Can 100Gb/s wavelengths be deployed using 10Gb/s engineering rules?

Ross Saunders[1], Gary Nicholl[2], Kevin Wollenweber[2] and Ted Schmidt[1]
[1] StrataLight Communications Inc., Los Gatos, California, USA, ross@stratalight.com
[2] Cisco Systems Inc., San Jose, California, USA, gnicholl@cisco.com

ABSTRACT

A key challenge set by carriers for 40Gb/s deployments was that the 40Gb/s wavelengths should be deployable over existing 10Gb/s DWDM systems, using 10Gb/s link engineering design rules. Typical 10Gb/s link engineering rules are:

1. Polarization Mode Dispersion (PMD) tolerance of 10ps (mean);
2. Chromatic Dispersion (CD) tolerance of ±700ps/nm;
3. Operation at 50GHz channel spacing, including transit through multiple cascaded [R]OADMs;
4. Optical reach up to 2,000km.

By using a combination of advanced modulation formats and adaptive dispersion compensation (technologies rarely seen at 10Gb/s outside of the submarine systems space), vendors did respond to the challenge and broadly met this requirement.

As we now start to explore feasible technologies for 100Gb/s optical transport, driven by 100GE port availability on core IP routers, the carrier challenge remains the same. 100Gb/s links should be deployable over existing 10Gb/s DWDM systems using 10Gb/s link engineering rules (as listed above).

To meet this challenge, optical transport technology must evolve to yet another level of complexity/maturity in both modulation formats and adaptive compensation techniques. Many clues as to how this might be achieved can be gained by first studying sister telecommunications industries, e.g. satellite (QPSK, QAM, LDCP FEC codes), wireless (advanced DSP, MSK), HDTV (TCM), etc.

The optical industry is not a pioneer of new ideas in modulation schemes and coding theory, we will always be followers. However, we do have the responsibility of developing the highest capacity "modems" on the planet to carry the core backbone traffic of the Internet. As such, the key to our success will be to analyze the pros and cons of advanced modulation/coding techniques and balance this with the practical limitations of high speed electronics processing speed and the challenges of real world optical layer impairments.

This invited paper will present a view on what advanced technologies are likely candidates to support 100GE optical IP transport over existing 10Gb/s DWDM systems, using 10Gb/s link engineering rules.

KEYWORD LIST

100Gb/s; 100GE; DWDM; IP; core networking; optical modulation schemes; adaptive compensation

INTRODUCTION

The market for 40Gb/s (OC-768 Packet-over-SONET, PoS) port interfaces on next generation IP routers is just beginning to be deployed in carrier networks. The first application is in the core IP network between major city populations. The expected growth in IP traffic due to converged IP networks, migration of voice traffic over to IP and future video over IP are driving the capacity growth projections at both traditional telecommunications carriers and MSOs.

As IP traffic growth continues, carriers will demand higher transmission speed between routers as inverse multiplexing IP traffic flows over multiple wavelengths is inefficient. This is due to different packet flow requirements and the hash algorithm used to mange the flows [1]. For this reason, just as IP network engineers are currently pushing the migration from 10Gb/s to 40Gb/s line rates, so in a few years they will push to adopt 100Gb/s from 40Gb/s at present. In addition, the IEEE is currently standardizing 100GE interfaces [2] and has already decided on a 100Gb/s MAC interface speed, so the adoption of Ethernet framing will likely act as a further catalyst to large scale adoption.

The development of DWDM-based systems is expected to continue with the wider scale adoption of ROADMs and further migration to multi-degree ROADM networks. With full band tunable lasers now standard on DWDM equipment and N-degree ROADMs just beginning to enter the market, it is expected that networks will migrate to some level of intelligent optical switching, for provisioning, protection and restoration purposes. One key technology piece required to enable dynamically switched optical networks is tunable optical filtering at the receiver. It is worth closely monitoring this as it is to be expected that the ability to filter out one wavelength from a group output of a WSS will be a strong requirement in next generation agile optical DWDM networks. The current standard for LH DWDM systems is to use 50GHz channel spacing. This tight channel spacing places severe restrictions on 100Gb/s support and requires the use of highly spectrally efficient modulation formats. Optical reach is another key attribute for DWDM core networking. Carriers are used to deploying 10Gb/s wavelengths, and now 40Gb/s wavelengths, with reach in excess of 1,500km. This greatly reduces the number of [relatively costly] OEO regenerators that must to be deployed, especially in long routes, such as coast-to-coast demands. To secure wide-scale adoption of 100Gb/s wavelengths, carriers will again expect/demand a similar type reach to minimize the “regen tax” in core IP network deployments. The required reach is clearly dependent on geography and network application, where metro/regional deployments and national deployments in smaller countries such as Japan, UK and Germany can live with shorter reach than the USA.

