Historical Trends in Wireline Communications

Behzad Razavi

60× improvement in speed in 20 years

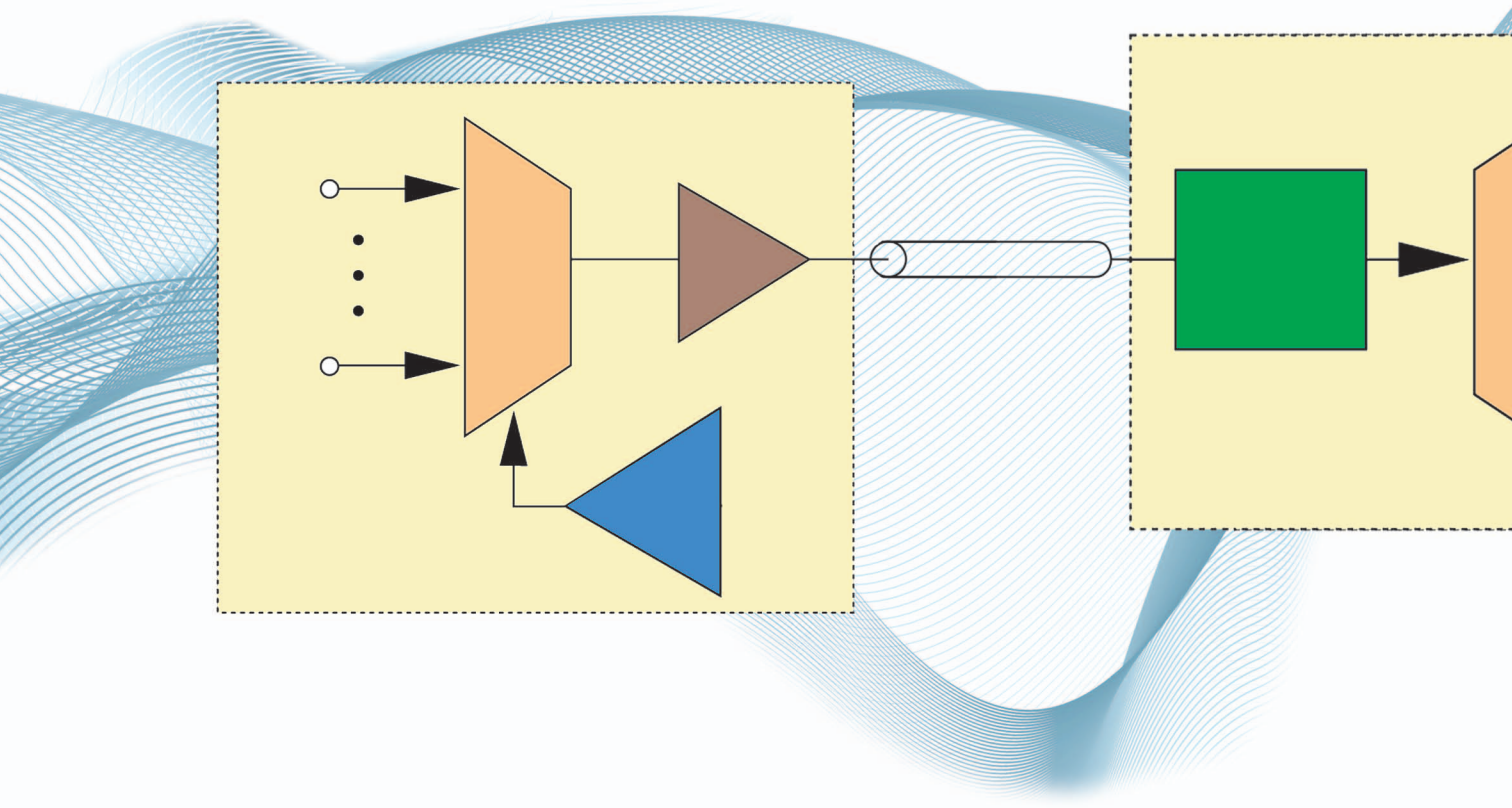

Communication over wires began decades before over the air—in fact, in binary form. The device named "the telegraph" by Claude Chappe was commercialized in the 1830s and served our civilization for more than one and a half centuries. Introduced by Alexander Graham Bell in the 1870s, the telephone is the first analog wireline communication system and also a long-lasting invention.

