# Elastic Optical Networking: A New Dawn for the Optical Layer?

*Ori Gerstel, Cisco, Israel*
*Masahiko Jinno, NTT, Japan*
*Andrew Lord, BT, United Kingdom*
*S. J. Ben Yoo, University of California, Davis, United States*

## ABSTRACT

Optical networks are undergoing significant changes, fueled by the exponential growth of traffic due to multimedia services and by the increased uncertainty in predicting the sources of this traffic due to the ever changing models of content providers over the Internet. The change has already begun: simple on-off modulation of signals, which was adequate for bit rates up to 10 Gb/s, has given way to much more sophisticated modulation schemes for 100 Gb/s and beyond. The next bottleneck is the 10-year-old division of the optical spectrum into a fixed “wavelength grid,” which will no longer work for 400 Gb/s and above, heralding the need for a more flexible grid. Once both transceivers and switches become flexible, a whole new elastic optical networking paradigm is born. In this article we describe the drivers, building blocks, architecture, and enabling technologies for this new paradigm, as well as early standardization efforts.

## INTRODUCTION

There continues to be ever increasing demand for bandwidth, to the point where, just as 10 Gb/s technology has reached maturity, service providers (SPs) are already installing higher bit rates, including 40 Gb/s and now 100 Gb/s per wavelength.[1] The 50 GHz International Telecommunication Union (ITU) wavelength grid divides the relevant optical spectrum range of 1530–1565 nm (the so-called C-band) into fixed 50 GHz spectrum slots, but it is likely that bit rates greater than 100 Gb/s will not fit into this scheme. This article examines the case for replacing this 10-year-old standard with a more flexible grid paradigm to meet the needs of future bandwidth demands.

Even if sufficiently broad spectrum is available, high-data-rate signals become increasingly difficult to transmit over long distances at high spectral efficiency (i.e., how much data rate can be supported for a limited spectral bandwidth) [1]; hence, it becomes beneficial for transceivers to be able to maximize spectral efficiency by adapting to the actual conditions of the network and data rate for each given traffic demand. In addition to the need to enhance the spectral efficiency, an increasing influence from large content providers, newly constructed data centers, and evolving peering relationships between providers are propelling the uncertainty and heterogeneity of the demands (the disparity between small and large demands) across the network.

To properly address this challenge, one needs flexible and adaptive networks equipped with flexible transceivers and network elements that can adapt to the actual traffic needs. Fortunately, the same technologies that are being considered for achieving very high bit rates at 100 Gb/s and beyond can also provide this added flexibility. The combination of adaptive transceivers, a flexible grid, and intelligent client nodes enables a new “elastic” networking paradigm [2], allowing SPs to address the increasing needs of the network without frequently overhauling it.

100-Gb/s-based transmission systems have been commercialized in the last two years. Since they are compatible with the 50 GHz ITU grid already deployed, the need for replacing the grid did not arise. Both the telecom and datacom industries are now considering a standard transmission data rate beyond 100 Gb/s, and 400 Gb/s is receiving a lot of attention. Unfortunately, the spectral width occupied by 400 Gb/s at standard modulation formats is too broad to fit in the 50 GHz ITU grid, and forcing it to fit by adopting a higher spectral efficiency modulation format would only allow short transmission distances. Figure 1a shows an existing ITU grid (top) vs. a flexible grid (bottom). The fixed grid does not support bit rates of 400 Gb/s and 1 Tb/s at standard modulation formats, as they overlap with at least one 50 GHz grid boundary.

Figure 1b shows different bit rate demands interconnecting node A with B, C, and D. The component that switches the channels arriving at B toward C or D is called a reconfigurable optical add-drop multiplexer (ROADM). If this device conformed to the ITU grid, it would not be able to switch the broader spectrum channels; as Fig. 1a shows, the optical spectrum coinciding with an ITU grid boundary (marked in black) will not be transmitted through the ROADM.

---

[1] *The optical layer typically uses a technology called dense wavelength division multiplexing (DWDM) to carry 40–80 such wavelengths on a fiber pair.*

