

Higher-Order Modulation for Client Optics

Chris Cole and Ilya Lyubomirsky, Finisar Corp.

Ali Ghiasi and Vivek Telang, Broadcom Corp.

ABSTRACT

Higher-order modulation is being studied for future 100G, 400G, and faster client optics. Amplitude, phase, and multi-subcarrier modulation approaches, used separately and in combination with each other, are being considered. Increased penalties due to higher-order modulation greater sensitivity to link impairments require use of DSP algorithms implemented in advanced CMOS nodes to close the link budget. DSP also enables programmability to reconfigure the optics for different applications and link conditions.

INTRODUCTION

The term *client optics* refers to intra- or short inter-data-center point-to-point fiber optic interfaces, with reaches from 100 m to 40 km. The term originally referred to point-to-point telecom interfaces that connect switch and transport equipment inside central offices. Client optics' bit rate is the product of five fundamental communication parameters:

- Symbols per second (signaling rate or Baud)
- Bits per symbol (modulation order)
- Number of fiber pairs (duplex physical channels)
- Number of wavelengths (optical frequency channels)
- Number of polarization states (optical spatial channels)

Multi-core fiber is being researched as a potential sixth parameter area. All deployed client optics use non-return to zero (NRZ) modulation, which encodes 1 b/symbol. The subject of this article is higher-order modulation (HOM), which encodes greater than 1 b/symbol, with at least 2 b/symbol considered useful.

OPTICAL INTERFACES

1G AND 10G CLIENT OPTICS

Mainstream deployed client optics operate at 1 Gb/s or 10 Gb/s, use NRZ modulation, one fiber pair, one wavelength, and one polarization state. Examples are IEEE 1000BASE-SX and

10GBASE-SR multimode fiber (MMF), and IEEE 1000BASE-LX and 10GBASE-LR single-mode fiber (SMF) standards, which support 10× bit rate increase by increasing the signaling rate from 1 GBaud to 10 GBaud.

40G AND 100G CLIENT OPTICS

To increase bit rate to 40 Gb/s and 100 Gb/s, both signaling rate and number of channels are increased [1]. Signaling rates for mainstream 40G and 100G optics are 10 GBaud and 25 GBaud, respectively. IEEE 40GBASE-SR4 and 100GBASE-SR4 MMF optics increase the number of fiber pairs to four, while retaining a single 850 nm wavelength. IEEE 40GBASE-LR4 and 100GBASE-LR4 SMF optics increase the number of 1310 nm band wavelengths to four, while retaining a single fiber pair. The IEEE 802.3bm Task Force (TF) is now defining the 100GBASE-SR4 MMF specification, and studying lower-cost and -power alternatives to the 100GBASE-LR4 SMF specification [2]. Alternative SMF optics approaches include multiple SMF pairs, forward error correction (FEC), and HOM. While TF consensus has not been reached for adoption of a new 100G SMF standard, the studied approaches are important for future SMF interfaces.

400G AND FASTER CLIENT OPTICS

IEEE 802.3 is likely to start a 400G Ethernet Study Group (SG) in March 2013. The anticipated focus is on a new medium access control (MAC) and physical coding sublayer (PCS) specification. To enable quick time to market and moderate development cost, the SG will consider first generation 400G optics approaches that reuse mainstream 100G 25 GBaud NRZ technology and simply increase the number of channels from 4 to 16. This results in the same bit per second port density, cost, and power as 100G. To improve these metrics, a follow-on IEEE 802.3 SG and TF will have to define second generation 400G optics with a lower number of channels. Some of the higher-density alternatives that may be considered are eight 50 GBaud NRZ wavelengths and four 50 GBaud 2 b/symbol HOM wavelengths.

Research is now ongoing into 1 Tb/s and higher bit rates, for which NRZ modulation requires an impractical number of channels. For example, a likely Ethernet follow-on bit rate to 400 Gb/s is 1.6 Tb/s, for which 25 GBaud and 50 GBaud NRZ modulation require 64 and 32 channels, respectively. To improve input/output (I/O) density and per bit power dissipation, reduction in the number of lanes through use of HOM is required [1]. Twisted-pair Ethernet copper (Cu) interfaces already use HOM for this, and to mitigate intersymbol interference (ISI), loss, and crosstalk. For example 100BASE-TX (100 Mb/s), 1000BASE-T (1 Gb/s), and 10GBASE-T (10 Gb/s) use 3-level, 5-level, and 16-level modulation, respectively. Insight gained from Cu interfaces is being applied to the study of HOM for client optics.

MODULATION ALTERNATIVES

GENERAL IMPLEMENTATION

When the per-lane bit rate of an interface is above the switching speed of digital integrated circuits (ICs), analog signal processing techniques that combine lower-rate tributaries into higher-rate lanes are required. The drawback of analog processing is the limitation on the complexity and precision of algorithms that can be implemented to deal with link impairments.

For client optics, digital IC switching speed is not the major per lane rate limitation. This makes digital signal processing (DSP) practical for dealing with penalties due to greater susceptibility of HOM to optical link impairments. These include ISI due to electro-optic (EO) component bandwidth and nonlinearities, laser relative intensity noise (RIN), fiber channel multipath interference (MPI), and receiver thermal and shot noise. A DSP-based intensity modulation direct detection (IMDD) client optics block diagram is shown in Fig. 1. The transmit (TX) path includes FEC encoder, TX DSP, clock multiplier, digital-to-analog converter (DAC), and linear TX optics. Two primary TX optics implementations are: 1) distributed feedback (DFB) direct modulation laser (DML), and 2) continuous wave (CW) DFB laser with Mach-Zehnder modulator (MZM). Historically, DML transmitters offer the lowest cost and power, while MZM transmitters offer the best performance. A TX optics implementation with intermediate cost and performance between the primary two is electro-absorption modulator laser (EML).