The notion of sending *data* over wires goes back to a system proposed by Shelford Bidwell in the 1880s, which he called the scanning photo telegraph, and it became the predecessor of the fax machine. In the 1940s, IBM used a modem to transmit data over phone wires, and in the 1960s, the fax machine was introduced by Xerox.

It was not until the 1970s that data rates higher than what modems could provide became commercially attractive. This decade would also witness a bifurcation in the transmission medium: the phone line or dedicated coax cables. In 1973, Robert Metcalfe and David Boggs proposed the idea of the Ethernet as a standard for connecting computers over the latter medium. (The coax cable would eventually be replaced with a type of twisted-pair line called the CAT-5 cable.) The digital subscriber line (DSL) patent was filed by Teltone Corporation in 1979 and the asymmetric

Digital Object Identifier 10.1109/MSSC.2015.2477016
Date of publication: 2 December 2015

1943-0582/15©2015IEEE

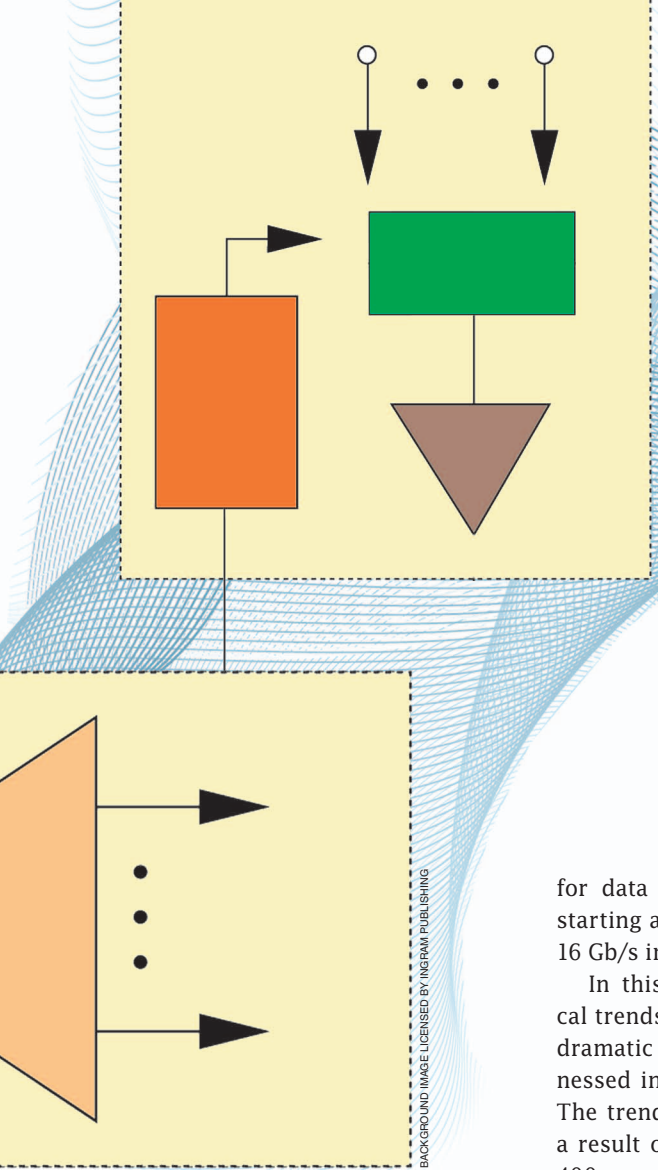

BACKGROUND IMAGE LICENSED BY INGRAM PUBLISHING

The need for greater speeds in wireline communications has arisen from both the increasing number of users and the higher data rates that they demand.

DSL patent by Bellcore in 1968, both for use over phone lines.

