

Optical Transceivers for High-Speed Data-Center Interconnects: CMOS Circuits, Integration, and Future Trends

Austin Rothschild

Stanford University, Department of Electrical Engineering

December 13, 2025

Contents

1	Introduction	2
2	Evolution of Optical Transceivers	3
2.1	Incoherent Links	3
2.2	Coherent Links	6
3	CMOS Circuits for Optical Interconnects and Transceivers	7
3.1	Transceiver Architectures	7
3.2	Receiver Systems	8
3.2.1	Limiting Amplifier Gain-Bandwidth	8
3.2.2	Transimpedance Amplifiers	9
3.2.3	Continuous Time Linear Equalization	10
3.2.4	Feedforward Equalization	10
3.2.5	Decision Feedback Equalization	11
3.3	Transmitter Systems	12
3.3.1	VCSEL Nonidealities	12
3.3.2	DAC Based Transmitters	13
4	State-of-the-Art in Optical Transceivers	14
4.1	Modern Receiver Architecture and Techniques	14
4.2	Modern Transmitter Architecture and Techniques	16
5	A Brief Look at Future Directions	17
5.1	Heterogeneous Integration and Co-Packaged Optics	17
6	Outlook and Conclusions	19

1 Introduction

Over the past decade, the rise of data-intensive computing, driven by machine learning, cloud services, and 5G communication networks, has placed interconnect bandwidth and energy efficiency at the forefront of system design. Modern data centers now operate at aggregate throughputs of tens of terabits per second per package, with individual electrical lanes scaling from 100 Gb/s to beyond 224 Gb/s. Electrical I/Os, long the foundation of intra- and inter-rack communication, are increasingly constrained by frequency-dependent conductor loss, aggressive equalization requirements, and thermal-density limits. As a result, electrical interconnects have reached a point where further gains in bandwidth or reach incur disproportionate increases in complexity, power consumption, and cost.

These challenges are exacerbated by the rapid expansion of data centers and scale-out networks to support machine-learning workloads, where interconnect power now represents a significant fraction of total system energy. According to the 2024 Report on U.S. Data Center Energy Use [1], data centers accounted for approximately 4.4% of total U.S. electricity consumption in 2023, with projections estimating this figure could rise to 12% by 2028. A substantial portion of this energy is consumed by the interconnect fabric itself, underscoring the urgent need for new approaches at the circuit, packaging, and system level to sustain scalable compute within realistic power budgets.

In contrast to electrical interconnects, optical communication systems have demonstrated sustained improvements in both aggregate data rate and normalized energy efficiency over the past two decades. As shown in Fig. 1, pluggable optical transceivers have increased data rates by nearly three orders of magnitude while simultaneously reducing energy per bit by approximately two orders of magnitude. These trends highlight the fundamental advantage of optical signaling for high-bandwidth communication, particularly as electrical links approach their scaling limits.




Figure 1: Energy consumption and data rates for optical transceivers over the past two decades [2].

Despite this remarkable progress, maintaining these scaling trends has become increasingly challenging as optical transceiver architectures approach fundamental technological limits. At the same time, data-center traffic growth has continued to outpace improvements in interface bandwidth. For example, between 2014 and 2021, Google’s data-center traffic increased by more than a factor of 36, while networking switch capacity grew by only a factor of 8—from 3.2 Tb/s (128×25 Gb/s) to 25.6 Tb/s

(256 × 100 Gb/s) [3]. This widening gap between traffic demand and interface capability motivates renewed innovation across the entire optical interconnect stack, including circuits, packaging, and system architecture.

Continued progress in electronics and hardware is therefore a key pillar in extending the scalability of optical interconnects. In this context, the optical transceiver serves as the critical interface between high-speed CMOS electronics and the optical domain, with its circuit-level design directly influencing achievable bandwidth, energy efficiency, and integration density. Accordingly, this paper reviews and discusses the key circuit technologies that enable high-speed optical communication, with emphasis on the fundamental non-idealities they impose and the techniques used to mitigate them. The focus is placed on short-reach, high-bandwidth applications relevant to modern data centers.

The remainder of this paper is organized as follows. Section 2 reviews prevalent signaling techniques in optical communication systems and the resulting implications for transceiver architecture. Section 3 examines the CMOS receiver and transmitter circuits that form the core of modern optical transceivers. Building on these foundations, Section 4 surveys a recent state-of-the-art implementation of high-speed CMOS circuits for optical communication at 224 Gb/s. Finally, Section 5 discusses emerging integration and architectural trends in short-reach optical interconnects, with particular emphasis on heterogeneous integration and co-packaged optics.

2 Evolution of Optical Transceivers

Over the past four decades, optical communication systems have undergone a remarkable evolution, driven by advances spanning materials science, device physics, signal processing, coding theory, networking, electronics, and photonics. Few technological fields integrate such a broad range of disciplines as tightly or as effectively. The resulting scalability and robustness of optical networks were underscored during the COVID-19 pandemic, when global fiber-optic infrastructure served as the critical backbone enabling widespread broadband connectivity. Early optical communication systems operated at data rates on the order of a few megabits per second, whereas modern systems are capable of transmitting multiple terabits per second over thousands of kilometers on a single optical fiber. This section briefly highlights the dominant signaling techniques that have emerged across different generations of optical communication systems, providing context for the transceiver architectures and circuit design choices discussed in the sections that follow.

2.1 Incoherent Links

For nearly three decades, optical communication systems have relied on intensity modulation with direct detection (IM-DD), which remains the dominant technique for short-reach data-center links, while coherent detection has become the de facto standard for long-haul and metro networks. A simplified view of an IM-DD optical link is shown in Fig. 5.

In the case of binary (2-PAM) IM-DD signaling, a two-level electrical drive current at the transmitter, $I_{TX}(t)$, is converted into an optical power waveform $P_{TX}(t) = \eta I_{TX}(t) + P_0$, where η denotes the electro-optic conversion efficiency of the laser and P_0 is the optical bias power. The optical signal propagates through a fiber channel—either single-mode (SMF) or multi-mode (MMF)—which introduces impairments such as pulse spreading due to chromatic dispersion. At the receiver, a photodiode converts the received optical power $P_{RX}(t)$ back into an electrical current for subsequent amplification and processing.