The receive (RX) path includes linear RX optics, analog-to-digital converter (ADC), clock recovery, RX DSP, and FEC decoder. The primary RX optics implementation uses direct-detection PIN photo-diode, which is inherently linear in converting optical power to electrical current, followed by a high dynamic range transimpedance amplifier (TIA), typically implemented as a variable-gain TIA in an automatic gain control (AGC) loop. The ADC often includes a variable gain amplifier (VGA) front-end stage to manage the ADC input dynamic range. The preferred RX DSP function for reducing ISI is an adaptive equalizer, shown in Fig. 2.

An alternate RX optics implementation referred to as coherent, mixes the received opti-

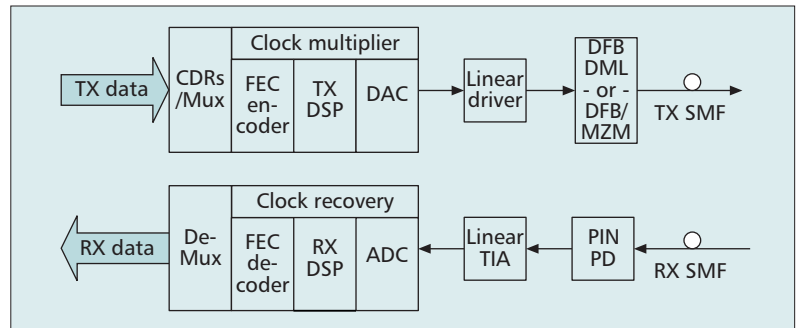


Figure 1. Higher-order modulation (HOM) duplex-SMF client optics block diagram.

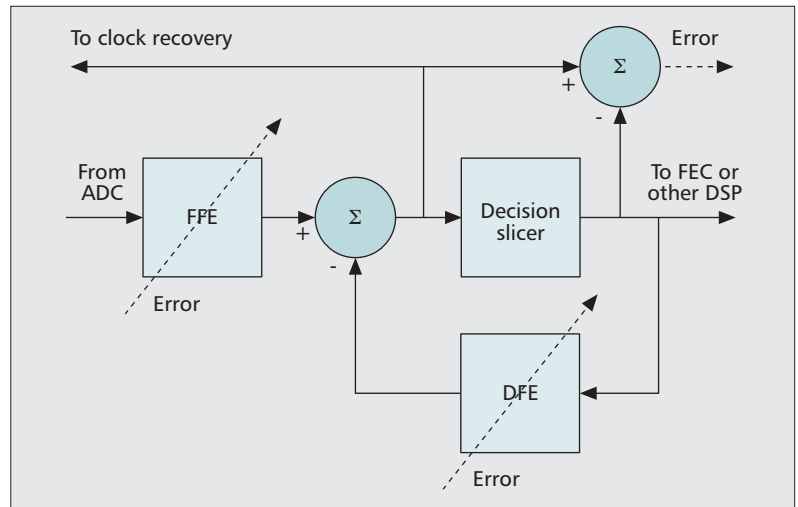


Figure 2. RX DSP adaptive equalizer block diagram.

cal signal with a local oscillator (LO) laser. The LO is either phase-locked to the received optical signal or free running with phase recovered in the RX DSP. Compared to direct detection, coherent detection has better sensitivity because of amplification by mixing with the LO, and more efficient equalization of chromatic dispersion (CD) because of preservation of phase information [3]. For client optics applications, direct detection has sufficient performance because link reaches are relatively short and CD is not a major impairment. Because coherent receivers have higher cost and power, direct detection receivers are the preferred approach.

AMPLITUDE MODULATION

In electrical baseband signaling, multilevel amplitude modulation is referred to as M -ary pulse amplitude modulation (PAM- M). For example, PAM-4 and PAM-8 refer to modulation with 4 and 8 amplitude levels, respectively. PAM- M also refers to multilevel intensity modulation in optics, even though this is not strictly accurate. Figure 3 shows a PAM-4 transmitter implementation using a DFB/MZM. The MZM is a segmented binary-weighted cascade implementing the DAC and TX optics blocks in one structure [4]. Figure 4 shows simulated eye diagrams for a realistic DFB/DML transmitter and an ideal DFB/MZM transmitter. The corresponding RX DSP includes an adaptive equalizer, for example, as shown in




Figure 3. PAM-4 TX DSP and binary-weighted DFB/MZM cascade block diagram.

Fig. 2, and inverse bit mapping. For IMDD client optics, PAM-4 is practical, and PAM-8 is the feasibility limit of modulation order because of penalties due to impairments [4].

PHASE MODULATION

Multiphase modulation of a carrier is referred to as M -ary phase shift keying (PSK- M). For example, PSK-4, also referred to as quadrature PSK (QPSK), and PSK-8 refer to modulation with 4 and 8 phases, respectively. PSK-8 is widely used

in wireless applications, for example, in the Enhanced Data GSM Environment (EDGE) mobile standard. An optics example is polarization multiplexed QPSK (PM-QPSK) used in 100G optical transport network (OTN) coherent systems.

