Chapter 2

Background

This chapter describes the basic principles of high-speed electrical and optical link design. It begins with an overview of the electrical circuits required to achieve high-speed communication over band-limited electrical channels, and then discusses optical channel advantages and optical source and detector properties. It ends with a brief review of the electrical circuit techniques commonly applied to interface with these optical devices.

2.1 High-Speed Electrical Links

High-speed point-to-point electrical links are commonly used in short distance chip-to-chip communication applications such as internet routers [6,7], multi-processor systems [14], and processor-memory interfaces [15-17]. In order to achieve high data rates, these systems employ specialized I/O circuitry that performs incident wave signaling over carefully designed controlled-impedance channels. As will be described later in this section, the electrical channel's frequency-dependent loss and impedance discontinuities become major limiters in data rate scaling. While traditionally simple binary non-return-to-zero (NRZ) pulse-amplitude-modulation (PAM-2) techniques have been used [18],

today's multi-Gb/s links require link designers to implement channel equalization [6-8] and consider more advanced modulation schemes [19,20].

This section begins by describing the three major link circuit components, the transmitter, receiver, and timing system. Next, it discusses the electrical channel properties that impact the transmitted signal. The section concludes by providing an overview of common equalization schemes and advanced modulation techniques that designers implement in order to extend data rates over the band-limited electrical channels.

2.1.1 Electrical Link Circuits

Figure 2.1 shows the major components of a typical high-speed electrical link system. Due to the limited number of high-speed I/O pins in chip packages and printed circuit board (PCB) wiring constraints, a high-bandwidth transmitter serializes parallel input data for transmission. Differential low-swing signaling is commonly used for common-mode noise rejection [21]. At the receiver, the incoming signal is sampled, regenerated to CMOS values, and deserialized. High-frequency clocks synchronize the data transfer and are generated by a frequency synthesis phase-locked loop (PLL) at the transmitter and recovered from the incoming data stream by a clock-and-data recovery (CDR) unit at the receiver.
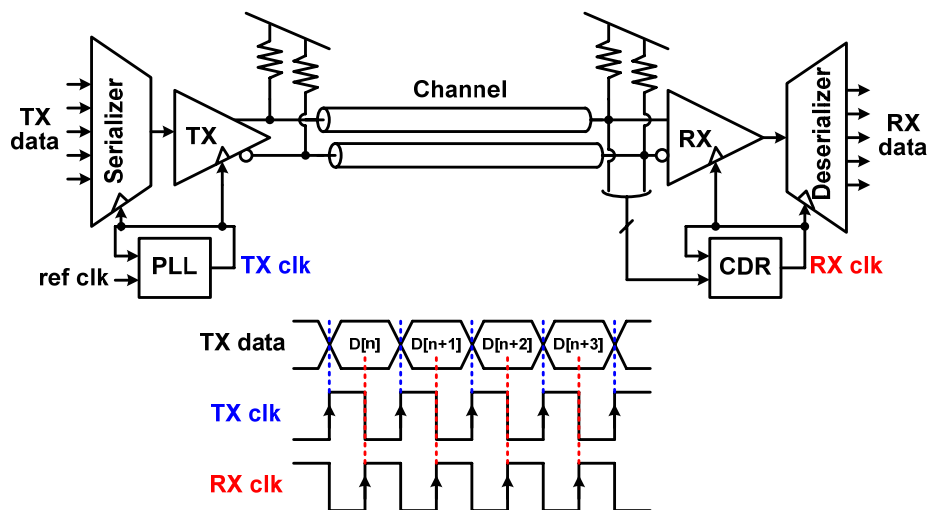

Figure 2.1: High-speed electrical link system

Transmitter

The transmitter must generate an accurate voltage swing on the channel while also maintaining proper output impedance in order to attenuate any channel-induced reflections. Either current or voltage-mode drivers, shown in Figure 2.2, are suitable output stages. Current-mode drivers typically steer current close to 20mA between the differential channel lines in order to launch a bipolar voltage swing on the order of ±500mV. Driver output impedance is maintained through termination which is in parallel with the high-impedance current switch. While current-mode drivers are most commonly implemented [22], the power associated with the required output voltage for proper transistor output impedance and the "wasted" current in the parallel termination led designers to consider voltage-mode drivers. These drivers use a regulated output stage to supply a fixed output swing on the channel through a series termination which is feedback controlled [23]. While the feedback impedance control is not as simple as parallel termination, the voltage-mode drivers have the potential to supply an equal receiver voltage swing at a quarter [24] of the common 20mA cost of current-mode drivers.
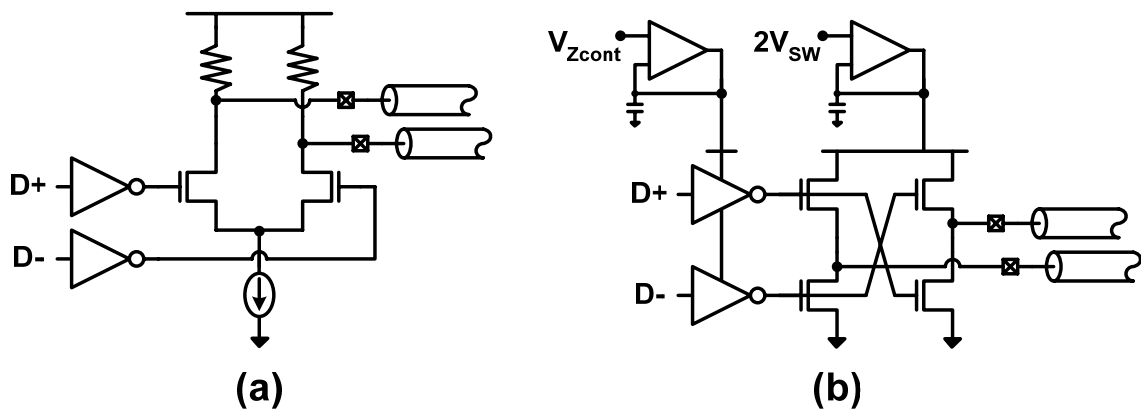

Figure 2.2: Transmitter output stages: (a) current-mode driver, (b) voltage-mode driver

Receiver

Figure 2.3 shows a high-speed receiver which compares the incoming data to a threshold and amplifies the signal to a CMOS value. This highlights a major advantage of binary differential signaling, where this threshold is inherent, whereas single-ended signaling requires careful threshold generation to account for variations in signal amplitude, loss, and noise [25]. The bulk of the signal amplification is often performed with a positive

feedback latch [26,27]. These latches are more power-efficient versus cascaded linear amplification stages since they don't dissipate DC current. While regenerative latches are the most power-efficient input amplifiers, link designers have used a small number of linear pre-amplification stages to implement equalization filters that offset channel loss faced by high data rate signals [15,28].
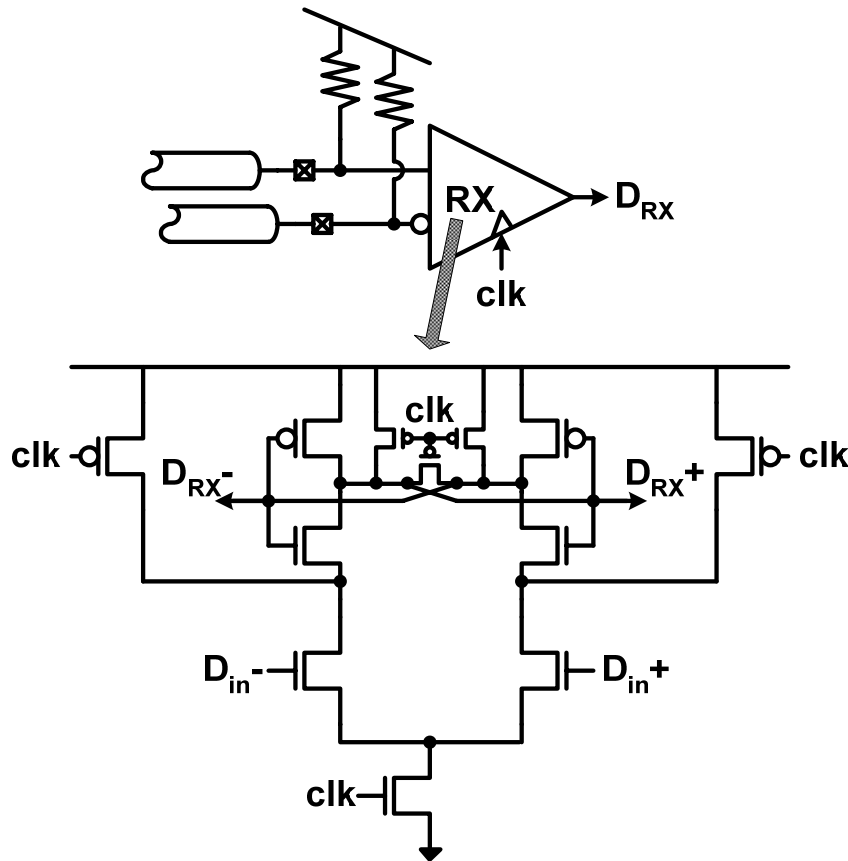