0163-6804/12/$25.00

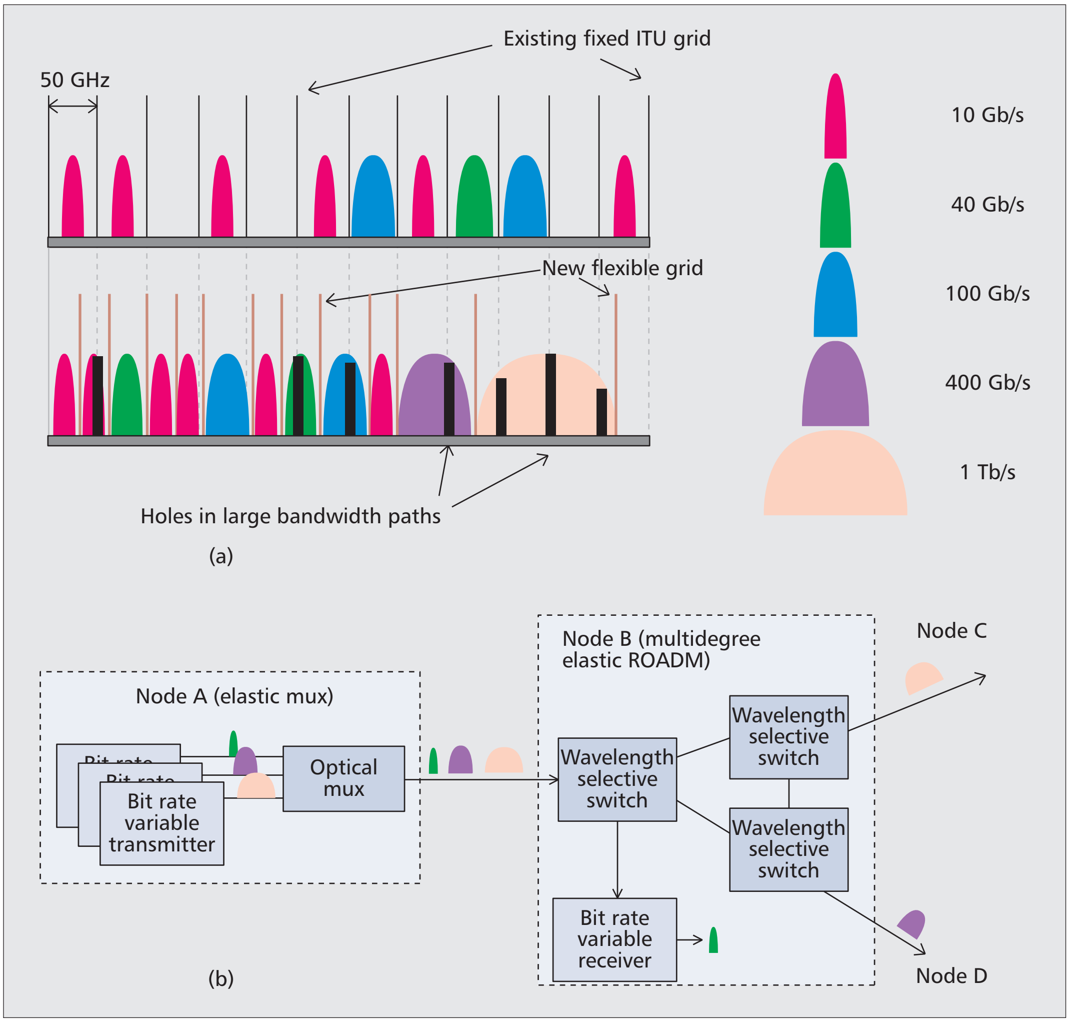

**Figure 1.** *a) Spectral widths of different bit rates relative to the ITU grid (for a given modulation format); b) three demands of bit rates 40 Gb/s, 1 Tb/s, and 400 Gb/s connect node A to B, C, and D, respectively.*

*Since 100Gb/s and higher bit rates must be supported by the same network, it makes sense to "properly size" the spectrum for each demand based on its bit rate and the transmission distance (instead of forcing all demands to use more spectrum).*

Therefore, in order to build a flexible network, a new kind of ROADM is required that allows flexible spectrum to be switched from the input to the output ports.

Figure 1 shows several features and will help us define some important terms. This new approach is called elastic optical networking (EON).[2] The term *elastic* refers to two key properties:

- The optical spectrum can be divided up flexibly.
- The transceivers can generate elastic optical paths (EOPs); that is, paths with variable bit rates.

These new transceivers are called bandwidth variable transceivers (BVTs). The main drivers for developing the EON paradigm are now listed.

**Support for 400 Gb/s, 1 Tb/s, and other high bit rate demands.** These can be carried on a fixed grid network by demultiplexing the demand to smaller ones such as 100 or 200 Gb/s, which can still fit in the fixed grid (this technique is known as inverse multiplexing, and although it will work, it will use up the spectrum more quickly than if the demand is carried in one contiguous EOP).

**Disparate bandwidth needs.** Since 100 Gb/s and higher bit rates must be supported by the same network, it makes sense to "properly size" the spectrum for each demand based on its bit rate and the transmission distance (instead of forcing all demands to use more spectrum).

**Tighter channel spacing.** Coherent signal detection, which is used in 100 Gb/s transmission and beyond [3], allows for closer spacing of channels. This technique is called Nyquist DWDM (or superchannels). Figure 1a shows how an EON allows channels to be spaced closer together, freeing up spectrum for other demands.

**Reach vs. spectral efficiency trade-off.** If an EOP is short in distance, the BVT can adjust to a modulation format occupying less optical spectrum, and the connection will still perform error-free due to the reduced impairments (e.g., added noise from optical amplifiers) on shorter-distance paths.

**Dynamic networking.** The optical layer can now respond directly to variable bandwidth demands from the client layer. An example would be IP over EON, in which the BVTs are adjusting their bandwidth in line with the IP layer demands.

---

[2] *Note that this concept has sometimes been called* flexgrid *or* flexible spectrum*. Other terms used in this context are* programmable transceivers *and* software defined optics*.*

*From a simple efficiency point of view, there is a significant advantage in building EONs, assuming there is sufficient traffic to merit it. Adding this to the other main drivers highlighted here provides sufficient benefits to undertake serious research into EONs.*

| Demand bit rate (Gb/s) | Modulation format | Channel bandwidth (GHz) | Fixed grid solution | Efficiency increase for EON |
|---|---|---|---|---|
| 40 | DP-QPSK | 25+10 | 1 50 GHz channel | 35 GHz vs. 50 = 43% |
| 100 | DP-QPSK | 37.5+10 | 1 50 GHz channel | 47.5 GHz vs. 50 = 5% |
| 100 | DP-16QAM | 25+10 | 1 50 GHz channel | 35 GHz vs. 50 = 43% |
| 400 | DP-QPSK | 75+10 | 4 100 Gb/s in 4 50 GHz channels | 85 GHz vs. 200 = 135% |
| 400 | DP-16QAM | 75+10 | 2 200Gb/s in 2 50 GHz channels | 85 GHz vs. 100 = 17% |
| 1000 | DP-QPSK | 190+10 | 10 100G in 10 50 GHz channels | 200 GHz vs. 500 = 150% |
| 1000 | DP-16QAM | 190+10 | 5 200Gb/s in 5 50 GHz channels | 200 GHz vs. 250 = 25% |

**Table 1.** *Efficiency improvement for flexible spectrum over a point-to-point link, assuming a 50 GHz grid for fixed DWDM and 10 GHz channel guard band and superchannels for EONs.*

The development of EON will require innovations in both hardware and software. New components will need to be developed, and they are often more complex than their fixed grid counterparts. Also challenging will be the control and management of the network, including setting up EOPs. To make this development worthwhile, it is important to quantify the benefits.

Table 1 compares fixed grid and EON for a 300 km point-to-point link,[3] and shows that the advantages of an EON over a fixed grid can be very significant, even up to 150 percent for 1 Tb/s. As transmission techniques improve, it is possible that the channel bandwidths will drop further, thus giving even larger benefits for EONs.

EONs always deliver efficiency improvements, but they are significantly different depending on the exact traffic scenario. It should be noted that 200 Gb/s dual polarization 16-quadrature amplitude modulation (QAM) can fit into the 50 GHz grid, and this makes it the most efficient fixed grid alternative (assuming inverse multiplexing). Similar modeling to Table 1 carried out over a network, rather than a point-to-point link, shows the effect of amplifying the efficiency benefits of EON due to the increased optical path blockage arising from multiple inversed multiplexed channels. This has the effect of increasing efficiency gains to well over 50 percent in many cases.