AMPLITUDE AND PHASE MODULATION

Combined amplitude and phase modulation, or complex modulation, is referred to as quadrature amplitude modulation (QAM). For example, QAM-16 has independent PAM-4 in-phase (I) and quadrature (Q) components. In IMDD systems, QAM requires a frequency shift, which can be implemented by mixing with a subcarrier. A practical alternative is carrierless amplitude and phase (CAP) modulation, which uses finite impulse response (FIR) filters to generate I and Q components [5, 6]. A block diagram of TX DSP for a CAP-16 system and the resulting output symbol constellation is shown in Fig. 5. Corresponding RX DSP has an adaptive equalizer, for example as shown in Fig. 2, I and Q FIR filters to recover the A_k and B_k symbols, and

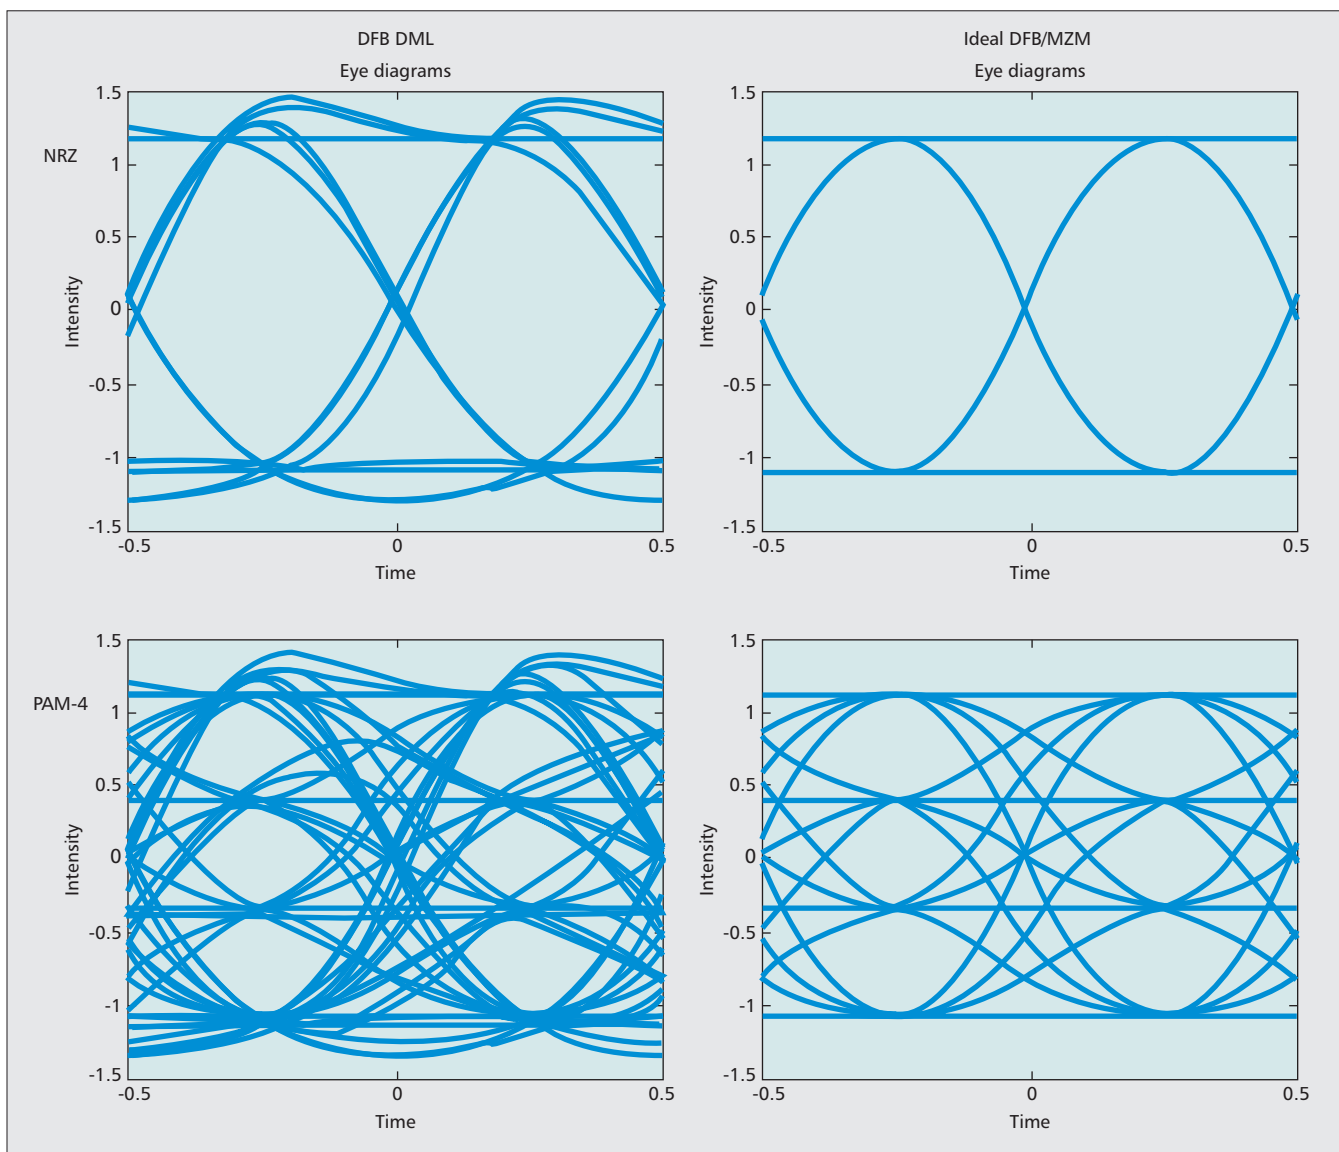


Figure 4. NRZ and PAM-4 DFB DML and ideal DFB/MZM TX eye diagrams.