Table 1: Evolution of Optical Interconnect Modules and Lane Rates

Rate (Gb/s)	Module	Electrical Lanes			Signal Rate (GBaud)
		# Lanes	Gb/s	Format	
1	SFP	1	1	NRZ	1
10	SFP	1	10	NRZ	10
40	QSFP	4	10	NRZ	4×10
100	QSFP-28	4	25	NRZ	4×25
200	QSFP-56	4	50	4PAM	4×25
400	QSFP-DD / OSFP	8	50	4PAM	4×50
800	QSFP-DD / OSFP	8	100	4PAM	8×50
1600	OSFP-XD	16	100	4PAM	4×100
3200	OSFP-XD	16	200	4PAM	$2 \times 4 \times 100$
3200	CPO	32	100	4PAM	32×50




Figure 2: Simplified view of intensity modulation with direct detection (IM-DD) in an optical communication system. [4]

IM-DD has remained the preferred architecture for short-reach links for several reasons. Its relative simplicity enables low-cost implementations, while the large symbol spacing associated with two-level signaling relaxes signal-to-noise ratio requirements and reduces reliance on complex, power-hungry equalization techniques. This inherent robustness to non-idealities in electro-optic and opto-electronic conversion has made IM-DD particularly attractive for data-center environments. Modern IM-DD systems predominantly employ four-level pulse-amplitude modulation (4-PAM), which increases link data rates without a proportional increase in analog bandwidth. 4-PAM has been widely adopted in data-center interconnect standards, including IEEE 400GbE (P802.3bs). As discussed in Section 3, 4-PAM operation imposes distinct circuit-level requirements compared to binary signaling. The evolution of IM-DD transceiver form factors is illustrated in Fig. 3, with corresponding module architectures and lane configurations summarized in Table 1. Historically, optical transceivers have been deployed as front-panel pluggable modules, simplifying system integration, replacement, and technology migration across different process nodes. This paradigm is increasingly evolving with the emergence of heterogeneous integration and co-packaged optics, discussed further in Section 5.

Table 2: Evolution of Optical Module Power and Energy Efficiency

Module	Year	Power (mW)	Rate (Gb/s)	Energy (pJ/bit)
1G SFP	2001	1.0	1	1000
10G SFP	2009	2.5	10	250
40G QSFP	2012	3.5	40	88
100G QSFP (CWDM4)	2015	3.5	100	35
200G QSFP (FR4)	2018	5.5	200	28
400G QSFP-DD (FR4)	2020	10	400	25
800G OSFP (FR4)	2022	16	800	20
CPO (OIF 2021.131.12)	2024	56	3200	18




Figure 3: The evolution of optical transceiver modules, arranged by year introduced (with CPO the most recent), over the past two decades [2].

As shown in Table 2, successive generations of optical modules have achieved substantial reductions in energy per bit under peak operating conditions, even as aggregate data rates have increased dramatically. However, these gains in energy efficiency have not translated into lower absolute power consumption; instead, total module power has increased by roughly an order of magnitude. Furthermore, increasing baud rates place stringent demands on the analog bandwidth of transmitter and receiver circuits. Together, these trends place a growing burden on high-speed circuit design, thermal management, and packaging strategies, motivating continued innovation in optical transceiver architectures.

2.2 Coherent Links




Figure 4: Block diagram of a standard digital coherent transceiver with common DSP functionality [5].

Coherent optical communication systems were originally developed for long-haul and metro networks, where maximizing spectral efficiency and receiver sensitivity is paramount. A block diagram of a coherent DSP based optical transceiver is displayed in Fig. 4. By encoding information in both the amplitude and phase of the optical field and employing a local oscillator at the receiver, coherent links enable higher dimensional modulation formats such as QPSK and high-order QAM, as well as polarization multiplexing. These capabilities allow coherent systems to achieve superior tolerance to fiber impairments, higher reach, and significantly improved optical signal-to-noise ratio (OSNR) efficiency compared to intensity-modulated direct-detection (IM-DD) links. However, coherent detection requires precise optical phase recovery and substantial DSP techniques to realize operation resulting in significantly higher power consumption and system complexity. At long-distances, where systems are limited by fiber impairments such as chromatic dispersion, polarization mode dispersion, and fiber nonlinearity, these DSP techniques are well worth the cost. However, short-reach links don't suffer from these effects and thus don't gain much from IQ based links beyond increased high-speed spectral efficiency.

As a result, despite their performance advantages, coherent links have historically been unattractive for short-reach data-center applications, where energy per bit, latency, and cost dominate over reach and spectral efficiency. In contrast, IM-DD architectures offer simpler transceivers with lower power consumption and have therefore remained the preferred solution for intra- and inter-rack links. Nevertheless, continued scaling of IM-DD systems toward higher lane rates increasingly relies on aggressive equalization and high-resolution mixed-signal circuits, eroding their traditional power and simplicity advantages. This has renewed interest in coherent systems with simplified or "coherent-lite" [6] architectures that reduce or eliminate DSP-intensive processing while retaining key benefits of coherent detection, as explored in prior work such as [7]. Enabled by advances in CMOS scaling, silicon photonics integration, and low-power mixed-signal design, these emerging coherent approaches may offer a viable path toward

energy-efficient optical interconnects for future data-center links in the kilometer-scale regime.

3 CMOS Circuits for Optical Interconnects and Transceivers

This section discusses the architecture and design philosophy of fundamental transmit and receiver circuit architectures and the impairments they introduce in an optical communication link. Timing and synchronization systems are omitted for brevity, as well as substantial discussion on optical devices (with the exception of VCSELs) with the focus on the key CMOS circuits that comprise these systems.

3.1 Transceiver Architectures

Optical transceiver and interconnect architectures are commonly partitioned into an electronic domain, responsible for waveform generation, timing, and signal conditioning, and a photonic domain, responsible for optical carrier generation, modulation, and detection. The simplified IM-DD interconnect architecture illustrated Fig. 5 reflects this separation: high-speed CMOS circuitry implements clock generation, M-PAM serialization, and amplification, while photonic devices such as the laser, modulator, and photodetector (PD) perform electro-optic and opto-electronic conversion. This division arises from fundamental differences in materials, scaling behavior, noise sensitivity, and thermal requirements between CMOS electronics and photonic devices, which historically favored fabrication and optimization in separate process technologies.