100GB/S STANDARDIZATION

IEEE Standardization Activities

Standardization of the next generation of Ethernet started in July 2006, when the IEEE agreed to form a Higher Speed Study Group (HSSG) to begin defining the project scope and objectives [11]. The initial focus of the group was on 100Gb/s only as the next rate for Ethernet. However a key finding by the HSSG was a divergence in bandwidth demand between the computing and networking industries. As a result the HSSG completed a Project Authorization Request (PAR) in July 2007 which included two new data rates, 40 Gb/s for server and storage applications and 100Gb/s for aggregation and core networking applications. The group also established physical layer objectives tailored towards the intended applications for each rate, including up to 100 meters for 40Gb/s and up to 40 kilometers for 100 Gb/s.

Assuming the PAR is approved by the IEEE Standards Board in the fourth quarter of 2007, the HSSG will become the IEEE P802.3ba task force and will begin work on the standard in early 2008. While no definitive standards development timeline has been agreed to, it is anticipated that the standard will be completed in 2010.

In the context of this paper it is the 100Gb/s data rate that is obviously of the most interest. As mentioned above the key driver for 100Gb/s is aggregation and core networking, and this will require transporting 100Gb/s ethernet interfaces between core routers/switches, over both metro and long-haul DWDM systems. It is expected that once 100Gb/s Ethernet becomes available, it will replace both Packet-over-SONET (POS) at OC-192 (10Gb/s) and OC-768 (40Gb/s) as the technology of choice in the core of carrier networks. In light of this the ITU have been tracking the work of the IEEE HSSG very closely, and at the June ITU meeting in Geneva it was agreed that the next level in the OTN hierarchy (OTU4) would be optimized for the transport of 100 Gb/s Ethernet (see below).

It is also worth pointing out at this point, that the primary application for 40 Gb/s Ethernet is seen as data center computing and storage. It is therefore anticipated that there will be little demand to extend these interfaces outside of the data center environment and to carry them over DWDM transport systems.

ITU Standardization Activities

In keeping with the recent activity of the IEEE HSSG in defining the next generation of ethernet, ITU SG15 has also been active in discussing the evolution of the G.709 OTN hierarchy and in particular how a 100Gb/s Ethernet client, once it has been defined by the IEEE, will be mapped and transported. The main topic of discussion has been around whether the next level in the OTN hierarchy (OTU4) should simply extend the existing OTN multiplexing structure and thus be optimized for the transport of SONET/SDH clients, or whether it should be modified slightly so that it is more optimized for the multiplexing and transport of ethernet clients, and in particular 100Gb/s.

The first thought was to extend the existing OTN hierarchy, as can clearly be seen from one the early proposals for a 130 Gb/s OTU4 rate, which is obviously excellent for transporting 3xOC-768 clients but not optimum for 100Gb/s ethernet. The mood of the group changed however, when at the June ITU meeting in Geneva contributions from a number of companies suggested that data, and in particular ethernet, would be driving the demand for next generation transport systems and that the new OTU4 rate should instead be optimized for the multiplexing and transport of 10Gb/s and 100Gb/s Ethernet clients. A number of technical proposals were made on how this could be achieved, but given that the IEEE has yet to define the exact frame format and encoded date rate for 100 Gb/s Ethernet no decisions could ultimately be made.

At the end of the June meeting ITU SG15 agreed that it would continue to track progress within the IEEE HSSG, and the resulting Task Force, and define the new OTU4 rate accordingly. Under the assumption that the IEEE decide to stick with 64B66B as the line code for 100Gb/s Ethernet (resulting in an encoded line rate of ~ 103 Gb/s), this means that the new OTU4 line rate will be approximately 112 Gb/s.

