Evolving Optical Transport Networks to 100G Lambdas and Beyond
Gaylord Hart
Infinera

Abstract

The cable industry is beginning to migrate to 100G core optical transport waves, which greatly improve fiber utilization while lowering transponder count for equivalent transport bandwidth. However, transporting 100G waves requires complex optical modulation to preserve performance and increase spectral efficiency. These modulation methods require several additional discrete optical components per lambda, so the migration to 100G waves alone does not fully address the scalability issues of increasing cost, power, space, and heat as bandwidth requirements continue to grow.

Optical networks are rapidly approaching the point where continual scaling of higher bit rate optical lambdas using discrete optical components is reaching its limits. Photonic integration, which combines multiple optical subsystems on a single IC, can efficiently support complex modulation schemes without increasing component counts and can offer significant improvements over discrete designs, providing a scalable path for future growth.

This paper provides an overview of complex optical modulation methods for 100G waves and higher. An analysis of discrete component solutions and photonic integration solutions is provided which demonstrates the impact of photonic integration in reducing component counts, heat, space, and optical coupling requirements.

INTRODUCTION

In 2014, the Internet is predicted to handle four times more traffic than it did in 2009 (see Figure 1, below). This is largely driven by the explosive growth of Internet video services, which are anticipated to be over fifty percent of all consumer Internet traffic by 2012, and by mobile data services, which are predicted to grow at a 108% compounded annual growth rate (CAGR) through 2014.
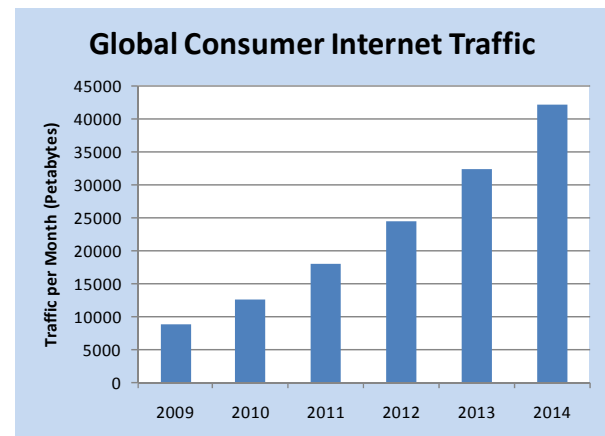

Figure 1—Global Consumer Internet Traffic

Cable subscription Video on Demand (VoD) traffic will also grow significantly over the next few years, more than doubling from 2011 to 2014 (see figure 2, below).
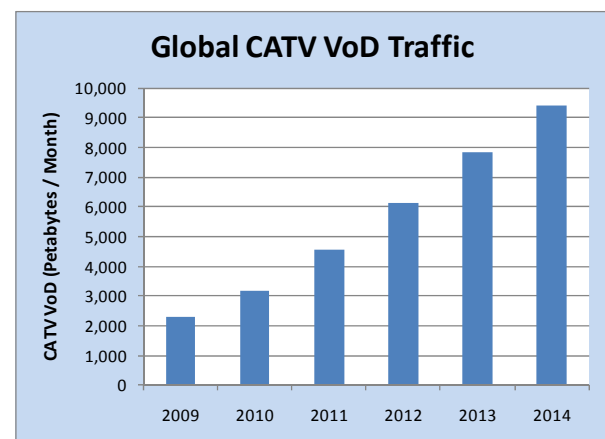

Figure 2—Global CATV VoD Traffic

The combined growth of high speed Internet access and VoD traffic alone will require an extensive expansion of transport capacity in cable networks. Currently, most high-speed optical transport is carried over 10 Gb/s waves using DWDM on a pair of fibers (Tx and Rx). Using 160 x 10G DWDM lambdas using double density spacing at 25 GHz, a single fiber pair can support 1.6 Tb/s capacity, or half that if 50 GHz spacing is used.

Typically, if additional bandwidth is required, another fiber pair is lit up in parallel to the existing lit fibers. This is relatively easy to accomplish if sufficient fiber pairs are available for growth, but lighting up new fibers requires additional chassis and common equipment to support the new fiber link, and incurs additional installation and operational costs associated with increased power, space, and cooling requirements. If additional fiber pairs are not available, considerable additional costs and time delays are usually incurred when new fiber must first be laid.

This linear approach to bandwidth growth, while workable in the short term, does not fundamentally lower the cost per bit for transport over time, nor does it improve scalability by making more efficient use of existing resources (e.g., fiber) or reducing resource consumption (e.g., space and cooling requirements). While fiber relief can be achieved by increasing the DWDM channel count per fiber or by expansion into the L-band, this does not sufficiently address the broader issues of scalability.