Figure 2.3: Receiver input stage with regenerative latch [26]

One issue with these latches is that they require time to reset or "pre-charge", and thus to achieve high data rates, often multiple latches are placed in parallel at the input and activated with multiple clock phases spaced a bit period apart in a time-division-demultiplexing manner [18,29], shown in Figure 2.4. This technique is also applicable at the transmitter, where the maximum serialized data rate is set by the clocks switching the multiplexer. The use of multiple clock phases offset in time by a bit period can overcome the intrinsic gate-speed which limits the maximum clock rate that can be efficiently distributed to 6-8 FO4 delays [30].
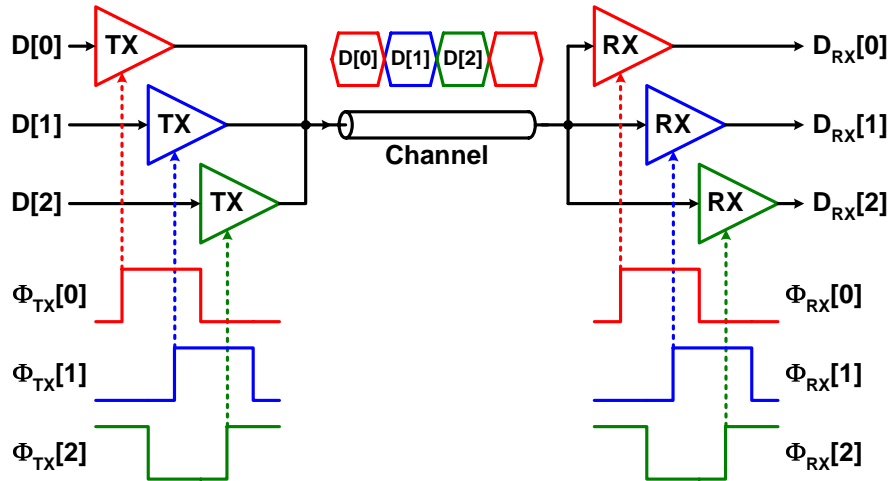
Figure 2.4: Time-division multiplexing link

Timing Circuits

High-precision low-noise clocks are necessary at both the transmitter and receiver in order to ensure sufficient timing margins at high data rates. Figure 2.5 show a PLL, which is often used at the transmitter for clock synthesis in order to serialize reduced-rate parallel input data and also potentially at the receiver for clock recovery. The PLL is a negative feedback loop which works to lock the phase of the feedback clock to an input reference clock. A phase-frequency detector produces an error signal which is proportional to the phase difference between the feedback and reference clocks. This phase error is then filtered to provide a control signal to a voltage-controlled oscillator (VCO) which generates the output clock. The PLL performs frequency synthesis by placing a clock divider in the feedback path, which forces the loop to lock with the output clock frequency equal to the input reference frequency times the loop division factor.
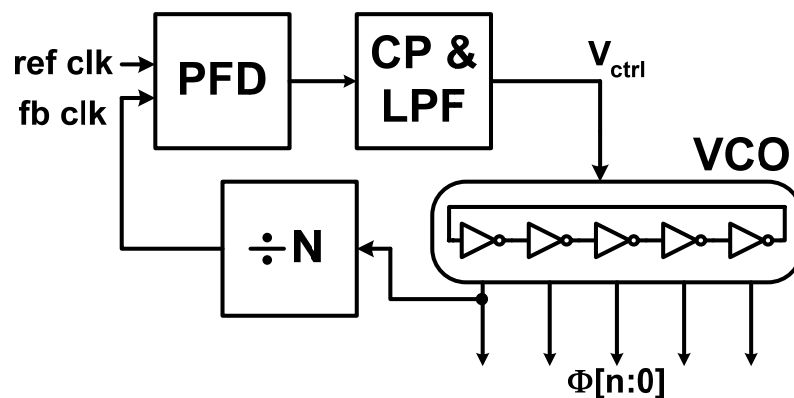
Figure 2.5: PLL frequency synthesizer

It is important that the PLL produce clocks with low timing noise, quantified in the timing domain as jitter and in the frequency domain as phase noise. Considering this, the most critical PLL component is the VCO, as its phase noise performance can dominate at the output clock and have a large influence on the overall loop design. LC oscillators typically have the best phase noise performance, but their area is large and tuning range is limited [31]. While ring oscillators display inferior phase noise characteristics, they offer advantages in reduced area, wide frequency range, and ability to easily generate multiple phase clocks for time-division multiplexing applications [18,29].

Also important is the PLL's ability to maintain proper operation over process variances, operating voltage, temperature, and frequency range. To address this, self-biasing techniques were developed by Maneatis [32] and expanded in [33,34] that set constant loop stability and noise filtering parameters over these variances in operating conditions.

At the receiver, clock recovery is required in order to position the data sampling clocks with maximum timing margin and also filter the incoming signal jitter. It is possible to modify a PLL to perform clock recovery with changes in the phase detection circuitry, as shown in Figure 2.6. Here the phase detector samples the incoming data stream to extract both data and phase information. As shown in Figure 2.7, the phase detector can either be linear [35], which provides both sign and magnitude information of the phase error, or binary [36], which provides only phase error sign information. While CDR systems with linear phase detectors are easier to analyze, generally they are harder to implement at high data rates due to the difficulty of generating narrow error pulse widths, resulting in effective dead-zones in the phase detector [37]. Binary, or "bangbang", phase detectors minimize this problem by providing equal delay for both data and phase information and only resolving the sign of the phase error [38]. In order to properly filter the input data jitter to prevent transfer onto the receiver clocks, the CDR bandwidth must be set sufficiently low, such that it has a hard time reducing the intrinsic phase noise of a ring VCO. Thus, while a PLL-based CDR is an efficient solution, generally one cannot optimally filter both VCO phase noise and input data jitter. This

motivates the use of dual-loop clock recovery [25], which provides two degrees of freedom to filter the two dominant clock noise sources.
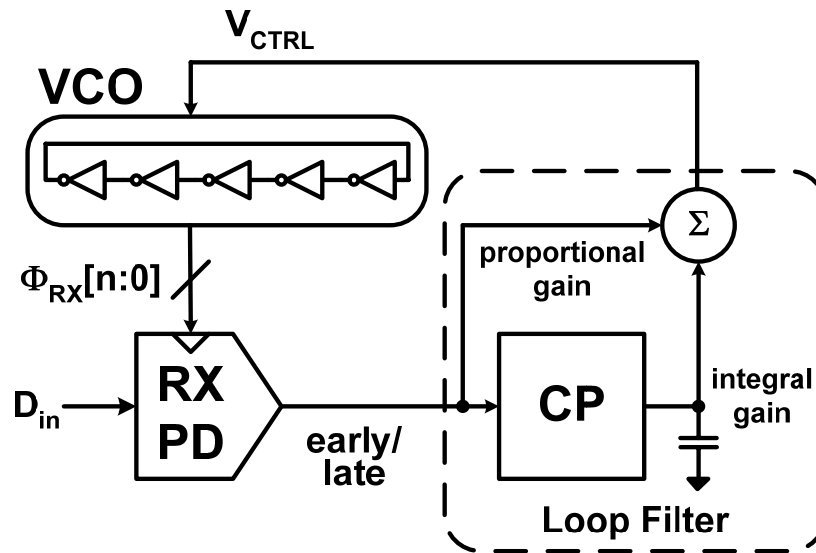

