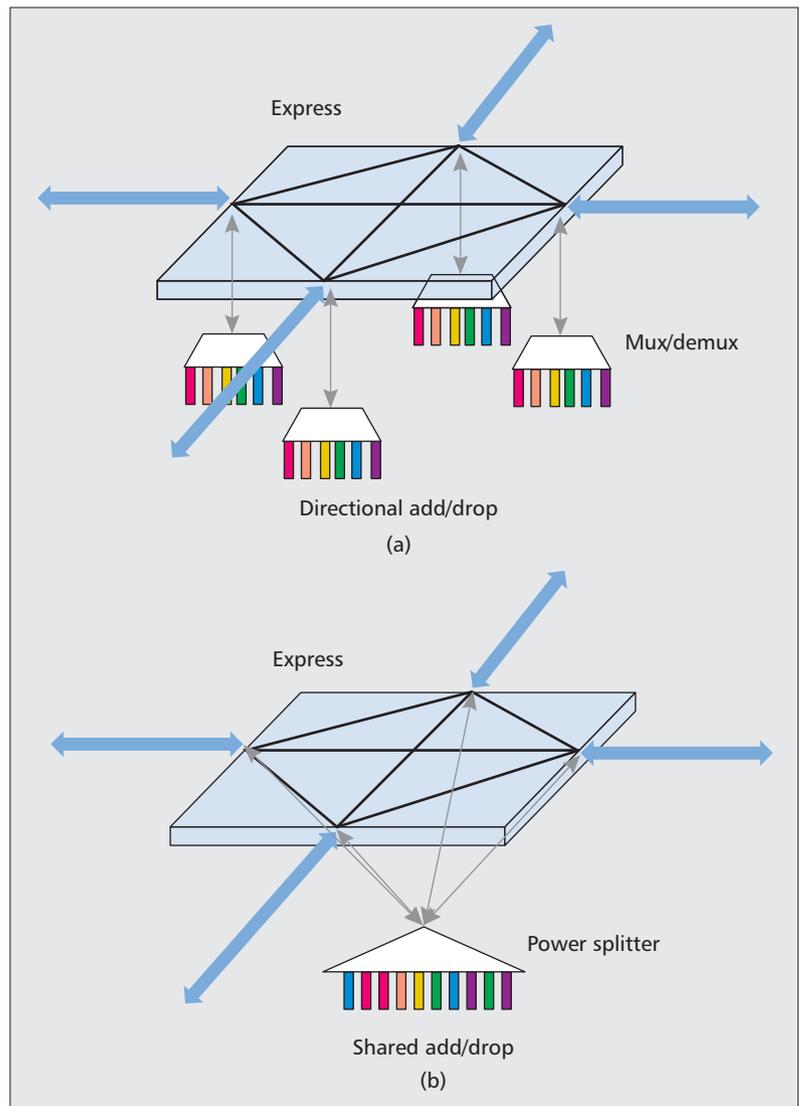


Figure 5. ROADM architecture illustration: a) today's color-specific, directional, and contentionless ROADM; b) the future colorless, directionless, and contentionless ROADM.

network-wide optimized bandwidth allocation. To optimize for the best throughput and performance, applications decide the communication patterns they need among data centers. With the intelligence of SDN, resource optimization, traffic control, and network planning can all be done in real time and in an automated fashion.

To make a software-defined packet-optical network a reality, real-time optical performance monitoring is needed in addition to intelligent control software. This topic of intelligent optical performance monitoring to enable "higher stability, reconfigurability and flexibility in a self-managed optical network" [14] was researched even before the current SDN framework was established. For example, an interesting research idea was to use an artificial neural network (ANN) model trained with parameters derived from asynchronous constellation diagrams to simultaneously identify levels of OSNR, CD, and polarization mode dispersion for QPSK signals [15]. Such schemes match well with, and will likely become useful in, a software-defined packet-optical WAN.