As cable operators' optical transport bandwidth requirements continue to grow at a compounded 50-100% per year, network growth cannot continue to scale on a linear basis simply by adding more optical transponders and add/drop multiplexers. Aside from equipment costs, the required space, power, cooling, and optical couplings rapidly become unmanageable when using a linear network expansion model to support geometric bandwidth growth.

SCALABILITY CONSIDERATIONS

A scalable transport solution accomplishes two things over time. It increases the transport capacity per fiber and lowers the total transport cost per Gb/s. The transport cost per Gb/s is comprised of several components, including power, space, hardware, and reliability costs. Additional factors which may contribute to these costs include bandwidth efficiency, the ability to minimize stranded bandwidth in the network, and the ability to simplify and accelerate operational process associated with the network. An ideal solution will make improvements across all these cost components as technology evolves and the network scales to higher transport bandwidth.

INCREASING TRANSPORT CAPACITY

To increase fiber transport capacity, two approaches are generally available: increase the DWDM channel count available on the fiber or pack more bits into the existing channels. Options for bandwidth expansion in the C-band are limited due to amplification requirements and the fact that most of the ITU-grid C-band is already used by existing transport systems. Expansion in the L-band is also possible, but this requires the use of special amplifiers. Consequently, most vendors have focused on increasing the capacity per DWDM channel.

Transport platforms supporting 40 Gb/s lambdas have been available for some years now, and 100 Gb/s lambdas have also recently been deployed. However, according to industry analyst Ovum, the cost per bit for 40G line side lambda transponders is not expected to drop below the 10G lambda cost until 2015, and the cost per bit for 100G line

side lambda transponders is not expected to drop below the 10G lambda cost until 2014. So for the near term, 10G transport costs will remain below 40G and 100G unless fiber exhaust is a factor.

The cable industry has recently begun migrating to 40 and 100 Gb/s transport waves, which has been accelerated by the ratification in June 2010 of the IEEE 802.3ba standard for 40 and 100 Gigabit Ethernet. 40G waves can be deployed using 25 GHz spacing, which means a 160 channel system can support 6.4 Tb/s per fiber pair. 100G waves are typically deployed using 50 GHz spacing, which means an 80 channel system can support 8 Tb/s per fiber, but with half the transponders as the 40G approach.

One important consideration for cable operators is the ability to seamlessly migrate to 40G and 100G transport lambdas on today's 10G networks without having to re-engineer or upgrade the network, including leaving existing dispersion compensation and amplifiers in place. To meet the dual objectives of increasing bandwidth per DWDM channel and operating these lambdas on existing 10G networks, most DWDM equipment uses higher order modulation methods for 40G and 100G lambdas to retain the current channel spacing and increase spectral efficiency.

HIGHER ORDER OPTICAL MODULATION

Most early DWDM systems used 1G or 2.5G lambdas which were modulated using NRZ OOK modulation (non-return-to-zero, on-off keying) applied to the laser light source. This is the simplest form of amplitude shift keying in which a one is represented by the presence of light and a zero by its absence (or vice versa). This modulation technique uses a minimal number of components, but is not very spectrally efficient. The migration

from 2.5G line side waves to 10G was fairly straight forward and did not require major changes or developments in modulation technology. The NRZ OOK modulation was simply speeded up to accommodate the higher line rate, which only required faster components. This is still the most common modulation used for 10G waves. Figure 3, below, shows a block diagram of the components required to transmit and receive a single NRZ OOK optical signal.
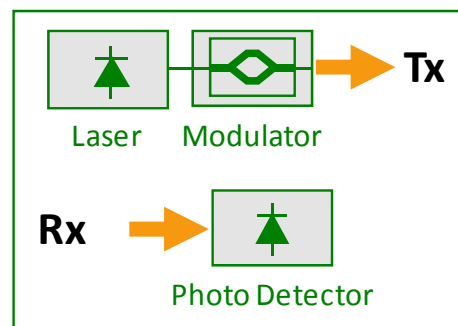

Figure 3—NRZ OOK Transmitter and Receiver Optical Components

As shown in the block diagram, the transmitter consists of a continuous wave laser and a simple OOK modulator. The receiver is even simpler, requiring only a photo detector.

One side effect of using OOK for 10G modulation is that the transport spectrum is wider than for 2.5G, while the bit period is shorter, which makes chromatic dispersion (CD) a more significant factor for 10G transport than for 2.5G. Depending on the transport distances encountered, it is common to use dispersion compensation at regular intervals along the fiber path to mitigate the effect of CD on 10G transport.