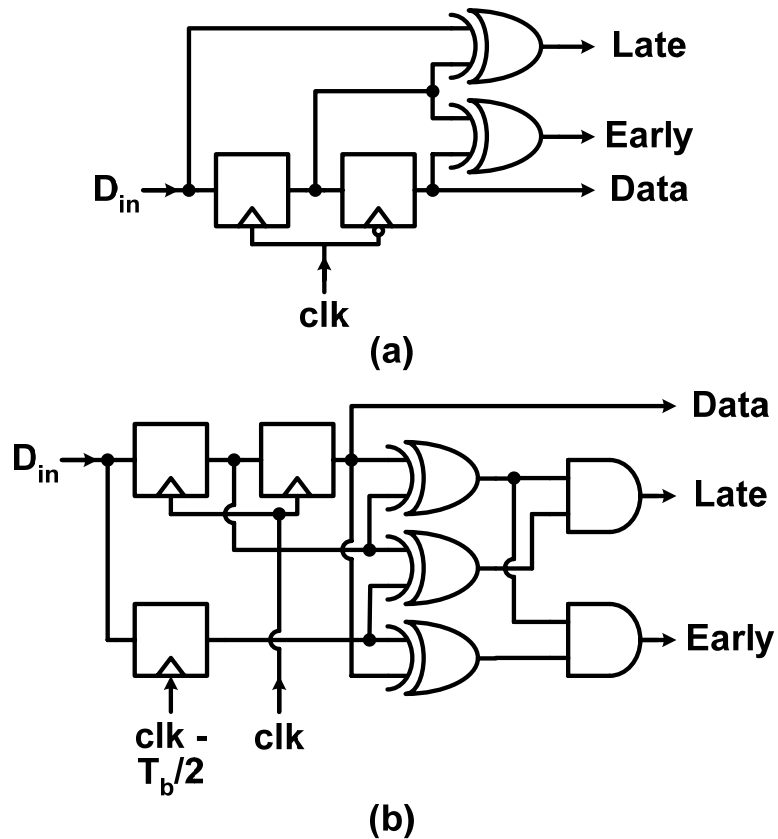Figure 2.6: PLL-based CDR system

Figure 2.7: CDR phase detectors: (a) linear [35], (b) binary [36]

While proper design of these high-speed I/O components requires considerable attention, CMOS scaling allows the basic circuit blocks to achieve data rates that exceed 10Gb/s [15,16]. However, as data rates scale into the low Gb/s, the frequency dependent loss of the chip-to-chip electrical wires disperses the transmitted signal to the extent that it is undetectable at the receiver without proper signal processing or channel equalization techniques. Thus, in order to design systems that achieve increased data rates, link designers must comprehend the high-frequency characteristics of the electrical channel, which are outlined next.

2.1.2 Electrical Channels

Electrical inter-chip communication bandwidth is predominantly limited by high-frequency loss of electrical traces, reflections caused from impedance discontinuities, and adjacent signal crosstalk, as shown in Figure 2.8. The relative magnitudes of these channel characteristics depend on the length and quality of the electrical channel which is a function of the application. Common applications range from processor-to-memory interconnection, which typically have short (<10") top-level microstrip traces with relatively uniform loss slopes [15], to server/router and multi-processor systems, which employ either long (~30") multi-layer backplanes [6] or (~10m) cables [39] which can both possess large impedance discontinuities and loss.

Dispersion

PCB traces suffer from high-frequency attenuation caused by wire skin effect and dielectric loss. As a signal propagates down a transmission line, the normalized amplitude at a distance x is equal to

$$\frac{V(x)}{V(0)} = e^{-(\alpha_R + \alpha_D)x} ,$$

(2.1)

where α_R and α_D represent resistive and dielectric loss factors [21]. The skin effect, which describes the process of high-frequency signal current crowding near the conductor surface, impacts the resistive loss term as frequency increases. This results in a resistive loss term which is proportional to the square-root of frequency
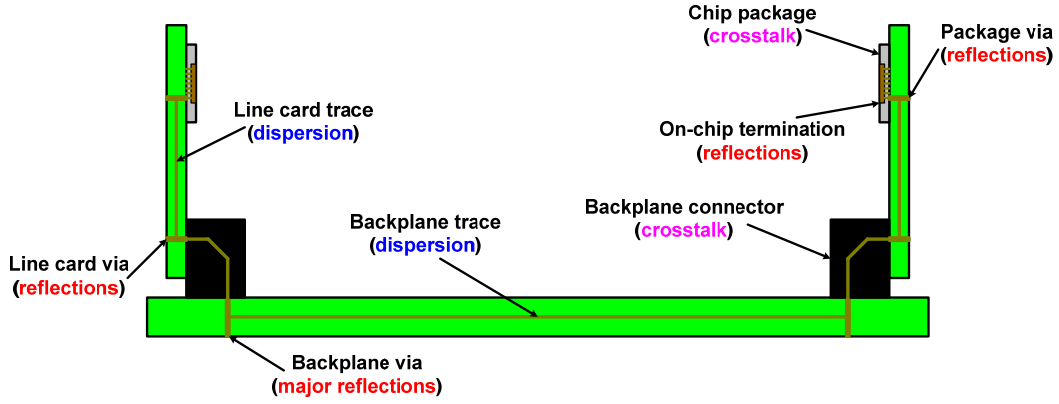
Figure 2.8: Backplane system cross-section

$$\alpha_R = \frac{R_{AC}}{2Z_0} = \frac{2.61 \times 10^{-7} \sqrt{\rho_r}}{\pi D 2 Z_0} \sqrt{f} \;,$$ (2.2)

where D is the trace's diameter (in), ρ_r is the relative resistivity compared to copper, and Z_0 is the trace's characteristic impedance [40]. Dielectric loss describes the process where energy is absorbed from the signal trace and transferred into heat due to the rotation of the board's dielectric atoms in an alternating electric field [40]. This results in the dielectric loss term increasing proportional to the signal frequency

$$\alpha_D = \frac{\pi \sqrt{\varepsilon_r} \tan \delta_D}{c} f \;,$$ (2.3)

where ε_r is the relative permittivity, c is the speed of light, and $\tan\delta_D$ is the board material's loss tangent [21].

Figure 2.9 shows how these frequency dependent loss terms result in low-pass channels where the attenuation increases with distance [41]. The high-frequency content of a pulses sent across such channel is filtered, resulting in an attenuated received pulse whose energy has been spread or dispersed over several bit periods, as shown in Figure 2.10(a). When transmitting data across the channel, energy from individual bits will now interfere with adjacent bits and make them more difficult to detect. This intersymbol interference (ISI) increases with channel loss and can completely close the received data eye diagram, as shown in Figure 2.10(b).
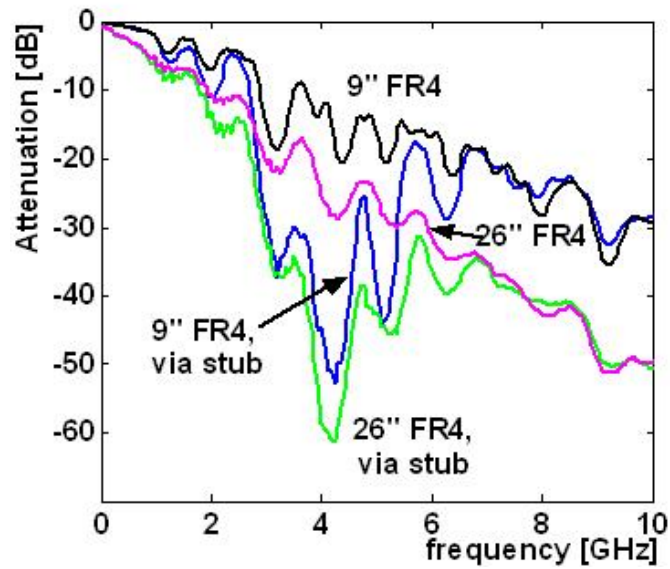
Figure 2.9: Frequency response of several backplane channels [41]

(a)   (b)
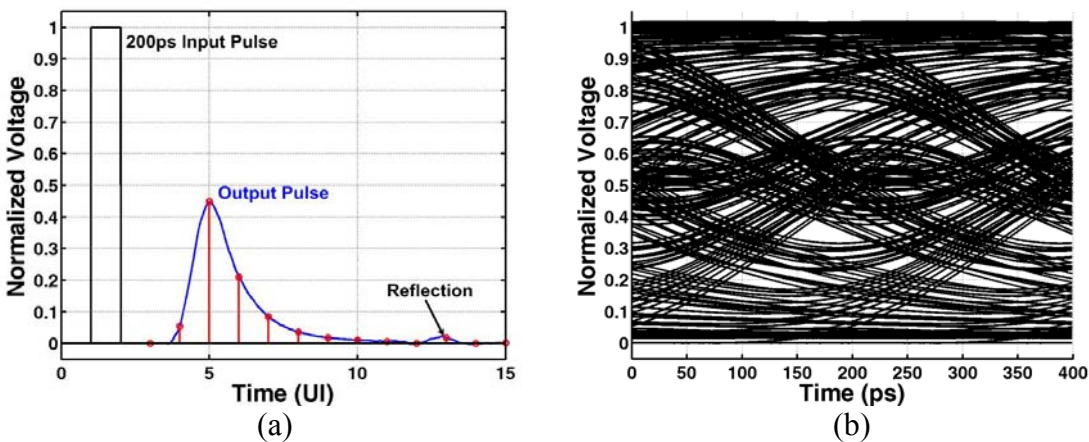
Figure 2.10: Backplane channel performance at 5Gb/s: (a) pulse response, (b) eye diagram