CORE IP NETWORKING DRIVERS

There is a shift in the market today towards a convergence of infrastructure components and integration of services as providers look to carry all of their traffic over a converged IP core. With the growth of broadband speeds to home users as well as a shift towards rich media content and high definition video, we are seeing increasing demands on the bandwidth of the network. Many providers are seeing faster growth in their networks than they have planned for, causing them to quickly plan the next generation of bandwidth growth. With existing technologies for running 40Gb/s links over existing 10Gb/s transport networks, we are seeing customers adopting 40Gb/s today, but even this will hit scaling limits. This is creating the market for a fast, spectrally efficient, long haul 100Gb/s interface that will allow carriers to scale their bandwidth. As carriers scale their networks to meet the growing demands, they have 3 main choices: overlay builds, bundled links, or faster interfaces which can run on their existing systems.

Overlay builds are obviously the most expensive choice. Providers are trying to converge networks and consolidate on IP, so building an overlay optical network to cope with this demand is undesirable. Leveraging the existing investment and amortizing it over a longer period of time helps their profitability. Any carriers who have channels available on their existing infrastructure would rather use that investment first, before making another one. If this were the only choice, 100Gb/s would not be a reality for many carriers due to the cost of a new build.

Before 40Gb/s interfaces were available on network equipment, many customers were forced to bundle links. This can be done at layer 1, 2, or 3, but all three methods have their advantages and disadvantages. One major issue with bundled links is that it is difficult to scale the number of connections as the bandwidth grows. This is one of the major factors driving 40Gb/s deployments today. As these networks continue to grow, bundling 20 links between two points to achieve 200Gb/s becomes difficult to manage. Many older devices were also only built to hash traffic across 8 or 16 links, so some links may run into scale limits due to HW limitations. Another issue with bundling is that bundling algorithms are based on hashing. Most hashing algorithms are based on some combination of source address and destination address to make sure the same traffic flows all hash to the same link. This can lead to inefficiencies and underutilization of links if there is not enough randomness in the traffic patters.

Increasing bandwidth demands will push many customers towards 100Gb/s, but next generation 100Gb/s interfaces must also be spectrally efficient to help wide scale adoption. Ideally, they will support the same 50GHz spacing that we see in today's DWDM systems at 10 and 40Gb/s. This will allow customers to easily integrate new higher speed links alongside their existing revenue generating circuits already in place. Another important characteristic will be to automatically compensate for differences in optical performance, allowing 100Gb/s to be deployed on the same engineered links as existing 10Gb/s links. Adding additional amplifiers or regenerators is not desirable. What we have seen with today's 40Gb/s technologies is that the ability to run across existing 10Gb/s systems and leverage the investment the provider has in their existing line system and the fiber in the ground has sped their adoption of the technology. The inherent advantage of this approach over inverse multiplexing is that it will allow the provider to grow to much higher bandwidths in their existing network before needing an overlay optical build and more fiber.

PRACTICAL NETWORKING CHALLENGES

Carriers around the world have deployed 10Gb/s DWDM systems, including more recent deployments with ROADMs and a few networks with 40Gb/s support. Carriers are largely comfortable deploying 10Gb/s DWDM systems and have learnt to accommodate the tighter dispersion tolerance requirements of 10Gb/s, versus 2.5Gb/s, wavelength deployments. Typical long haul 10Gb/s deployments include distributed per span dispersion compensation and the use of 50GHz channel spacing and associated add/drop and MUX/DEMUX optical filter technology. The rough “Rules of Thumb” comparing the challenges of 100Gb/s vs. 10Gb/s and 100Gb/s vs. 40Gb/s design are summarized in Table.1.

