Signal Processing for High-Speed Links

Naresh Shanbhag, Andrew Singer, and Hyeon-min Bae

Abstract The steady growth in demand for bandwidth has resulted in data-rates in the 10s of Gb/s in back-plane and optical channels. Such high-speed links suffer from impairments such as dispersion, noise and nonlinearities. Due to the difficulty in implementing multi-Gb/s transmitters and receivers in silicon, conventionally, high-speed links were implemented primarily with analog circuits employing minimal signal processing. However, the relentless scaling of feature sizes exemplified by Moore's Law has enabled the application of sophisticated signal processing techniques to both back-plane and optical links employing mixed analog and digital architectures and circuits. As a result, over the last decade, signal processing has emerged as a key technology to the advancement of low-cost, high data-rate, back-plane and optical communication systems. In this chapter, we provide an overview of some of the driving factors that limit the performance of high-speed links, and highlight some of the potential opportunities for the signal processing and circuits community to make substantial contributions in the modeling, design and implementation of these systems.

1 Introduction

Demand for high-bandwidth data networks continues to grow, requiring the deployment of a hierarchy of communication networks spanning 1000s of kilometers (ultra long-haul (ULH) optical), through 100s of kilometers (metro optical links), to 10s of kilometers (wireless cellular), through the last mile access networks (DSL and

Naresh Shanbhag
University of Illinois at Urbana-Champaign, Urbana, e-mail: shanbhag@illinois.edu

Andrew Singer
University of Illinois at Urbana-Champaign, Urbana, e-mail: acsinger@illinois.edu

Hyeon-min Bae
KAIST, Daejeon, South Korea, e-mail: hmbae@ee.kaist.ac.kr

S.S. Bhattacharyya et al. (eds.), *Handbook of Signal Processing Systems*,
DOI 10.1007/978-1-4419-6345-1_4, © Springer Science+Business Media, LLC 2010

HFC, and wireless), through 10s of meters (Ethernet LAN), to less than a meter (back-plane and chip-to-chip I/O)[1]. Interestingly, the structure of data aggregation and distribution in this network hierarchy has resulted in the highest data rates in its two extremes, ULH and metro optical (longest) and back-plane and I/O (shortest), compared to any other type of communication link. Indeed, today, both back-plane and optical links support data rates in excess of 10Gb/s. We refer to such links as high-speed links. Also interesting is the fact that high-speed links have, in the last decade, seen an extensive application of signal processing techniques, and VLSI architectures in the design of high-speed transceivers simultaneously for both link types. This chapter discusses this recent trend and the role of signal processing and integrated circuits in the design of high-speed links.

While the data rates for both back-plane and optical links exceed those of all other digital communications media, in some respects, they are among the least sophisticated, using simple on-off keying (OOK) or non-return to zero (NRZ), and baseband comparators for symbol-by-symbol clock and data recovery (CDR). For example, until recently, optical links did not employ any form of equalization, and back-plane links to this day do not incorporate any physical layer forward error-correction (FEC). In this respect, these links are quite primitive in that they haven't leveraged the tremendous advances in statistical signal processing and communication techniques that have enabled a number of important advances in both wired and wireless communications such as those in digital voice-band modems, cellular technologies, such as GSM and CDMA, broadband enabling technologies, such as cable modems and DSL modems, and OFDM technology in use in a host of wireless digital transmission standards.

The reason that simple modulation, transmitter and receiver structures were suitable for high-speed links is because at lower data-rates and short distances, both the back-plane and optical links act as infinite bandwidth channels. As a result, optical links have been the transmission media of choice in backbone networks and are rapidly making in-roads into customer premises, enterprise networks, while back-plane links dominate system-level interconnect such as those in storage area networks. However, as data rates in these links rise above a few Gb/s into the 10 Gb/s range, these links began to exhibit intersymbol interference (ISI) or dispersion in addition to noise.

Figure 1 shows block diagrams of a high-speed back-plane link and an optical link. Much of the processing in both links is similar even though significant differences exist. Serial data $d[k]$ in the back-plane transmitter in Fig. 1(a) is clocked using a transmit phase-locked loop (PLL) and sent to a pre-emphasis filter/driver to pre-equalize channel ISI. The back-plane channel is typically a copper trace, typically few tens of inches, running over an FR4 dielectric through vias or stubs connecting different metal layers and connectors linking line cards to the board. At the receiver, the channel output is processed by a receive amplifier, which filters out-of-band noise, and amplifies signal for subsequent processing. The receiver recovers a symbol-rate clock from the data using one of several clock recovery techniques. A linear equalizer (LE) or decision feedback equalizer (DFE) may further reduce ISI. The equalized output is sliced to make a decision. In an optical transmitter, an

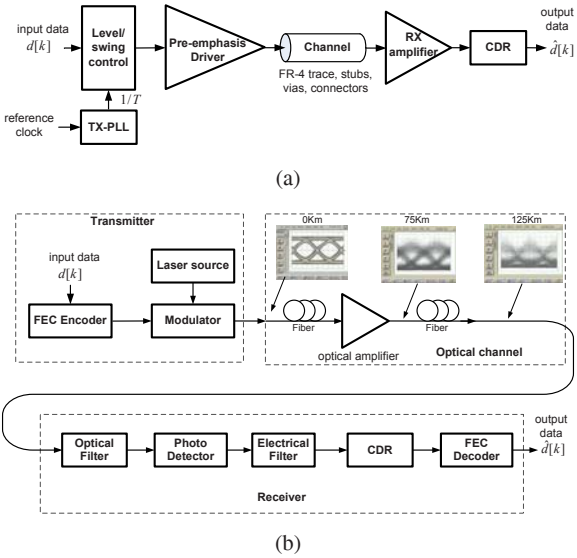

Fig. 1 High-speed link configurations for: (a) a back-plane serial link, and (b) an optical fiber link.

information sequence $d[k]$ is used to modulate the intensity of a laser source. For long-haul links (typically greater than 200 km), the optical signal will propagate along multiple spans of fiber with repeated optical amplification to overcome attenuation. Also shown in Fig. 1(b) are data *eye diagrams*, which illustrate the transmitted and received symbol patterns as they would appear on an oscilloscope at various stages through the optical fiber. Note how an *open* transmit eye from which correct decisions can easily be made using a simple slicing operation, begins to *close* due to dispersion and noise, as fiber length increases. Eye diagrams for back-plane links similarly indicate significant eye closure at the receiver input.

As a result, high-speed links have recently began to employ sophisticated signal processing and communication techniques such as equalization, and FEC. In back-plane links, FEC is only recently being deployed [2, 3], while equalization [4, 5] was introduced within the last decade. For optical links, electronic dispersion compensation (EDC) [6] based on signal processing techniques have emerged as a technology of great promise for OC-192 (10 Gb/s) data rates and beyond. In addition, the unrelenting progress of semiconductor technology exemplified by Moore's Law has provided an implementation platform that is cost-effective, low-power and high-performance. The next several decades of high-speed links will likely be dominated by the signal processing, architectures and circuit techniques that are used both for signal transmission, modulation and coding, as well as for equalization and decoding at the receiver.

The design of signal processing-enhanced back-plane and optical communication links presents unique challenges spanning algorithmic issues in transmitter and receiver design, mixed-signal analog front-end design, and VLSI architectures for

implementing the digital signal processing back-end. Though these are tethered applications, reducing power is important due to stringent power budgets imposed by the systems, e.g., boards, transponders and line-cards, in which these devices reside. For example, high-speed I/O links have a power budget of a few tens of mW/Gb/s and a target $BER <= 10^{-15}$, making their design extremely challenging. Thus, a cost-effective solution, i.e., a solution that meets the system performance specifications within the power budget, requires joint optimization of the signal procession algorithms, VLSI architectures, analog and digital integrated circuit design.

This chapter will provide an overview of some of the driving factors governing the development of current and next generation high-speed links, with a particular emphasis on the tools, models, and methods by which the signal processing and circuits communities can play an integral role. The remainder of this chapter is structured as follows. In Section 2, system models describing the functionality of link components are developed to enable simulation and modeling of end-to-end back-plane and fiber links. Section 3 describes the application of signal processing techniques for combating dispersion, and noise, in high-speed links. Section 4 provides a case study of an OC192 EDC-based receiver and the chapter concludes with some remarks on emerging areas for research in high-speed links in Section 5.

2 System Models

In this section, we briefly discuss some of the salient characteristics in high-speed links of practical signal processing interest. We pay particular attention to models that can be used for system design, performance simulation and experimentation. These are the essential building blocks from which a signal processing-based development can make a substantive, practical impact on current and future backplane and optical links. The discussion is centered around the link components, viz., the transmitter, channel and receiver.

2.1 The Transmitter

Both the back-plane and optical links primarily employ binary signaling techniques such as NRZ modulation. Unlike an optical link, the back-plane link is carrierless, i.e., the output of a back-plane channel is the baseband pulse $s(t) = \sum_k d[k]g(t-kT)$, where $d[k]$ is the information sequence, and $g(t)$ is the baseband pulse (Fig. 2(a)), typically NRZ, that captures the frequency response of the transmit driver (see Fig. 1(a)). Furthermore, unlike traditional communication links which are average power limited, the back-plane link is intrinsically peak-power limited. This is because back-plane transmit drivers are integrated in modern day low-voltage ($1V$) CMOS technologies. Both optical and back-plane transmitters also include a trans-

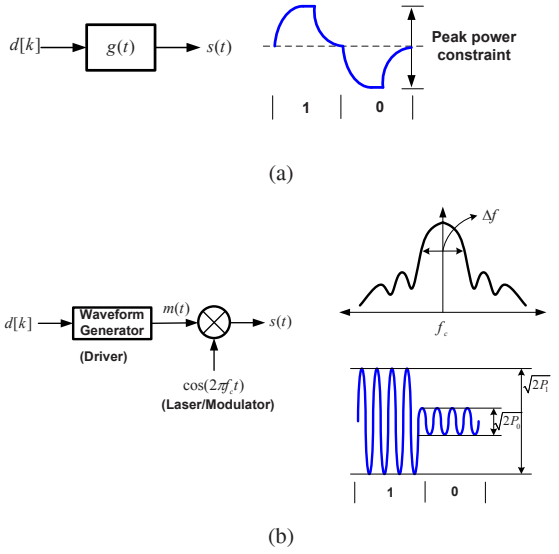

Fig. 2 Transmitter functional model for: (a) a back-plane channel, and (b) an optical channel.

mit PLL that generates a transmit clock. The jitter on the transmit clock plays a strong role in determining the *BER* in the back-plane receiver.