Reflections

Signal interference also results from reflections caused by impedance discontinuities. If a signal propagating across a transmission line experiences a change in impedance Z_r relative to the line's characteristic impedance Z_0, a percentage of that signal equal to

$$\frac{V_r}{V_i} = \frac{Z_r - Z_0}{Z_r + Z_0} \tag{2.4}$$

will reflect back to the transmitter. This results in an attenuated or, in the case of multiple reflections, a time delayed version of the signal arriving at the receiver. The most common sources of impedance discontinuities are from on-chip termination mismatches and via stubs that result with signaling over multiple PCB layers. Figure 2.9 shows that the capacitive discontinuity formed by the thick backplane via stubs can cause severe nulls in the channel frequency response.

Crosstalk

Another form of interference comes from crosstalk, which occurs due to both capacitive and inductive coupling between neighboring signal lines. As a signal propagates across the channel, it experiences the most crosstalk in the backplane connectors and chip packages where the signal spacing is smallest compared to the distance to a shield. Crosstalk is classified as near-end (NEXT), where energy from an aggressor (transmitter) couples and is reflected back to the victim (receiver) on the same chip, and far-end (FEXT), where the aggressor energy couples and propagates along the channel to a victim on another chip. NEXT is commonly the most detrimental crosstalk, as energy from a strong transmitter ($\sim 1V_{pp}$) can couple onto a received signal at the same chip which has been attenuated ($\sim 20mV_{pp}$) from propagating on the lossy channel. Crosstalk is potentially a major limiter to high-speed electrical link scaling, as in common backplane channels the crosstalk energy can actually exceed the through channel signal energy at frequencies near 4GHz [6].

2.1.3 Channel Equalization and Advanced Modulation Techniques

The previous subsection discussed interference mechanisms that can severely limit the rate at which data is transmitted across electrical channels. As shown in Figure 2.9(b), frequency dependent channel loss can reach magnitudes sufficient to make simple NRZ binary signaling undetectable. Thus, in order to continue scaling electrical link data rates, designers have implemented systems which compensate for frequency dependent loss or equalize the channel response. This subsection discusses how the equalization circuitry is often implemented in high-speed links, and other approaches for dealing with these issues.

Equalization Systems

In order to extend a given channel's maximum data rate, many communication systems use equalization techniques to cancel intersymbol interference caused by channel distortion. Equalizers are implemented either as linear filters (both discrete and continuous-time) that attempt to flatten the channel frequency response, or as non-linear filters that directly cancel ISI based on the received data sequence. Depending on system data rate requirements relative to channel bandwidth and the severity of potential noise sources, different combinations of transmit and/or receive equalization are employed.

Transmit equalization, implemented with an FIR filter, is the most common technique used in high-speed links [42]. This TX "pre-emphasis" (or more accurately "de-emphasis") filter, shown in Figure 2.11, attempts to invert the channel distortion that a data bit experiences by pre-distorting or shaping the pulse over several bit times. While this filtering could also be implemented at the receiver, the main advantage of implementing the equalization at the transmitter is that it is generally easier to build high-speed digital-to-analog converters (DACs) versus receive-side analog-to-digital converters. However, because the transmitter is limited in the amount of peak-power that it can send across the channel due to driver voltage headroom constraints, the net result is that the low-frequency signal content has been attenuated down to the high-frequency level, as shown in Figure 2.11.
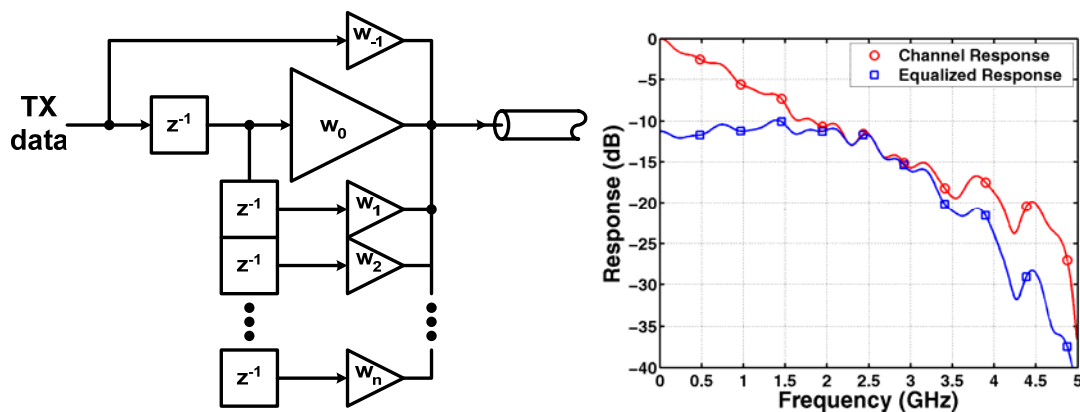

Figure 2.11: TX equalization with an FIR filter

Figure 2.12 shows a block diagram of receiver-side FIR equalization. A common problem faced by linear receive side equalization is that high-frequency noise content and

crosstalk are amplified along with the incoming signal. Also challenging is the implementation of the analog delay elements, which are often implemented through time-interleaved sample-and-hold stages [43] or through pure analog delay stages with large area passives [44,45]. Nonetheless, one of the major advantage of receive side equalization is that the filter tap coefficients can be adaptively tuned to the specific channel [43], which is not possible with transmit-side equalization unless a "back-channel" is implemented [46].
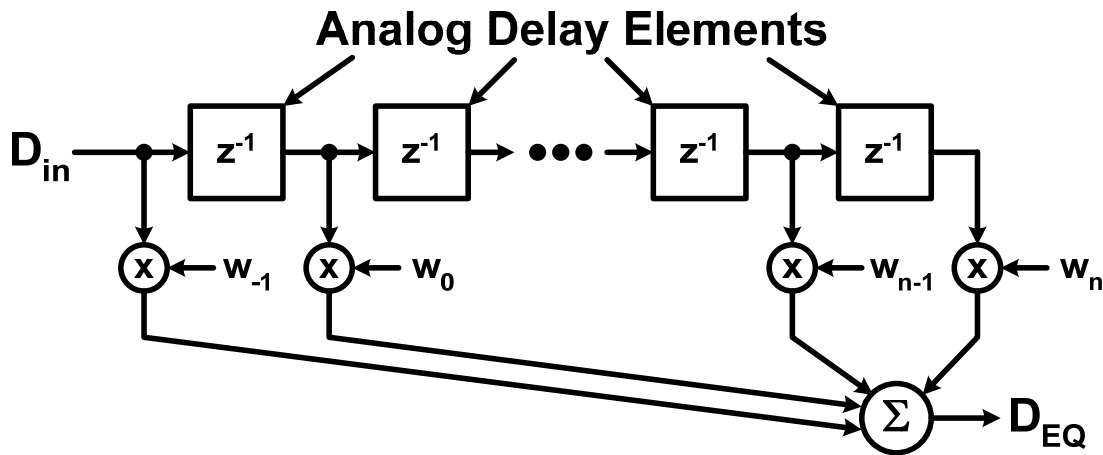
Figure 2.12: RX equalization with an FIR filter

Linear receiver equalization can also be implemented with a continuous-time amplifier, as shown in Figure 2.13. Here, programmable RC-degeneration in the differential amplifier creates a high-pass filter transfer function which compensates the low-pass channel. While this implementation is a simple and low-area solution, one issue is that the amplifier has to supply gain at frequencies close to the full signal data rate. This gain-bandwidth requirement potentially limits the maximum data rate, particularly in time-division demultiplexing receivers.