Figure 5: Simplified block diagram of a high-speed optical interconnect implementing IM-DD. In modern IM-DD links, $M = 4$.

Recent efforts in co-packaged optics (CPO) seek to reduce the energy and bandwidth penalties associated with long electrical interfaces by integrating photonic engines in close proximity to the switch or compute ASIC. By shortening high-speed electrical paths, CPO enables lower I/O power and higher bandwidth density, but introduces new challenges in thermal management, testability, laser integration, and system-level co-design, making electronic–photonic integration an active and rapidly evolving research area briefly mentioned in section 5. In this section, fundamental CMOS components comprising the transmitter and receiver blocks are described.

3.2 Receiver Systems

The basic optical receiver architecture is shown below in Fig. 6. It consists of a front-end photodetector (PD) to convert the incident optical signal into a current, a low-noise transimpedance amplifier (TIA) stage to convert the PD current into an amplified voltage signal. Depending on the topology of the TIA, an equalizer may follow in order to help compensate the bandwidth limiting effects of the TIA. This also allows the use of more narrow-band front-end TIAs in order to reduce the total integrated RX noise, with bandwidth compensation provided by the post TIA equalizer. This approach greatly relaxes some of the tradeoffs of gain, power, noise, bandwidth, and linearity that exist in the analog TIA design, and is a common approach in TIA implementation. A cascade of limiting amplifiers (LA) ¹ follow the front-end in order to convert the variable-amplitude electrical signal at the output of the TIA/Equalizer stage into a clean, constant-amplitude digital level waveform suitable to be processed by the clock and data recovery (CDR) system, and their selection of gain and number of stages is described in detail next.

2




Figure 6: Block diagram of an optical receiver front-end circuit [8].

3.2.1 Limiting Amplifier Gain-Bandwidth

A cascade of LAs are used in order to distribute the gain across stages, which allows for lower overall stage gain and thus higher bandwidth per stage. If each gain stage in the LA chain is a single pole amplifier, the total multi-stage LA transfer function will be:

$$\frac{v_{out}}{v_{in}} = \left(\frac{A_{vs}}{1 + \frac{s}{\omega_{3dBs}}} \right)^n = A_{vs}^n \frac{1}{\left(1 + \frac{s}{\omega_{3dBs}} \right)^n} \quad (1)$$

where A_{vs} is the per-stage voltage gain, and ω_{3dBs} is the per-stage 3-dB bandwidth frequency. The total cascaded 3-dB bandwidth ω_{3dBtot} is given as:

$$\left(1 + \left(\frac{\omega_{3dBtot}}{\omega_{3dBs}} \right)^2 \right)^n = 2 \quad (2)$$

$$\omega_{3dBtot} = \omega_{3dBs} \sqrt{2^{1/n} - 1}$$

If each LA stage supports maximum per-stage GBW_s due to technology constraints, then $\omega_{3dBs} = \frac{GBW_s}{A_{vs}}$. Thus if a high bandwidth is needed, the per stage gain is reduced and more stages are used. However, from Eq. 2, the bandwidth will compress with cascaded stages. Therefore, there exists an optimum number of stages for maximizing potential gain bandwidth. To obtain the optimal number of stages and per-stage gain, the following approximation is made:

¹Also referred to as post-amplifiers or automatic gain control (AGC) amps.

²CDRs remain outside the scope of this review, however, they are a crucial system in any optical RX.

$$\omega_{3dBtot} = \frac{GBW_s}{G_{tot}^{1/n}} \sqrt{2^{1/n} - 1} \approx \frac{GBW_s}{G_{tot}^{1/n}} \sqrt{\frac{1}{n} \ln 2} \quad (3)$$

Taking the reciprocal of this expression and minimizing the natural log of the denominator yields an expression for the optimal number of stages:

$$\begin{aligned} \frac{d}{dn} \left(\ln \left(\left(\frac{\sqrt{n}}{GBW_s \sqrt{\ln 2}} \right) G_{tot}^{1/n} \right) \right) &= \frac{d}{dn} \left(\frac{1}{2} \ln(n) + \frac{1}{n} \ln(G_{tot}) - \ln(GBW_s \sqrt{\ln 2}) \right) = 0 \\ \frac{1}{2n} - \frac{1}{n^2} \ln(G_{tot}) &= 0 \\ \frac{1}{n} \ln(G_{tot}) &= \frac{1}{2} \end{aligned} \quad (4)$$

Thus, the optimum number of LA stages in the cascade is:

$$n_{opt} = 2 \ln(G_{tot}) \quad (5)$$

which corresponds to an optimal per-stage gain of:

$$\begin{aligned} A_{vs,opt}^{2 \ln(G_{tot})} &= G_{tot} \\ 2 \ln(G_{tot}) \ln(A_{vs,opt}) &= \ln(G_{tot}) \\ A_{vs,opt} &= \sqrt{e} = 1.65 \end{aligned} \quad (6)$$

From a power and noise perspective, the number of stages is usually chosen less than the theoretical optima. Typically, LA stages range between $n = 2$ to $n = 6$ in published works and implementations. Additionally, the bandwidth of each LA stage is frequently extended using inductive techniques such as shunt peaking, series peaking, and T-coil peaking using on-chip CMOS inductors [9].

3.2.2 Transimpedance Amplifiers

Optical receiver (RX) design faces a fundamental tradeoff between gain, noise, bandwidth, linearity, and power efficiency, and the choice of TIA topology attempts to balance these tradeoffs.³ The shunt-feedback inverter based topology is a commonly implemented TIA block as it offers high linearity, high transconductance per unit bias current, high swing, and self-biasing. The TIA in Fig. 7 displays the simplified architecture.




Figure 7: TIA circuit diagram (left) and associated small right model (right) with noise sources shown in red [11].