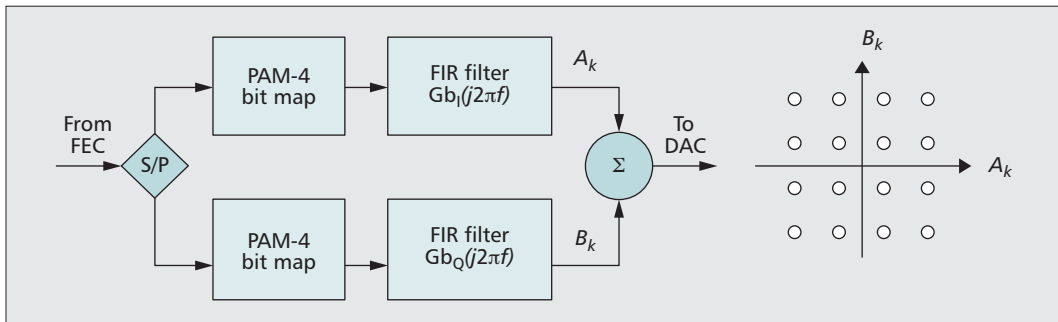


Figure 5. CAP-16 TX DSP and symbol constellation.

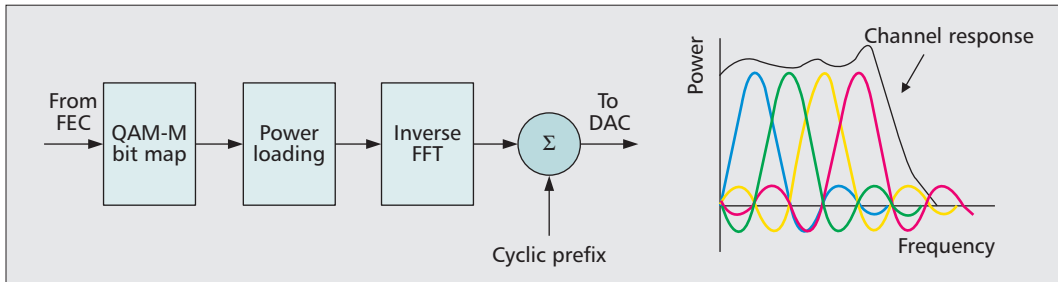


Figure 6. DMT QAM-M TX DSP and power spectrum.

inverse bit mapping. An optics example is dual-wavelength PM-QAM-16 currently in development for 400G OTN coherent systems.

MULTI-SUBCARRIER MODULATION

Modulation of multiple subcarriers is commonly referred to as discrete multitone (DMT,) with QAM of the subcarriers the preferred approach. A block diagram of TX DSP implementation based on the inverse fast Fourier transform (FFT) and the resulting output spectrum is shown in Fig. 6 [7, 8]. Corresponding RX DSP includes DMT symbol synchronization, cyclic prefix removal, FFT, frequency domain equalization, symbol detection, and inverse bit mapping. Since each subcarrier only samples a fraction of the channel frequency response, as shown in Fig. 6, DMT is robust against amplitude and phase ripples and requires a simple equalizer, for example, one tap. The drawback is high peak-to-average power ratio (PAPR), which requires clipping and higher average optical power [9]. An optics example is orthogonal frequency-division multiplexing (OFDM) being researched for use in OTN systems, and refers primarily to coherent implementations.

PERFORMANCE

SPECTRAL EFFICIENCY

HOM has higher spectral efficiency, specified in bits per second per Hertz than NRZ. Ideal NRZ Nyquist signaling at bit rate R has a symbol rate $R_S = R$ and a two-sided bandwidth $BW = R$, resulting in spectral efficiency of 1 b/s/Hz. PAM-4 modulation at same bit rate halves the symbol rate and two-sided bandwidth, doubling the spectral efficiency to 2 b/s/Hz. PAM-8 triples the spectral efficiency. The ideal CAP spectrum can be modeled as an ideal PAM spectrum shifted by an $R_S/2$ subcarrier. Therefore, at the same

symbol rate, CAP requires twice the bandwidth of PAM. However, if CAP modulation is complex, CAP-M² has the same bit rate and spectral efficiency as PAM-M, for example, CAP-16 is the same as PAM-4. With the same bit rate and subcarrier modulation, DMT requires the same bandwidth as CAP, and therefore has the same spectral efficiency as CAP. This ignores bandwidth increase due to DMT cyclic prefix overhead, which is a reasonable approximation for client optics because only a short prefix is required (e.g., two samples). For all formats, DAC and ADC minimum (Nyquist) sampling rate F_S in samples per second is equal to the two-sided bandwidth BW . The above examples are listed in Table 1.

Ideal Nyquist signal sampling requires a brick wall anti-aliasing filter frequency response or, equivalently, infinite duration $\sin(x)/x$ filter impulse response. A practical approach is to increase bandwidth and use a shaped frequency response like raised cosine. Example baseband components of PAM and CAP spectrum with roll-off factor of $\alpha = 0.5$ are shown in Fig. 7, also showing the resulting $(1 + \alpha)$ increase in two-sided bandwidth BW , minimum sampling rate F_S , and implicit CAP subcarrier. Figure 7 also shows a nearly rectangular DMT spectrum with close to ideal Nyquist bandwidth achieved by using a high number of subcarriers at the expense of high PAPR. The sampling rate F_S increases by $(1 + \beta)$ to enable practical anti-aliasing. Another approach is to use lower sampling rate F_S , including down to Nyquist sampling, at the expense of increased aliasing penalty.

Since intensity modulation spectrum is symmetric, Fig. 7 shows only the positive frequency response. In contrast, coherent modulation has distinct positive and negative frequency response, or equivalently a complex impulse response,

PSK-8 is widely used in wireless applications, for example in the Enhanced Data GSM Environment (EDGE) mobile standard. An optics example is Polarization Multiplexed QPSK (PM-QPSK) used in 100G Optical Transport Network (OTN) coherent systems.