Thus, from a simple efficiency point of view, there is a significant advantage in building EONs, assuming there is sufficient traffic to merit it. Adding this to the other main drivers highlighted above provides sufficient benefits to undertake serious research into EONs. That research is currently in full flow and this article will summarize its main aspects.

## SOME NOVEL EON CONCEPTS

To gain a better understanding of EON and the flexibilities it provides, let us review several concepts that are not typically part of fixed DWDM networks.

### THERE'S MORE THAN ONE WAY TO SKIN A DEMAND

In fixed DWDM networks, there is typically one way to implement a given demand: the wavelength bit rate is fixed, the optical reach is fixed, and the spectrum is fixed. So the demand can occupy less than a full wavelength (wasting capacity, like demand C in Fig. 2c) or more than one (requiring multiple wavelengths, as in Fig. 2a). An EON allows for multiple choices when implementing a demand:

- A given demand can be assigned a modulation format that gives sufficient performance to reach the required distance, while minimizing the spectral bandwidth occupied by the optical path; see demands D and E in Fig. 2d.
- Today the ratio between the amount of forward error correction (FEC) and payload is fixed, but it could be made adaptive in EONs to enable greater distances to be reached when the required bandwidth is lower; see [4] for details.
- Whenever a connection passes through an ROADM, the ROADM acts as a filter that reduces the optical bandwidth for the channel. When this happens over and over, the resulting bandwidth may be too narrow, affecting the quality of the signal and limiting the reach. With EONs, the spectrum allocated for longer EOPs can be increased to account for such bandwidth narrowing, increasing the number of ROADMs the connection can go through.

If the demand is too large to be handled by a single optical channel, "superchannels" can be constructed; these contain multiple very closely spaced channels, which traverse the network as a single entity, but can be demultiplexed at the receiver. Fig. 2b illustrates this (compared to the fixed network case shown in Fig. 2a). This concept should not be confused with "virtual concatenation," which allows a single demand to span several independent channels, in which case guard bands will still be needed. Thus, it is possible that a single demand will be virtually concatenated over multiple superchannels.

---

[3] *These results are reach sensitive: higher order modulation formats will not be feasible when the required reach is more than a few hundred kilometers.*

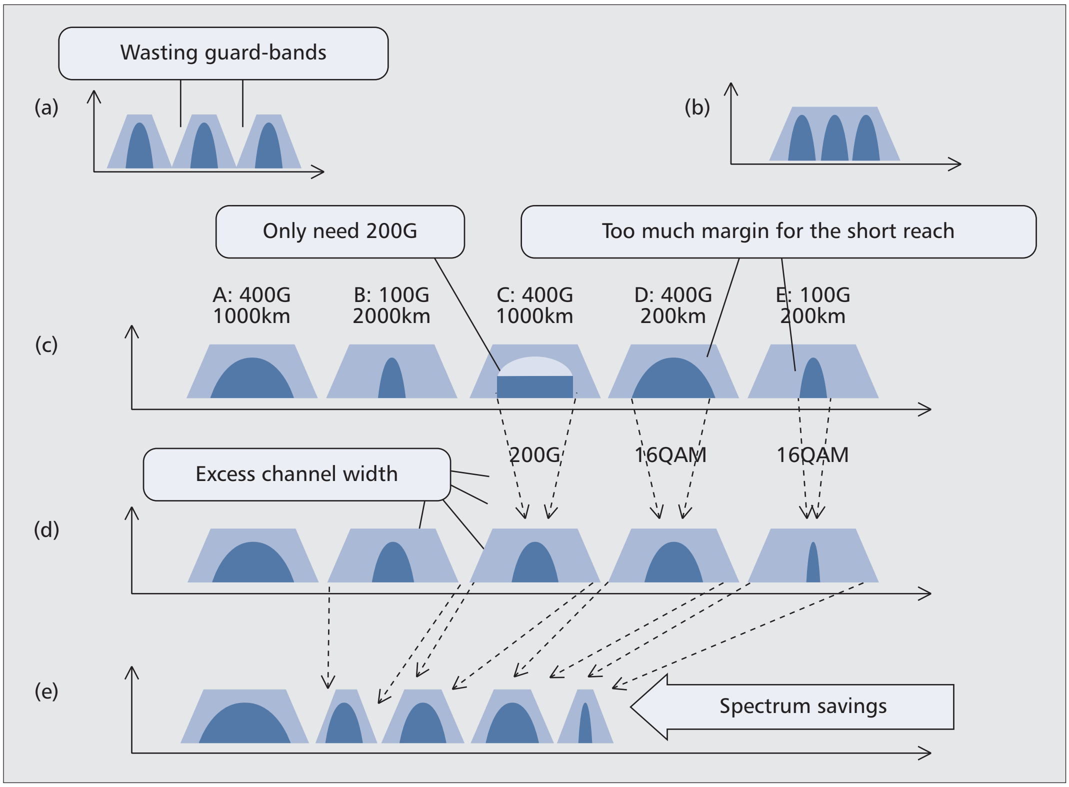

**Figure 2.** *a) Fixed grid maintains strict guard bands between optical paths that implement a 300 Gb/s demand; b) the channels for the demand can be grouped tightly into a superchannel and transported as one entity; c) five demands and their spectrum needs on a 100 GHz fixed grid, assuming quaternary phase shift keying (QPSK) modulation; d) the same demands, with adaptive modulation optimized for the required bit rate and reach; e) the same demands, with additional flexible spectrum.*

*Since the spectrum slice in an EON varies, it will create fragmentation in the spectrum over time, much like a computer hard disk becomes fragmented. A related but different problem is the need to grow superchannels over time, causing them to use up free spectrum between them and at some point limiting their ability to expand.*

## SLICEABLE TRANSCEIVER

This concept is about "slicing" a single BVT into several "virtual transceivers" that serve separate EOPs (in Fig. 3b, a sliceable 400 Gb/s transceiver is sliced into three EOPs: 100 Gb/s, 100 Gb/s and 200 Gb/s). This flexibility is key to the economic justification of EONs since it is hard to justify "wasting" a 400G BVT on, say, a 150 Gb/s EOP alone (as in Fig. 3a). One could use instead a standard rigid 400 Gb/s transceiver and electrical subwavelength grooming to fill up the remaining 250 Gb/s, but this introduces another layer and eliminates some of the cost gains of EONs [5]. This flexibility seems feasible in next generation BVT designs, as is the ability to use a different modulation format for each virtual transceiver. While such a feature may challenge the transceiver design, the added cost seems fairly modest.