Figure 2(b) describes the modulation process for an optical link. The transmitter output signal is given by

$$s(t) = \sum_k d[k]g(t - kT)cos(2\pi f_c t) \tag{1}$$

where f_c is the carrier frequency of the optical source, typically around 193 THz for 1550 nm wavelength[1] Note that the transmit signal spectrum for an NRZ transmitter has finite bandwidth. The optical signal is hard to attenuate completely at such high data rates. Thus, one measure of the quality of the laser is the *extinction ratio* E_r (see Fig. 2)

$$E_r = 10\log_{10}\frac{P_1}{P_0}. \tag{2}$$

where P_1 and P_0 are the powers in a transmitted *one* and *zero*, respectively. Extinction ratios of 8 dB are typical for directly modulated lasers employed in short reach, while long-haul applications employ externally modulated lasers using Mach-Zehnder modulators [7] with an extinction ratio of around 15 dB. Unlike back-plane links where higher signal to noise ratio (*SNR*) usually implies a higher quality signal and a lower bit-error rate *BER*, for optical communications, a higher optical signal-

[1] The center frequency is 228 THz for 1310nm wavelength, which can be obtained using the relationship $f\lambda = c$.

to-noise ratio (*OSNR*) may not provide a lower *BER*, if the extinction ratio is poor. The power in the "signal" component of the *OSNR* includes that in both the *ones* and the *zeros*, while the *BER* of an optical link will be a function of the difference in these optical powers.

2.2 The Channel

Channel modeling is a critical first step in developing transmitter and receiver algorithms, architectures and circuits. Back-plane channel models tend to be simpler than those for fiber due to the presence of non-linearities in the latter. This subsection describes the back-plane and fiber channel models.

2.2.1 Back-plane Channel Model

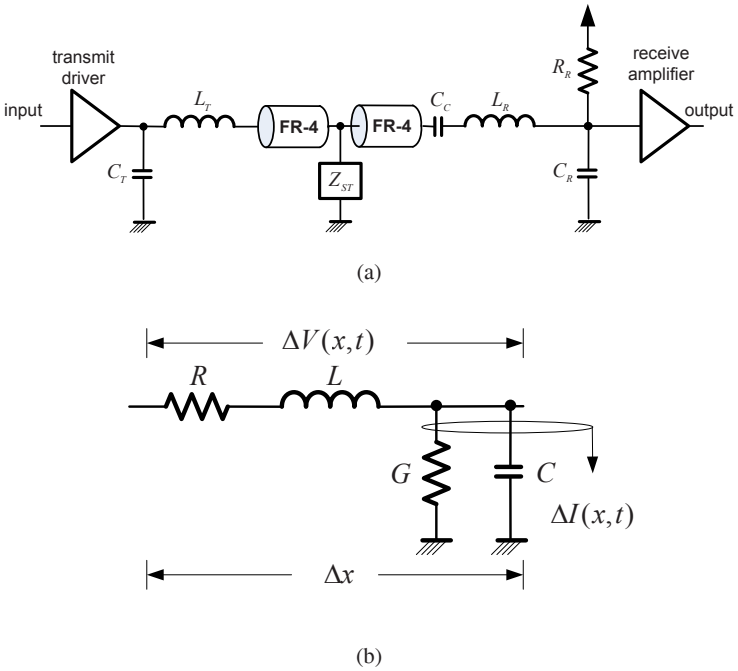

Fig. 3 Back-plane channel model: (a) a circuit model of the link, and (b) a lossy transmission line element.

A complete back-plane channel begins at the output node of the transmit driver, and ends at the input node of the receive amplifier. Starting from the transmitter,

a simplified channel model (see Fig. 3(a)) includes the transmit side pad capacitance C_T, bond wire inductance L_T, a transmission line model of the FR-4 trace, impedance Z_T representing discontinuities due to vias and unterminated stubs, and ending with coupling capacitor C_C (in some cases) at the receiver, receive termination resistor R_T (either off-chip or on-chip), bond wire inductance L_R, and receive side pad capacitance C_R. The transmission line model for the FR-4 trace can be described by Telegrapher's equations shown below, which in turn can be derived via the application of KCL and KVL to the RLGC circuit in Fig. 3(b), as follows:

$$\frac{\partial V(x,t))}{\partial x} = RI(x,t) + L\frac{\partial I(x,t)}{\partial t} \tag{3}$$

$$\frac{\partial I(x,t)}{\partial x} = GV(x,t) + C\frac{\partial V(x,t)}{\partial t} \tag{4}$$

where $V(x,t)$ and $I(x,t)$ are the voltage and current signals, respectively, at time t from an initial time $t = 0$, and distance x from the source, and R, L, G, and C, are the series resistance, series inductance, shunt conductance and shunt capacitance per unit length. The elements R and G result in loss in the transmitted signal energy. In addition, the back-plane channel suffers from frequency-dependent loss mechanisms at frequencies above few GHz owing to *skin effect* and *dielectric losses*. Skin effect arises due to crowding of high-frequency current towards the surface of the conductor leading to the resistance R becoming frequency-dependent, i.e., $R = R(f)$ above a certain frequency. The skin depth δ is given by

$$\delta = \frac{1}{\sqrt{\pi f \mu \sigma}} \tag{5}$$

where μ and σ are the permeability and resistivity, respectively, of the conductor. For $f \geq f_s$, where f_s is the frequency at which $\delta = r_w$ (r_w is the radius of the FR-4 trace), skin-effect kicks in and the resistance $R(f) \approx R(0)(0.5(r_w/\delta) + 0.26)$. Otherwise, for $f < f_s$, $R(f) = R(0)$. Dielectric loss occurs due to absorption of electromagnetic energy in the dielectric surrounding the conductor. Dielectric loss is specified in terms of the *loss tangent* $\tan\delta$, which is the ratio of the resistive to the reactive component of the electromagnetic energy in the dielectric material, and is proportional to f.

These loss mechanisms result in the back-plane channel being a band-limited low-pass channel, whose frequency response is given by:

$$H(f) = \exp(-\gamma l) \tag{6}$$

where $\gamma(f)$ is the frequency dependent propagation coefficient that captures the skin-effect and dielectric-loss effects, and l is the trace length. Typical values of the parameters for a FR-4 trace are $\tan\delta = 0.035$, and $\varepsilon_r = 4.3$. Figure 4(a) shows that a typical 20" FR-4 channel exhibits a loss of $20dB$ at around $4GHz$, which is close to the $-3dB$ bandwidth of an NRZ pulse. This indicates that the back-plane channel suffers from severe ISI ranging over approximately 15-to-20 symbol-periods. In

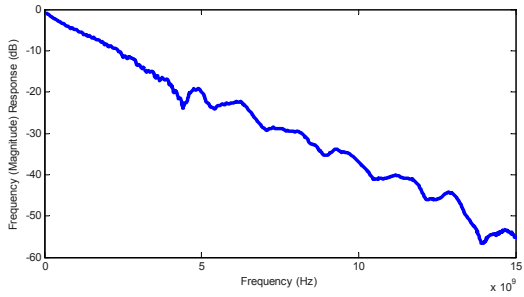

(a)

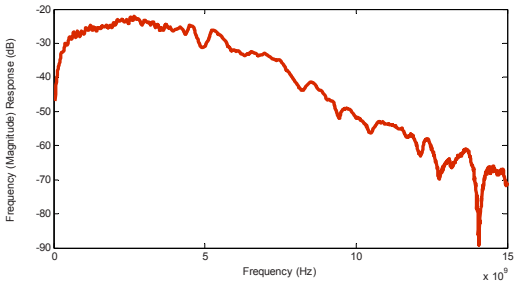

(b)

Fig. 4 Time and frequency domain responses for a 20" FR-4 back-plane channel: (a) data channel, and (b) cross-talk channel.

addition, impedance discontinuities due to vias, stubs, and connectors, cause reflections, which result in notches in the frequency response and further aggravate the interference problem. The back-plane channel also experiences far-end cross-talk from neighboring traces as shown in Fig. 4(b).

2.2.2 Optical Fiber Channel Model

Optical fiber comes in two types: single-mode fiber (SMF) and multi-mode fiber (MMF). Single-mode fiber (SMF) permits only a single optical mode (see Fig. 5(a)) to propagate thereby enabling greater propagation distances before pulse spreading occurs. Thus, SMF is employed in long reach applications and some short reach ones [8]. Longer reach applications use 1550 nm wavelengths as this represents the wavelength of minimum loss due to absorption.

Propagation of the optical signal over SMF fiber is governed by the nonlinear Schrödinger equation (NLSE)[8]. Figure 5(a) shows that the propagating wave is supported by two orthogonal axes of polarization, or *polarization states* within the fiber. A fraction ρ^2 of the light energy couples into one of the polarization states

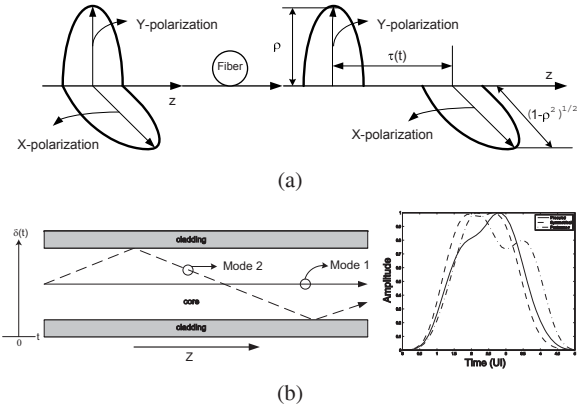

Fig. 5 Optical fiber: (a) signal propagation single-mode fiber, and (b) in multi-mode fiber. The outputs shown are three sample impulse responses developed by the IEEE 802.3aq working group to model modal dispersion.

while the remainder $(1-\rho^2)$ couples into the other polarization state. These two states experience different group velocities, resulting in the pulse arrival times at the receiver being separated by $\tau(t)$. This differential group delay experienced by these two states gives rise to polarization mode dispersion (PMD) causing pulse spreading at the output of the photodetector in the receiver. A length L fiber with attenuation α and PMD results in the following time-varying impulse response:

$$h(t,\tau) = e^{-\frac{\ln(10)\alpha L}{20}}(\rho\delta(t-\tau_0) + \sqrt{1-\rho^2}\delta(t-\tau_0-\tau(t))), \tag{7}$$

where τ_0 is the baseline delay of the fiber and $\tau(t)$ is the instantaneous DGD of the fiber due to first-order PMD. The mean DGD of the fiber, expressed in $ps/\sqrt{km}$, gives a measure of the expected delay between the two polarization axes. A typical value of the mean DGD is $1ps/\sqrt{km}$ [9]. A good model of the variability of PMD over time is captured by [10]:

$$f_\tau(\tau) = \frac{32\tau^2}{\pi^2\tau_m^3}\exp\left\{-\frac{4\tau^2}{\pi\tau_m^2}\right\}$$

where $f_\tau(\tau)$ is the marginal distribution, and τ_m is the mean DGD. We see that the instantaneous DGD of a fiber will exceed three times its mean DGD with probability 4.2×10^{-5}.

Chromatic dispersion (CD) occurs due to the dependence of the index of refraction on frequency. This gives rise to a difference in the velocities with which the frequencies within the transmitted signal spectrum propagate resulting in pulse spreading in time. The resulting equivalent baseband linear time-invariant model for the channel is given by

$$H_{\mathrm{CD}}(f) = \exp\left\{-j\pi D\frac{\lambda^2}{c}Lf^2\right\}, \tag{8}$$

where D is the dispersion parameter at wavelength λ, L is the fiber length, f is the equivalent baseband frequency. Typical fiber parameter values are listed in Table 1.