Increasing the feedback resistance R_F helps lower input-referred noise at the cost of reduced front-end bandwidth and thus additional ISI. This bandwidth reduction can be compensated using various

³Modern short-reach links have increasingly moved towards ADC based receivers that may often entirely omit a first-stage TIA to save power if the link has sufficient noise margin performance [10].

equalization techniques, which are later described. It can be shown that the transfer function of this TIA is:

$$Z_{TIA}(s) = \frac{-R_a(g_m R_f - 1 - sC_{gd}R_f)}{1 + g_m R_a + sK_1 + s^2K_2}, \quad (7)$$

where

$$K_1 = C_{IN}(R_a + R_f) + C_a R_a + C_{gd}R_f(1 + g_m R_a) \quad (8)$$

and

$$K_2 = R_f R_a (C_{IN}C_a + C_{IN}C_{gd} + C_{gd}C_a). \quad (9)$$

It is readily observed from the TIA transfer function that there are several bandwidth limiting parasitic capacitance's, namely the sum of the photodetector capacitance, C_{PD} , the chip pad capacitance, C_{PAD} , and inverters C_{gs} . Additionally, the C_{gd} capacitance is important due to the Miller effect, which can become the dominant pole in the case of large transistor sizes and feedback resistances, both of which are employed to lower the TIA's input referred noise. To compensate the frequency response of TIA's transfer function in Eq. 7, several receive techniques may be used in isolation or simultaneously and are briefly discussed. While each of these bandwidth extension techniques can provide good performance, at some point package level optimization must be explored to fundamentally reduce bandwidth limiting parasitics in the TIA front end. Techniques such as flip-chip co-packaging of the TIA and PD have become common practices to minimize parasitics and also expand I/O density in favor of traditional wire-bond based solution [12].

3.2.3 Continuous Time Linear Equalization

A first-stage TIA may have its bandwidth limitation compensated by a second-stage continuous time linear equalizer (CTLE) in order to help reduce intersymbol interference (ISI). An ideal unity-gain CTLE stage completely recovers the bandwidth of the RX front-end by compensating the TIA and channel frequency responses. The ideal CTLE has a transfer function of:

$$H_{CE}(s) = \frac{1 + g_m R_a + sK_1 + s^2K_2}{(1 + g_m R_a) \left(1 + \frac{s}{\chi Q 2\pi f_{TIA}} + \frac{s^2}{(\chi 2\pi f_{TIA})^2} \right)} \quad (10)$$

where f_{TIA} is the 3-dB bandwidth of the TIA, and Q is the quality factor of the TIA. The zeros of an ideal CTLE perfectly cancel the poles of the TIA, and the poles of the CTLE are χ times higher with respect to those of the TIA. The product of Eq. 7 and Eq. 10 gives the equalized transfer function, resulting in a flat gain across frequency at the expense of attenuation. The basic analog CTLE can easily be realized using resistive and capacitive degeneration in a differential pair in order to synthesize one zero and two poles, with additional inductive peaking techniques employed for wideband operation [9]. In practice, this is primarily used as a first order compensation approach, and pairs with a combination of the techniques described next.

3.2.4 Feedforward Equalization

A feed-forward equalizer (FFE) can be employed after the TIA, which consists of multiple taps that produce scaled, delayed versions of the input signal as shown in Fig 8a. Taking a linear combination of these scaled and shifted versions of the signals allows for reduction in the pre-cursor and post-cursor ISI, as shown in Fig. 8b.




Figure 8: FFE block diagram (a). Basic 3-tap FFE operation (b) [11].

The tap weights can be computed using an adaptive algorithm (such as adaptive LMS) which minimizes the observed error between the received output of the FFE system and a known training sequence. This training sequence is often performed on startup to initially calibrate a link, and may be periodically refreshed in order to deal with temperature variation, offset, and drift which may arise over time and require a recalibration. A drawback of the FFE is noise-enhancement, which arises as the FFE amplifies noise in the same frequency bands that it boosts the signal. As such, high speed optical links commonly employ techniques that avoid noise enhancement such as decision feedback equalization (DFE).

3.2.5 Decision Feedback Equalization

The DFE based receiver uses a finite impulse response (FIR) feedback system as shown in Fig. 9a with each tap of an M-tap DFE used to eliminate an associated post-cursor. Decision feedback equalization mitigates ISI by subtracting an estimate of symbol-dependent interference formed from previously detected symbols at the receiver input. Because this subtraction is performed using reconstructed, noise-free symbol decisions rather than by amplifying the received waveform, a DFE cancels ISI without introducing noise enhancement.




Figure 9: An FIR-DFE (a). Basic FIR-DFE operation (b) [11].

The DFE operation is shown in Fig. 9b using $M = 2$ taps. In the ideal case, all post-cursors can be removed by a DFE, and only the precursors remain and limit the eye opening. Further transmit equalization (also known as TX pre-emphasis) can be used to compensate precursor ISI using a bidirectional link training phase. While the DFE avoids noise enhancement, it has strict timing requirements in the feedback loop, as the DFE output slicer must traverse the feedback path back to the decision device's input in one unit interval (UI). To relax the stringent feedback latency of high-speed DFEs, loop unrolling pre-computes multiple candidate feedback paths corresponding to possible symbol decisions, enabling parallel evaluation while preserving the subtractive interference-cancellation behavior of the DFE.

3.3 Transmitter Systems

Transmitter circuits constitute an important role in the electrical-optical interface in high-speed optical transceivers and play a decisive factor in setting data rate, energy efficiency, and modulation accuracy. As shown in Fig. 5, the TX path consists of an M-PAM serializer followed by a CMOS driver. The CMOS driver output interfaces with an electro-optic modulator, such as a Mach-Zehnder modulator or a ring-resonator. The modulator imprints the data onto a continuous-wave optical carrier supplied by the laser; this separation enables higher bandwidths, improved linearity, and greater system robustness (especially thermal performance) than direct laser modulation.