Figure 2.13: Continuous-time equalizing amplifier

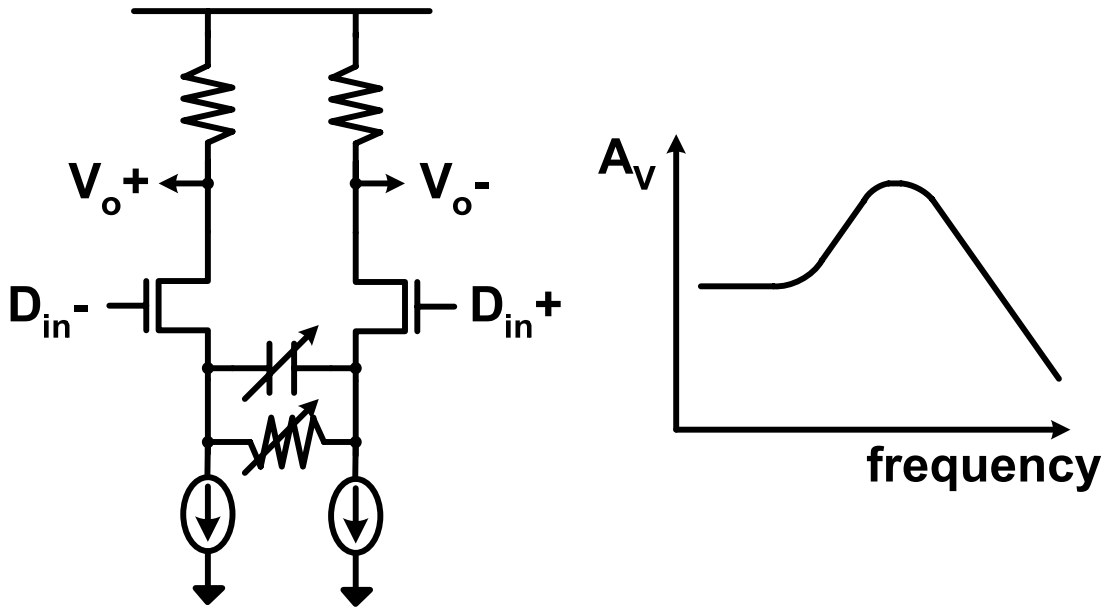The final equalization topology commonly implemented in high-speed links is a receiver-side decision feedback equalizer (DFE). A DFE, shown in Figure 2.14, attempts to directly subtract ISI from the incoming signal by feeding back the resolved data to control the polarity of the equalization taps. Unlike linear receive equalization, a DFE doesn't directly amplify the input signal noise or cross-talk since it uses the quantized input values. However, there is the potential for error propagation in a DFE if the noise is large enough for a quantized output to be wrong. Also, due to the feedback equalization structure, the DFE cannot cancel pre-cursor ISI. The major challenge in DFE implementation is closing timing on the first tap feedback since this must be done in one bit period or unit interval (UI). Direct feedback implementations [6] require this critical timing path to be highly optimized. While a loop-unrolling architecture eliminates the need for first tap feedback [47], if a multiple tap implementation is required the critical path simply shifts to the second tap which has a timing constraint also near 1UI [8].
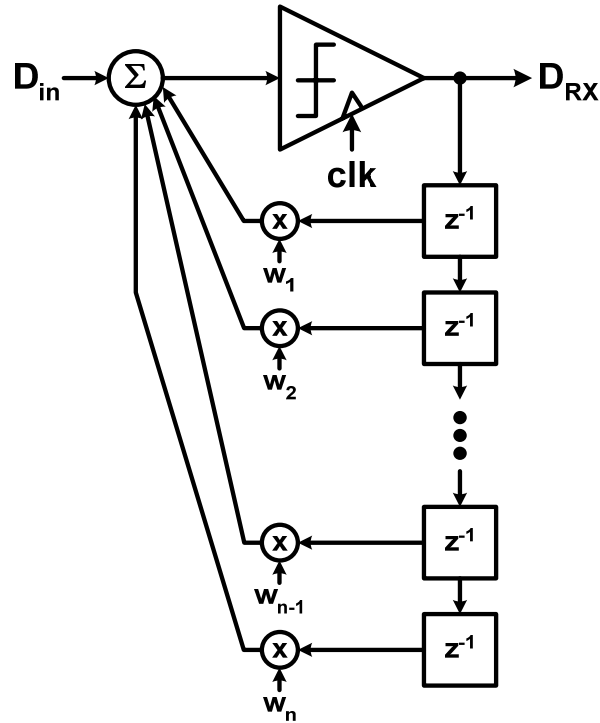
Figure 2.14: RX equalization with a DFE

Advanced Modulation Techniques

Modulation techniques which provide spectral efficiencies higher than simple binary signaling have also been implemented by link designers in order to increase data rates over band-limited channels. Multi-level PAM, most commonly PAM-4, is a popular modulation scheme which has been implemented both in academia [48] and industry [49,50]. Shown in Figure 2.15, PAM-4 modulation consists of two bits per symbol, which allows transmission of an equivalent amount of data in half the channel bandwidth. However, due to the transmitter's peak-power limit, the voltage margin between symbols is 3x (9.5dB) lower with PAM-4 versus simple binary PAM-2 signaling. Thus, a general rule of thumb exists that if the channel loss at the PAM-2 Nyquist frequency is greater than 10dB relative to the previous octave, then PAM-4 can potentially offer a higher signal-to-noise ratio (SNR) at the receiver. However, this rule can be somewhat optimistic due to the differing ISI and jitter distribution present with PAM-4 signaling [51]. Also, PAM-2 signaling with a non-linear DFE at the receiver further bridges the performance gap due to the DFE's ability to cancel the dominant first post-cursor ISI without the inherent signal attenuation associated with transmitter equalization [7].
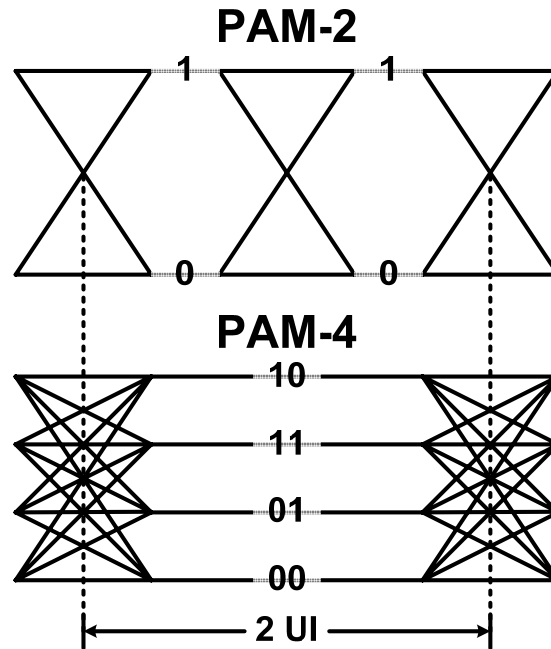
PAM-2

Figure 2.15: Pulse amplitude modulation – simple binary PAM-2 (1bit/symbol) and PAM-4 (2bits/symbol)

Another more radical modulation format under consideration by link researchers is the use of multi-tone signaling. While this type of signaling is commonly used in systems such as DSL modems [52], it is relatively new for high-speed inter-chip communication applications. In contrast with conventional baseband signaling, multi-tone signaling breaks the channel bandwidth into multiple frequency bands over which data is transmitted. This technique has the potential to greatly reduce equalization complexity relative to baseband signaling due to the reduction in per-band loss and the ability to selectively avoid severe channel nulls. Typically, in systems such as modems where the data rate is significantly lower than the on-chip processing frequencies, the required frequency conversion in done in the digital domain and requires DAC transmit and ADC receive front-ends [53,54]. While it is possible to implement high-speed transmit DACs [55], the excessive digital processing and ADC speed and precision required for multi-Gb/s channel bands results in prohibitive receiver power and complexity. Thus, for power-efficient multi-tone receivers, researchers have proposed using analog mixing techniques combined with integration filters and multiple-input-multiple-output (MIMO) DFEs to cancel out band-to-band interference [20].

Serious challenges exist in achieving increased inter-chip communication bandwidth over electrical channels while still satisfying I/O power and density constraints. As discussed, current equalization and advanced modulation techniques allow data rates near 10Gb/s over severely band-limited channels. However, this additional circuitry comes with a power and complexity cost, with typical commercial high-speed serial I/O links consuming close to 20mW/Gb/s [56,39] and research-grade links consuming near 10mW/Gb/s [15,17]. The demand for higher data rates will only result in increased equalization requirements and further degrade link energy efficiencies. While there has been recent work on reducing link power [23,28,57], these implementations have focused on moderate data rates over relatively tame channels. This approach will require extremely dense I/O architectures over optimized electrical channels that will ultimately be limited by the chip bump/pad pitch and crosstalk constraints. These issues motive investigation into the use of optical links for chip-to-chip applications, discussed in the next section.