Similarly, higher bit rate transmission is more susceptible to other impairments as well, and many of these effects increase with the square of the bit rate. Two of these limiting

factors are polarization mode dispersion (PMD) and optical signal to noise ratio (OSNR), which must be taken into account in optical network design for optimal BER and reach performance.

Using NRZ OOK modulation above 10G line rates, these optical impairments begin to severely impact DWDM transport performance. Because these impairments are related to symbol rate rather than bit rate, most DWDM vendors are using higher order modulation methods for 40G and 100G wave transport, which effectively reduces the symbol rate while increasing the bit rate, and thus minimizes the negative impact of these impairments. Higher order modulation also is spectrally more efficient, which means the information capacity of the fiber is improved as well.

There are many higher order modulation formats which can be used for optical transport, but the most commonly used today rely on some form of phase shift keying (PSK) to translate bits into optical phase states for transport across the network. Table 1, below, provides the bits per symbol for common optical modulation schemes. The more bits per symbol encoded in transmission, the more efficient the modulation format is. As seen in the table below, BPSK encodes one bit per symbol, so the baud rate is the same as the bit rate. QPSK, which is twice as efficient as BPSK, encodes two bits per symbol, so the baud rate is half the bit rate.

Modulation Format	Bits / Symbol
BPSK	1
QPSK	2
8 PSK	3
8 QAM	3
16 QAM	4

Table 1—Bits per Symbol for Common Modulation Formats

Further improvements in spectral efficiency and reduction of symbol rate can be achieved using polarization multiplexing on the optical signal. In this case, the transmitting laser's output is split into two signals, and the polarization of one of these is shifted 90 degrees before being modulated. Each polarized signal is then modulated separately and then combined at the transmitter output for transport over a single fiber. A block diagram of the polarization multiplexed transmitter and receiver is provided below in Figure 4.
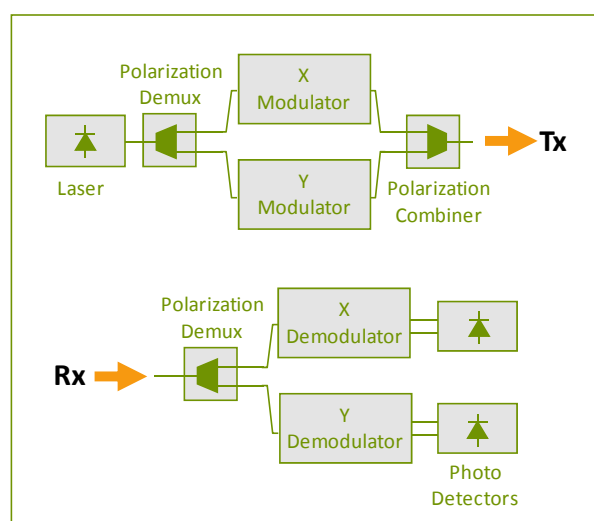

Figure 4—Polarization Multiplexed Transmitter and Receiver

On the receive side, the incoming optical signal is split into the original two polarized signals using a polarization demultiplexer, and then each is independently demodulated. The two recovered bit streams are then combined to reproduce the original data stream used to modulate the transmit laser. Polarization multiplexing may be used with any of the common modulation formats, and for a given transport data rate it effectively increases the bits per symbol for a given modulation format by a factor of two.

Most DWDM transport equipment uses QPSK modulation for 40G and 100G

transport waves. While QPSK is roughly twice the complexity of BPSK and requires about twice as many components to implement, it represents a good tradeoff between spectral efficiency and OSNR performance. For even better performance, polarization multiplexed QPSK (PM-QPSK) is typically used, though at the expense of greater complexity and increased component count.

Further transmission performance enhancements can be achieved if coherent detection is used at the receiver. In differential QPSK (DPQSK), detection is performed by measuring changes in the phase of the received signal rather than the absolute phase itself. Because a copy of the originating reference signal is not required, DQPSK receivers are much simpler to implement. However, DQPSK receivers provide a lower level of performance which translates into higher BER or shorter reach when compared to coherently detected QPSK. This difference is usually sufficient to justify the use of coherent detection, and the Optical Internetworking Forum (OIF) has standardized on coherent PM-QPSK for 100G wave transport.

(sometimes referred to as synchronous detection) requires a local copy of the transmitter's original CW carrier to perform direct phase detection. In the optical domain, this requires a separate local laser at the receiver (equivalent to a local oscillator in the RF domain) which is phase locked to the incoming optical signal to create a copy of the original optical carrier.