Modulation for bit rate R	Symbol rate R_s (Baud)	Bandwidth* BW (Hz)	Spectral efficiency (b/s/Hz)
NRZ	R	R	1
PAM-4	R/2	R/2	2
PAM-8	R/3	R/3	3
CAP-16	R/4	R/2	2
DMT QAM-16	N.A.	R/2	2

* Also minimum DAC and ADC sampling rate F_s in samples per second.

Table 1. Ideal Nyquist modulation symbol rate, two-sided bandwidth and spectral efficiency.

which doubles spectral efficiency. Again, because coherent modulation and detection implementation has higher cost and power, IMDD is the preferred approach for client optics.

SIGNAL-TO-NOISE RATIO

The higher spectral efficiency of HOM has a drawback. For fixed transmit power or peak amplitude, the separation between signal levels decreases with increasing modulation order. For example, an electrical domain NRZ signal with 1 V peak-to-peak (PP) amplitude has a signal separation of 1 V. A PAM-4 signal with the same amplitude has a signal separation of only 333 mV, a decrease of 9.5 dB, as can be seen in Fig. 4. This translates into more stringent required signal-to-noise ratio (SNR) at the receiver decision slicer. As the number of modulation levels increases, the required SNR to achieve the same bit error ratio (BER) also increases.

A key figure of merit in optical systems is average optical power. Figure 8 shows a plot of BER vs. received average optical power for single-channel 100 Gb/s NRZ, PAM-4, CAP-16, and DMT QAM-16. For reference, 25 Gb/s NRZ is also included. All transmitters are ideal linear power DML, with PAPR and clipping effects included for CAP and DMT. The receivers are direct detection and limited by thermal noise. Link impairments such as RIN and MPI are not included.

The ideal model is useful for comparing modulation performance. NRZ incurs a 3 dB optical power penalty to increase bit rate from 25 Gb/s to 100 Gb/s. HOM alternatives incur at least a 7 dB penalty for the same quadrupling of bit rate. Therefore, to support applications that currently use NRZ optics, HOM implementations require higher SNR, for example, by including FEC.

IMPAIRMENTS

RIN

Laser RIN manifests itself in the receiver as an intensity-dependent noise source (i.e., higher noise for higher-level symbols). IMDD laser RIN requirements, as a function of PAM- M

order derived using Monte Carlo simulations [10], are listed in Table 2. Because SNR penalty is dominated by the noise of the highest-level symbol, the required RIN is approximately inversely proportional with bandwidth, and M^2 for large M . For example, RIN has to be improved (reduced) by ~ 3 dB with each bandwidth doubling, and by at least 6 dB with each additional symbol bit [11].

MPI

MPI results from reflections at optical interfaces in cascaded patch cords with connectors having finite reflectivity. Each patch cord acts as a Fabry-Perot Etalon cavity, and cascaded patch cords act as a set of external cavities. The reflected optical waves interfere at the receiver in a time varying manner. MPI impacts higher-level symbols more severely [4]. IMDD optical return loss (ORL) requirements, as a function of PAM- M order derived using Monte Carlo simulations [10], are listed in Table 2.

QUANTIZATION

The effective number of bits (ENOB) determines the electrical signal-to-noise distortion ratio (SNDR). The SNDR is given by

$$\text{SNDR} = 6.02 \times \text{ENOB} + 1.76$$

For high-speed DACs and ADCs that are suitable for 100 Gb/s applications, the practical ENOB limit is 6 bits. The resulting SNDR is ~ 35 dB, which becomes the HOM SNR limit.

For comparison, NRZ electrical I/O SNDR is up to 50 dB.

DSP IMPLEMENTATIONS

To meet the size and power requirements of client optics, the data conversion, DSP, and FEC blocks have to be implemented in advanced CMOS nodes like 40 nm, 28 nm, or 20 nm.

DATA CONVERSION

HOM TX waveform reconstruction requires DAC sampling rate equal to or greater than the signaling rate. An example DAC design for 100G OTN coherent transmitters has nominal 8 bits at 65 GS/s and dissipates 0.75 W in 40 nm complementary metal oxide semiconductor (CMOS) [12]. Modulation-specific approaches can reduce the DAC and TX optics implementation complexity, for example, as shown in Fig. 3 for PAM.

The primary high-speed ADC architectures are flash and successive approximation register (SAR). Historically, flash has been the preferred approach because of circuit simplicity and no error floors. An example flash ADC has nominal 6 bits at 10 GS/s, ENOB of ~ 5 bits, ~ 5 GHz single-side bandwidth, and dissipates 330 mW in 65 nm CMOS [13]. Estimated dissipation in 20 nm CMOS is 100 mW. SAR architectures use less power for equivalent performance. An example SAR ADC for 100G OTN coherent receivers has nominal 8 bits at 65 GS/s, ENOB of ~ 6 bits, ~ 20 GHz single-side bandwidth, and dissipates 1.5 W in 40 nm CMOS [12].