## FLEXIBLE CLIENT INTERCONNECT

Since different EOPs may carry different payload bit rates, one key question is how to connect the client (e.g., an IP router) to a BVT. From a protocol perspective, extensions of the optical transport network (OTN) and Ethernet standards seem appropriate, similar to the concept of ODUflex in OTNs. This client interface is further challenged if the BVT is sliceable, and the bandwidth for each virtual interface must expand/contract over time. In this case, the interface between the client and the BVT becomes a flexible pool of channels that can be grouped in different ways. This also affects the client line card architecture. A simpler option is to integrate the BVT into the client box, and avoid the need for a client interface altogether.

## HITLESS SPECTRUM REALLOCATION

Since the spectrum slice in EON varies, it will create fragmentation in the spectrum over time, much like a computer hard disk becomes fragmented. A related but different problem is the need to grow superchannels over time, causing them to use up free spectrum between them and at some point limiting their ability to expand. These two problems will require periodically reallocating the spectrum in a dynamic fashion. The required mechanism will remove stranded fragments of spectrum between EOPs and redistribute them to allow for further growth. In the context of a superchannel comprising different channels, this hitless spectrum shift can be done by adding an adjacent channel to the EOP in the desired direction of the shift and remapping the data sent over the channel that will be released to the new channel using either a transport mechanism called the Link Capacity Adjustment Scheme (LCAS) or an IP layer mechanism called link bundling. Both mechanisms allow for such changes to occur in a hitless manner. Note that this mechanism requires coordinating the changes in spectrum at the client device, the transmitter, and the optical switching device, but this can be done slowly and hence seems achievable.

One of the main challenges in an EON is how to determine the minimum necessary spectral resources and adaptively allocate them to an optical channel with minimum guard band assigned between channels.

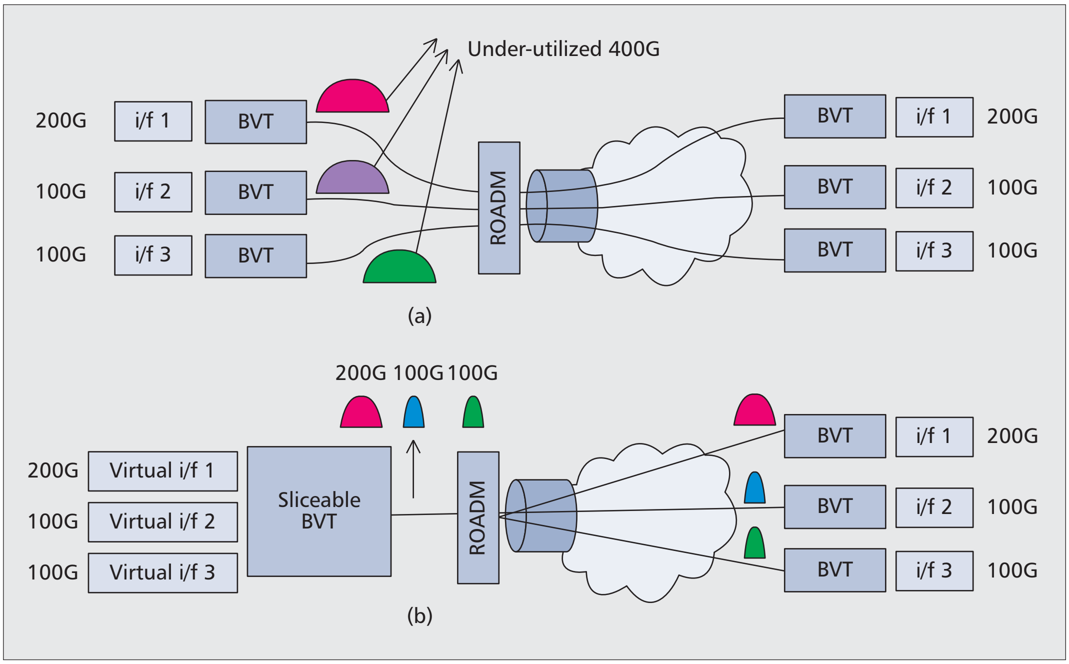

**Figure 3.** *Support for three-subwavelength demand using fixed or sliceable transceivers (the ROADMs are assumed to support EON): a) fixed transceivers; b) sliceable transceivers.*

### MANAGING A CONNECTION PER DEMAND INSTEAD OF MANAGING WAVELENGTH

In fixed DWDM networks, each wavelength must be managed individually: it is switched individually through ROADMs, consumes the same amount of spectrum whether it is part of a larger demand or not, and appears as a separate entity to management systems. If multiple wavelengths are used to realize a larger demand, their bundling occurs at the client layer without really affecting the network. This is not the case with EONs, where most demands will be mapped 1:1 to a single EOP: if the demand is smaller than a BVT, the BVT will be (ideally) sliced to carry it independently, and if the demand is larger than a BVT, multiple channels will be bundled into a single superchannel that implements the demand. Either way, the demand will be switched as one entity in ROADMs and managed as a single entity in management systems. A side effect of this is that the number of managed entities is dramatically reduced, at least for large IP networks, in which a large number of channels are routinely bundled. This has a positive impact on the scale of ROADMs: no longer does an ROADM have to add and drop many dozens of separate connections; instead, it will only need to add/drop a handful of EOPs (typically the nodal degree of a large IP core router is just 2–5, which means that the number of EOPs for a dual core router site will be typically less than 10).

## FROM BUILDING BLOCKS TO ARCHITECTURE

The EON allows for very close mapping between the client and transport layers. This is especially true for a client layer that consumes significant bandwidth like an IP core network. The following description provides a vision for a fully automated EON, based on an intelligent impairment-aware control plane and clients that have the intelligence to exploit a configurable optical network.