Table 1 Typical parameter values for single mode fiber used in long haul communications.

parameter	symbol	value
wavelength	λ	1550 nm
propagation velocity	c	2.997925×10^8 m/s
fiber attenuation	α	0.2×10^{-3} dB/m
dispersion	D	17 ps/nm-km
second order dispersion	$\beta_2 = -Dc(10^{-6})/(2\pi f_c^2)$	20×10^{-27} s^2/m
dispersion slope	β_3	1.1×10^{-40} s^3/m
center frequency	f_c	193.1×10^{12} Hz
core refractive index	n_2	2.6×10^{-20} $\mathrm{m}^2\mathrm{W}^{-1}$
core area	A	90×10^{-12} m^2
nonlinearity	$\gamma = 2\pi n_2/\lambda A$	5×10^{-3} $\mathrm{m}^{-1}\mathrm{W}^{-1}$

A simple model for signal propagation in MMF captures the coupling of the modulated laser source onto a set of N_m propagating modes, each of which propagates a slightly different distance through the fiber (see Fig. 5(b)). If the detector area is sufficiently large, then the output of a square-law detector yields a sum of powers of the modal arrivals since the modes are spatially orthogonal. The following finite-length impulse response model for MMF propagation can be obtained when modal noise and mode coupling is ignored:

$$h(t) = \sum_{k=0}^{N_m-1} e^{-\ln(10)\alpha L/20} h_k \delta(t - \tau_0 - \tau_k), \tag{9}$$

where $\sum_{k=0}^{N_m-1} |h_k|^2 = 1$, N_m is the number of propagating modes, an amount equal to $|h_k|^2$ of the incident power is coupled into the kth mode, τ_k is the differential modal delay with respect to the overall baseline delay of τ_0 induced by the wavelength of the source, α is the attenuation, and L is the fiber length. In the presence of modal coupling, the terms h_k would become time-varying and can be modelled as Rayleigh fading under suitable assumptions on mode orthogonality [11].

2.3 Receiver Models

The receiver processing in both back-plane and optical links begins with a filter that removes out-of-band noise. In back-plane systems, the receive filter is an electrical component, which also serves as an amplifier, and an equalizer. In an optical link, the optical signal is optically filtered and then passed on to a device for optical-to-

electrical (O/E) conversion, which generates an electrical current signal. We refer to this device as the O/E detector in order to distinguish it from statistical detectors, such as slicers, that appear later in the receive chain. A transimpedance amplifier (TIA) which converts the detector current into a voltage signal for subsequent processing.

Both links have synchronization sub-systems referred to as a clock-recovery unit (CRU), which recovers a baud-rate clock from the receiver signal. The output clock of a CRU is employed to sample the output of the receive filter (or the TIA in case of an optical link), using a decision-making latch or slicer. The combination of the CRU and slicer is referred to as a clock-data recovery (CDR). A CDR, including the front-end filter (TIA in case of optical links), implements the bulk of processing in optical and electrical receivers. However, a CDR employs little or no signal processing and thus fails to achieve the requisite *BER* in the presence of ISI. As a result, modern day receiver for both link types, include additional components such as equalizers (both digital and analog), ADCs, in addition to a CDR. In the optical context, such signal processing enhanced links are called EDC-based links.

2.3.1 O/E Detectors

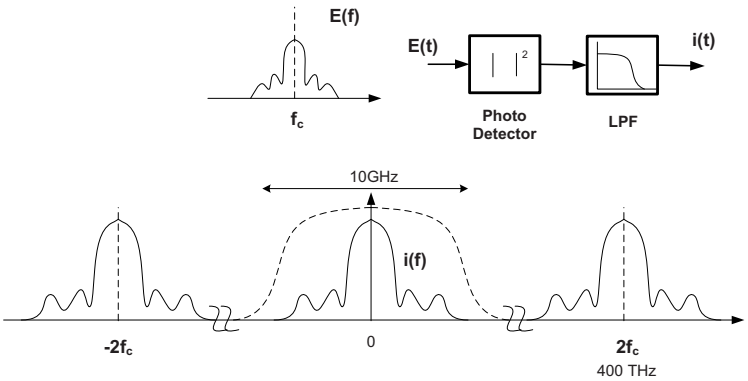

Fig. 6 A block diagram model for a photodetector receiver is shown as an ideal square-law detector followed by a low pass filter.

Typical optical detectors include the avalanche photodiode (APD) and the PIN diode, each of which produce an electrical current proportional to the incident optical power. Intensity modulation can be modeled as $E(t) = m(t)\cos(2\pi f_c t)$, with f_c around 200 THz, as in Fig. 6. A photodetector in the absence of noise and dispersion, can be viewed as a bandlimited square-law device, approximately producing an output current

$$i(t) = \mathscr{R} \int_{t-\tau}^{t} |E(t)|^2 dt = \mathscr{R} \int_{t-\tau}^{t} \frac{|m(t)|^2}{2} dt$$

where $\mathscr{R}$ is the responsivity of the detector (typically between 0.7 A/W and 0.85 A/W for APD and PIN receivers), owing to the finite integration bandwidth of the detector, where $1/\tau << f_c$ is typically on the order of 10 GHz, for a 10 Gb/s link, for example. The double-frequency term induced by the quadratic detector is therefore filtered out, leaving only the baseband demodulated intensity signal $m^2(t)$, in the absence of dispersion. Since $m(t) > 0$ for intensity modulation, the data can be completely recovered from $m^2(t)$ by a simple CDR, in the absence of noise and dispersion.

2.3.2 Analog Front-End (AFE)

The AFE refers to the linear signal processing blocks in the receive path. The back-plane AFE consists of a variable gain amplifier (VGA) for signal amplification. Usually, the receive amplifier also provides some high-frequency peaking in order to compensate of the channel roll-off, and limits out-of-band noise. In EDC-based optical links, the AFE consists of a cascade of the TIA and the VGA.

The TIA is a linear amplifier that takes an input current $i_{in}(t)$ and generates an output voltage $v_o(t)$ where

$$v_o(t) = -i_{in}(t) R_f \frac{A_v}{1 + A_v} \tag{10}$$

where $R_f \frac{A_v}{1+A_v}$ is the transimpedence (in Ohms) and A_v is the open loop gain of the TIA amplifier kernel. Typical values of R_f reside in the range $800\,\Omega$ to $2k\,\Omega$. The 3dB bandwidth of the TIA $f_{c,tia}$ is typically designed to be approximately 70% of the symbol rate in order to achieve a balance between ISI and SNR. The performance of the TIA is also specified in terms of its sensitivity.

A VGA is typically used to boost the receive signal level if subsequent processing will include more than a simple CDR. Linearity of the VGA (for back-plane), and the detector-TIA-VGA combination (for an optical link) is important if an ADC is used, in order to preserve the information. The VGA can be modeled as a transconductance stage with an output load resistor R_L and an output capacitance C_L. The low-frequency gain of such a stage is given by

$$A_v = -g_m R_L \tag{11}$$

where g_m is the transconductance of the transistor, and its 3dB bandwidth is given by

$$f_{c,vga} = \frac{1}{2\pi R_L C_L}. \tag{12}$$

2.3.3 Analog-to-Digital Converter (ADC)

An ADC is typically modeled as an ideal sampler followed by a memoryless quantizer. Such models do not account for the many non-idealities that exist in practice especially in the multi-GHz regime. A practical ADC consists of a band-limited front-end, a sampler and a quantizer, and may exhibit hysteresis and other pattern-dependent behavior. Non-idealities of the ADC are often aggregated in terms of the effective number of bits (ENOB) of the ADC, which attempts to quantify the resulting quantization noise as if it were additive and independent of the input. At 10 GS/s, ADCs with 3-to-6 bits of resolution have been produced for both optical and back-plane channels using SiGe bipolar processes[12, 13] (optical only), and more recently in CMOS for back-plane [14], optical [15], and both [16], illustrating the strong emergence of signal processing enhanced solutions for both link types.

2.3.4 Clock-Recovery Unit

Design of a high-performance CRU is critical in high-speed links. Unlike other communication links such as DSL or wireless, the CRU in back-plane and optical can dominate the complexity of the receiver. This is very true for a CDR-based receiver. The CRU generates symbol timing information and the receiver clock from the received analog signal in the presence of dispersion, noise, and input jitter.

A PLL forms the basis of the operation of a CRU. A PLL can be viewed as an extremely narrowband filter with its passband centered at a frequency equal to the symbol-rate. A traditional PLL is referred to as a *linear PLL* because the PLL func-

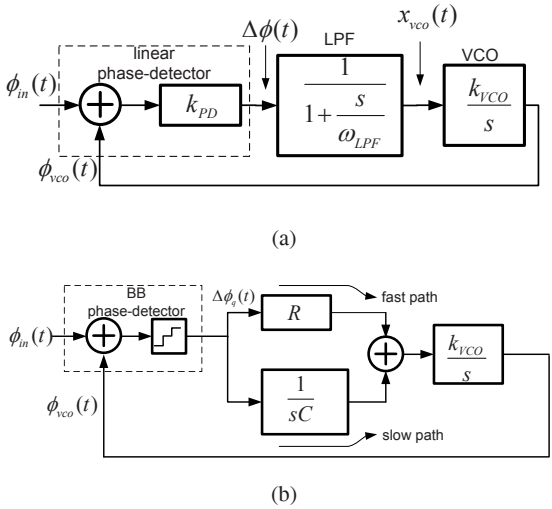

Fig. 7 Phase-locked loop (PLL): (a) linear, and (b) binary or bang-bang PLL.

tionality in the phase domain can be easily described using linear feedback system theory. The linear PLL in Fig. 7(a) employs a linear phase-detector (PD) to compute the instantaneous phase-error $\Delta\phi(t)$ between the input phase $\phi_{in}(t)$ and the clock phase $\phi_{vco}(t)$. The phase error is integrated using a low-pass filter (LPF) whose output is used to control the input voltage, and hence the frequency, of a voltage-controlled oscillator (VCO). At high-speeds, the linear PD is unable to generate narrow pulses that represent small phase-errors. Thus, most high-speed links employ a *binary* or a *bang-bang* (BB) PLL (see Fig. 7(b)). The operation of a BB-PLL is best understood by considering a 1^{st} order loop, which is obtained by eliminating the slow-path in Fig. 7(b). In a 1^{st}-order BB-PLL, a quantized, typically 3-level, phase-error $\Delta\phi_q(t)$ is generated using a BB phase-detector (BB-PD). The BB-PD ensures that the VCO frequency switches between two extremes $f_{nom} \pm f_{bb}$, also called the *lock range*, where f_{nom} is the nominal data frequency, which is made equal to the symbol-rate, and f_{bb} is the frequency step. Thus, unlike in a linear PLL, the VCO in a BB-PLL operates at two discrete frequencies separated by $2f_{bb}$. A 1^{st}-order BB-PLL locks as long as the average data frequency lies in the lock range. Otherwise, the VCO starts to hunt for the right frequency and in the process generates hunting jitter, whose amplitude is proportional to the lock range. A separate frequency-acquisition loop (not shown) is usually implemented in order to acquire the symbol frequency. Even then, high-frequency jitter wander, as dictated by the jitter tolerance requirements, can push the BB-PLL out of lock resulting in hunting jitter. This presents an intrinsic limitation of a 1^{st}-order BB-PLL in that its lock range and its output jitter amplitude are tied together via a single parameter f_{bb}. The 2^{nd}-order BB-PLL is able to achieve the conflicting goals of having a large lock range and on minimizing output jitter amplitude, by processing the quantized phase-error through two paths: (1) a direct/proportional path for fast updates, and (2) an integral path for slow updates. The proportional path is able to track high-frequency jitter such as those due to data, while the slow path manages to track the high-frequency jitter wander thereby resulting is small output jitter.