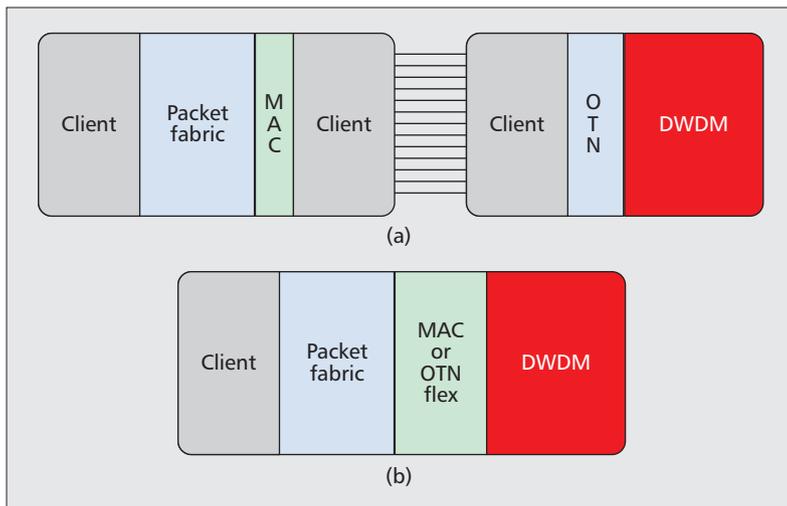


Figure 6. Flexible Ethernet MAC proposal: a) today's fixed Ethernet client-to-OTN DWDM line mapping; b) proposed packet and DWDM integration with flexible client-to-line mapping.

CONCLUSION

Scaling large inter-data-center networks on existing fiber plants, from both the capacity and operational perspectives, is the primary challenge that Google and other large network operators and carriers will face in the coming decade. It is exciting to see various technology advances in digital coherent systems for capacity scaling and reach enhancement. There are research areas such as all-optical regeneration for coherent and PSK systems that have promise for further scaling. In addition, packet-optical integration and SDN enable us to fully unlock the potential of the technology advances on the transmission side. In summary, the continued growth and increasing criticality of inter-data-center networks are likely to be the key drivers for advances in optical communications over the next decade.

REFERENCES

- [1] R.-J. Essiambre and R. W. Tkach, "Capacity Trends and Limit of Optical Communication Networks," *Proc. IEEE*, vol. 100, no. 5, Mar. 2012, pp. 1035–55.
- [2] R. J. Essiambre *et al.*, "Capacity Limit of Optical Fiber Networks," *J. Lightwave Tech.*, vol. 28, no. 4, Feb. 2010, pp. 662–701.
- [3] Y. Emori, K. Tanaka, and S. Namiki, "100 nm Bandwidth Flat-Gain Raman Amplifiers Pumped and Gain-Equalized by 12-Wavelength-Channel WDM Laser Diode Unit," *Electron. Lett.*, vol. 35, no. 16, Aug. 1999, pp. 1355–56.
- [4] C. Lam *et al.*, "Fiber Optic Communication Technologies: What's Needed for Data Center Network Operations," *IEEE Commun. Mag.*, vol. 48, no. 7, July 2010, pp. 32–39.
- [5] K. Sugihara *et al.*, "A Spatially-Coupled Type LDPC Code with an NCG of 12 dB for Optical Transmission Beyond 100 Gb/s," *OFC 2013*, paper OM2B.4, Anaheim, CA, 2013.
- [6] T. Sugihara *et al.*, "Collaborative Signal Processing with FEC in Digital Coherent Systems," *OFC 2013*, paper OM2B.3, Anaheim, CA, 2013.
- [7] H. Zhang *et al.*, "30.58 Tb/s Transmission over 7230 km Using PDM Half 4D-16QAM Coded Modulation with 6.1 b/s/Hz Spectral Efficiency," *OFC 2013*, paper OTu2B.3, Anaheim, CA, 2013.
- [8] E. Ip and J. M. Kahn, "Compensation of Dispersion and Nonlinear Impairments Using Digital Backpropagation," *J. Lightwave Tech.*, vol. 26, no. 20, Oct. 2008, pp. 3416–25.
- [9] Z. Tao *et al.*, "Complexity-Reduced Digital Nonlinear Compensation for Coherent Optical System," *Proc. SPIE*, vol. 8647, Jan. 2013.