- Every client link is mapped to a single EOP, without an extra layer of multiplexing. This mapping is resource-conserving, assuming sliceable BVTs that can serve multiple client demands.
- When a new link is needed, an EON considers the best implementation of the request, in terms of modulation format, FEC, and spectrum, yielding the lowest-cost solution. This decision is based on adaptive behavior that adjusts itself based on actual link conditions and control plane knowledge of impairments in the network,[4] while still taking into account margins for network aging.
- When the link requires more (or less) bandwidth, the EOP is extended (or contracted) hitlessly in the most efficient way to accommodate the demand.
- The operator need not have good visibility of the traffic forecast and the link conditions, but just needs to provide sufficient BVTs per site and connect them in a flexible manner to the clients (if they are not already integrated into the clients). Since these transceivers are sliceable, they can be used effectively even for client links that do not require significant amount of bandwidth.
- When a failure occurs, some client links will carry the load from the failed links. Their bandwidth should be extended automatically. Some of the failed links may be rerouted automatically in the EON, by exploiting possibly different modulation, FEC and spectrum to overcome the longer reach of a protection path. This may imply less available bandwidth at the client layer, but if it is intelligent enough, it could choose to drop low-priority traffic or reroute some of the

---

[4] *By contrast, today's DWDM systems have little knowledge about the actual link characteristics. They rely on laborious manual measurements of link characteristics and provisioning this data into offline planning tools. This approach does not allow for adjustments over time and must assume inaccurate data. It therefore requires significant margins in assessing link quality, resulting in more expensive networks.*

traffic over other paths (via traffic engineering).

- A similar behavior occurs when transmission conditions change due to fiber plant aging. Each EOP will extend itself to accommodate the new conditions. Again, such behavior will benefit greatly from an intelligent client that is able to move traffic over less busy routes.
- Periodically, there may be a need for the above described hitless spectrum reallocation process.

In such a network, the planning process is significantly reduced: no longer do you need to measure the fiber plant, enter the results into the planning tool, provide an accurate forecast on traffic growth, select different transponders to satisfy different connection reaches, and repeat the process when conditions change. Instead, the network will adjust itself to optimally meet the demands with the given transmission conditions, and readjust as needed if demands or conditions change.

## SPECTRUM ALLOCATION IN A STATIC AND DYNAMIC NETWORK

One of the main challenges in an EON is how to determine the minimum necessary spectral resources and adaptively allocate them to an optical channel with minimum guard band assigned between channels. The minimum necessary spectrum is determined by the conditions along the optical path that result in the delivered signal-to-noise ratio (SNR), and linear and nonlinear distortions, while guaranteeing the required data rate and optical reach. This problem is exacerbated in a dynamic environment, where connection must expand/contract without affecting traffic.

The problem of calculating a route and wavelength for an optical channel in conventional transparent wavelength-routed optical networks is called the routing and wavelength assignment (RWA) problem maintaining the same wavelength along the path. Adaptive spectral allocation in EONs introduces another constraint to wavelength-continuity in terms of the available spectrum on each fiber link along the route. This is a more general routing and *spectrum* assignment (RSA) problem.

In order to explain how to determine the route, spectrum range, and modulation format for an optical channel, it is useful to define some terminology. The frequency range an optical channel is allowed to occupy is defined as a frequency slot (FS). The full width of the FS can be flexibly adjusted. From a practical viewpoint to simplify the network design, it is useful to quantize the usable fiber spectrum resources into contiguous frequency units with an appropriate width of, for example, 12.5 GHz. This is a building block for the FS, which is called a frequency slot unit (FSU) in this article. The RSA problem can be divided into two stages. The first stage is to create a route list with a sufficient number of free contiguous FSUs from source to destination (*corridor width*), and determine the modulation format for each source-destination pair at a given bit rate; see Fig. 4a for more details. The second stage allocates contiguous FSUs to the route as shown in Fig. 4b.

There are several proposed algorithms to select contiguous FSUs from an unused spectral resource pool to minimize spectrum fragmentation and to retain as much as possible contiguous spectrum for future utilization. The simplest one is the “First Fit” algorithm, which searches the available contiguous FSUs are from lower to higher-numbered FSU, thus favoring the lowest available contiguous FSUs.

When considering a dynamic network, it is important to leave sufficient available spectrum around a demand, to allow it to grow either up or down in spectrum. However, leaving such room for growth impacts the promised savings of the superchannel concept. While some solutions have been suggested, they only serve to delay the need to reallocate the spectrum at a future point in time.

## ENABLING TECHNOLOGIES

Key enabling technologies for the EON are flexible bandwidth transmitters and receivers (BTVs) that can scale up to terahertz bandwidth and beyond, and flexible spectrum selective switches (flex WSSs) that can multiplex and switch variable spectral bands.

### FLEXIBLE SPECTRUM SELECTIVE SWITCHES

An ROADM is typically constructed out of several interconnected WSS devices and amplifiers (Fig. 1b). While mature WSS technologies were specific to the ITU grid, newer WSS technologies have been productized recently and allow switching almost arbitrary spectrum slices(in 3.125~6.250 GHz steps), enabling elastic ROADMs. These devices are based on one of several technologies: optical micro electromechanical systems (MEMS), liquid crystals on silicon (LICOS), or silica planar lightwave circuits (PLCs) [6]. Due to the low-cost overhead for these devices, we expect them to become common in next-generation ROADMs, and to be deployed irrespective of whether the larger EON vision will materialize, simply because they provide assurance that the network will be robust against future uncertainty without paying a significant premium.

*An ROADM is typically constructed out of several interconnected WSS devices and amplifiers. While mature WSS technologies were specific to the ITU grid, newer WSS technologies have been productized recently and allow switching almost arbitrary spectrum slices (in 3.125~6.250 GHz steps), enabling elastic ROADMs.*

### TRANSPONDER (BVT)

Multicarrier solutions such as coherent wavelength-division multiplexing (CoWDM) [7], coherent optical orthogonal frequency-division multiplexing (CO-OFDM) [8], Nyquist-WDM, as well as dynamic optical arbitrary waveform generation (OAWG) [9] have been proposed as possible transponder implementations for EONs. These solutions rely on the generation of many low-speed subcarriers to form broadband data waveforms using lower-speed modulators so that terabit-per-second data can be generated using lower speed electronics at < 40 Gbaud.