In 1988, the AT&T StarLAN system achieved a data rate of 10 Mb/s over the twisted pair, which paved the way for the ratification by IEEE of the 10Base-T Ethernet standard in 1990. The rising demand for higher speeds led to the standardization of 100Base-T Ethernet in 1995, gigabit Ethernet in 1999, and 10GBase-T in 2006. In parallel with the Ethernet development, two other standards emerged as well: 1) the fibre channel (originally conceived for optical fibers and later applied to copper media) for data storage networks, starting with a speed of 132 Mb/s and now reaching 32 Gb/s, and 2) the peripheral component interconnect (PCI) for data buses within a computer, starting at 33 Mb/s and approaching 16 Gb/s in PCI-Express.

In this article, we study historical trends, inflection points, and the dramatic progress that we have witnessed in wireline communications. The trend charts produced here are a result of surveying approximately 400 papers published in the past three decades. We consider mainly bipolar/BiCMOS and CMOS implementations and refer to the former as "bipolar" for the sake of brevity. We also focus on communication in copper media.

Generic Wireline System

To appreciate the significance of the developments in this field, it is helpful to first consider a basic wireline communication system. Shown in Figure 1, such a system consists of a transmitter (TX), a propagation medium (the channel), and a receiver (RX). The TX employs a multiplexer (MUX) to serialize parallel data streams $D_1,\ldots,D_n$ and a driver capable of delivering sufficient voltage swings to the channel. A phase-locked loop (PLL) generates the clock frequencies necessary for the MUX. (The MUX output is preferably retimed at full rate by a flip-flop.) The RX incorporates an equalizer to compensate for the channel's imperfections, a clock and data recovery (CDR) circuit to recover the clock and retime the received data, and a demultiplexer (DMUX). (In modern systems, the TX also shoulders part of the equalization task.)

The system of Figure 1 faces design tradeoffs among three parameters: the loss and impedance discontinuities of the channel, the maximum

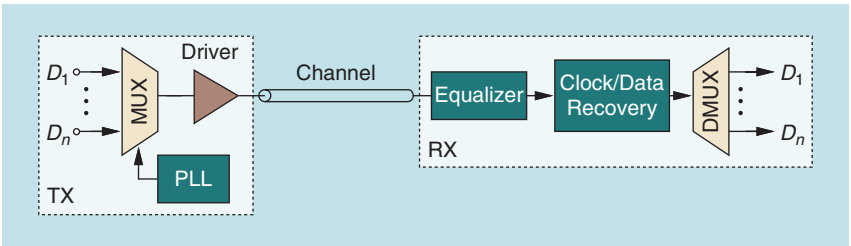

FIGURE 1: A generic wireline communication system.

data rate that can be communicated, and the TX and RX power consumption. For example, we wish to transport data at a rate of 40 Gb/s through a channel that has 25 dB of loss at 20 GHz while the transceiver (TRX) must consume less than 1 mW/Gb/s (40 mW).

This brief perspective reveals that today's high-performance, high-speed systems have grown from extensive work on a multitude of building blocks: MUXs and DMUXs, PLLs and frequency dividers, equalizers, and CDR circuits. We therefore study the history with the aid of the evolution of some of these building blocks.

It is important to note that the voltage swing delivered by the TX driver in Figure 1 to the copper channel has not scaled much with technology due to signal-to-noise ratio limitations at the receive end. For a given channel characteristic impedance, e.g., 50 Ω, this means that the driver draws a weakly scalable power. This issue becomes more severe at rates higher than a few gigabits per second as the driver must incorporate back-termination, on-chip resistors to suppress secondary reflections. Figure 2 depicts an example where the R_T resistors reduce reflections returning from the receive end due to the mismatch between R_{RX} and R_0. In this case, the driver consumes a power of

$$P_{dr} = \frac{V_{swing}}{R_0 \| R_T} V_{DD}, \tag{1}$$

where V_{swing} denotes the peak differential voltage swing delivered to the channel. If the driver is modified to generate rail-to-rail swings, (1) can be reduced to $V_{swing}^2/(R_0 \| R_T)$.