Coherent detection is usually implemented with an ASIC which integrates the A/D conversion for the optical detector outputs as well as a Digital Signal Processor (DSP), AGC controller, and other functions on the ASIC. Using digital signal processing in the coherent receiver enables other features as well, including electronic dispersion compensation (EDC). EDC is far more tolerant than fiber based dispersion compensation and can readily provide compensation for +/-50,000 ps/nm without the use of bulky dispersion compensation modules and their inherent attenuation. Similarly, PMD performance is greatly improved when coherent detection is used with digital signal processing, and a DGD tolerance of 200 ps peak can be achieved for 100G waves.
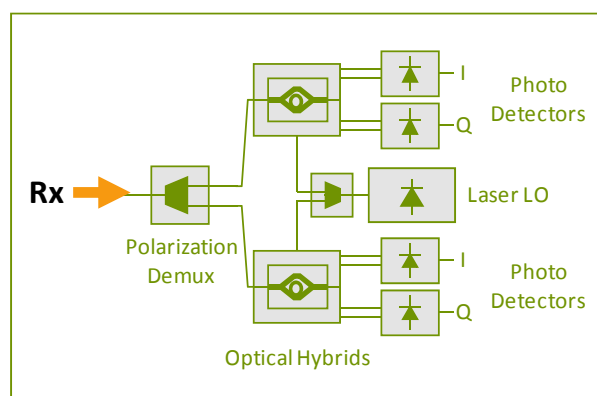
Figure 5—PM-QPSK Coherent Receiver Block Diagram

Figure 6—QM-PSK Transmitter Block Diagram

Figure 5, above, shows a block diagram of a PM-QPSK receiver. Coherent detection
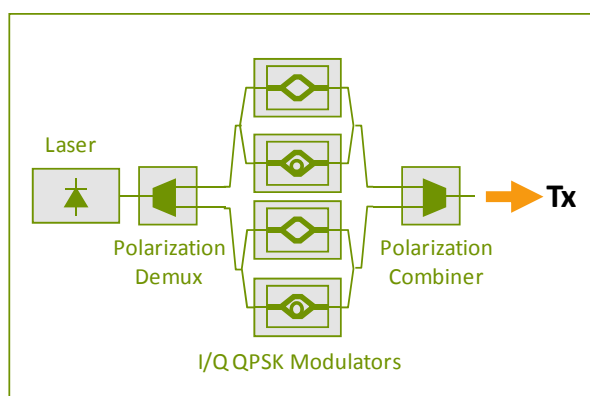
Figure 6, above, shows a block diagram of a PM-QPSK transmitter. The source laser

signal is split into two polarized sources 90 degrees out of phase. These are then QPSK modulated, creating I and Q modulation output components, and these are then combined for transport on the fiber.

One additional benefit of DSP-based coherent detection is the ability to provision the modulation format in software. Thus the transmitter and receiver pair can be designed to be operated using either BPSK or QPSK, for example, and the operating mode may be selected by the operator based upon link requirements. BPSK supports longer reach at a lower data rate for ULH applications, where QPSK is optimized for higher data rates with some compromise in reach.

PHOTONIC INTEGRATION

As one can see comparing the number of components required for NRZ OOK transmission of a single lambda in Figure 3, above, with the number required for PM-QPSK in Figures 5 and 6, also above, it is apparent that a significant number of additional optical components and fiber connections are required per lambda to support 40G and 100G wave transmission. If one considers transmission of 80 x 100G or 160 x 40G lambdas per fiber, the number of required additional components and fiber couplings is quite large. This has a direct and negative impact on reliability, power consumption, space, and cooling requirements, all of which lead to higher transport costs and a lower degree of scalability.

In practice, some of these discrete components may be combined in a single package using either hybrid construction (where two or more discrete components are placed in the same package) or small-scale integration (where two or more components are monolithically manufactured and packaged together). However, this approach offers only limited relief since it does not address the entire receiver and transmitter subsections or the DWDM system as a whole, which requires multiple lambdas as well as mux/demuxes.

Ideally, as optical transport networks scale in bandwidth by using higher order modulation techniques, new technology should be available to mitigate the effects of the increased complexity and component counts. Fortunately, this is indeed the case. Modern photonic integration allows multiple optical subsystems to be monolithically manufactured on an Indium Phosphide (InP) chip using large-scale integration. It is possible today to put all the optical components necessary to support multiple lambdas, including the mux/demux functions, on a single photonic integrated circuit (PIC) and address integration at the system level rather than the component level.