Design parameter	100Gb/s vs. 10Gb/s	100Gb/s vs. 40Gb/s
OSNR requirement	10dB higher	4dB higher
CD tolerance	100x smaller	6.25x smaller
DGD tolerance	10x smaller	2.5x smaller
PMD-limited reach	100x shorter	6.25x shorter
Optical bandwidth	10x larger	2.5x larger

Table 1. Rules of Thumb for 10Gb/s vs. 40Gb/s vs. 100Gb/s Optical System Design

To enable 100Gb/s deployment using 10Gb/s link engineering rules, the following design criteria must be met:

1. Optical reach of ≥ 1,500km;
2. Support 50GHz DWDM channel spacing;
3. No change to existing DWDM common equipment;
4. Non traffic-affecting upgrades on existing 2.5Gb/s, 10Gb/s and 40Gb/s circuits;
5. Must not induce significant linear or nonlinear crosstalk penalty on existing DWDM channels;
6. Power per channel must be ≤ +2dBm;
7. Chromatic dispersion tolerance of ±700ps/nm;
8. Polarization Mode Dispersion tolerance of 10ps (mean DGD);
9. Ability to express signal through ≥ 4 ROADMs at 50GHz DWDM channel spacing;
10. Ability to express signal through ≥ 24 ROADMs at 100GHz DWDM channel spacing;
11. Automated dispersion tuning/tracking, if applicable;
12. Full band tunable lasers on 50GHz ITU grid;
13. Should support Digital Wrapper (OTU4) frame structure for fault management, alarm signaling and performance monitoring;
14. Must be bit-for-bit transparent for the 100GE client interface, no flow control or packet loss at transport layer;
15. Must be “plug and play” and installable by existing field technicians, no “PhDs with screwdrivers” required for link turn-up in the field;

If the above criteria are fully met at 100Gb/s line rate, then significant savings are realized by the carrier related to: (i) **capital preservation**, 100Gb/s retrofit upgradability increases the usual life of existing fiber/DWDM plant; (ii) **increased spectral efficiency**, 100Gb/s upgrade on an 80 channel system increases the capacity of the fiber/DWDM by an order of magnitude to 8Tb/s; (iv) **new revenue opportunity** by supporting 100GE services.

The downside of designing a 100Gb/s solution to be deployable over existing 10Gb/s systems is that the additional engineering design constraints add cost/complexity versus a complete new DWDM system that could be engineered from the bottom-up and optimized for multi-channel 100Gb/s DWDM. An example of this is that a 100Gb/s solution based on DQPSK would have extremely attractive cost points and is highly achievable, but due to it’s spectral width it would require 100GHz DWDM channel spacing and a reach of 500-600km. For applications with limited reach and/or channel count, e.g. metro/regional, European PTT, Japan, etc. this solution may offer better price points. It is therefore expected that different 100Gb/s designs will give better cost points for specific applications, rather than “one size fits all”.

DESIGN TRADEOFFS AND MODULATION SCHEME SELECTION

To meet the challenges highlighted in Table 1 and still enable 100Gb/s retrofits into existing 10Gb/s DWDM systems, different modulation schemes must be considered versus Non Return-to-Zero On-Off Keying (NRZ OOK), which is still the standard used in 10Gb/s deployments. Given that optical engineers are novices at using advanced modulation and signal processing techniques, a logical first step is to take a look at other communications and IT industries where limited spectrum and SNR of other media has led to innovation in coding schemes. Examples of novel modulation schemes used in other industries include: (i) [Gaussian] Minimum-Shift-Keying (MSK) in wireless GSM for high spectral efficiency; (ii) Trellis Coded Modulation (TCM) [3] for bandwidth-efficient increased coding gain in HDTV; (iii) QPSK/QAM for high spectral efficiency/SNR performance in satellite communications; (iv) Orthogonal Frequency Division Multiplexing (OFDM)/Discrete Multi-Tone (DMT) in WiMAX/DSL to optimize performance over varying spectral characteristics; (v) advanced FEC Low Density Parity Checking (LDPC)/turbo codes approaching Shannon's limit in Satellite digital video broadcasting and WiMAX for improved SNR sensitivity; (vi) use of Digital Signal Processing (DSP) [4] in many digital communications fields to compensate for line/Tx/Rx bandwidth impairments, noise filtering, channel filtering and adaptive compensation.

Many lessons can be learnt for looking at the practical design challenges in each of these communications fields, identifying why particular modulations schemes or processing techniques were chosen and then evaluating the practical implementation at the very high data rates required for carrying optical backbone internet traffic. Many modulation techniques also tradeoff spectral efficiency vs. OSNR sensitivity. Figure 1 shows this tradeoff for 100Gb/s optical transport.