2.2 High-Speed Optical Links

The primary motivation for an I/O architecture modification as radical as optical signaling is the magnitude of potential bandwidth offered with an optical channel. Conventional optical data transmission is analogous to wireless AM radio, where data is transmitted by modulating the optical intensity or amplitude of the high-frequency optical carrier signal. In order to achieve high fidelity over the most common optical channel – the glass fiber, high-speed optical communication systems typically use infrared light from source lasers with wavelengths ranging from 850-1550nm, or equivalently frequencies ranging from 200-350THz. Thus, the potential data bandwidth is quite large since this high optical carrier frequency exceeds current data rates by over three orders of magnitude. Moreover, because the loss of typical optical channels at short distances varies only fractions of dBs over wide wavelength ranges (tens of nanometers) [58], there is the potential for data transmission of several Tb/s without the requirement of channel equalization. This simplifies design of optical links in a manner similar to non-channel limited electrical links. However, optical links do require additional circuits that

Bibliography

[1] B. Landman and R. L. Russo, "On a Pin vs. Block Relationship for Partitioning of Logic Graphs," *IEEE Transactions on Computers*, vol. C-20, no. 12, Dec. 1971, pp. 1469-1479.

[2] N. McKeown, Router Capacity: Raw Data (http://www.stanford.edu/~nickm/RouterCapacity.xls)

[3] G. E. Moore, "No Exponential is Forever: But "Forever" Can Be Delayed!," *IEEE International Solid-State Circuits Conference*, Feb. 2003.

[4] S. L. Graham, M. Snir, and C. A. Patterson, *Getting Up to Speed: The Future of Supercomputing*, The National Academies Press, 2005.

[5] Semiconductor Industry Association (SIA), International Technology Roadmap for Semiconductors 2006 Update, 2006.

[6] R. Payne *et al.*, "A 6.25-Gb/s Binary Transceiver in 0.13-μm CMOS for Serial Data Transmission Across High Loss Legacy Backplane Channels," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 12, Dec. 2005, pp. 2646-2657.

[7] J. F. Bulzacchelli *et al.*, "A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 12, Dec. 2006, pp. 2885-2900.

[8] B. S. Leibowitz *et al.*, "A 7.5Gb/s 10-Tap DFE Receiver with First Tap Partial Response, Spectrally Gated Adaptation, and 2nd-Order Data-Filtered CDR," *IEEE International Solid-State Circuits Conference*, Feb. 2007.

[9] G. A. Keeler *et al.*, "The Benefits of Ultrashort Optical Pulses in Optically Interconnected Systems," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 9, no. 2, Mar.-Apr. 2003, pp. 477-485.

[10] J. J. Liu *et al.*, "Multichannel Ultrathin Silicon-on-Sapphire Optical Interconnects," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 9, no. 2, Mar.-Apr. 2003, pp. 380-386.

[11] D. V. Plant *et al.*, "256-Channel Bidirectional Optical Interconnect Using VCSELs and Photodiodes on CMOS," *Journal of Lightwave Technology*, vol. 19, no. 8, Aug. 2001, pp. 1093-1103.

[12] D. Agarwal and D. A. B. Miller, "Latency in Short Pulse based Optical Interconnects," *IEEE Lasers and Electro-Optics Society Annual Meeting*, Nov. 2001.

[13] A. Emami-Neyestanal *et al.*, "A 1.6Gb/s, 3mW CMOS Receiver for Optical Communication," *IEEE Symposium on VLSI Circuits*, June 2002.

[14] R. Nair *et al.*, "A 28.5GB/s CMOS Non-Blocking Router for Terabits/s Connectivity between Multiple Processors and Peripheral I/O Nodes," *IEEE International Solid-State Circuits Conference*, Feb. 2001.

[15] B. Casper *et al.*, "A 20Gb/s Forwarded Clock Transceiver in 90nm CMOS," *IEEE International Solid-State Circuits Conference*, Feb. 2006.

[16] J. Jaussi *et al.*, "A 20Gb/s Embedded Clock Transceiver in 90nm CMOS," *IEEE International Solid-State Circuits Conference*, Feb. 2006.

[17] E. Prete, D. Scheideler, and, A. Sanders, "A 100mW 9.6Gb/s Transceiver in 90nm CMOS for Next-Generation Memory Interfaces," *IEEE International Solid-State Circuits Conference*, Feb. 2006.

[18] C.-K. Yang and M. Horowitz, "A 0.8-μm CMOS 2.5Gb/s Oversampling Receiver and Transmitter for Serial Links," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 12, Dec. 1996, pp. 2015-2023.

[19] J. L. Zerbe *et al.*, "Equalization and Clock Recovery for a 2.5-10-Gb/s 2-PAM/4-PAM Backplane Transceiver Cell," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 12, Dec. 2003, pp. 2121-2130.

[20] A. Amirkhany *et al.*, "Analog Multi-Tone Signaling for High-Speed Backplane Electrical Links," *IEEE Global Telecommunications Conference*, Nov. 2006.

[21] W. Dally and J. Poulton, *Digital Systems Engineering*, Cambridge University Press, 1998.

[22] K. Lee *et al.*, "A CMOS serial link for fully duplexed data communication," *IEEE Journal of Solid-State Circuits*, vol. 30, no. 4, Apr. 1995, pp. 353-364.

[23] K.-L. J. Wong *et al.*, "A 27-mW 3.6-Gb/s I/O Transceiver," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 4, Apr. 2004, pp. 602-612.

[24] C. Menolfi *et al.*, "A 16Gb/s Source-Series Terminated Transmitter in 65nm CMOS SOI," *IEEE International Solid-State Circuits Conference*, Feb. 2007.

[25] S. Sidiropoulos and M. Horowitz, "A Semidigital Dual Delay-Locked Loop," *IEEE Journal of Solid-State Circuits*, vol. 32, no. 11, Nov. 1997, pp. 1683-1692.

[26] J. Montanaro *et al.*, "A 160MHz, 32b, 0.5W CMOS RISC Microprocessor," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 11, Nov. 1996, pp. 1703-1714.

[27] A. Yukawa *et al.*, "A CMOS 8-bit high speed A/D converter IC," *IEEE European Solid-State Circuits Conference*, Sep. 1988.

[28] R. Palmer *et al.*, "A 14mW 6.25Gb/s Transceiver in 90nm CMOS for Serial Chip-to-Chip Communications," *IEEE International Solid-State Circuits Conference*, Feb. 2007.

[29] J. Kim and M. Horowitz, "Adaptive-Supply Serial Links with Sub-1V Operation and Per-Pin Clock Recovery," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 11, Nov. 2002, pp. 1403-1413.

[30] M. Horowitz, C.-K. Yang, and S. Sidiropoulos, "High-Speed Electrical Signaling: Overview and Limitations," *IEEE Micro*, vol. 18, no. 1, Jan.-Feb. 1998, pp. 12-24.

[31] A. Hajimiri and T. H. Lee, "Design Issues in CMOS Differential LC Oscillators," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 5, May 1999, pp. 717-724.

[32] J. G. Maneatis, "Low-jitter Process-Independent DLL and PLL based in Self-Biased Techniques," *IEEE Journal of Solid-State Circuits*, vol. 31, no. 11, Nov. 1996, pp. 1723-1732.

[33] S. Sidiropoulos *et al.*, "Adaptive Bandwidth DLLs and PLLs using Regulated Supply CMOS Buffers," *IEEE Symposium on VLSI Circuits*, June 2000.

[34] J. G. Maneatis *et al.*, "Self-Biased High-Bandwidth Low-Jitter 1-to-4096 Multiplier Clock Generator PLL," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 11, Nov. 2003, pp. 1795-1803.

[35] C. R. Hogge, Jr., "A Self Correcting Clock Recovery Circuit," *Journal of Lightwave Technology*, vol. 3, no. 6, Dec. 1985, pp. 1312-1314.

[36] J. D. H. Alexander, "Clock Recovery From Random Binary Signals," *Electronics Letters*, vol. 11, no. 22, Oct. 1975, pp. 541-542.

[37] Y. M. Greshichchev and P. Schvan, "SiGe Clock and Data Recovery IC with Linear-Type PLL for 10-Gb/s SONET Application," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 9, Sep. 2000, pp. 1353-1359.

[38] Y. M. Greshichchev *et al.*, "A Fully Integrated SiGe Receiver IC for 10-Gb/s Data Rate," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 12, Dec. 2000, pp. 1949-1957.