CRUs in optical links have much more stringent requirements than those in backplane channels. This is because in long haul applications, the CRU needs to satisfy SONET (Synchronous Optical NETwork) specifications on jitter transfer, jitter tolerance and jitter generation. Jitter manifests itself as a random spread in the zero-crossings of a signal. The input jitter J_{in} is a function of the input *SNR* (SNR_{in}), for which an approximate relationship is given by

$$J_{in} = k\frac{T_{tr}}{\sqrt{SNR_{in}}} \tag{13}$$

where T_{tr} is the transition period, and k is a dimensionless constant.

A typical CRU uses a wideband PLL for tracking the jitter in the input signal so that the correct sampling phase is used for the data. However, the resulting recovered clock at the output of the PLL will contain jitter. A second narrowband PLL is used to clean-up the jitter in the clock at the output of the first PLL. This requires the

use of a first-in-first-out (FIFO) buffer at the output of the sampler to align data and output clock.

2.4 Noise Models

High-speed links suffer from various noise sources. In addition to ISI, these noise sources fundamentally limit the achievable *BER*.

2.4.1 Back-plane Noise Models

Noise in back-plane links is dominated by [17]: (1) jitter in the recovered clock, and (2) thermal noise from the AFE and the input termination (see Fig. 3(a)). Jitter refers to the random variation in the zero-crossings of a nominally periodic signal. Clock jitter in the receiver is caused by a number of sources such as jitter in the transmit PLL, noise in the channel and power supply, jitter generated by the receive PLL in the CRU. Jitter analysis [18] is an important aspect of back-plane link design. It can be shown that the sampled received signal in the presence of jitter and ISI is given by:

$$y[n] = \sum_k d[k]h_g[n-k] - \sum_k v[k]q[k]h_{t0}[n-k] - s[n]\sum_k v[k]h_{t0}[n-k] \tag{14}$$

where $h_g[k]$ is the sampled pulse-response (composite of transmit shaping $g(t)$ and channel impulse response $h(t)$), $h_{t0}[k] = h(t0+nT)$ is the sampled channel impulse response with a sampling offset of $t0$, $v[k] = d[k] - d[k-1]$ is the data transition, $q[k]$ is the jitter in the transmit clock, and $s[n]$ is the jitter in the recovered clock. The first term in (14) represents ISI, while the second and third terms represent the impact of jitter. Note: jitter appears as additional noise source in the received signal, and that jitter appears only when there is a data transition (non-zero $v[k]$).

Thermal noise in back-plane links are generated by the termination resistors (R_T in Fig. 3(a)), the sampler, and the receive amplifier. The voltage noise PSD (in V^2/Hz) is given by:

$$N_o = 4kT(R_T + R_{sw} + \frac{\alpha}{g_m}) \tag{15}$$

where R_T, R_{sw}, g_m and α are the termination resistance, sampler resistance, input stage transconductance, and a technology dependent parameter, respectively. Assuming typical values of $R_T = 50\Omega$, $R_{sw} = 200\Omega$, $g_m = 1mA/V$ and $\alpha = 1$, we get $N_o = 2\times 10^{-17}V^2/Hz$.

In addition, the slicer imposes a minimum input voltage requirement referred to as *slicer resolution*, needed to resolve the input signal. The slicer resolution is a function of static offsets that occur in differential circuit structures due to transistor mismatch. Typical values of uncorrected offsets are in the range of $10mV$.

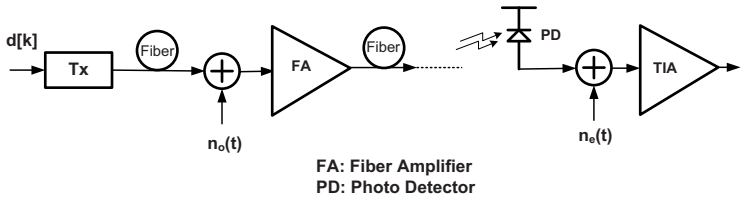

Fig. 8 Block diagram of an optical link model showing primary sources of noise.

2.4.2 Optical Fiber Noise Models

In amplified links, the dominant source of noise as observed at the receiver is amplified spontaneous emission noise $n_o(t)$, or ASE noise (see Fig. 8). In unamplified links, the dominant source of noise is the electrical receiver noise $n_e(t)$, which encompasses all sources of noise in the receiver, including that of the detector, the TIA and and subsequent electronics.

The OSNR is computed as the ratio of the signal power in the transmitted optical signal to the ASE noise power in a reference bandwidth at a reference frequency adjacent to the transmitted signal. The typical noise reference bandwidth used is 0.1 nm, which is approximately 12.5 GHz for 1550 nm transmission.[2] As such, to simulate a given OSNR level, the noise variance used in a T_s-sampled discrete-time simulation should be

$$\sigma^2_{\text{ase}} = E\{|E(t)|^2\}10^{-\text{OSNR}/10}\frac{W_{\text{samp}}}{W_{\text{ref}}},$$

where $W_{\text{samp}} = 1/T_s$, is the sampling rate of the numerical simulation, $W_{\text{ref}} = 12.5$ GHz for a 0.1 nm noise resolution bandwidth and OSNR is the target OSNR. Since the signal will propagate along two orthogonal axes of polarization, to adequately model the ASE noise, a separate noise simulation can be run for each polarization axis with the same target OSNR and the results added after detection. For directly detected optical signals, additive Gaussian noise in the optical domain, becomes signal-dependent and non-Gaussian in the electrical domain. The conditional probability density function for the detector output $r(t)$ given the value of noise-free channel output $y(t)$ can be shown to be a signal-dependent non-central chi-square distribution with $2M$ degrees of freedom [19].

For modeling or numerical simulation, each of the component models given in this section can be combined to create an end-to-end system model. The system model can be employed to design transmitter and receiver algorithms, VLSI architectures and integrated circuits for high-speed links, in order to meet system specifications.

[2] This is computed by noting that $f = c/\lambda$ and therefore $|\Delta f| = (c/\lambda^2)|\Delta\lambda| = 12.478GHz$, for $\lambda = 1550$ nm and $|\Delta\lambda| = 0.1nm$.

3 Signal Processing Methods

In the last several years, a number of efforts have emerged to bring signal processing into back-plane and optical communcation links. In this section, we will explore some of the basic models and structures that have been proposed and brought into practice that rely on digital signal processing methods to improve the performance of optical links.

3.1 Modulation Formats

NRZ is the dominant modulation format for both back-plane and optical links. However, other more spectrally-efficient techniques have also been explored. In back-plane, modulation formats such as 4-PAM [20] and more recently duobinary (DB) [21, 22] have been explored. Duobinary has also been explored for optical links [23]. Duobinary modulation refers to the use of a controlled amount of ISI in the transmitted signal or in the composite channel response. In back-plane channels, duobinary is implemented by modulating the amplitude while in optical links, it is implemented using phase modulation. The transmitted power spectrum of DB-modulated signals is roughly half that of NRZ modulation resulting in reduced ISI induced by the channel.

3.2 Adaptive Equalization

In practice, the channel in high-speed links are unknown. Therefore, adaptive equalization techniques are needed that can learn the channel parameters via a process of training. Many of the adaptive equalization techniques employed in high-speed links tend to be simplified versions of those in DSL and wireless. A key difference being that the high data-rates in high-speed links force the designer to explore a rich mix of analog and digital equalization approaches instead on a predominantly digital approach as in other communication links. The baseband output $r(t)$ of a high-speed link can be written in terms of the transmitted bit sequence $d[n]$ as

$$r(t) = \sum_{k=-\infty}^{\infty} d[k]h_g(t-kT) + w(t),$$

where $h_g(t) = g(t) * h(t)$ is the convolution of the transmit pulse $g(t)$ with the channel impulse response $h(t)$ assuming that the latter is time-invariant. It is well-known [24] that the optimal receive filter for such a linear modulation over a linear channel is a matched filter to the transmit pulse $h_g(t)$ and that a set of sufficient statistics for the optimal detection of the sequence $d[n]$ (for any criteria of optimality) can be

obtained by sampling the output of this matched filter at the symbol-synchronous intervals, $t = nT$.

An attractive means of implementing such a matched filter when the channel response is not known precisely in advance, is the use of an adaptive fractionally-spaced equalizer (FSE), whose output can be written

$$y[n] = y(nT) = \sum_{k=0}^{N} c_k r(nT - k\tau), \tag{16}$$

where the the delay-line spacing of the N taps in the equalizer, τ, is typically a fraction of the symbol period, i.e. $\tau = T/M_f$, where $M_f = 2$ is typical. The coefficients of the tapped delay line, or feed-forward equalizer (FFE), can be adapted using a variety of methods such as the least-mean square (LMS) algorithm, followed by subsequent symbol-spaced sampling. In principle, $y[n]$ can exactly reproduce the T-spaced outputs of the optimal continuous-time matched filter $h_g(nT - t)$ when the highest frequency in $r(t)$ is less than $1/2\tau$, i.e., $\tau-$spaced sampling satisfies the sampling theorem. FSEs have the added benefit of having performance that is relatively insensitive to the sampling phase of the receiver, which is an important feature in high-speed links, where the recovered clock is often frought with jitter.

FSEs with $T/2$-spacing have been implemented with reasonable success for both back-plane [25] and MMF channels [26]. These tend to cover a span of 3-to-7 T-spaced symbols. A block diagram of a τ-spaced FFE is shown in Fig. 9(a). FFE structures are typically implemented with analog tapped delay-line circuits in CMOS, where the cost (and therefore the power consumption) of a transceiver need to be kept to a minimum. Challenges include the realization of accurate, low-loss delay-lines, such that the spacing τ between adjacent taps are equal and equal to a fraction of the bit-period, and that the signal does not decay too substantially as it is passed from the input to the end of the delay line. For symbol decisions $\hat{d}[n] = 1$ for $y[n] > \eta$ and $\hat{d}[n] = 0$ for $y[n] \leq \eta$ and error signal $e[n] = (\hat{d}[n] - y[n])$, the coefficient update algorithms using LMS, for step-size μ,

$$c_k[n] = c_k[n-1] + \mu r(nT - k\tau)e[n], \quad k = 0, 1, \ldots, N-1$$

need not be overly complex, since the resolution with which the multiplications (16) can be carried out is limited.