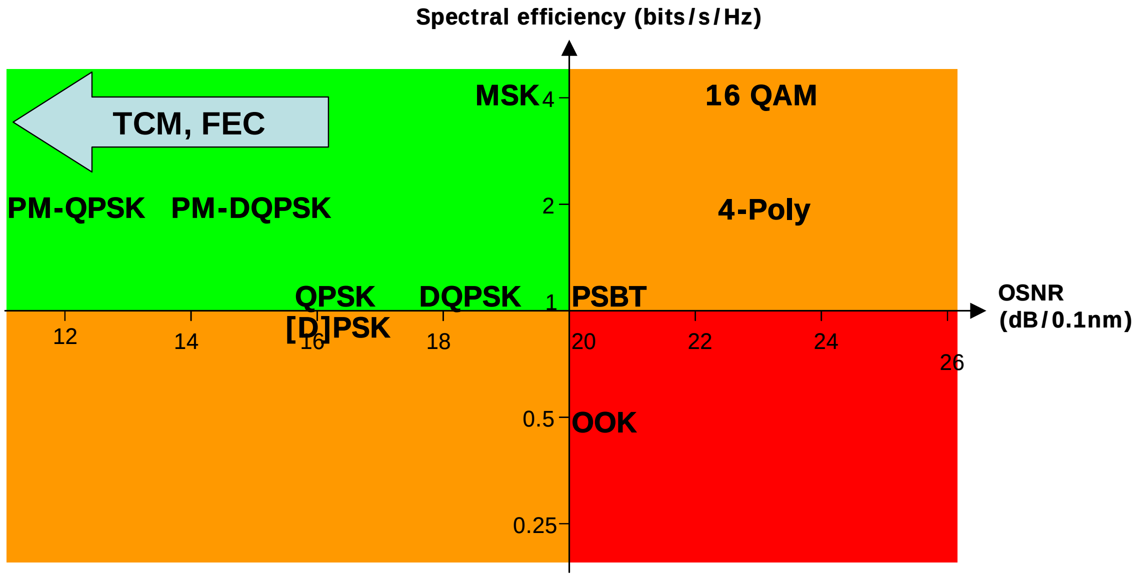

Figure 1 - Spectral Efficiency vs. OSNR Sensitivity Tradeoffs

To meet the requirements of 100Gb/s wavelengths retrofitting into existing 10Gb/s DWDM systems, modulation formats in the green quadrant are needed to meet the spectral efficiency and OSNR sensitivity criteria. Note this is a simplistic model that does not include fiber nonlinear propagation distortion and assumes 8.5dB FEC Net Equivalent Coding Gain (NECG). In this first-order analysis, it is clearly evident that OOK format does not come close to supporting the required level of spectral efficiency (it would only support 200GHz DWDM channel spacing) or OSNR (reach would be around 400km). Neither [D]PSK or [D]QPSK meets the spectral efficiency requirement and would only support 100GHz channel spacing, with [D]PSK suffering a higher penalty per ROADM that [D]QPSK due to the larger spectral width of the signal. 8-PSK (octal, 8 states in signal-space diagram) can support 50GHz channel spacing but the reduced Euclidean distance between signal states results in a corresponding reduction in OSNR vs. [Q]PSK, so the reach would be in the order of 400km. 4-Poly (4 level polybinary coding, integrate over 4 bits) again supports 50GHz channel spacing but has even worse OSNR sensitivity than 8-PSK due to the multi-level slicer and reduced amplitude margin available in the Rx. 16-QAM (16 level Quadrature Amplitude Modulation) and MSK offer even higher spectral efficiency than is really needed (i.e. could actually support 100Gb/s channels at 25GHz channel spacing) but the cost is too high in OSNR sensitivity. Polarization Multiplexing (PM) offers a doubling of the spectral efficiency (supports 100Gb/s channels at 50GHz DWDM channel spacing) and good OSNR performance for PM-[D]QPSK [5,6]. In addition to modulation scheme choice, TCM and more advanced FEC algorithms [7] that approach Shannon's Limit can also improve the OSNR sensitivity and hence reach to achieve transmission distances in excess of 1,500km.