Despite the similarity of multicarrier signal generation between CO-OFDM, CoWDM, Nyquist-WDM, and OFDM transmitters, there are fundamental differences with respect to their operation principles and capabilities. Figure 5a shows the spectrum of multibanded CO-OFDM [8]. Here, each modulator generates many low-bandwidth ($\Delta f$) subcarriers to form each band.

*In OAWG, the coherent combination of many spectral slices generated in parallel enables the creation of a continuous output spectrum. In this case, arbitrary bandwidth, single- and multi-carrier channels can be generated, in which each channel can be in a different modulation format.*

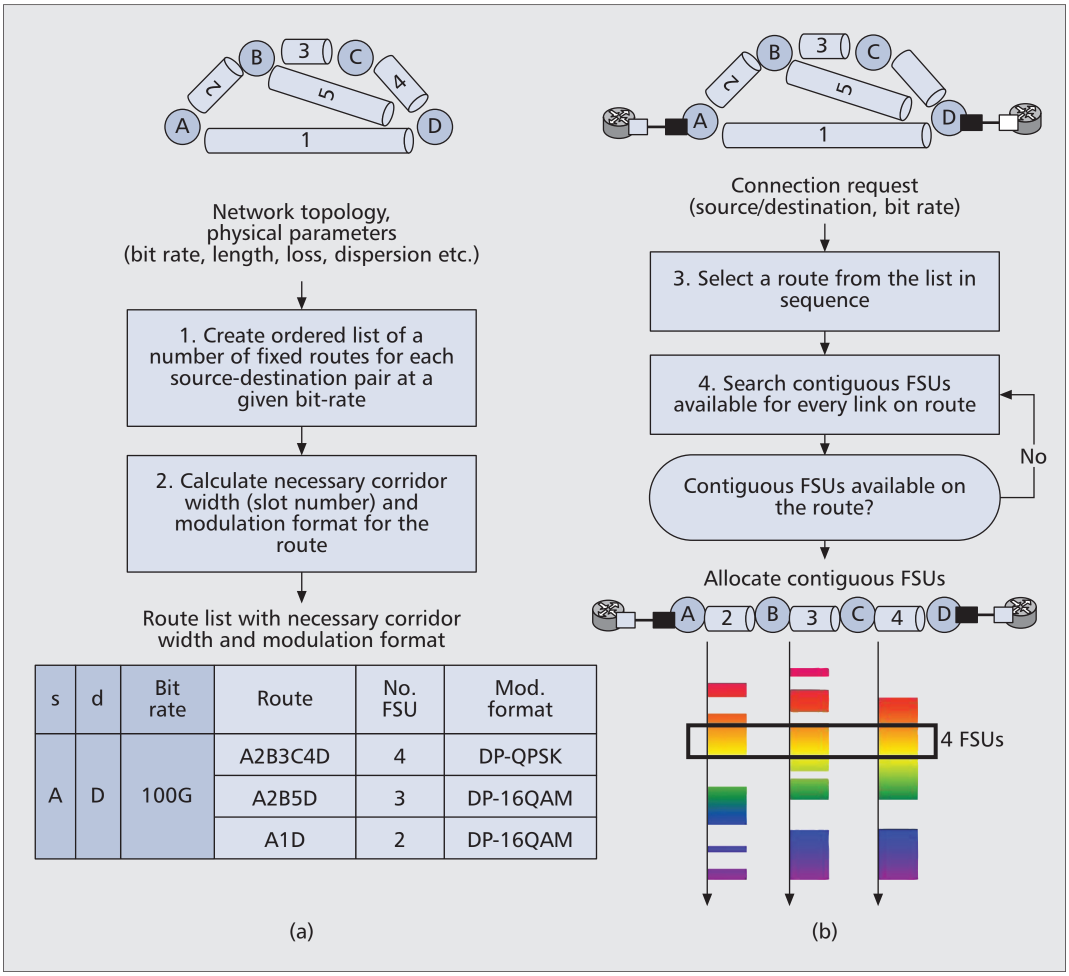

| s | d | Bit rate | Route | No. FSU | Mod. format |
|---|---|---|---|---|---|
| A | D | 100G | A2B3C4D | 4 | DP-QPSK |
| | | | A2B5D | 3 | DP-16QAM |
| | | | A1D | 2 | DP-16QAM |

**Figure 4.** *The spectrum assignment process: a) creating a route list with necessary corridor width and modulation format; b) allocating contiguous FSUs on the route.*

Orthogonality is maintained between bands by ensuring a spacing of $\Delta fG = m \times \Delta f$ between the outside subcarriers of adjacent bands for integer values of $m$. Figure 5b shows a CoWDM spectrum that consists of several subcarriers. To maintain orthogonality between subcarriers, the CoWDM subcarrier symbol rate is set equal to the subcarrier frequency spacing. This restricts each modulator to generating only an integer number of subcarriers. Additionally, for both OFDM and CoWDM, the chromatic dispersion tolerance scales with the subcarrier data rate instead of the total bit rate [8]. However, the use of lower-bandwidth subcarriers (~100 MHz) in OFDM causes high peak-to-average power ratios, which increases susceptibility to nonlinear impairments. The larger subcarrier bandwidths (~10–40 GHz) typically used in CoWDM have a lower peak-to-average-power ratio with comparable performance to isolated single-carrier systems [7].

Nyquist-WDM attempts to minimize the spectral utilization of each channel and reduce the spectral guard bands required between WDM channels generated from independent lasers. Using aggressive optical prefiltering with spectral shape approaching that of a Nyquist filter with a square spectrum minimizes the channel bandwidth to a value equal to the channel baud rate. The channels are then packed closely together such that the subcarrier spacing is equal to or slightly larger than the baud rate. However, significant power penalties arise from setting the channel frequency spacing equal to the baud rate, which requires FEC to achieve error-free performance. Hence, a trade-off exists between spectral efficiency and intercarrier interference [10].

In OAWG, the coherent combination of many spectral slices generated in parallel enables the creation of a continuous output spectrum [9]. In this case, arbitrary-bandwidth single- and multicarrier channels can be generated, in which each channel can be in a different modulation format (Fig. 5c). The versatility of this signal generation technique enables customization of generated waveforms over the total operational bandwidth of the transmitter. This enables avoiding large peak-to-average-power ratios and incorporating precompensation for impairments such as chromatic dispersion. OAWG also removes the restriction that the modulator bandwidth must be a multiple or sub-multiple of any generated channel or subcarrier bandwidth [9].