FFEs are not particularly well-suited to channels with deep spectral nulls such as those in back-plane channels with discontinuities. For such channels, the DFE shown in Fig. 9, is a natural extension of the FFE. A DFE typically consists of an FFE with an additional linear (or non-linear) filter used to process past symbol decisions $\hat{d}[n]$ in order to cancel post-cursor ISI. Assuming past symbol decisions are correct, then the FFE portion of a DFE converts as much of the pre-cursor ISI as possible, within its degrees of freedom, into post-cursor ISI, the effects of which will be subtracted off by the feedback filter. The output of the DFE with a linear feedback filter is given by

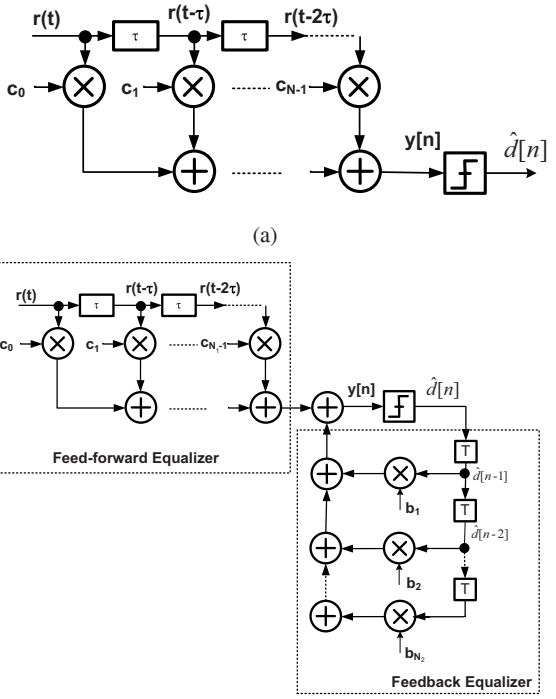

Fig. 9 Linear equalization techniques: (a) block diagram of a fractionally-spaced feedforward linear equalizer (FFE), and (b) block diagram of a fractionally-spaced decision feedback equalizaer (DFE).

$$y[n] = \sum_{k=0}^{N_1-1} c_k[n] r(nT - k\tau) - \sum_{j=1}^{N_2} b_j[n] \hat{d}[n-j], \tag{17}$$

whose LMS coefficient updates are given by

$$c_k[n] = c_k[n-1] + \mu r(nT - k\tau) e[n], \quad k = 0, 2, \ldots, N_1 - 1 \tag{18}$$

$$b_j[n] = b_j[n-1] + \mu \hat{d}[n-j] e[n], \quad j = 1, 2, \ldots, N_2. \tag{19}$$

3.3 Equalization of Nonlinear Channels

Unlike back-plane and MMF channels, the SMF channel is non-linear. This is due to the presence of a direct (square-law) O/E detector in the receiver front-end. As a result, there is no discrete-time equivalent baseband channel, and adaptive equalization techniques, such as the FFE or DFE discussed previously, do not provide much benefit for SMF applications. Some success for small amounts of CD has been

achieved using sub-optimal techniques such as DFEs with state-dependent thresholds [27] or with nonlinear feedback taps, i.e. where the feedback filter contains not only prior detected bits, but also pairwise products [26].

In order to develop optimum detection techniques, the SMF channel is described as a nonlinear continuous-time channel with finite memory whose output at the detector can be written,

$$r(t) = S(t; d[n], d[n-1], \ldots, d[n-L_c]) + n(t), \tag{20}$$

where $n(t)$ comprises optical noise in an amplified link or electrical noise in an unamplified span of fiber, and $S(t; d[n], d[n-1], \ldots, d[n-L_c])$ describes the noise-free, state-dependent, component of the signal, which is a function of the L_c adjacent transmitted symbols. Sufficient statistics for optimal sequence detection [28] for such continuous-time channels with memory can be obtained from a baud-sampled (with perfect symbol timing) bank of matched filters. For data rates at or above 10 Gb/s, processing such matched filters at line rate, together with an MLSE engine poses significant implementation challenges.

3.4 Maximum Likelihood Sequence Estimation (MLSE)

The optimum receiver in terms of minimizing the BER for a nonlinear channel with memory, such as the optical fiber channel, is given by the maximum a-posteriori probability (MAP) detector and can be computed using the BCJR algorithm [3], which is a two-pass algorithm, requiring a forward and backward recursion over the entire data sequence. The MLSE algorithm provides the optimal detector in terms of minimizing sequence error rate (SER) rather than BER. The MLSE algorithm provides nearly the same performance as MAP detection for uncoded transmissions, or when symbol detection is carried out separately from FEC decoding, using only a single forward-pass through the data.

The Viterbi algorithm [30] can be used to recursively solve for the SER-optimal transmitted sequence consistent with the observations given a channel model at the receiver that can describe the probability density governing the observations $r[n]$. Specifically, we can define the set of possible states at time index k as the 2^{L_c} vectors $\mathbf{s}[k] = [d[k-1], \ldots, d[k-L_c]]$, where $d[n] \in \{0,1\}$. In addition, we define the set of valid state transitions from time index k to time index $k+1$ as the 2^{L_c+1} possible vectors of the form $\mathbf{b}[k] = [d[k], d[k-1], \ldots, d[k-L_c]]$. Using this notation, the noise-free outputs of our channel model $x[n] = f(d[n], \mathbf{s}[n])$ and channel observations $r[n] = x[n] + w[n]$, where the noise-free outputs, the channel observations, and the noise samples, $w[n]$, can be either scalars (in the case of baud-sampled statistics), or vectors (in the case of over-sampled channel outputs or state-dependent matched filter outputs).

Given a sequence of channel observations $\mathbf{r}[k] = \{r[k], r[k-1], \ldots, r[k-L_r+1]\}$ of length L_r, the channel when viewed as a 2^{L_c}-state machine can exhibit $2^{L_c+L_r}$

possible state sequences $\tilde{\mathbf{s}}[k] = \{\mathbf{s}[k], \mathbf{s}[k-1], \ldots, \mathbf{s}[k-L_r+1]\}$ (owing to the shift-structure of adjacent states). Each valid state sequence corresponds to a *path* in the trellis representation of the channel (see Fig. 10). The MLSE algorithm seeks a particular state sequence or a path $\hat{\mathbf{s}}[k]$ for which the probability of the channel observations sequence $\mathbf{r}[k]$ given $\hat{\mathbf{s}}[k]$ is maximum, i.e.

$$\hat{\mathbf{s}}[k] = \arg\max_{\tilde{\mathbf{s}}[k]} P(\mathbf{r}[k]|\tilde{\mathbf{s}}[k]) = \arg\max_{\tilde{\mathbf{s}}[k]} P(\mathbf{r}[k]|\tilde{\mathbf{s}}[k])P(\tilde{\mathbf{s}}[k]) = \arg\max_{\tilde{\mathbf{s}}[k]} P(\mathbf{r}[k], \tilde{\mathbf{s}}[k]) \tag{21}$$

assuming that each state sequence is equally likely. Assuming further that the noise samples are conditionally independent, we can write

$$P(\mathbf{r}[k], \tilde{\mathbf{s}}[k]) = \prod_{k=0}^{L_r} P(\mathbf{s}[k+1]|\mathbf{s}[k]) \prod_{k=0}^{L_r} P(\mathbf{r}[k]|\mathbf{s}[k+1], \mathbf{s}[k]),$$

and define the *branch metric*

$$\lambda(b[k]) = -\ln P(\mathbf{s}[k+1]|\mathbf{s}[k]) - \ln P(\mathbf{r}[k]|\mathbf{b}[k] = \{\mathbf{s}[k+1], \mathbf{s}[k]\})$$

to quantify the likelihood that a given state transition occurred. Then, we obtain the log probability of any state sequence, or path, as a sum of branch metrics along the state transitions in that path,

$$L(\tilde{\mathbf{s}}[k]) = -\ln(P(\mathbf{r}[k], \tilde{\mathbf{s}}[k]) = \sum_{k=0}^{L_r} \lambda(\mathbf{b}[k]),$$

which is also referred to as the path metric. The label $\hat{S}(s[k])$ is used to designate the shortest path (the path with the smallest path metric) through the trellis ending in state $s[k]$, which is known as a *survivor path* since all longer paths also ending in state $s[k]$ are discarded. The path metric of the survivor path ending in state $s[k]$ is labeled $M(s[k])$, i.e. $M(s[k]) = L(\hat{S}(s[k]))$.

This is best visualized using a trellis, as shown in Fig. 10, where all permissable/valid state transitions $\mathbf{b}[k]$ have equal probability. The Viterbi algorithm is now a dynamic programming approach to finding the path through the trellis of smallest accumulated branch metric.

MLSE implementation in high-speed links brings up a number of practical considerations. First, the link needs to be ADC-based. This constraint is no longer a major issue for data-rates in the 10Gb/s range today though power considerations may preclude their use in many back-plane applications in the short-term. Second, since the data is continuously transmitted, it will be necessary, for latency reasons, to output bit decisions as the data is processed. This is called the *sliding window* version of the Viterbi algorithm, which requires a *finite look-back* window of length D. It is desirable to make the look-back window as long as possible, in order to reduce the probability that more than one of the survivor paths (see Fig. 10) will exist at time $k-D$ when processing at time k. However, the storage requirements and complexity of the algorithm per output symbol will increase linearly in D. We

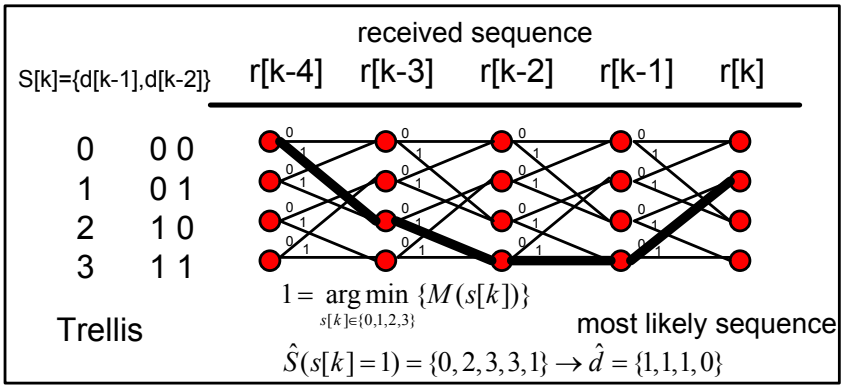

Fig. 10 Maximum likelihood sequence estimation (MLSE): A trellis for a 4-state ($L_c = 2$) MLSE algorithm with $L_r = 5$ is shown. The trellis diagram illustrates valid state sequences with solid lines as paths through the trellis. Above each line connecting two states is the corresponding symbol that would have been transmitted for that transition.

have found that setting D equal to a small multiple of the channel memory L_c works well in practice, and setting $D > 3L_c$ does not improve performance significantly. For D in this range and when more than one survivor path exists at $k - D$, selecting the one with the smallest path metric at k appears to work well in practice.