Coherent detection has to date never been implemented in a commercial DWDM product, but it certainly offers some important performance advantages, at the expense of added cost/complexity. The performance advantage for coherent vs. direct detection PSK formats is shown in Figure 2. For Binary PSK, the OSNR advantage for implementing coherent vs. direct detection is less than 1dB, typically around 0.7dB at the pre-FEC BER levels of interest. Therefore the added cost/complexity of synchronous coherent detection with binary PSK is not really worth the 0.7dB performance enhancement. For Quatenary PSK (QPSK), coherent detection offers a significant 2.5dB OSNR sensitivity improvement over direct detection Differential QPSK (DQPSK]. As the doubling in spectral efficiency offered by quaternary PSK is extremely attractive for 100Gb/s implementation, it is now a tougher decision to de-select coherent detection as the 2.5dB OSNR improvement is also critical to meet the > 1,500km reach requirement. In addition to the 2.5dB OSNR improvement, coherent detection offers other advantages such as: (i) electromagnetic phase information passes into the electrical domain, so analogue and/or digital filters can be used for PMD and CD compensation; (ii) a single channel can be filtered out from many DWDM channels by Local Oscillator tuning (tunable laser), rather than requiring a tunable optical filter and (iii) polarization demultiplexing is a minor increment, given polarization diversity is already anecessity for coherent detection.

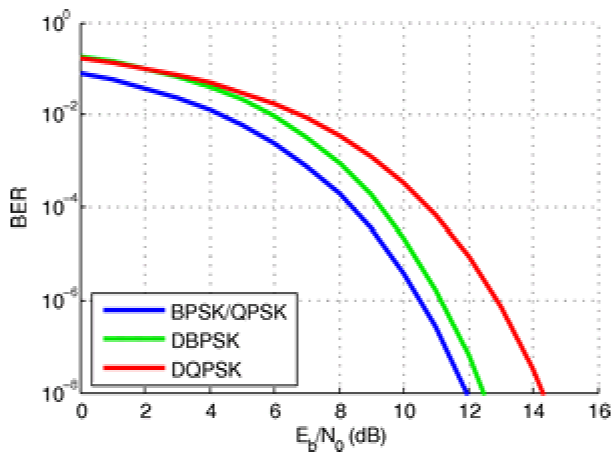

Figure 2 – PSK/DPSK/DQPSK SNR Comparison

Another important constraint that cannot be ignored is the cost factor. The different design implementations have very different levels of complexity that in turn impacts development cost, component Bill of Material (BOM) and manufacturability/yield. The cost/complexity of each modulation scheme must be closely investigated on both the electrical and optical level. There are often tradeoffs involved here, as for example formats with an increased number of optical levels (i.e. coding more bits per baud) results in increased optical complexity in encoding/decoding multiple levels, but reduced electronics complexity due to the lower baud rate required for the electronic components, such as driver amplifiers, multiplexers, etc. A summary of the design complexity tradeoffs for different modulation formats is shown in Figure 3.