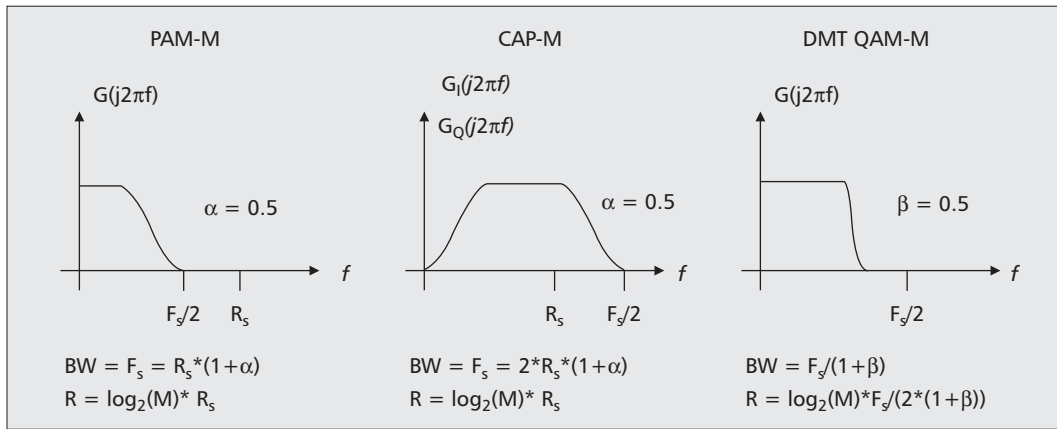


Figure 7. PAM-M, CAP-M and DMT QAM-M baseband component positive spectra, two-sided bandwidth BW , sampling rate F_s , and bit rate R .

FEC

Mainstream NRZ client optics do not use FEC. In contrast, HOM client optics require FEC coding gain to close the link budget.

The IEEE 802.3bj Task Force is now specifying Reed-Solomon (RS) FEC codes for the 100G Ethernet backplane (100GBASE-KR4 and 100GBASE-KP4) and Cu cabling (100GBASE-CR4) standards. The NRZ based KR4 specification has 0 percent overhead (in-line) RS code with 5.8 dB gain, 105 ns latency, and estimated 32 mW dissipation. The PAM-4-based KP4 specification has 3 percent overhead RS code with 7.6 dB gain, 240 ns latency, and estimated 75 mW dissipation in 20 nm CMOS.

KR4 and KP4 FECs can be used for client optics. The KR4 FEC is under consideration for 32G Fibre Channel (FC) and 100GBASE-SR4 MMF optics, as well as multi-channel SMF optics. 40G and 100G OTN SMF client optics already use G.709 RS FEC. KR4 and KP4 FECs can also be used with HOM formats like PAM-4 and CAP-16. Higher modulation order requires greater coding gain, which is possible with higher overhead and block size RS codes, increasing latency. Other codes, such as Bose-Chaudhuri-Hocquenghem (BCH), offer greater coding gain with lower latency. Even more powerful codes combine modulation and soft-decision decoding such as those used in OTN coherent systems.

EQUALIZATION

Figure 2 shows an RX DSP-based adaptive equalizer with feed-forward equalizer (FFE) and decision feedback equalizer (DFE) blocks. An FFE is characterized by the number of taps and their spacing; either T-spaced (symbol rate) or T/2-spaced (fractional). An FFE approximates the channel matched filter response and equalizes the precursor portion of the ISI. A DFE is characterized by the number of taps and cancels the post-cursor portion of the ISI. The major DFE implementation challenge is closing the feedback loop in one symbol period at high signaling rate, which requires interleaving and look-ahead techniques. A major DFE drawback is error propagation when erroneous decisions are fed back, which is especially problematic for some FEC codes. An example optical interface

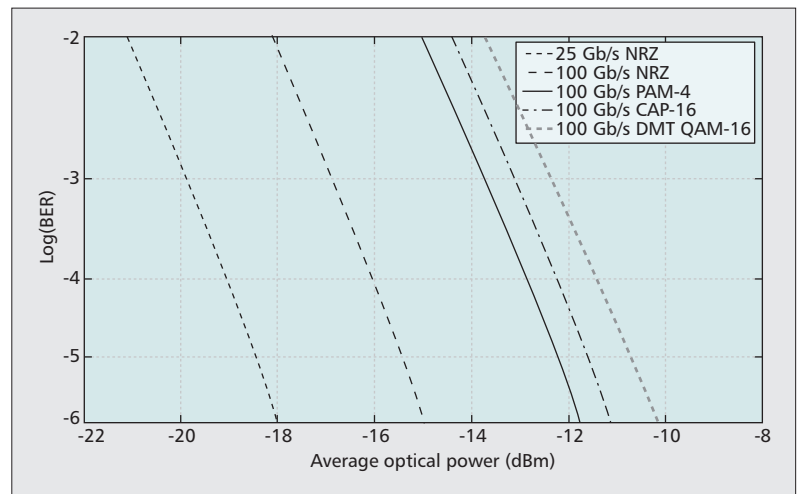


Figure 8. NRZ, PAM-4, CAP-16, and DMT QAM-16 bit error rate.

specified with an adaptive equalizer is IEEE 10GBASE-LRM, used for 10 GBaud transmission over legacy MMF, requiring a 14-tap T/2 FFE and 5-tap DFE. Another example is 100G OTN coherent receiver using FEC error feedback to optimally adapt RX parameters. This can be applied in client optics to set some parameters if low and zero BER operation is properly handled.

PROGRAMMABILITY

Programmable optics change interface parameters to enable multiple application support, link reconfiguration, and link optimization.

BAUD

Mainstream 1G and 10G client optics are single-channel NRZ, leaving only the signaling rate as a variable parameter. This is used in dual-mode MAC/PCS application-specific ICs (ASICs) that support both 1000BASE and 10GBASE, by operating the I/O at 1 Gb/s and 10 Gb/s, respectively. Unfortunately, in the 1000BASE mode this idles most of the ASIC processing bandwidth. Because lower-rate optics are a fraction of the cost of higher-rate optics, 1000BASE applications use

Modulation for 100 Gb/s	Required RIN for 0.5 dB penalty (dB/Hz)	Required ORL for 0.1 dB penalty (dB)
NRZ	-126	-18
PAM-4	-135	-23
PAM-8	-143	-26

Table 2. PAM-M laser RIN and ORL requirements for FEC-based optics at $BER = 1e-5$ [10].

separate, dedicated 1G optics. In contrast, FC storage standards support $1\times$, $2\times$, and $4\times$ rates in the same interface, for example, 2, 4, and 8 Gb/s [14]. This is possible because the $4\times$ rate FC optics are sufficiently close in cost to $2\times$ and $1\times$ rate FC optics.