As illustrated in Fig. 10, the conventional Viterbi decoder would employ the current received sample $r[n]$ to compute the branch metrics, and update the four path metrics that correspond to the likelihood of the best path (survivor path) ending at each of the four states in one symbol period. The survivor paths are stored in memory. The path metric update step is followed by the trace-back step where one of the survivor paths is traced back to a predetermined depth, D, and a decision is made. The survivor path metric update is performed using an add-compare-select (ACS) operation, which is recursive and therefore also difficult to implement at optical line rates.

A probabilistic model for the observations $r[k]$ given the state transitions $\mathbf{b}[k]$, i.e., given the transmitted bits $d[k], \ldots, d[k - L_c]$, is required for branch metric computations. In a practical applications, this model must be determined adaptively from the channel observations $r[n]$ without the aid of any training, i.e. without any knowledge of the true bit sequence $d[n]$. One such model for a baud-sampled receiver uses an adaptive system identification algorithm based on a Volterra series expansion of the channel impairments [12, 31, 32]

$$r[n] = c_0 + \sum_{k=0}^{L_c} a_k d[n-k] + \sum_{k=0}^{L_c} \sum_{\ell=0, \ell \neq k}^{L_c} b_{k,\ell} d[n-k] d[n-\ell] + \cdots, \qquad (22)$$

where, $r[n] = r(nT)$, are baud-sampled receiver outputs, $d[n] \in \{0, 1\}$ are assumed transmitted channel symbols, and whose parameters c_0, a_k and $b_{k,\ell}$ can be adaptively

estimated using the LMS algorithm.

$$c_0[n] = c_0[n-1] + \mu e[n] - \nu(c_0[n] - c_0[0]), \tag{23}$$

$$a_k[n] = a_k[n-1] + \mu \hat{d}[n-k]e[n] - \nu(a_k[n] - a_k[0]), \;\; 0 \le k \le L \tag{24}$$

$$b_{k,\ell}[n] = b_{k,\ell}[n-1] + \mu \hat{d}[n-k]\hat{d}[n-\ell]e[n] - \nu(b_{k,\ell}[n] - b_{k,\ell}[0] \;\; 0 \le k, \ell \le L \tag{25}$$

where, $e[n] = (r_{ADC}[n] - r[n])$, for ADC outputs $r_{ADC}[n]$, $\hat{d}[n]$ are past decided bits from the MLSE algorithm, and ν is a step-size parameter for the regularization terms $c_0[0], a_k[0]$ and $b_{k,\ell}[0]$, which are used to incorporate prior knowledge and to mitigate accumulation of finite-precision effects. In a practical implementation, LMS updates can be performed at a large fraction of the line rate to reduce power, which also enables the use of MLSE outputs, or perhaps FEC-corrected outputs, when making LMS updates. Given the model coefficents, (22) can be used to pre-compute the noise-free outputs for each state transition in the trellis. The branch metrics can then be computed using either a parametric or non-parametric model for the noise statistics within each state. We have found that a suitably chosen parametric model for the noise statistics [12], e.g., state-dependent Gaussian, works well in practice, requires relatively few parameters to be estimated, and enables the use of Euclidean branch metrics, which simplifies the architecture.

4 Case Study of an OC-192 EDC-based Receiver

Designing EDC-based receivers at OC-192 rates, i.e., 9.953 Gb/s to 12.5 Gb/s, poses numerous signal processing and mixed-signal circuit design challenges. In this section, we illustrate some of the issues that arise in implementing EDC using mixed-signal ICs by describing the design of a maximum likelihood sequence estimation (MLSE) receiver [12] with adaptive non-linear channel estimation. Key issues include: the design of a baud-sampled ADC, clock recovery in the presence of dispersion, non-linear channel estimation in the absence of a training sequence, and the design of a high-speed Viterbi equalizer. This is the first reported design of a complete MLSE receiver design meeting the specifications of OC-192/STM-64 long-haul(LH), ultra long-haul (ULH), and metro fiber links at rates up to 12.5 Gb/s.

4.1 MLSE Receiver Architecture

The MLSE receiver was designed as a two-chip solution in two different process technologies that were chosen to match the functionality to be implemented. The analog signal conditioning, sampling and clock recovery functions were implemented in an analog front-end (AFE) IC in a 0.18 μm, 3.3 V, 75 GHz, SiGe BiCMOS process. The digital MLSE equalizer and the non-linear channel estimator

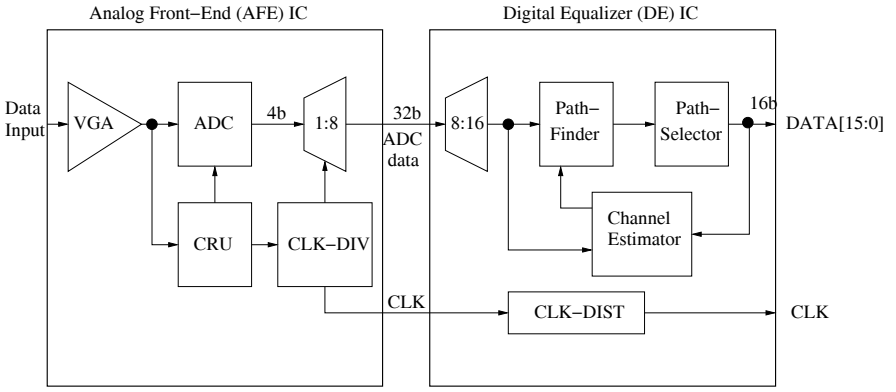

Fig. 11 Block diagram of MLSE-based receiver.

were implemented in a digital equalizer (DE) IC in a $0.13\,\mu$m, 1.2 V CMOS process.

The architecture of the MLSE receiver is shown in Fig. 11. The received signal is at a line rate of between 9.953 Gb/s for uncoded links and up to 12.5 Gb/s for FEC-based links. The received signal is amplified by a VGA then sampled by a 4-bit flash ADC. A dispersion-tolerant CRU recovers a line rate clock for the ADC. The 4-bit line rate ADC samples are demuxed by a factor of 1:8 to generate a 32-bit interface to the DE IC, which implements a 4-state MLSE algorithm, i.e. it assumes a channel memory of 3 symbol periods. The digital equalizer includes a blind, adaptive channel estimator of the form in (22) and a data decoding unit on-chip that requires no training.

4.2 MLSE Equalizer Algorithm and VLSI Architecture

Designing an MLSE equalizer at OC-192 line rates (9.953 Gb/s-12.5 Gb/s) is extremely challenging because of the high data rates involved. Conventional high-speed Viterbi architectures often employ parallel processing or higher radix processing [33]. Parallel/block processing architectures tend to suffer from edge-effects while higher-radix architectures have been shown to achieve speed-ups of up to a factor of 2, which is not sufficient for this application.

Instead, the architecture in [12, 34] (see Fig. 12) incorporates a number of techniques that enable the MLSE computations to complete within the bounds imposed by the high OC-192 line rates. First, the channel trellis is traversed in a time-reversed manner (see Fig. 12(a)). Doing so eliminates the trace-back step and the survivor memory associated with conventional approaches thereby providing a considerable power saving while enhancing throughput.

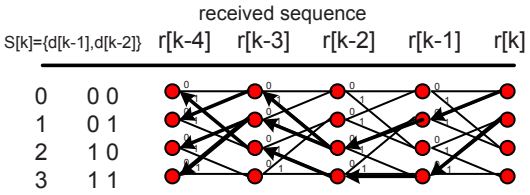

(a)

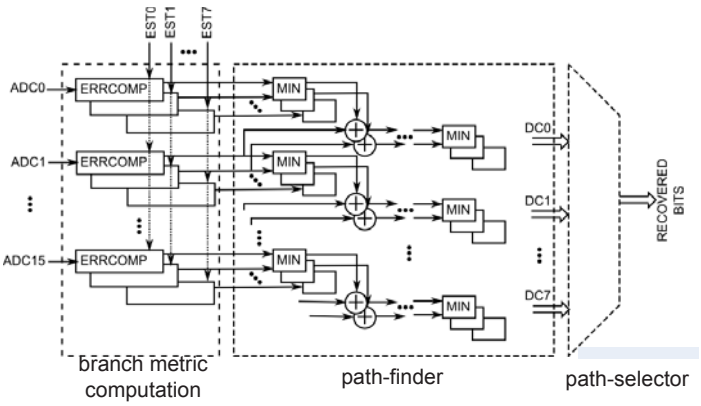

(b)

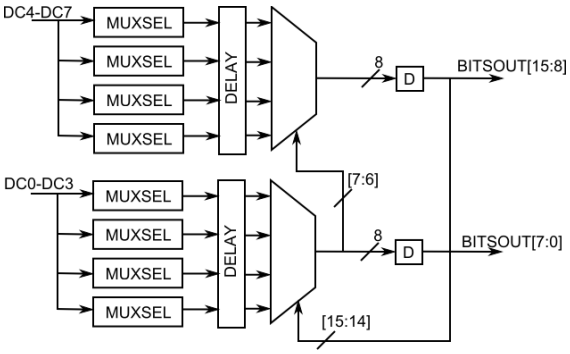

(c)

Fig. 12 High-speed MLSE architecture: (a) time-reversed trellis, in which a cold-start at $r[k]$ leads to 4 possible reverse-survivor paths ending at states $s[k-D]$, (b) VLSI architecture, and (c) path-selector.

Second, the ACS bottleneck associated with conventional approaches is eliminated by restricting the survivor path memory to a reasonably small number (6) of symbol periods. This number is obtained by noting that the channel memory spans roughly 3 symbol periods, and the structure of the channel trellis. With a fixed survivor depth D, and time-reversed traversal of the trellis, it becomes possible to implement the ACS computations in a fully feed-forward architecture. This feed-forward ACS unit is referred to as the *path-finder* block in Fig. 12. A feed-forward architecture [35] can be easily pipelined in order to meet an arbitrary speed requirement.

Third, in order to avoid edge effects, past decisions are employed to select among the four possible survivor paths. Though this choice results in a recursive stage referred to as the *path-selector* (see Fig. 11) consisting of a series of multiplexers it is not difficult to meet the speed requirements. This is because the path-selector is small and localized part of the architecture and hence easily amenable to circuit optimizations.

The three innovations described above provide a favorable trade-off between BER with ease of implementation. Though these do not reduce frequency at which the architecture needs to be clocked, which remains at up to 12.5 GHz, they do result in an architecture that can be easily transformed into one that can operate at a lower clock-frequency as described next.

The fourth innovation is to make multiple (N_d) decisions along the chosen survivor path in each clock cycle resulting in a clock frequency that needs to be $1/N_d$ times the line rate. Fifth, the architecture is unrolled [35] in time by a factor of N_u resulting a further reduction of clock frequency by a factor of N_u. Thus, the final clock frequency is given by

$$f_{clk} = \frac{R}{N_u N_d} \tag{26}$$

where R is the line rate. Algorithmically, unfolding [35] does not ease the throughput problem as it expands the number of computations in the critical path by the same factor as it lengthens the clock period. This fact does not present a problem in this case because the critical path is localized to the path-selector, which as mentioned earlier, is a simple cascade of multiplexers and hence can be custom designed to meet the speed requirements. A side benefit of making multiple decisions and employing unfolding is that the equalizer doubles up as a $1 : N_d N_u$ demultiplexer as well. Most receivers need a 1:16 demux functionality which is usually implemented using a separate demux chip. By choosing $N_d = 2$ and $N_u = 8$, the DE chip demultiplexes and decodes the data simultaneously.