CO-OFDM, CoWDM, Nyquist-WDM, and OAWG are all capable of adopting various modulation formats with elastic spectrum assignments as well as generation of Tb/s superchannels. The transmitters for the three technologies utilize a

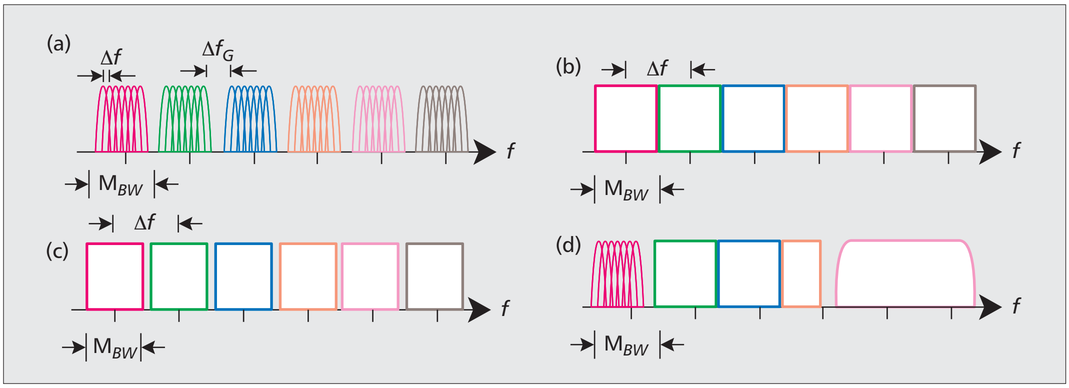

**Figure 5.** *Comparison between CO-OFDM, CoWDM, Nyquist-WDM, and OAWG waveform generation: a) CO-OFDM generates many low-speed subcarriers using an inverse fast Fourier transform (IFFT) to ensure orthogonality; b) CoWDM operates by combining many orthogonal subcarriers together to form a seemingly random waveform; c) Nyquist-WDM combines many independently generated channels together with a minimum guard band; d) OAWG operates over a continuous gapless spectrum, enabling the generation of any combination of single-carrier or multicarrier waveforms. The tick marks indicate locations of modulated comb lines. Df: subcarrier frequency spacing; DfG: frequency spacing between CO-OFDM bands; MBW: modulator bandwidth.*

*This is obviously just a starting point: standards are needed to define a flexible client interface into a BVT, and control plane standards are needed to signal the setup of an EOP as well as its modification over time due to increased/decreased demands and changing transmission conditions.*

similar structure based on many modulators in parallel at low speeds (< 40 Gbaud) modulating as many optical subcarriers. In case of CO-OFDM, it also utilizes electronic subcarriers. The receivers for CO-OFDM, CoWDM, Nyquist-WDM, and OAWG also require many coherent receivers in parallel at low speeds (< 40 Gbaud) and as many optical subcarriers. While CO-OFDM, CoWDM, and Nyquist-WDM require standard optical subcarrier frequency spacing and fixed symbol rates, OAWG can mix and match variety of optical subcarrier frequency spacings, and the OAWG waveforms can include data waveforms in both single carrier modulation formats and multicarrier modulation formats such as CoWDM, CO-OFDM, and Nyquist-WDM. Hence OAWG can interoperate with legacy DWDM as well as CoWDM, CO-OFDM, Nyquist-WDM systems. Such interoperability can help seamless network upgrades from today's networks to future EON. On the other hand, Nyquist-WDM may be the first technology to be deployed for superchannel applications due to its closest resemblance to WDM.

## FROM RESEARCH TO COMMERCIALIZATION

In the past two years, EON research has transitioned from theory to experimentation. For example, a recent field trial of EON based OFDM transmission has demonstrated over 620 km distance with 10G/40G/100G/555G with defragmentation [11], and an EON network testbed with real-time automated adaptive control plane has also been demonstrated [12]. We expect more substantial testbeds to be built in the near term.

Early standardization initiatives have also started, such as revisiting the notation of frequency resources assigned to an optical channel. The conventional ITU grid, which is a reference set of frequencies used to denote the allowed nominal central frequencies of optical channels, is anchored to 193.1 THz, and supports various channel spacings of 12.5 GHz, 25 GHz, 50 GHz, and 100 GHz. In conventional optical networks, when a nominal central frequency is assigned to an optical channel, a fixed frequency range between plus and minus half the channel spacing from the central frequency is implicitly assigned to the channel. Now, ITU Study Group 15 is introducing a flexible DWDM grid into its Recommendation G.694.1, where the allowed frequency slots have a nominal central frequency (THz) defined by $193.1 + n \times 0.00625$ ($n$ is a positive or negative integer including 0), and a slot width defined by $12.5$ GHz $\times\, m$ ($m$ is a positive integer). Depending on the application, a subset of the possible slot widths and central frequency defined in the flexible DWDM grid can be selected.

This is obviously just a starting point: standards are needed to define a flexible client interface into a BVT, control plane standards are needed to signal the setup of an EOP as well as its modification over time due to increased/decreased demands and changing transmission conditions. Network management interfaces should also be standardized. Further research is still needed on more flexible modulation formats and FEC as well as the dynamic behavior of the network, such as allocation of spectrum under changing transmission conditions. EON will also greatly benefit from research and development on BVT [10] and WSS [5] component technologies, particularly from photonic integration of those technologies leading to advanced WSS such as high-port-count flex WSS and > 1 Tb/s BVT integrated circuits.

## SUMMARY

In this article we describe one of the most exciting future directions of optical networks. We start by motivating the transition from fixed grid networks to elastic optical networks. Once the fixed spectrum grid is broken down, many new technologies that are just now crossing the boundary between research and product development will enable a much more flexible network. This process has already started with

*Impairment-aware control planes will be essential for orchestrating the adaptation of the network. Whether many optical networks will become fully elastic is too early to tell, but we certainly seem to be at an interesting juncture in which very significant elasticity seems quite feasible.*

elastic ROADMs, which will not carry significant extra cost and will enable the EON. Then transceivers will become more flexible, fueled by the need to support both high and low demands depending on the required reach. These transceivers may eventually allow for adaptive use of resources, flexible use of spectrum, and a flexible relationship between client technologies such as IP and the optical layer. Impairment-aware control planes will be essential for orchestrating the adaptation of the network. It is too early to tell whether many optical networks will become fully elastic, but we certainly seem to be at an interesting juncture in which very significant elasticity seems quite feasible.