[39] Y. Moon *et al.*, "A Quad 6Gb/s Multi-rate CMOS Transceiver with TX Rise/Fall-Time Control," *IEEE International Solid-State Circuits Conference*, Feb. 2006.

[40] H. Johnson and M. Graham, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, 1993.

[41] V. Stojanović and M. Horowitz, "Modeling and Analysis of High-Speed Links," *IEEE Custom Integrated Circuits Conference*, Sep. 2003.

[42] W. Dally and J. Poulton, "Transmitter equalization for 4-Gbps signaling," *IEEE Micro*, vol. 17, no. 1, Jan.-Feb. 1997, pp. 48-56.

[43] J. Jaussi *et al.*, "8-Gb/s Source-Synchronous I/O Link With Adaptive Receiver Equalization, Offset Cancellation, and Clock De-Skew," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 1, Jan. 2005, pp. 80-88.

[44] H. Wu *et al.*, "Integrated Transversal Equalizers in High-Speed Fiber-Optic Systems," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 12, Dec. 2003, pp. 2131-2137.

[45] D. Hernandez-Garduno and J. Silva-Martinez, "A CMOS 1Gb/s 5-Tap Transversal Equilzer Based on Inductorless 3rd-Order Delay Cells," *IEEE International Solid-State Circuits Conference*, Feb. 2007.

[46] A. Ho *et al.*, "Common-mode Backplane Signaling System for Differential High-Speed Links," *IEEE Symposium on VLSI Circuits*, June 2004.

[47] K. K. Parhi, "Pipelining in Algorithms with Quantizer Loops," *IEEE Transactions on Circuits and Systems*, vol. 38, no. 7, July 1991, pp. 745-754.

[48] R. Farjad-Rad *et al.*, "A 0.3-μm CMOS 8-Gb/s 4-PAM Serial Link Transceiver," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 5, May 2000, pp. 757-764.

[49] J. Zerbe *et al.*, "1.6Gb/s/pin 4-PAM signaling and circuits for a multidrop bus," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 5, May 2001, pp. 752-760.

[50] J. Stonick *et al.*, "An adaptive pam-4 5-Gb/s backplane transceiver in 0.25-μm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 38, no. 3, Mar. 2003, pp. 436-443.

[51] V. Stojanović, "Channel-Limited High-Speed Links: Modeling, Analysis and Design," *PhD. Thesis*, Stanford University, Sep. 2004.

[52] G. Cherubini *et al.*, "Filter Bank Modulation Techniques for Very High-Speed Digital Subscriber Lines," *IEEE Communications Magazine*, vol. 38, no. 5, May 2000, pp. 98-104.

[53] S. B. Weinstein and P. M. Ebert, "Data Transmission by Frequency-Division Multiplexing Using the Discrete Fourier Transform," *IEEE Transactions on Communications*, vol. 19, no. 5, Oct. 1971, pp. 628-634.

[54] B. Hirosaki, "An Orthogonally Multiplexed QAM System using the Discrete Fourier Transform," *IEEE Transactions on Communications*, vol. 29, no. 7, July 1981, pp. 982-989.

[55] A. Amirkhany *et al.*, "A 24Gb/s Software Programmable Multi-Channel Transmitter," *IEEE Symposium on VLSI Circuits*, June 2007.

[56] R. Farjad-Rad *et al.*, "0.622-8.0Gbps 150mW Serial IO Macrocell with Fully Flexible Preemphasis and Equalization," *IEEE Symposium on VLSI Circuits*, June 2003.

[57] G. Balamurugan *et al.*, "A Scalable 5-15Gbps, 14-75mW Low Power I/O Transceiver in 65nm CMOS," *IEEE Symposium on VLSI Circuits*, June 2007.

[58] L. Kazovsky, S. Benedetto, and A. Wilner, *Opitcal Fiber Communication Systems*, Artech House, 1996.

[59] D. A. B. Miller, "Rationale and Challenges for Optical Interconnects to Electronic Chips," *Proc. IEEE*, vol. 88, no. 6, June 2000, pp. 728-749.

[60] H. A. Willebrand and B. S. Ghuman, "Fiber Optics Without Fiber," *IEEE Spectrum*, vol. 38, no. 8, Aug. 2001, pp. 40-45.

[61] H. Thienpont *et al.*, "Plastic Microoptical Interconnection Modules for Parallel Free-Space Inter- and Intra-MCM Data Communication," *Proc. IEEE*, vol. 88, no. 6, June 2000, pp. 769-779.

[62] M. Gruber, R. Kerssenfischer, and J. Jahns, "Planar-Integrated Free-Space Optical Fan-Out Module for MT-Connected Fiber Ribbons," *Journal of Lightwave Technology*, vol. 22, no. 9, Sep. 2004, pp. 2218-2222.

[63] [AMP NETCONNECT: XG Optical Fiber System](http://www.ampnetconnect.com/documents/XG_Optical_Fiber_Systems_WhitePaper.pdf) (http://www.ampnetconnect.com/documents/XG_Optical_Fiber_Systems_WhitePaper.pdf)

[64] A. Narasimha *et al.*, "A Fully Integrated 4x10Gb/s DWDM Optoelectronic Transceiver in a standard 0.13μm CMOS SOI," *IEEE International Solid-State Circuits Conference*, Feb. 2007.

[65] D. Bossert *et al.*, "Production of high-speed oxide confined VCSEL arrays for datacom applications," *Proc. SPIE*, vol. 4649, June 2002, pp. 142-151.

[66] M. Grabherr *et al.*, "Fabricating VCSELs in a high tech start-up," *Proc. SPIE*, vol. 4942, Apr. 2003, pp. 13-24.

[67] D. Vez *et al.*, "10 Gbit/s VCSELs for datacom: devices and applications," *Proc. SPIE*, vol. 4942, Apr. 2003, pp. 29-43.

[68] J. Jewell *et al.*, "1310nm VCSELs in 1-10Gb/s commercial applications," *Proc. SPIE*, vol. 6132, Feb. 2006, pp. 1-9.

[69] M. A. Wistey *et al.*, "GaInNAsSb/GaAs vertical cavity surface emitting lasers at 1534nm," *Electronics Letters*, vol. 42, no. 5, Mar. 2006, pp. 282-283.

[70] D. Wiedenmann *et al*., "Design and Analysis of Single-Mode Oxidized VCSEL's for High-Speed Optical Interconnects," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 5, no. 3, May-June 1999, pp. 503-511.

[71] N. Suzuki *et al*., "1.1-μm-Range InGaAs VCSELs for High-Speed Optical Interconnections," *IEEE Photonics Technology Letters*, vol. 18, no. 12, June 2006, pp. 1368-1370.

[72] L. A. Coldren and S. W. Corzine, *Diode Lasers and Photonic Integrated Circuits*, Wiley-Interscience, 1995.

[73] Y. Liu *et al*., "Numerical Investigation of Self-Heating Effects of Oxide-Condfined Vertical-Cavity Surface-Emitting Lasers," *IEEE Journal of Quantum Electronics*, vol. 41, no. 1, Jan. 2005, pp. 15-25.

[74] K. W. Goossen, "Fitting Optical Interconnects to an Electrical World: Packaging and Reliability Issues of Arrayed Optoelectronic Modules," *IEEE Lasers and Electro-Optics Society Annual Meeting*, Nov. 2004.

[75] M. Teitelbaum and K. W. Goossen, "Reliability of Direct Mesa Flip-chip Bonded VCSEL's," *IEEE Lasers and Electro-Optics Society Annual Meeting*, Nov. 2004.

[76] C. Helms *et al.*, "Reliability of Oxide VCSELs at Emcore," *Proc. SPIE*, vol. 5364, June 2004, pp. 183-189.

[77] D. Kucharski *et al.*, "A 20Gb/s VCSEL Driver with Pre-Emphasis and Regulated Output Impedance in 0.13μm CMOS," *IEEE International Solid-State Circuits Conference*, Feb. 2005.

[78] S. Palermo and M. Horowitz, "High-Speed Transmitters in 90nm CMOS for High-Density Optical Interconnects," *IEEE European Solid-State Circuits Conference*, Feb. 2006.

[79] A. Kern, A. Chandrakasan, and I. Young, "18Gb/s Optical I/O: VCSEL Driver and TIA in 90nm CMOS," *IEEE Symposium on VLSI Circuits*, June 2007.