3.3.1 VCSEL Nonidealities

In short reach links, the TX laser is often a vertical-cavity surface-emitting laser (VCSEL) due to their low cost and low power consumption. However, VCSEL's have three main non-idealities which limit further improvement in optical link speed as shown in Fig. 10. In particular, the VCSEL's light-current conversion efficiency, η , is a function of the bias current [13]. As the bias current of the VCSEL increases, η decreases. Secondly, the VCSEL bandwidth nonlinearly varies with bias current, with increased bias current leading to larger bandwidth. These two effects have little impact for simple two-level (NRZ) systems, however, 4-PAM links must deal with these effects. The VCSEL naturally produces eye heights with different spacings depending on bias current as shown in Fig. 10. Additionally, the widths of all the eye levels are different as a result of bandwidth nonlinearity in applied bias current. These effects lead to distortion of the output optical eye diagram. Finally, the VCSEL produces a rising edge that is faster than the falling edges, leading to an asymmetric output response that degrades with decreasing bias current and increasing data rates [14]. As a result, the VCSEL eye has significant skew, narrowing the sampling window of 4-PAM signals, putting strain on the RX equalizer and sampler performance.




Figure 10: VCSEL nonidealities [15].

To deal with these nonidealities, the use of DSP and a high-speed digital-to-analog converter (DAC) are commonly employed to generate pre-distorted TX waveforms that can compensate these effects. Such architectures have been enabled by CMOS process node scaling and ultra-high speed DACs that implement parallel time-interleaving architectures in order to achieve data rates exceeding 224 Gb/s.

3.3.2 DAC Based Transmitters

As lane speeds have advanced beyond 100 Gb/s toward 224 Gb/s and higher, traditional current-mode driver architectures have increasingly been supplemented—or replaced—by DAC-based transmitters that enable precise multi-level waveform generation and transmitter-side equalization [16]. As previously discussed, the transmitter nonidealities introduced by the VCSEL and downstream components (optical channel and receiver) pose significant obstacles for reliable optical links and must be compensated.

Fig. 15 shows two popular methods for system level implementation of the TX circuit in high-speed wireline and optical links. The analog based TX approach implements both data modulation and feedforward equalization, with FFE tap delays achieved using custom digital delay elements such as latches clocked with multiple phases of half or quarter-rate clocks. A primary drawback of the analog approach is the increased complexity with the number of TX FFE taps, and increased capacitive loading on the analog TX output that results. As such, DAC-based TX architectures have become the standard. Analog design requirements are relaxed in the DAC based approach, as DSP is used to implement modulation and equalization.




Figure 11: Analog based TX (a) and DAC based TX (b) [16].

The architecture of the high-speed TX DAC has two fundamental driver topologies shown in Fig. 12.




Figure 12: Tailless current-mode logic cell (a) and source-series terminated (SST) cell [16].

Both of these driver topologies comprise relatively simple unit cells that are instantiated in parallel, with their current outputs summed in proportion to the digital input code. For example, a CML cell may form one of many identical unit elements in a thermometer-coded DAC, where its bias current is selectively steered to the output node when D_{in} is asserted, or diverted to a dummy load when de-asserted. This operating principle, known as current steering [17], enables data transitions to be realized by redirecting current into a low-impedance load rather than by slewing the output voltage across a large swing. As a result, the output node remains biased at an approximately constant common-mode

voltage, and only a small differential voltage is developed across the load, avoiding the large-signal RC charging that fundamentally limits voltage-mode drivers. This small-signal operation makes current-steering architectures the preferred choice for high-bandwidth DAC-based transmitters. For instance, a 6-bit DAC may employ 64 such unit cells operating in parallel, with their currents summed into a matched $50\ \Omega$ load to synthesize a multi-level or pre-emphasized drive waveform at hundreds of gigahertz. Unit cells may also be binary weighted, or segmented, which combines thermometer and binary weighting approaches for different output bits to optimize DAC linearity.

In contrast, the source-series-terminated (SST) driver cell operates as a tunable transconductance or effective resistive element, producing an output voltage approximately proportional to the input signal and enabling highly linear modulation. This linearity advantage is particularly important at higher resolution, as current-mode logic (CML) based DACs exhibit a code-dependent output impedance that scales with the number of active unit cells, resulting in third-order harmonic distortion [18], which can be characterized by

$$HD_3 = \left[2^{N-2} \frac{R_L}{R_u} \right]^2 \quad (11)$$

By maintaining a more constant output impedance across codes, SST-based linear drivers suppress this distortion mechanism and can achieve superior intrinsic linearity compared to current-steering architectures. However, this improvement in linearity comes at the expense of bandwidth, since SST drivers operate in a voltage-mode regime that requires large-signal slewing of the output node and repeated charging and discharging of load and parasitic capacitance. Consequently, the effective $R_{\text{out}}C_L$ time constant becomes a dominant bandwidth limitation, making SST topologies less scalable to the highest baud rates than current-steering drivers, which operate in a small-signal regime around a fixed common-mode voltage. Despite this trade-off, SST-based linear drivers are increasingly adopted in linear-drive optical interfaces, such as linear pluggable optics (LPO), where reduced DSP complexity, lower power consumption, and high linearity outweigh the ultimate bandwidth limitations [19; 20].

4 State-of-the-Art in Optical Transceivers

A modern example from Intel [21] highlights the most advanced circuit architectures and topologies that enable the highest recorded speeds and best in class pJ/bit performance in leading edge process nodes, providing a case-study into the modern designs that are pushing next generation electrical and optical circuits. This design is representative of a broader industry shift toward ADC- and DAC-based architectures at extreme lane rates.

4.1 Modern Receiver Architecture and Techniques

A dedicated ISSCC publication from this group in 2023 [22] focuses on the RX implementation, with the block diagram shown in 13. This work presents a 224-Gb/s PAM4 ADC-based SERDES receiver implemented in a 5-nm FinFET process, achieving 1.41 pJ/b analog RX energy efficiency while supporting channels with up to 38 dB insertion loss at Nyquist. The receiver architecture combines a hybrid analog front-end (AFE) with a 64-way time-interleaved 6-bit SAR ADC, followed by digital FFE and DFE, and a DCO-based CDR. The design demonstrates that carefully partitioning equalization between the analog and digital domains enables both high bandwidth scalability and energy efficiency at data rates beyond 200 Gb/s.