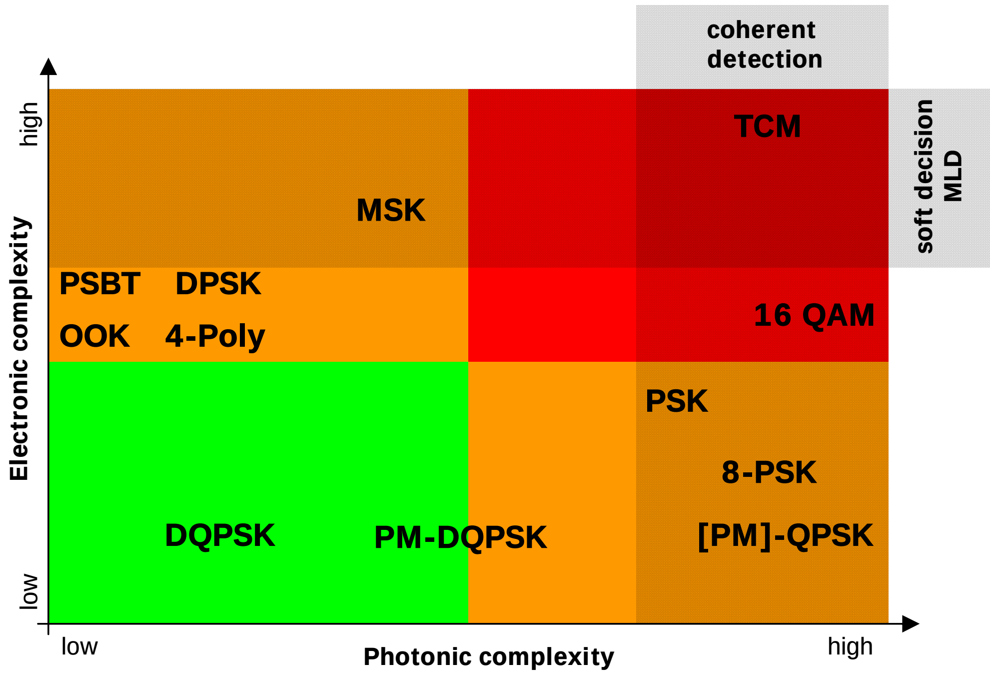

Figure 3 – Implementation Complexity Tradeoffs

For the options illustrated in Figure 3, it is again preferred that we select a format in the green quadrant, as these options offer the easiest implementation and lowest cost. It is interesting to note that OOK (or any format that codes only 1 information bit/baud) does not qualify for low electronics complexity. This is due to the immaturity of qualified electronic/electro-optic processes that can support 100Gb/s TDM line rate for MUX/DEMUX, driver amplifier, modulator and photodetector functions. As mentioned previously, DQPSK is highly attractive from a low electronics/photonics implementation complexity/cost perspective. Unfortunately it does not meet all the market requirements (in particular 50GHz channel spacing) as described in Figure 2, but it may still make sense in applications such as metro/regional networks where 100GHz channel spacing is acceptable. Coherent detection adds significantly to the complexity/cost, but the gains in spectral efficiency/OSNR are considerable and coherent detection schemes are a likely candidate format to meet all the requirements for 100Gb/s wavelength deployment using 10Gb/s link engineering rules.

ADAPTIVE DISPERSION COMPENSATION SCHEMES

To meet the stringent link engineering rules imposed by a network designed for 10Gb/s, modulation scheme choice alone may not be sufficient. The use of adaptive dispersion compensators enables compensation of both CD and PMD and with proper design, can increase the CD and PMD tolerance to the equivalent of [uncompensated] 10G DWDM NRZ OOK wavelengths. The adaptive compensators can be implemented in the optical domain and in both analogue or digital implementation in the electronic domain. At 40Gb/s line rates, optical Tunable Dispersion Compensators (TDCs) and PMD Compensators (PMDCs) are used to increase the CD/PMD tolerance to the same levels as 10Gb/s. TDCs are often implemented in Fiber Bragg Grating (FBG) and Gires-Tournois Etalon (GTE) form [8]. PMD compensator designs have been implemented in the electrical domain [9] and optical domain [10], see Figure 4.

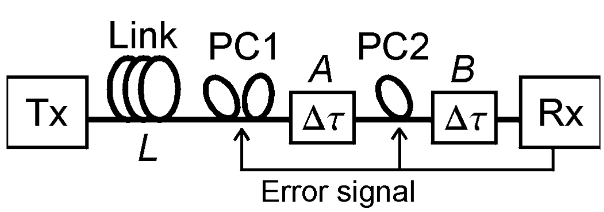

Figure 4 – Optical Polarization Mode Dispersion Compensator

In a coherent detection scheme, after digital quantization, a digital adaptive Finite Impulse Response (FIR) filter [4] can be implemented where the tap weight coefficients are adjusted in real-time to compensate for the time-varying PMD-induced pulse distortion.

SUMMARY

A summary table of the pros and cons of the different 100Gb/s modulation schemes considered is shown in Table 2. As can be seen, it is possible to meet the reach and spectral efficiency requirements necessary to deploy 100Gb/s wavelengths using 10Gb/s DWDM design rules, but there is a penalty to pay in terms of added cost/complexity. In order to meet the CD and PMD tolerance levels, optical and/or electrical domain adaptive dispersion compensation technology is needed, which does again incur additional cost.