4.3 Measured Results

The MLSE receiver was tested with up to 200 km of SMF. The MLSE receiver achieves a BER of 10^{-4} at an OSNR of 14.2 dB with 2200 ps/nm of dispersion as shown in Fig. 13(a). A received power of -14 dBm is used throughout the testing.

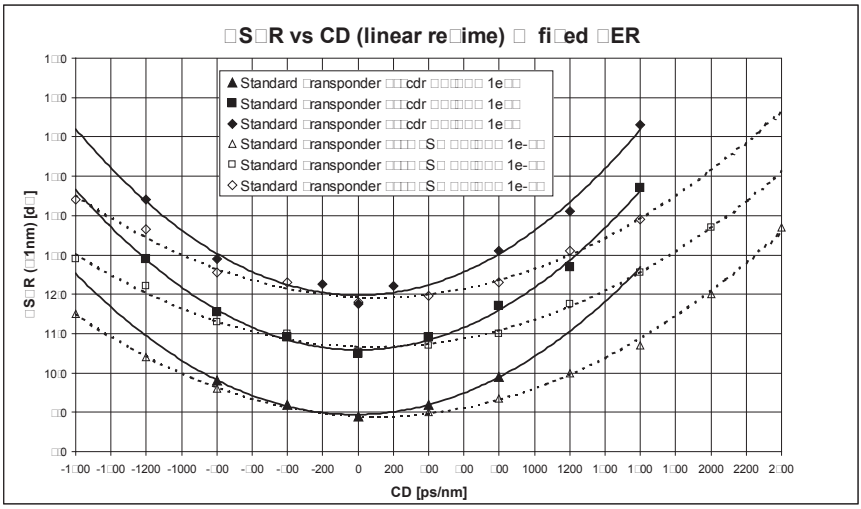

(a)

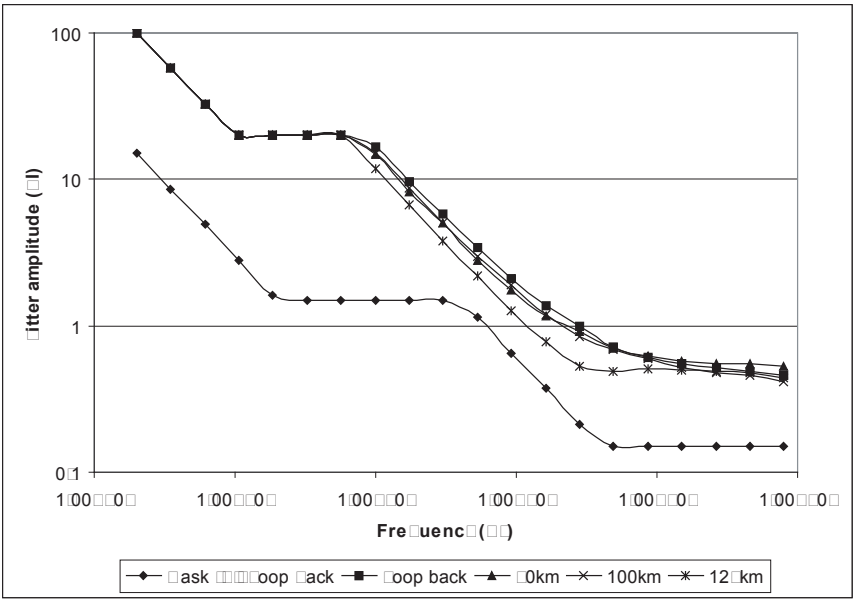

(b)

Fig. 13 MLSE receiver test results: a) chromatic dispersion test, and b) SONET jitter tolerance.

The MLSE receiver reduces the OSNR penalty of a standard CDR by more than 2 dB at CD of 1600 ps/nm and BER of 10^{-4}. The penalty for a standard CDR increases rapidly beyond 1600 ps/nm of CD. Figure 13(a) also shows that the MLSE receiver does not have a penalty in the back-to-back (i.e. 0 km) configuration.

The receiver has been shown to provide an error free (BER $< 10^{-15}$) post-FEC output, with a pre-FEC BER of 10^{-3} at 10.71 Gb/s with 2000 ps/nm of dispersion. It can compensate for 60 ps of instantaneous differential group delay (DGD) with a 2 dB OSNR penalty (BER=10^{-6}). The channel estimator in the MLSE engine adapts at a rate of 30 MHz and thus is easily able to track variations in DGD, which are typically less than 1 KHz.

The MLSE receiver satisfies the SONET jitter tolerance specifications with 2200 ps/nm of dispersion (see Fig. 13(b)). A HP OMNIBER OTN is used for this measurement. The measured output clock jitter is $< 0.5\,\text{ps}_{\text{rms}}$. The output BER is kept at 10^{-3} and the CD is varied up to 2200 ps with a $2^{31}-1$ PRBS. The fixed BER is achieved by adjusting the OSNR while maintaining the received optical power at -14 dBm. The PLL output clock jitter is less than $0.64\,\text{ps}_{\text{rms}}$ across test conditions. The PLL maintains locked without cycle slips with up to 1300 consecutive identical digits (CID) at a BER of 10^{-2}, with 125 km SMF-28 optical fiber link.

The MLSE receiver illustrates the benefits of jointly addressing signal processing algorithm design, VLSI architectures, mixed-signal integrated circuit implementation, as well as electromagnetic signal integrity issues in packaging and circuit board design.

5 Advanced Techniques

In this section, we will describe a few of the directions in which new signal processing advances are being made in high-speed links of the future. In particular, we focus on the application of FEC to reduce power in back-plane links, and the use of coherent detection in optical links.

5.1 FEC for Low-power Back-Plane Links

As described earlier, modern state-of-the-art back-plane links today rely exclusively upon an equalization-based transceiver to compensate for ISI and noise in order to achieve $BER < 10^{-15}$ [5, 20]. This is highly unusual because most other communication links today, such as DSL, wireless, optical, disk drives, and others, employ some form of FEC in order to achieve the requisite *BER*. In such links, the equalizer achieves a *BER* in the range of 10^{-3}-to-10^{-4} at its output, and the FEC to further reduces the *BER* by orders-of-magnitude to 10^{-12}-to-10^{-15}. Therefore, in the absence of FEC, back-plane links are designed to be ISI-dominated with high $SNR > 30dB$, and hence consume more power than necessary. Though digital techniques such as

FEC tend to consume more power than analog, scaling of feature sizes favor digital designs more than analog. Thus, it is expected that in nanoscale process technologies, a digital heavy FEC-based I/O link [2], and even perhaps an ADC-based link, might be an attractive approach to reduce power.

In an (n,k,d_{min}) code, a block of k data/information bits/symbols (dataword) are mapped a block of n code bits/symbols (codeword) where $n > k$. Such a code is said to have a code-rate of $r = \frac{k}{n}$. FEC results in the line-rate being greater than the data rate by a factor of $1/r$. Thus, a coded link suffers from increased ISI and hence exhibit an *ISI penalty*. The ability to correct errors due to FEC results in a *coding gain*, where the coding gain is the reduction in channel *SNR* to achieve the same *BER*. For back-plane links, we expect that the coding gain for specific types of codes will offset the ISI penalty with an acceptable latency, and thereby result in a reduced power link.

In order to capture the trade-offs inherent in FEC, consider the specific example of a (327, 265, 15) ($r = 0.81$) BCH code with $d_{min} = 15, t = 7$ employed over a link comprising of a 20" FR4 channel supporting NRZ shaped $10Gb/s$ data and a receiver with an LE and a DFE.

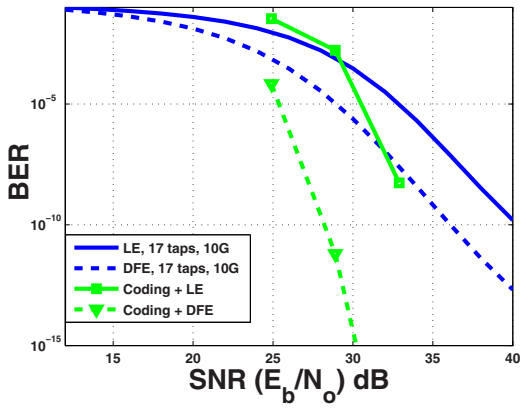

(a)

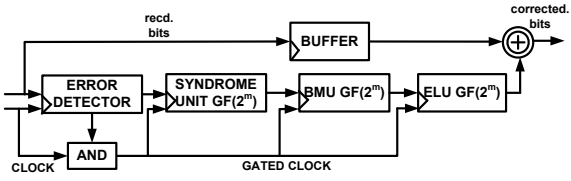

(b)

Fig. 14 BCH codes in back-plane links: (a) performance of a (327, 265, 15) BCH code over a 20" FR-4 channel at 10 Gb/s, and (b) a low-power BCH decoder architecture

Figure 14 shows the performance of a BCH code. We observe that the *BER* reduces by six orders-of-magnitude and ten orders-of-magnitude with an LE and a DFE, respectively. Furthermore, the crossover point between coded and uncoded link, due to the ISI penalty, occurs between $25dB$-to-$30dB$ and around $24dB$ for an LE and a DFE, respectively. In addition, we find that at a $BER = 10^{-10}$, FEC provides a coding gain of $7dB$-to-$8dB$ for both LE and DFE, and that the coding gain increases as the *BER* reduces. It is expected that at the desired $BER = 10^{-15}$, a coding gain of at least $10dB$ is feasible. The coding gain or *SNR* slack can be budgeted across the transmit driver, and the receiver circuitry in order to save power.

The use of FEC in multi-Gb/s links raises interesting implementation questions especially those related to power consumption. For example, as the pre-FEC error-rate is quite small, e.g., 10^{-4}, a power-efficient decoder can be designed by incorporating a low-complexity error-detector (see Fig. 14(b)), which is constantly operating but which is used to gate the subsequent processing in the decoder. In addition to FEC, the back-plane channel can benefit from other signal processing techniques such as the use of MLSE detector, as was demonstrated for an optical link. A channel shortening equalizer will need to precede the MLSE detector as the back-plane channel is known to have a large memory. The availability of high-speed ADCs opens up the application of a number of low-power techniques including algorithmic noise-tolerance (ANT) [36], where statistical signal processing techniques are employed to detect and correct errors in signal processing architectures.

5.2 Coherent Detection in Optical Links

In section 2.1, we described a basic model for both directly modulated and externally modulated optical sources for intensity modulations, such as NRZ or RZ signalling. We also described a more spectrally efficient method, based on a form of optical duobinary signalling, in which the phase of the carrier is additionally modulated by either 0 or π radians, depending on the state of the precoder. Attempts to achieve even greater spectral efficiency in optical transmitters invariably require innovations in both the optical devices involved as well as in the transmit and receive signal processing. Recently, a number of additional phase-shift keyed modulation formats have been demonstrated, including the differential QPSK format shown in Fig. 15 (a) [37]. The transmitter requires a digital precoder, and an optical encoder with dual Mach-Zehnder modulators. Receiver decoding is performed optically, using an optical delay-and-add structure to avoid the need for generating a coherent local oscillator.