A central contribution is the hybrid CTLE-based AFE, which uses both inductive peaking and source degeneration to shape the frequency response and closely match channel loss profiles as highlighted in

Fig. 14. By performing a large fraction ($\approx 60\text{--}70\%$) of the total equalization in the analog domain, the design significantly reduces ADC dynamic range requirements, allowing the ADC resolution to be limited to 6 bits, which is critical for power efficiency at 112-GS/s sampling rates




Figure 13: Block diagram of the Intel 224 Gb/s ADC based 4-PAM RX [22].

On the data-conversion side, the receiver employs a loop-unrolled (LU) SAR ADC architecture with extensive time interleaving. Each SAR bit has a dedicated comparator that retains its decision, eliminating reset logic and reducing conversion latency and power. Combined with carefully engineered clock distribution and low-jitter phase interpolation, the ADC sustains Nyquist-rate operation at 56 GHz with manageable jitter and thermal noise. The measured results demonstrate pre-FEC BERs as low as $1.9 \cdot 10^{-11}$ for moderate-loss channels and $6 \cdot 10^{-7}$ at 38 dB loss, validating the robustness of the architecture across a wide channel envelope




Figure 14: Analog front-end (AFE) block diagram including input network, stage 1 CTLE, stage 2/3 limiting amplifiers (VGA) and simulated hybrid CTLE transfer function[22].

In this architecture, a conventional TIA is omitted in favor of a CTLE-based front end, as sensitivity margins in short-reach links permit a linear, wideband voltage-mode interface better suited for high-speed ADC sampling. This approach avoids the bandwidth and power penalties of ultra-wideband TIAs

while enabling flexible DSP-based equalization [10].

4.2 Modern Transmitter Architecture and Techniques

Complementing the highly digital receiver architecture, the associated transmitter adopts a DAC-based design philosophy that mirrors the same shift toward digitally assisted waveform synthesis and equalization. The associated high-speed DAC based transmitter block of the 224 Gb/s Intel transceiver is described in [23], first reported in the 2022 IEEE Journal of Solid State Circuits (JSSC). The block-diagram is highlighted in Fig. 15. In this design, the digital back end performs pulse shaping, equalization, and coefficient scaling, while the DAC converts the digitally equalized symbols into a wideband analog waveform at the output pad. This architecture avoids the heavy capacitive loading and clocking overhead associated with large analog FFE tap banks, which become increasingly inefficient as the number of taps grows at higher data rates, as was previously described. As noted in the paper, a CML current-steering driver topology is used for the DAC unit cells.




Figure 15: Block diagram of the Intel 224 Gb/s DAC based 4-PAM TX [23].

In the Intel transmitter, the DAC is segmented into binary and thermometer-coded slices to balance linearity, mismatch tolerance, and power consumption, achieving the resolution required for 4-PAM while maintaining sub-2 pJ/bit energy efficiency at 224 Gb/s. Thermometer coding of the MSBs reduces differential nonlinearity, while binary-coded LSBs limit area and switching overhead. The segmentation approach is highlighted in Fig. 16. As noted by the authors, the transition from binary to thermometer coding of the DAC was based on Monte Carlo simulations measuring the drive current of each DAC slice to achieve less than 0.5 dB worst-case differential nonlinearity (DNL) in the presence of device mismatch.




Figure 16: DAC segmentation and slice architecture [23].

The paper indicates that state-of-the-art 224-Gb/s 4-PAM transmitters operate near the 2 pJ/b analog regime, with digital FFE adding a few tenths of pJ/b; at 224 Gb/s this corresponds to roughly a half-watt per lane, making I/O power density and thermal constraints first-order design drivers.

As demonstrated by these transmitter and receiver implementations, sustaining bandwidth scaling beyond 100-Gbaud per lane requires increasingly sophisticated circuit architectures that carefully balance linearity, bandwidth, and energy efficiency. DAC-based transmitters rely on segmented current-steering drivers, digital FFE, and multi-phase clocking to generate high-fidelity 4-PAM waveforms, while ADC-based receivers employ linear CTLE front ends and DSP-assisted equalization to maintain scalability at extreme data rates. Together, these designs reflect a broader trend in which circuit complexity migrates from isolated analog blocks to tightly co-designed analog-digital systems. Continued reductions in pJ/bit will depend on advances in wideband data converters, clock generation, and packaging-aware circuit optimization.

5 A Brief Look at Future Directions

5.1 Heterogeneous Integration and Co-Packaged Optics

Heterogeneous integration refers to the assembly of multiple dies or components fabricated in different process technologies and material systems within a single package. In the context of optical interconnects, this paradigm enables CMOS electronic circuits, silicon-photonic devices, and often III–V laser sources to be integrated while preserving independent optimization of each domain. As optical and electrical interconnect requirements continue to diverge, heterogeneous integration has emerged as a key enabler for overcoming the limitations of monolithic integration in advanced systems.

In contemporary data-center architectures, fiber-optic links interface with compute boards through pluggable optical transceiver modules, which then communicate with the compute package via board-level electrical interconnects, which implement some of the circuits and equalizers previously described. This final electrical segment increasingly imposes a bandwidth and energy bottleneck at high data rates. More tightly integrated optical interconnect architectures seek to address this limitation by shortening the electrical reach between electronic and photonic components and relocating optical interfaces closer to the compute fabric.

One prominent manifestation of this trend is co-packaged optics (CPO), in which optical transceiver engines are integrated within the same package as a switch or compute ASIC. By minimizing high-speed electrical interfaces, CPO reduces I/O power consumption and enables significantly higher bandwidth density compared to traditional pluggable solutions [24; 25; 26; 19]. As illustrated in Fig. 17 and Fig. 18, heterogeneous integration can be realized across a spectrum of packaging approaches, including 2D integration on a shared silicon-photonics interposer, 2.5D integration using organic substrates and interposers, and 3D stacking that vertically integrates photonic and electronic dies. Progress along this integration continuum enables progressively shorter electrical interfaces, improved signal integrity, and enhanced energy efficiency. While 2D and 2.5D approaches offer favorable near-term trade-offs in manufacturability and yield, 3D integration represents a longer-term path toward fully integrated optical I/O, positioning heterogeneous integration and co-packaged optics as central elements in the future evolution of data-center interconnects.