The developing standards work in IEEE and ITU for 100Gb/s will help define the interfaces and OTN link management functions that are necessary to allow interoperability between next-generation IP routers/Ethernet switches and the OTN-compatible DWDM optical transport equipment. The technology to enable 100Gb/s deployment using 10Gb/s link engineering rules is feasible and realizable, but the challenge will be to make this available at a cost point that is acceptable to drive wide scale adoption for use in next-generation internet backbones.

	OOK	PSBT	DPSK	DQPSK	QPSK	PM-DQPSK	PM-QPSK
Spectral Efficiency	0.4 bits/s/Hz	1 bits/s/Hz	0.8 bits/s/Hz	1.6 bits/s/Hz	1.6 bits/s/Hz	3 bits/s/Hz	3 bits/s/Hz
OSNR sensitivity	20dB/0.1nm	20dB/0.1nm	17dB/0.1nm	18dB/0.1nm	15.5dB/0.1nm	18dB/0.1nm	15.5dB/0.1nm
PMD tolerance	1ps	1ps	1ps	2ps	2ps	2.5ps	2.5ps
CD tolerance	15ps/nm	50ps/nm	12ps/nm	35ps/nm	35ps/nm	140ps/nm	140ps/nm
Analogue electronics complexity	HIGH	HIGH	HIGH	MEDIUM	MEDIUM	LOW	LOW
Digital electronics complexity	LOW	LOW	LOW	LOW	HIGH	MEDIUM	HIGH
Optical complexity	LOW	MEDIUM	MEDIUM	MEDIUM	HIGH	MEDIUM	HIGH
Reach estimate	400km	400km	800km	700km	1,000km	700km	1,000km
Cost estimate	0%	+10%	+20%	+50%	+70%	+90%	+110%

Table 2 - 100Gb/s Modulation Scheme Comparison

N.B. [1] [Q]PSK schemes assume coherent detection with Rx ADC/DSP; [2] Assumes 8.5dB NECG FEC; [3] No adaptive dispersion compensation assumed, dispersion tolerances are intrinsic per modulation scheme

REFERENCES

[1] Stewart, Joe, Cisco Systems Inc., " OC-768: How? When ? Where ?", OFC 2006 workshop presentation, Anaheim, USA.

[2] IEEE 802.3 40Gb/s and 100Gb/s Ethernet HSSG
http://grouper.ieee.org/groups/802/3/hssg/public/july07/par_0707.pdf

[3] Proakis, John G., "Digital Communications", 4th Ed., 2001, McGraw-Hill, ISBN 0-07-232111-3.
[4] Proakis, John, G. "Digital Signal Processing: Principles, Algorithms and Applications", 4th Ed., 2007, Prentice-Hall, ISBN 0-13-187374-1.
[5] Taylor, Michael "Coherent Detection Method using DSP for Demodulation of Signal and Subsequent Equalization of Propagation Impairments", p 674-676, IEEE Photonics Technology Letters, vol. 16(**2**), February 2004.
[6] Laperle, Charles et al., "Wavelength Division Multiplexing (WDM) and Polarization Mode Dispersion (PMD) Performance of a Coherent 40Gbit/s Dual-Polarization Quadrature Phase Shift Keying (DP-QPSK) Transceiver", Paper PDP16, OFC 2007, Anaheim, USA.
[7] Djordjevic, Ivan, et al., "Generalized Low Density Parity-Check Codes for Optical Communications Systems", JLT 23(5), May 2005
[8] F. Gires, and P. Tournois (1964). "Interferometre utilisable pour la compression d'impulsions lumineuses modulees en frequence". C. R. Acad. Sci. Paris 258: 6112–6115.
[9] Hong, Jiang et al., "SiGe Equalizer IC for PMD Mitigation and Signal Optimization of 40Gbit/s Transmission", Paper OWO2, OFC 2005, Anaheim, USA.
[10] Heffner, Brian et al., "43Gbit/s Adaptive Polarization Mode Dispersion Compensator Field Trial", NFOEC 2007, Anaheim, USA.
[11], IEEE 802.3 Higher Speed Study Group, http://grouper.ieee.org/groups/802/3/hssg/