- [10] L. Li *et al.*, "Implementation Efficient Nonlinear Equalizer Based on Correlated Digital Backpropagation," *OFC 2011*, paper OWW3, Los Angeles, CA, 2011.
- [11] X. Wei, "Power-Weighted Dispersion Distribution Function for Characterizing Nonlinear Properties of Long-haul Optical Transmission Links," *Optics Letters*, vol. 31, no. 17, Sept. 2006, pp. 2544–46.
- [12] E. Ciaramella, "Wavelength Conversion and All-Optical Regeneration: Achievements and Open Issues," *J. Lightwave Tech.*, vol. 30, no. 4, Feb. 2012, pp. 572–82.
- [13] Y. Tian *et al.*, "112-Gb/s DP-QPSK Transmission over 7,860-km DMF Using Phase-Conjugated Copy and Digital Phase-Sensitive Boosting with Enhanced Noise and Nonlinearity Tolerance," *OFC 2013*, paper OTu2B.5, Anaheim, CA, 2013.
- [14] Z. Pan, C. Yu, and A. E. Willner, "Optical Performance Monitoring for the Next Generation Optical Communication Networks," *Optical Fiber Tech.*, vol. 16, no. 1, Jan. 2010, pp. 20–45.
- [15] J. A. Jargon *et al.*, "Optical Performance Monitoring of QPSK Data Channels by Use of Neural Networks Trained with Parameters Derived from Asynchronous Constellation Diagrams," *Optics Express*, vol. 18, no. 5, Mar. 2010, pp. 4931–38.

BIOGRAPHIES

XIAOXUE ZHAO (wendyzhao@google.com) joined Google as a networking hardware engineer in 2008, where she worked on high-speed system design and optical interconnect technologies. She joined the network architecture group at Google in 2009, and is currently a staff network architect. Her recent focuses include data center network architecture, metro and long-haul DWDM system design, as well as DWDM mesh network planning, automation, and optimization. She has numerous journal and conference publications, and holds more than 10 patents in optical devices, subsystems, and data center architecture. She received her B.S. degree from Peking University in China, and her Ph.D. degree in electrical engineering from the University of California at Berkeley.

VIJAY VUSTRIKALA currently is an optical network architect at Google, where he is focused on developing solutions for scaling Google's long-haul and metro optical networks. Prior to Google, he was at Infinera, Motorola, and Sycamore Networks in senior marketing, business development, and architecture roles and covered applications ranging from core to access networks. He has published extensively, spoken at numerous industry events, and holds several patents in optical devices and systems. He obtained a Ph.D. from the University of Maryland, College Park in the area of optoelectronic integration, and a B.S.E.E. from the Indian Institute of Technology (IIT), Madras.

BIKASH KOLEY is the principal architect and manager of the Network Architecture team at Google. He is responsible for scaling optimization and automation of the network layer of Google's cloud infrastructure. He also oversees the network technology road map at Google in order to support all present and future Google services. Prior to Google, he was CTO of Qstreams Networks, a company he co-founded. He also spent several years at Ciena Corporation in various technical roles developing DWDM and Ethernet technologies. He is a frequent speaker at conferences and industry forums, and is an active participant in various networking standard bodies. He has received eight patents related to various optical and networking technologies. He received a B.Tech. from IIT, and M.S. and Ph.D. degrees from the University of Maryland at College Park, all in electrical engineering.

VALEY KAMALOV is a staff optical transport engineer at Google, where he is focused on the performance of the optical network. Prior to joining Google, he spent seven years at Nokia Siemens Networks designing optical networks. He received B.S. and Ph.D. degrees from Moscow University in quantum electronics, and his Sc.D. degree in chemical physics from the Russian Academy of Sciences.

TAD HOFMEISTER has been a network architect at Google since 2011, with a focus on metro and long-haul DWDM networking. Prior to Google, he was an architect, system engineer, and hardware designer for several optical transport and packet processing companies including Ciena, Matisse Networks, OpVista, and Applied Signal Technology. He earned M.S. and Ph.D. degrees in electrical engineering from Stanford University, and B.S. degrees from Columbia University and Bates College.