Figure 17: 2D, 2.5D, and 3D integration techniques [2].




Figure 18: Evolving approaches to packaging of optical transceivers [2].

6 Outlook and Conclusions

Optical interconnects have become a foundational technology for sustaining bandwidth and energy scaling in modern data-center systems. As electrical I/O approaches fundamental limits imposed by loss, equalization complexity, and power density, optical transceivers provide a scalable path forward by decoupling bandwidth from physical distance and enabling continued reductions in energy per bit. Over the past two decades, advances in CMOS scaling, mixed-signal circuit design, and photonic integration have enabled optical links to increase aggregate data rates by orders of magnitude while significantly improving normalized energy efficiency.

This paper reviewed the circuit-level building blocks that underpin this progress, with a focus on short-reach optical transceivers relevant to data-center environments. On the receiver side, the evolution toward ADC-based receivers with linear CTLE front ends reflects a broader shift toward digitally assisted equalization. This transition enables bandwidth scalability beyond 100 Gbaud per lane while maintaining energy efficiency. On the transmitter side, DAC-based architectures employing segmented current-steering or linear driver topologies enable precise multi-level waveform synthesis and transmitter-side equalization at extreme data rates, but place increasing demands on power delivery, clocking, and thermal management. Together, these designs illustrate how sustaining performance at 200 Gb/s and beyond requires tightly co-designed analog and digital circuits rather than isolated block-level optimization.

Looking forward, continued scaling of optical interconnect bandwidth will increasingly depend not only on circuit innovation but also on advances in system architecture and packaging. Heterogeneous integration and co-packaged optics represent a critical inflection point in this evolution, enabling photonic and electronic subsystems fabricated in disparate technologies to be integrated with minimal electrical reach and high bandwidth density. Despite these advantages, challenges related to thermal management, yield, laser integration, and system test remain active areas of research. At the same time, emerging coherent-lite approaches suggest that selectively reintroducing aspects of coherent detection, while aggressively simplifying signal processing, may offer attractive operating points for future short- to medium-reach links as intensity-modulated direct-detection architectures approach diminishing returns.

Ultimately, the trajectory of optical transceivers underscores a central theme: future gains in bandwidth density and energy efficiency will arise from holistic co-design across circuits, photonics, packaging, and system architecture. As data-center workloads continue to scale in both size and intensity, optical transceivers will remain a critical interface where innovations at the circuit and integration level directly shape the performance, efficiency, and feasibility of next-generation computing systems.

References

- [1] A. Shehabi, S. J. Smith, A. Hubbard, A. Newkirk, N. Lei, M. A. Siddik, B. Holecek, J. G. Koomey, E. R. Masanet, and D. A. Sartor, “2024 united states data center energy usage report,” Lawrence Berkeley National Laboratory, Technical Report LBNL-200163, Dec. 2024. [Online]. Available: <https://escholarship.org/uc/item/32d6m0d1>
- [2] R. Nagarajan, L. Ding, R. Coccioli, M. Kato, R. Tan, P. Tumne, M. Patterson, and L. Liu, “2.5d heterogeneous integration for silicon photonics engines in optical transceivers,” *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 29, no. 3: Photon. Elec. Co-Inte. and Adv. Trans. Print., pp. 1–10, 2023.

- [3] E. Agrell, M. Karlsson, F. Poletti, S. Namiki, X. V. Chen, L. A. Rusch, B. Puttnam, P. Bayvel, L. Schmalen, Z. Tao, F. R. Kschischang, A. Alvarado, B. Mukherjee, R. Casellas, X. Zhou, D. van Veen, G. Mohs, E. Wong, A. Mecozzi, M.-S. Alouini, E. Diamanti, and M. Uysal, “Roadmap on optical communications,” *Journal of Optics*, vol. 26, no. 9, p. 093001, jul 2024. [Online]. Available: <https://doi.org/10.1088/2040-8986/ad261f>
- [4] M. Chagnon, “Optical communications for short reach,” *Journal of Lightwave Technology*, vol. 37, no. 8, pp. 1779–1797, 2019.
- [5] K. Zhong, X. Zhou, J. Huo, C. Yu, C. Lu, and A. P. T. Lau, “Digital signal processing for short-reach optical communications: A review of current technologies and future trends,” *Journal of Lightwave Technology*, vol. 36, no. 2, pp. 377–400, 2018.
- [6] Ciena Corporation, “Coherent comes to the data center—meet ciena’s 1.6 tb/s coherent-lite pluggable,” <https://www.ciena.com/insights/blog/2024/coherent-comes-to-the-data-center-meet-cienas-1-6-tbs-coherent-lite-pluggable>, 2024, accessed: 2025-11-06.
- [7] J. K. Perin, A. Shastri, and J. M. Kahn, “Design of low-power DSP-free coherent receivers for data center links,” *Journal of Lightwave Technology*, vol. 35, no. 21, pp. 4650–4662, 2017.
- [8] D. Li, G. Minoia, M. Repossi, D. Baldi, E. Temporiti, A. Mazzanti, and F. Svelto, “A low-noise design technique for high-speed cmos optical receivers,” *IEEE Journal of Solid-State Circuits*, vol. 49, no. 6, pp. 1437–1447, 2014.
- [9] S. Shekhar, J. S. Walling, and D. J. Allstot, “Bandwidth extension techniques for CMOS amplifiers,” *IEEE Journal of Solid-State Circuits*, vol. 41, no. 11, pp. 2424–2439, 2006.
- [10] S. Palermo, S. Hoyos, A. Shafik, E. Z. Tabasy, S. Cai, S. Kiran, and K. Lee, “Cmos adc-based receivers for high-speed electrical and optical links,” *IEEE Communications Magazine*, vol. 54, no. 10, pp. 168–175, 2016.
- [11] B. Radi, D. Abdelrahman, O. Liboiron-Ladouceur, G. Cowan, and T. C. Carusone, “Optimal optical receivers in nanoscale CMOS: A tutorial,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 69, no. 6, pp. 2604–2609, 2022.
- [12] I. A. Young, E. Mohammed, J. T. S. Liao, A. M. Kern, S. Palermo, B. A. Block, M. R. Reshotko, and P. L. D. Chang, “Optical I/O technology for tera-scale computing,” *IEEE Journal of Solid-State Circuits*, vol. 45, no. 1, pp. 235–248, 2010.
- [13] G. Belfiore, R. Henker, and F. Ellinger, “40 gbit/s 4-pam vcsel driver circuit with power roll-off compensation,” in *2018 IEEE 61st International Midwest Symposium on Circuits and Systems (MWSCAS)*, 2018, pp. 747–750.
- [14] D. Kucharski, Y. Kwark, D. Kuchta, D. Guckenberger, K. Kornegay, M. Tan, C.-K. Lin, and A. Tandon, “A 20 gb/s vcsel driver with pre-emphasis and regulated output impedance in 0.13 μm cmos,” in *ISSCC. 2005 IEEE International Digest of Technical Papers. Solid-State Circuits Conference, 2005.*, 2005, pp. 222–594 Vol. 1.
- [15] C. Zhang, F. Chen, L. Wang, L. Wang, and C. P. Yue, “Recent advances of high-speed short-reach optical interconnects for data centers,” *IEEE Open Journal of the Solid-State Circuits Society*, vol. 5, pp. 86–100, 2025.