## REFERENCES

[1] M. Jinno *et al.*, "Distance-Adaptive Spectrum Resource Allocation In Spectrum-sliced Elastic Optical Path Network," *IEEE Commun. Mag.*, 2010, pp. 138–45.

[2] M. Jinno *et al.*, "Spectrum-Efficient and Scalable Elastic Optical Path Network: Architecture, Benefits, and Enabling Technologies," *IEEE Commun. Mag.*, vol. 47, 2009, pp. 66–73.

[3] K. Roberts *et al.*, "100G and Beyond with Digital Coherent Signal Processing," *IEEE Commun. Mag.*, vol. 48, July 2010, pp. 62–69.

[4] G.-H. Gho *et al.*, "Rate-Adaptive Coding for Optical Fiber Transmission Systems," *J. Lightwave Tech.*, vol. 29, 2011, pp. 222–33.

[5] O. A. Gerstel, "Flexible Use of Spectrum and Photonic Grooming," *Photonics in Switching*, 2010, p. PMD3.

[6] T. A. Strasser and J. L. Wagener, "Wavelength-Selective Switches for ROADM Applications," *IEEE J. Sel. Topics in Quantum Electronics*, vol. 16, Sept.–Oct. 2010, pp. 1150–57.

[7] P. Frascella *et al.*, "Unrepeatered Field Transmission of 2 Tbit/s Multi-Banded Coherent WDM Over 124 km of Installed SMF," *Opt. Express*, vol. 18, 2010, pp. 24745–52.

[8] W. Shieh *et al.*, "Coherent Optical OFDM: Theory and Design," *Optics Express*, vol. 16, Jan 2008, pp. 841–59.

[9] D. J. Geisler *et al.*, "Bandwidth Scalable, Coherent Transmitter Based on the Parallel Synthesis of Multiple Spectral Slices Using Optical Arbitrary Waveform Generation," *Optics Express*, vol. 19, Apr. 2011, pp. 8242–53.

[10] X. Liu and S. Chandrasekhar, "High Spectral-Efficiency Transmission Techniques for Systems Beyond 100 Gb/s," *Sig. Proc. Photonic Commun.*, 2011, paper # SPMA1.

[11] N. Amaya *et al.*, "Gridless Optical Networking Field Trial: Flexible Spectrum Switching, Defragmentation and Transport of 10G/40G/100G/555G over 620-km Field Fiber," *ECOC '11*, Geneva, Switzerland, 2011.

[12] D. J. Geisler *et al.*, "The First Testbed Demonstration of a Flexible Bandwidth Network with a Real-Time Adaptive Control Plane," *ECOC '11*, Geneva, Switzerland, 2011.

## BIOGRAPHIES

ORI GERSTEL [M'95, SM'05, F'08] (ogerstel@cisco.com) is a principal engineer at Cisco, where he is responsible for identifying opportunities for integration of routers and transport technologies, such as the IP over DWDM architecture. Prior to that he was in charge of Cisco's Optical Advanced Technology team, and is the key inventor behind the advanced capabilities of Cisco's DWDM product. Before joining Cisco in 2002, he was a senior systems architect for Nortel Networks' MEMS-based photonic crossconnect product. Until 2000, he was the systems and software architect for the Optical Networking Group at Tellabs, where he architected the first commercial mesh DWDM system (TITAN 7100). Prior to that, he performed early optical networking research at IBM Research. He authored over 70 papers in international conferences and journals, and over 35 patents on optical networks. He has served on conference committees, such as OFC and INFOCOM, and has been the technical co-chair of Broadnets and IPoP. He also serves as an editor for several international journals such as *JOCN* and has been teaching short courses at OFC. He holds a Ph.D. degree from the Technion, Israel.

MASAHIKO JINNO [M'91] (jinno.masahiko@lab.ntt.co.jp) is a senior research engineer, supervisor, at Nippon Telegraph and Telephone Corporation (NTT) Network Innovation Laboratories. He received his B.E. and M.E. degrees in electronics engineering from Kanazawa University, and his Ph.D. degree in engineering from Osaka University in 1984, 1986, and 1995, respectively. He joined NTT in 1986. He was a guest scientist at the National Institute of Standards and Technology (NIST), Boulder, Colorado, during 1993-1994. He received the Young Engineer Award from the Institute of Electronics, Information, and Communication Engineers (IEICE) in 1993, and the Best Paper Awards for the 1997, 1998, and 2007 Optoelectronics and Communications Conferences.

ANDREW LORD (andrew.lord@bt.com) is an optical specialist at British Telecom where he is responsible for directing research into optical core transport and networks. He received his B.A. Hons degree in Physics from Oxford University in 1985, and joined BT after graduating. He worked on all implementations of DWDM technology since then, including large subsea systems such as TAT14 as well as national and Ministry of Defense DWDM networks. He has helped lead the EU collaborative projects NOBEL and STRONGEST and currently serves on the technical subcommittees of several conferences including OFC and ECOC, where in 2011 he led workshops on EONs. He is a member of the Institute of Physics.

S. J. BEN YOO [S'82, M'84, SM'97, F'07] (sbyoo@ucdavis.edu) received his B.S. degree in electrical engineering with distinction, his M.S. degree in electrical engineering, and his Ph.D. degree in electrical engineering with minor in physics, all from Stanford University, California, in 1984, 1986, and 1991, respectively. He currently serves as a professor of electrical engineering at the University of California at Davis (UC Davis). His research at UC Davis includes optical switching devices, systems, and networking technologies for the future computing and communications. Prior to joining UC Davis in 1999, he was a senior research scientist at Bellcore, leading technical efforts in optical networking research and systems integration. He participated ATD/MONET testbed integration and a number of standardization activities including GR-2918-CORE, GR-2918-ILR, GR-1377-CORE, and GR-1377-ILR on dense WDM and OC-192 systems. He is a Fellow of the Optical Society of America (OSA), and is a recipient of the DARPA Award for Sustained Excellence (1997), the Bellcore CEO Award (1998), the Outstanding Mid-Career Research Award (UC Davis, 2004), and the Outstanding Senior Research Award (UC Davis, 2011).