Rate variability is possible with 40G and 100G client optics since the mainstream interfaces are both four channels. A dual-mode MAC and PCS ASIC with 10 Gb/s and 25 Gb/s capable I/O can be reconfigured between 40GBASE and 100GBASE modes. Similar to reconfiguration between 100BASE and 10GBASE, this idles most of the ASIC processing bandwidth in the 40GBASE mode.

NUMBER OF CHANNELS

Because 40G and 100G Ethernet interfaces are multi-channel, rate change can be based on channel reconfiguration. The advantage over just varying the per-channel rate is that ASIC processing bandwidth is fully utilized at all rates and not mostly idled at lower rates.

The IEEE 100GBASE-SR10 (MMF) interface is specified for a 10G vertical cavity surface emitting laser (VCSEL) array, and can be directly connected to $10\times$ 10G electrical I/O [15]. The array elements and electrical I/O can be reconfigured into four element groups to support two 40GBASE-SR4 interfaces, or into individual elements to support ten 10GBASE-SR interfaces. This requires the MAC to be reconfigured between 100GBASE, 40GBASE, and 10GBASE modes, and the PCS to be reconfigured between SR10, SR4, and SR modes, respectively.

To support sending 10GBASE-R bits over $4\times$ 25G electrical I/O, the Optical Interface Forum (OIF) developed the Multi-Link Gearbox (MLG) Implementation Agreement (IA) [16]. OIF is standardizing a second IA to support sending 40GBASE-R bits. The two specifications enable reconfiguring 10G VCSEL array optics with $4\times$ 25G electrical I/O into one 100GBASE-SR10 interface, two 40GBASE-SR4 interfaces, and ten 10GBASE-SR interfaces. Also $8\times$ 25G electrical I/O optics can be reconfigured into two SR10, five SR4, and twenty SR interfaces. Another approach is to use the 40GBASE physical medium attachment (PMA) specification to bit multiplex two 10G virtual lanes over one 20G electrical I/O lane. Reconfiguration from one 100G to two 40G interfaces then additionally requires changing the electrical I/O from $4\times$ 25G to $4\times$ 20G.

MLG and PMA specifications can similarly

be used to reconfigure $4\times$ 25G electrical I/O to support LR and LR4 (SMF) interfaces, although separate optics are required because fixed (not programmable) wavelengths are used.

MODULATION ORDER

The earliest programmable communication interfaces are voice-band data modems, which adapt their data rate to the maximum supported by the link. For example, the 1200 b/s v.22 interface specified by the International Telecommunication Union (ITU) uses QPSK modulation. The next rate interface, 2400 b/s v.22bis, uses QAM-16 modulation. Initially, the link starts at 1200 b/s to adopt the receive equalizer to enable operation with the higher modulation order required for 2400 b/s. This protocol is used by all higher-rate voice-band modem standards. TX and RX DSP implementations enable the rate programmability.

Digital subscriber loop (DSL) interfaces use DMT modulation, with adaptive variable rate through selection of the optimum bit rate for each subcarrier based on SNR. DSP programmability of DMT modulation optimizes the link bit rate in the face of impairments like nulls caused by wire taps.

HOM optics can similarly support rate variability. For example, changing modulation between NRZ and PAM-4, or between QPSK and QAM-16, while holding the other link parameters constant, changes the rate by a factor of two. DMT modulation offers the most programming flexibility because its adjustable parameters are modulation order and power of each subcarrier, and the number of subcarriers. Just as when signaling rate is reduced, ASIC processing bandwidth is partially idled when modulation order is reduced. The cost of the highest-rate optics is incurred when used at lower rates.

CONCLUSIONS

Mainstream 1G and 10G client optics use one channel and NRZ modulation, while mainstream 40G and 100G optics use four channels and NRZ modulation. The industry is now studying higher-order amplitude, phase, and multi-subcarrier modulation for future 100G, 400G, and faster client optics. Modulation approaches under study, like PAM, CAP, and DMT, have fewer channels and higher spectral efficiency than NRZ, but require higher SNR because of greater sensitivity to impairments like laser RIN and MPI. To support the same link budget applications supported by NRZ, DSP and FEC implemented in advanced CMOS nodes are required. Programmability of the signaling rate, number of channels, and modulation order enables reconfiguring the optics to support different applications and link conditions.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Julie Eng, senior vice president of Transceiver Engineering, Finisar Corp., and Lorenzo Longo, vice president and general manager of High Speed Interconnect Products (HSIP), Broadcom Corp., for their support.