- [16] T. Dickson, Z. Deniz, M. Cochet, J. Bulzacchelli, M. Kessel, P. A. Francese, T. Morf, J. Proesel, H. Ainspan, M. Braendli, and M. Meghelli, "Digital-to-analog converters for 100+ gb/s wireline transmitters: Architectures, circuits, and calibration," in *2024 IEEE Custom Integrated Circuits Conference (CICC)*, 2024, pp. 1–8.
- [17] B. Razavi, "The current-steering dac [a circuit for all seasons]," *IEEE Solid-State Circuits Magazine*, vol. 10, no. 1, pp. 11–15, 2018.
- [18] C.-H. Lin, F. M. L. van der Goes, J. R. Westra, J. Mulder, Y. Lin, E. Arslan, E. Ayranci, X. Liu, and K. Bult, "A 12 bit 2.9 gs/s dac with im3 –60 dbc beyond 1 ghz in 65 nm cmos," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 12, pp. 3285–3293, 2009.
- [19] Marvell Technology, Inc., "Marvell introduces 1.6 tb/s LPO chipset to enable optical short-reach, scale-up compute fabric interconnects," <https://www.marvell.com/company/newsroom/marvell-introduces-1-6-tbps-lpo-chipset.html>, Dec. 2024, accessed: 2025-11-06.
- [20] G. de Valicourt, P. Pupalaikis, R. Giles, M. Lamponi, L. Elsinger, S. Liu, B. Sawyer, J. Proesel, S. T. Le, E. Ho, K. Liu, Z. Zhu, S. Corteselli, L. Alloine, C. Daunt, M. Ferriss, B. Rahmani, F. Warning, A. Bruno, S. Abbaslou, M. Zaman, Z. Pan, G. Fischer, G. Reichert, J. George, U. Alakusu, J. Binkley, I. Martinez, K. van Acoleyen, F. Fesharaki, S. Paul, V. Karra, T. Nguyen, N. Machineni, A. Sullivan, D. Assumpcao, and P. Winzer, "1.6-tb/s low-power linear-drive high-density optical interface for machine learning/AI," *Optics Express*, vol. 33, no. 7, pp. 15 338–15 354, 2025. [Online]. Available: <https://opg.optica.org/oe/abstract.cfm?URI=oe-33-7-15338>
- [21] A. Khairi, A. Laufer, I. Radashkevich, Y. Krupnik, J. Kim, T. Warshavsky Grafi, A. Balankutty, Y. Sabag, Y. Segal, U. Viobnik, M. P. Li, I. Levin, Y. Ben Ezra, and A. Cohen, "Beyond 200-gb/s PAM4 ADC and DAC-based transceiver for wireline and linear optics applications," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 4, pp. 265–276, 2024.
- [22] A. Khairi, Y. Krupnik, A. Laufer, Y. Segal, M. Cusmai, I. Levin, A. Gordon, Y. Sabag, V. Rahinski, I. Lotan, G. Ori, N. Familia, S. Litski, T. W. Grafi, U. Viobnik, D. Lazar, Y. Horwitz, A. Balankutty, S. Kiran, S. Palermo, P. M. Li, F. O'Mahony, and A. Cohen, "A 1.41-pj/b 224-gb/s pam4 6-bit adc-based serdes receiver with hybrid afe capable of supporting long reach channels," *IEEE Journal of Solid-State Circuits*, vol. 58, no. 1, pp. 8–18, 2023.
- [23] J. Kim, S. Kundu, A. Balankutty, M. Beach, B. C. Kim, S. T. Kim, Y. Liu, S. K. Murthy, P. Wali, K. Yu, H. S. Kim, C.-C. Liu, D. Shin, A. Cohen, Y. Segal, Y. Fan, M. P. Li, and F. O'Mahony, "A 224-gb/s DAC-based PAM-4 quarter-rate transmitter with 8-tap FFE in 10-nm FinFET," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 1, pp. 6–20, 2022.
- [24] Broadcom Inc., "Co-packaged optics (CPO)—highest power efficiency and bandwidth density," <https://www.broadcom.com/info/optics/cpo>, 2025, accessed: 2025-11-06.
- [25] Ayar Labs, Inc., "Enabling next-gen AI infrastructure with co-packaged optics and advanced packaging," <https://ayarlabs.com/resources/enabling-next-gen-ai-infrastructure-with-co-packaged-optics-and-advanced-packaging/>, Sep. 2025, presentation at 2025 TSMC OIP Ecosystem Forum. Accessed: 2025-11-06.

- [26] NVIDIA Corporation, “NVIDIA announces spectrum-X photonics, co-packaged optics networking switches to scale AI factories to millions of GPUs,” <https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-Announces-Spectrum-X-Photonics-Co-Packaged-Optics-Networking-Switches-to-Scale-AI-Factories-to-Millions-of-GPUs/default.aspx>, Mar. 2025, accessed: 2025-11-06.