Commercially deployed PICs for 10G transport have been available for NRZ OOK modulation since 2005. These currently support 10 x 10 Gb/s DWDM wavelengths on a pair of PICs less than 5 mm square each (complete TX and RX systems, including mux/demux functions, 100 Gb/s per PIC). PICs which support 5 x 100 Gb/s DWDM wavelengths using PM-QPSK modulation and coherent detection have already been demonstrated in long haul networks and are planned for commercial availability in DWDM systems in 2012.

PICs provide the unique benefit of integrating not only the discrete optical components on the IC, but the optical connections between them as well. Table 2, below, provides a summary of the currently available 10 x 10G PICs, including the number of functions integrated on the Tx and Rx PICs, the number of equivalent discrete component "gold box" packages eliminated, and the number of fiber connections eliminated.

10 x 11G NRZ OOK PICs	100G PICs
InP monolithic ICs (Tx & Rx)	2
Integrated optical functions	62
"Gold box" replacements	20-50
Fiber coupling reduction	>100

Table 2—100G PIC Large-Scale Integration

Modeling an 80 x 10G lambda terminal configuration node (800 Gb/s), photonic integration, when compared against discrete transponder based solutions, can save up to 45% on power costs while reducing the required rack space from three 7 foot bays to one 7 foot bay.

The benefits of photonic integration only increase as more complex modulation methods are used to provide more bandwidth per lambda. Table 3, below, shows the corresponding integration summary for 500G PICs supporting 5 x 100G lambdas using PM-QPSK. This represents a fivefold increase in bandwidth and a fourfold improvement in fiber coupling and discrete component reduction when compared to today's 100G PICs.

5 x 114G PM-QPSK PICs	500G PICs
InP monolithic ICs (Tx & Rx)	2
Integrated optical functions	>600
"Gold box" replacements	>100
Fiber coupling reduction	>400

Table 3—500G PIC Large-Scale Integration

Moore's law indicates significant improvements are yet possible in future generations of PIC technology, and 10 x 100G Terabit PICs have already been produced and tested successfully in the lab. Current models predict PIC bandwidth will double about every three years, keeping pace with bandwidth growth while lowering the overall cost per bit.

EVOLUTION OF THE TERABIT PIC AND MULTI-CARRIER SUPER-CHANNELS

As the industry migrates to Terabit optical transport, there will be an ever-increasing need to squeeze more bandwidth from fiber and yet to retain flexibility on how that bandwidth is used. Modulation formats will continue to evolve, with 8QAM and 16QAM being next in line for extending transport bandwidth. At the same time, it is desirable to retain the flexibility to support multiple modulation formats on the same fiber, which will allow seamless migration to higher bandwidth and allow suitable tradeoffs to be made for optical reach versus total fiber capacity.

Current DWDM systems are based on an ITU channel grid which provides for lambda spacing of 100, 50, or 25 GHz. Inherent in this spacing are dead zones to allow for optical filtering of individual wavelengths, and these dead zones limit channel density, which in return can result in up to 50% of the available fiber spectrum being unusable.

In future DWDM transport systems, it will make more sense to move beyond the current ITU channel plan and implement multi-carrier super-channels that eliminate the dead zones for the carriers within the super-channel, but which preserve a guard band at the edges of each super-channel for filtering purposes. Multiple sub-rate carriers can readily be implemented within the super-channel using PIC technology which also allows multiple modulation formats and flexible channel spacing to be supported and provisioned in software. Such DWDM systems should enable transport capacities of up to 25 Tb/s per fiber, well beyond the capacity of today's systems.

SUMMARY AND CONCLUSIONS

As transport bandwidth requirements grow at a 50-70% CAGR, DWDM transport systems will have to migrate to complex modulation methods such as PM-QPSK to conserve fiber spectrum and increase the transport capacity per lambda. Complex modulation requires substantially more optical components per lambda than the current 10G transport lambdas require. This increase in components is not sustainable in the long run due to the cost, space, heat, power, and reliability requirements associated with discrete component implementations of DWDM optical transport systems.

Photonic integration, which combines onto a single VLSI chip the entire optical subsystems needed to transport multiple lambdas using complex modulation, provides a scalable technology that enables future migration to Terabit PICs which will lower the overall transport cost per bit while reducing space, power, and cooling required per bit.

NOTES

Source for bandwidth growth forecasts: "Cisco Visual Networking Index: Forecast and Methodology, 2009-2014," June 2. 2010.

"Optical networks volumes and revenue history and forecast, 2006-2015," Ovum, September 21, 2010.

Author Contact Information:

Gaylord Hart
Infinera
303.840.9381
ghart@infinera.com