REFERENCES

- [1] C. Cole, "Beyond 100G Client Optics," *IEEE Commun. Mag.*, Feb. 2012, pp. 58-66.
- [2] IEEE 802.3bm 40 Gb/s and 100 Gb/s Fiber Optic Task Force and Next Generation Optical Ethernet Study Group, (<http://www.ieee802.org/3/bm/index.html>)
- [3] E. Ip *et al.*, "Coherent Detection in Optical Fiber Systems," *Optics Express*, vol. 16, no. 2, Jan. 2008, pp. 753-91.
- [4] A. Ghiassi *et al.*, "Investigation of 100GbE Based on PAM-4 and PAM-8," IEEE 802.3bm Fiber Optic Task Force, Geneva, Switzerland, Sept. 2012.
- [5] J. D. Ingham *et al.*, "40 Gb/s Carrier-less Amplitude and Phase Modulation for Low-Cost Optical Data Communication Links," *OFC '11*, OThZ3, Los Angeles, CA, Mar. 2011.
- [6] J. L. Wei *et al.*, "Feasibility of 100G Ethernet Enabled by Carrier-less Amplitude/Phase Modulation and Optical OFDM," *Euro. Conf. Opt. Commun.*, Amsterdam, The Netherlands, Sept. 2012, p. 6.05.
- [7] J. A. C. Bingham, "Multicarrier Modulation for Data Transmission: An Idea Whose Time Has Come," *IEEE Commun. Mag.*, May 1990, pp. 5-14.
- [8] T. Tanaka *et al.*, "50 Gb/s Class Transmission in Single Mode Fiber using Discrete Multi-tone Modulation with 10G Directly Modulated Laser," *OFC '12*, OTh4G, Los Angeles, CA, Mar. 2012.
- [9] E. Vanin, "Performance Evaluation of Intensity Modulated Optical OFDM System with Digital Baseband Distortion," *Optics Express*, vol. 19, no. 5, Feb. 2011 pp. 4280-93.
- [10] G. Nicholl *et al.*, "Update on Technical Feasibility for PAM Modulation," IEEE 802.3 Next Generation Optical Ethernet Study Group, Kona, HI, Mar. 2012.
- [11] I. Lyubomirsky *et al.*, "100 Gb/s SMF PMD Alternatives Analysis," IEEE 802.3bm Fiber Optic Task Force, San Antonio, TX, Nov. 2012.
- [12] I. Dedic, "High Speed CMOS DSP and Data Converters," *OFC '11*, OTuN1, Los Angeles, CA, Mar. 2011.
- [13] J. Cao *et al.*, "A 500mW Digitally Calibrated AFE in 65nm CMOS for 10Gb/s Serial Links over Backplane and Multimode Fiber," *ISSCC Dig. Technical Papers*, Feb. 2009, pp. 370-71.
- [14] "Official Roadmap," Fibre Channel Industry Association, <http://www.fibrechannel.org/roadmaps>.
- [15] C. Cole *et al.*, "100GbE- Optical LAN Technologies," *IEEE Applications & Practice*, Dec. 2007.
- [16] S. Trowbridge *et al.*, "Multi-link Gearbox Implementation Agreement," Optical Interface Forum doc. no. OIF2011.381.02, Apr. 2012, pp. 12-19.

BIOGRAPHIES

CHRIS COLE [SM] (chris.cole@finisar.com) is a director at Finisar Corp., Sunnyvale, California. He received a B.S. degree in aeronautics and astronautics, and B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology. At Hughes Aircraft Co. (now Boeing SDC) and then M.I.T. Lincoln Laboratory, he contribut-

ed to multiple imaging and communication satellite programs such as Milstar. Later, he consulted on telecom ASIC design for Texas Instruments DSP Group and Silicon Systems Inc. (now Maxim.) He was one of the architects of the Sequoia coherent imaging ultrasound platform at Acuson Corp. (now Siemens Ultrasound), where he also managed hardware and software development groups. As a principal consultant with the Parallax Group he carried out signal processing analysis and product definition for several imaging and communication systems. At BBN, a Finisar acquisition, he developed 10 Gb/s and 40 Gb/s optical transceivers. He is now managing the development of 100 Gb/s and 400 Gb/s optical standards and transceivers.

ALI GHIAZI (aghiasi@broadcom.com) is an associate technical director at HSIP Group, Broadcom Corp., Santa Clara, California. He received B.S. and M.S. degrees in electrical engineering from North Dakota State University, and a Ph.D. degree in electrical engineering from the University of Minnesota. At 3M Corp. he worked on advanced optical transceivers. Then at IBM Corp. he was responsible for the development of ESCON and 1G Fibre Channel optics. At Sun Microsystems (now Oracle), he developed optical and cluster interconnect. Now at Broadcom, he is working on 100 Gb/s and 400 Gb/s interconnect, higher-order modulation for optics, and silicon photonics. He is a founder of the XFP and SFP+ MSAs.

ILYA LYUBOMIRSKY (ilya.lyubomirsky@finisar.com) is a senior principal scientist at Finisar Corp., Sunnyvale, California. He received a B.S. in electrical engineering and a B.S. in mathematics from the University of Maryland, and M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology. At ONI Systems Inc., he contributed to system design of optical metro networks. Next, he developed ultra-dense WDM systems at Ciena Corp., and researched long-haul fiber optic PIC-based communication systems at Infinera Corp. He then joined the Electrical Engineering faculty at the University of California Riverside, where he taught and researched high-speed optical communication systems. Now at Finisar, he is researching and developing 100 Gb/s and 400 Gb/s optical transceivers.

VIVEK TELANG (vtelang@broadcom.com) is a senior director of engineering at HSIP Group, Broadcom Corp., Irvine, California. He received a B.S. degree in electrical engineering from the Indian Institute of Technology, and M.S. and Ph.D. degrees in electrical engineering from the University of Notre Dame. At the Tellabs Research Center, he worked on wireline communication systems including voice-band modems and xDSL transceivers. Then at Cirrus Logic, he developed 100M Ethernet, T1E1, and xDSL transceivers. At Cicada Semiconductor (now Vitesse) he was responsible for the design of a 1000BASE-T system. He designed 10GBASE-KR and SFP+ transceivers at Broadcom, where he is now developing 25 Gb/s and 50 Gb/s SerDes. He is the author of numerous patents and technical publications.

Mainstream 1G and 10G client optics use one channel and NRZ modulation, while mainstream 40G and 100G optics use four channels and NRZ modulation. The industry is now studying higher-order amplitude, phase, and multiple-subcarrier modulation for future 100G, 400G, and faster client optics.