However, there are benefits of using a coherent receiver when additional signal processing is used for EDC. While the received signal after direct detection is proportional to the received optical power, if a coherent receiver were used, then the received electrical signal would be proportional the electromagnetic field in the fiber. As a result, CD and PMD might appear as linear distortions, enabling linear equalization techniques to substantially outperform their counterparts when used

after direct (square-law) detection. Four photo detectors are presently required to recover the in-phase and quadrature (I and Q) components of the electromagnetic field from both polarizations of the received optical signal. If the length of fiber to be compensated were known in advance, techniques for coherent moduation combined with linear pre-equaliation [38] can be used. One advantage of electronic pre-equalization is that direct detection can still be used at the receiver, in exchange for substantial signal processing followed by digital to analog conversion at the modulator.

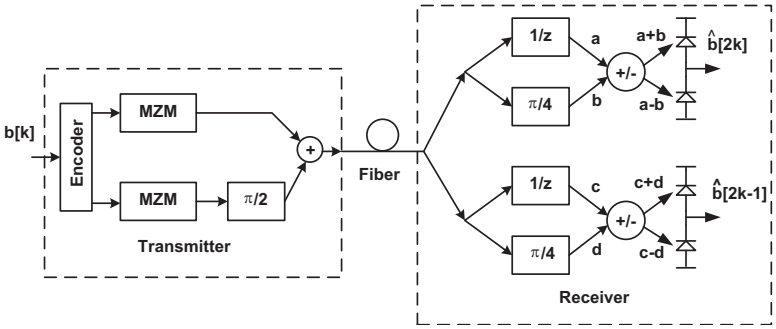

Fig. 15 Block diagram for a differential QPSK transmitter and receiver.

In addition, techniques that exploit *polarization diversity* [39] and *modal diversity* [40, 41] are being explored, whereby PMD and modal dispersion are viewed as opportunities for enhancing date rates, rather than impediments. This is analogous to what is being done in wireless systems today. Thus, as was the case for back-plane links, there exists a wide range of unexplored signal processing techniques in the context of optical links.

6 Concluding Remarks

Signal processing techniques have found their way into the design of high-speed links, such as back-plane and optical links, today. Thanks to the tremendous advances in semiconductor technology, equalization in back-plane links and EDC-based optical products are making their way into commercial systems. We hope that this chapter has provided an overview of the general concepts, and will enable further exploration into this fascinating and challenging topic.

7 Acknowledgements

The authors thank their colleagues at the University of Illinois at Urbana-Champaign, and Finisar Corporation, Inc., for their contributions.

References

1. L. D. Paulson, "The ins and outs of new local I/O trends," *IEEE Computer Magazine*, July 2003.
2. R. Narasimha and N. R. Shanbhag, "Forward error-correction for high-speed I/O," *Proceedings of the 42th Annual Asilomar Conference on Signals, Systems, and Computers*, October 2008.
3. N. Blitvic, M. Lee, and V. Stojanović, "Channel coding for high-speed links: A systematic look at code performance and system simulation," *IEEE Transactions on Advanced Packaging*, pp. 268–279, May 2009.
4. M.-J. Lee, W. J. Dally, and P. Chiang, "Low-power area-efficient high-speed i/o circuit techniques," *IEEE Journal of Solid-State Circuits*, pp. 1591–1599, November 2000.
5. J. E. Jaussi, G. Balamurugan, D. R. Johnson, B. Casper, A. Martin, J. Kennedy, N. Shanbhag, and R. Mooney, "8-Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation, and clock de-skew," *IEEE Journal of Solid-State Circuits*, pp. 80–88, January 2005.
6. A. Singer, N. R. Shanbhag, and H.-M. Bae, "Electronic dispersion compensation," *IEEE Signal Processing Magazine*, pp. 110–130, November 2008.
7. G. E. Keiser, *Optical Fiber Communications*. McGraw-Hill International Editions: Electrical Engineering Series, 2000.
8. G. P. Agrawal, *Fiber-Optic Communication Systems*. Wiley-Interscience, 2002.
9. J. Peters, A. Dori, and F. Kapron, "Bellcore's fiber measurement audit of existing cable plant for use with high bandwidth systems," *Proc. of the NFOEC'97*, September 1997.
10. M. Karlsson, "Probability density functions of the differential group delay in optical fiber communication systems," *IEEE Journal of Lightwave Technology*, vol. 19, pp. 324–332, 2001.
11. G. Papen and R. Blahut, *Lightwave Communication Systems*. 2005. draft.
12. H. M. Bae, J. Ashbrook, J. Park, N. Shanbhag, A. C. Singer, and S. Chopra, "An MLSE receiver for electronic-dispersion compensation of OC-192 links," *Journal of Solid State Circuits*, vol. 41, pp. 2541–2554, November 2006.
13. A. Farbert, S. Langenbach, N. Stojanovic, C. Dorschky, T. Kupfer, C. Schulien, J.-P. Elbers, H. Wernz, H. Griesser, and C. Glingener, "Performance of a 10.7 Gb/s receiver with digital equaliser using maximum likelihood sequence estimation," *ECOC'2004 Proceedings*, pp. PD–Th4.1.5, 2004.
14. M. Harwood, et al., "A 12.5 Gb/s SerDes in 65nm CMOS using a baud-rate ADC with digital receiver equalization and clock recovery," *IEEE International Solid-State Circuits Conference*, February 2007.
15. O. Agazzi, et al., "A 90nm CMOS DSP MLSD transceiver with integrated AFE for electronic dispersion compensation of multi-mode optical fibers at 10 Gb/s," *IEEE International Solid-State Circuits Conference*, 2008.
16. J. Cao, et al., "A 500mW digitally calibrated AFE in 65nm CMOS for 10Gb/s serial links over backplane and multimode fiber," *IEEE International Solid-State Circuits Conference*, February 2009.
17. V. Stojanović, A. Amirkhany, and M. A. Horowitz, "Optimal linear precoding with theoretical and practical data rates in high-speed serial link back-plane communications," *IEEE Int. Conf. Communications*, 2004.

18. G. Balamurugan and N. R. Shanbhag, "Modeling and mitigation of jitter in multi-gbps source-synchronous I/O links," *Proceedings of the 21st International Conference on Computer Design*, October 2003.
19. P. Humblet and M. Azizoglu, "On the bit error rate of lightwave systems with optical amplifiers," *Journal of Lightwave Technology*, vol. 9, pp. 1576–1582, November 1981.
20. J. Stonick, G.-Y. Wei, J. L. Sonntag, and D. K. Weinlader, "An adaptive PAM-4 5-Gb/s backplane transceiver in 0.25μm CMOS," *IEEE Journal of Solid-State Circuits*, pp. 436–443, March 2003.
21. K. Yamaguchi and et al., "12Gb/s duobinary signaling with x2 oversampled edge equalization," *IEEE International Solid-State Circuits Conference*, February 2005.
22. J. Lee, M.-S. Chen, and H.-D. Wang, "Design and comparison of three 20-Gb/s backplane transceivers for duobinary, PAM4, and NRZ data," *IEEE Journal of Solid-State Circuits*, pp. 2120–2133, September 2008.
23. T. Franck, P. B. Hansen, T. N. Nielsen, and L. Eskildsen, "Duobinary transmitter with low intersymbol interference," *IEEE Photonics Technology Letters*, pp. 597–599, April 1998.
24. S. Qureshi, "Adaptive equalization," *Proceedings of the IEEE*, vol. 73, pp. 1349 – 1387, September 1985.
25. M. BiChan and A. C. Carusone, "A 6.5 Gb/s backplane transmitter with 6-tap FIR equalizer and variable tap spacing," *IEEE Custom Integrated Circuits Conference*, 2008.
26. Q. Yu and A. Shanbhag, "Electronic data processing for error and dispersion compensation," *Journal of Lightwave Technology*, vol. 24, pp. 4514 – 4525, December 2006.
27. J. Winters and R. Gitlin, "Electrical signal processing techniques in long-haul fiber optic systems," *IEEE Transactions on Communications*, pp. 1439–1453, September 1990.
28. S. Benedetto, E. Bigliere, and V. Castellani, *Digital Transmission Theory*. Englewood Cliffs, NJ: Prentice Hall, 1987.
29. L.R. Bahl et al., "Optimal decoding of linear codes for minimizing symbol error rate," *IEEE Trans. on Information Theory*, vol. 20, pp. 284–287, March 1974.
30. G. Forney, "Maximum-likelihood sequence estimation of digital sequences in the presence of intersymbol interference," *IEEE Transactions on Communications*, vol. 18, May 1972.
31. W. Sauer-Greff, M. Lorang, H. Haunstein, and R. Urbansky, "Modified Volterra series and state model approach for nonlienar data channels," *Proc. IEEE Signal Processing '99*, pp. 19–23, 1999.
32. O. Agazzi, M. Hueda, H. Carrer, and D. Crivelli, "Maximum likelihood sequence estimation in dispersive optical channels," *Journal of Lightwave Technology*, vol. 23, pp. 749–763, February 2005.
33. P. J. Black and T. H. Meng, "A 140-Mb/s, 32-state, radix-4 Viterbi decoder," *IEEE Journal of Solid-State Circuits*, vol. 27, pp. 1877–1885, Dec. 1992.
34. R. Hegde, A. Singer, and J. Janovetz, "Method and apparatus for delayed recursion decoder," *US Patent*, no. 7206363. filed June, 2003, issued April, 2007.
35. K. K. Parhi, *VLSI Digital Signal Processing Systems: Design and Implementation*. Wiley, 1999.
36. R. Hegde and N. R. Shanbhag, "Soft digital signal processing," *IEEE Transactions on VLSI Systems*, pp. 813–823, December 2001.
37. R. Griffin, "Integrated DQPSK transmitters," *Optical Fiber Communication Conference, 2005. Technical Digest*, vol. OFC/NFOEC Volume 3, p. OWE3, March 2005.
38. M. E. Said, J. Stich, and M. Elmasry, "An electrically pre-equalized 10Gb/s duobinary transmission system," *Journal of Lighwave Technology*, no. 1, pp. 388–400, 2005.
39. C. Laperle, B. Villeneuve, Z. Zhang, D. McGhan, H. Sun, and M. OŠSullivan, "Wavelength division multiplexing (WDM) and polarization mode dispersion (PMD) performance of a coherent 40Gbit/s dual-polarization quadrature phase shift keying (DP-QPSK) transceiver," *Post Deadline Papers OFC/NFOEC 2007*, no. PDP16, 2007.
40. H. Stuart, "Dispersive multiplexing in multimode optical fiber," *Science*, vol. 289, pp. 281–283, 2000.
41. A. Tarighat, R. Hsu, A. Shah, A. Sayed, and B. Jalali, "Fundamentals and challenges of optical Multi-Input Multiple-Output multimode fiber links," *IEEE Communications Magazine*, pp. 1–8, May 2007.