[80] H. C. Neitzert *et al.*, "InGaAs-InP Superlattice Electroabsorption Waveguide Modulator," *IEEE Photonics Technology Letters*, vol. 7, no. 8, Aug. 1995, pp. 875-877.

[81] H. Fukano *et al.*, "Very-Low-Driving-Voltage Electroabsorption Modulators Operating at 40Gb/s," *Journal of Lightwave Technology*, vol. 24, no. 5, May 2006, pp. 2219-2224.

[82] D. A. B. Miller *et al.*, "The Quantum Well Self-Electrooptic Effect Device: Optoelectronic Bistability and Oscillation, and Self-Linearized Modulation," *IEEE Journal of Quantum Electronics*, vol. QE-21, no. 9, Sep. 1985, pp. 1462-1476.

[83] A. L. Lentine *et al.*, "High-Speed Optoelectronic VLSI Switching Chip with >4000 Optical I/O Based on Flip-Chip Bonding of MQW Modulators and Detectors to Silicon CMOS," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 2, no. 1, Apr. 1996, pp. 77-84.

[84] G. A. Keeler, "Optical Interconnects to Silicon CMOS: Integrated Optoelectronic Modulators and Short Pulse Systems," *PhD. Thesis*, Stanford University, Dec. 2002.

[85] N. C. Helman *et al.*, "Misalignment-Tolerant Surface-Normal Low-Voltage Modulator for Optical Interconnects," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 11, no. 2, Mar.-Apr. 2005, pp. 338-342.

[86] N. C. Helman, "Optoelectronic Modulators for Optical Interconnects," *PhD. Thesis*, Stanford University, May 2005.

[87] D. A. B. Miller, "Quantum Well Optoelectronic Switching Devices," *International Journal of High Speed Electronics*, vol. 1, no. 1, Jan. 1990, pp. 19-46.

[88] A. V. Krishnamoorthy and D. A. B. Miller, "Scaling Optoelectronic-VLSI Circuits into the 21st Century: A Technology Roadmap," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 2, no. 1, Apr. 1996, pp. 55-76.

[89] A. V. Krishnamoorthy *et al.*, "Ring oscillators with optical and electrical readout based on hybrid GaAs MQW modulators bonded to 0.8μm silicon VLSI circuits," *Electronics Letters*, vol. 31, no. 22, Oct. 1995, pp. 1917-1918.

[90] M. Jutzi *et al.*, "40-Gbit/s Ge-Si Photodiodes," *Silicon Monolithic Integrated Circuits in RF Systems*, Jan. 2006.

[91] M. Yang *et al.*, "A High-Speed, High-Sensitivity Silicon Lateral Trench Photodetector," *IEEE Electron Device Letters*, vol. 23, no. 7, July 2002, pp. 395-397.

[92] M. R. Reshotko, D. L. Kencke, and B. Block, "High-speed CMOS compatible photodetectors for optical interconnects," *Proc. SPIE*, vol. 5564, Oct. 2004, pp. 146-155.

[93] S. S. Mohan *et al.*, "Bandwidth Extension in CMOS with Optimized On-Chip Inductors," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 3, Mar. 2000, pp. 346-355.

[94] C. Kromer *et al.*, "A 100-mW 4x10Gb/s Transceiver in 80-nm CMOS for High-Density Optical Interconnects," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 12, Dec. 2005, pp. 2667-2679.

[95] S. Palermo, A. Emami-Neyestanak, and M. Horowitz, "A 90nm CMOS 16Gb/s Transceiver for Optical Interconnects," *IEEE International Solid-State Circuits Conference*, Feb. 2007.

[96] Avalon Photonics datacom VCSELs (http://www.avap.ch/an1.pdf)

[97] J. G. Proakis and M. Salehi, *Communication Systems Engineering*, Prentice Hall, 2001.

[98] A. Emami-Neyestanak *et al.*, "CMOS Transceiver with Baud Rate Clock Recovery for Optical Interconnects," *IEEE Symposium on VLSI Circuits*, June 2004.

[99] R. Pierret, *Semiconductor Device Fundamentals*, Pearson Higher Education, 2003.

[100] T. K. Woodward *et al.*, "Modulator-Driver Circuits for Optoelectronic VLSI," *IEEE Photonics Technology Letters*, vol. 9, no. 6, June 1997, pp. 839-841.

[101] B. Serneels *et al.*, "A High-Voltage Output Driver in a 2.5-V 0.25-μm CMOS Technology," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 3, Mar. 2005, pp. 576-583.

[102] A.-J. Annema, G. J. G. M. Geelen, and P. C. de Jong, "5.5-V I/O in a 2.5-V 0.25-μm CMOS Technology," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 3, Mar. 2001, pp. 528-538.

[103] S. Hara *et al.*, "Broad-Band Monolithic Microwave Active Inductor and Its Application to Miniaturized Wide-Band Amplifiers," *IEEE Transactions on Microwave Theory and Techniques*, vol. 36, no. 12, Dec. 1988, pp. 1920-1924.

[104] R. Ho *et al.*, "Applications of On-Chip Samplers for Test and Measurement of Integrated Circuits," *IEEE Symposium on VLSI Circuits*, June 1998.

[105] J. E. Roth *et al.*, "An optical interconnect transceiver at 1550nm using low voltage electroabsorption modulators directly integrated to CMOS," Accepted for publication in *Journal of Lightwave Technology*.

[106] M. J. M. Pelgrom, A. C. J. Duinmaijer, and A. P. G. Welbers, "Matching Properties of MOS Transistors," *IEEE Journal of Solid-State Circuits*, vol. 24, no. 5, Oct. 1989, pp. 1433-1440.

[107] M.-J. E. Lee, W. J. Dally, and P. Chiang, "Low-Power Area-Efficient High-Speed I/O Circuit Techniques," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 11, Nov. 2000, pp. 1591-1599.

[108] B. Nikolić *et al.*, "Improved Sense-Amplifier-Based Flip-Flop: Design and Measurements," *IEEE Journal of Solid-State Circuits*, vol. 35, no. 6, June 2000, pp. 876-884.

[109] A. X. Widmer and P. A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM Journal of Research and Development*, vol. 27, no. 5, Sep. 1983, pp. 440-451.

[110] A. Emami-Neyestanak, "Design of CMOS Receivers for Parallel Optical Interconnects," *PhD. Thesis*, Stanford University, Aug. 2004.

[111] L. Sun and T. A. Kwasniewski, "A 1.25-GHz 0.35-μm Monolithic CMOS PLL Based on a Multiphase Ring Oscillator," *IEEE Journal of Solid-State Circuits*, vol. 36, no. 6, June 2001, pp. 910-916.

[112] B. Razavi, *Design of Analog CMOS Intregrated Circuits*, McGraw-Hill, 2000.

[113] E. Alon *et al.*, "Replica Compensated Linear Regulators for Supply-Regulated Phase-Locked Loops," *IEEE Journal of Solid-State Circuits*, vol. 41, no. 2, Feb. 2006, pp. 413-424.

[114] H. Ebenhoech, "Make IC digital frequency comparators," *Electronic Design*, vol. 15, no. 14, July 1967, pp. 62-64.

[115] M. Mansuri, D. Liu, and C.-K. K. Yang, "Fast Frequency Acquisition Detectors for GSamples/s Phase-Locked Loops," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 10, Oct. 2002, pp. 1331-1334.

[116] P. Larsson, "A 2-1600-MHz CMOS Clock Recovery PLL with Low-Vdd Capability," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 12, Dec. 1999, pp. 1951-1960.

[117] H. Lee *et al.*, "Improving CDR Performance via Estimation," *IEEE International Solid-State Circuits Conference*, Feb. 2006.

[118] L. Wu and W. C. Black Jr., "A Low-Jitter Skew-Calibrated Multi-Phase Clock Generator for Time-Interleaved Applications," *IEEE International Solid-State Circuits Conference*, Feb. 2001.

[119] K. Yamaguchi *et al.*, "2.5GHz 4-phase Clock Generator with Scalable and No Feedback Loop Architecture," *IEEE International Solid-State Circuits Conference*, Feb. 2001.

[120] S. A. Blokhin *et al.*, "Vertical-Cavity Surface-Emitting Lasers Based on Submonolayer InGaAs Quantum Dots," *IEEE Journal of Quantum Electronics*, vol. 42, no. 9, Sep. 2006, pp. 851-858.

[121] Y.-H. Kuo *et al.*, "Quantum-Confined Stark Effect in Ge/SiGe Quantum Wells on Si for Optical Modulators," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 12, no. 6, Nov.-Dec. 2006, pp. 1503-1513.