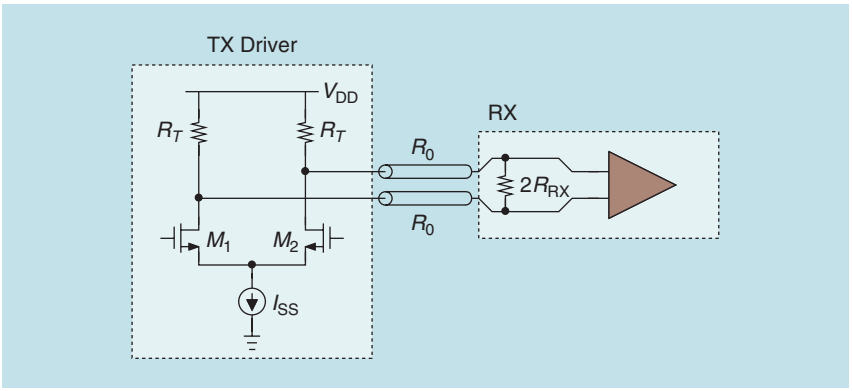

FIGURE 2: The problem of output driver power dissipation.

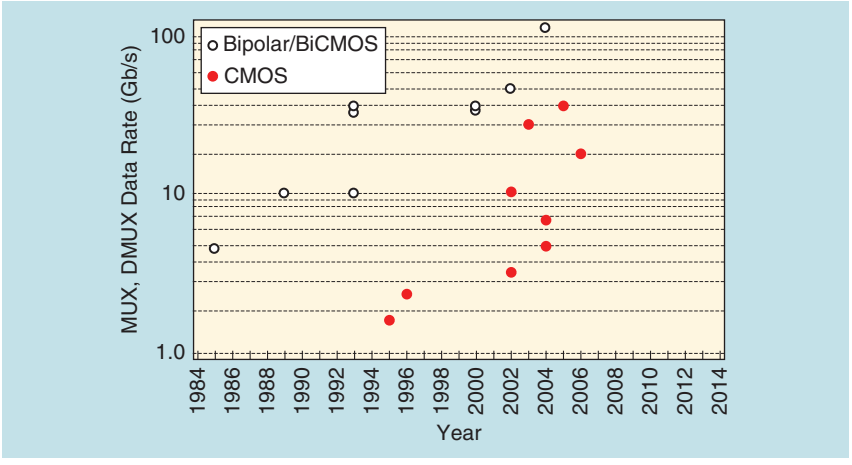

FIGURE 3: The speed of MUX and DMUX circuits versus time.

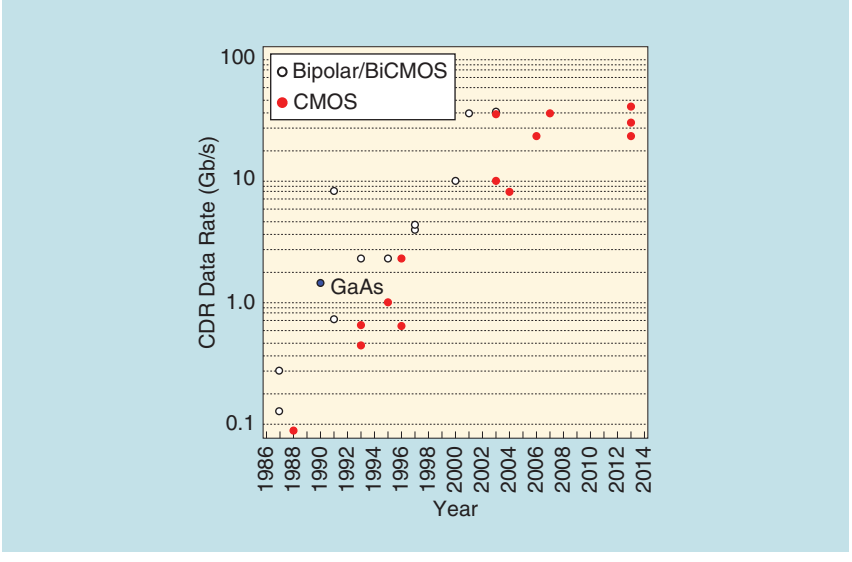

FIGURE 4: The speed of CDR circuits versus time.

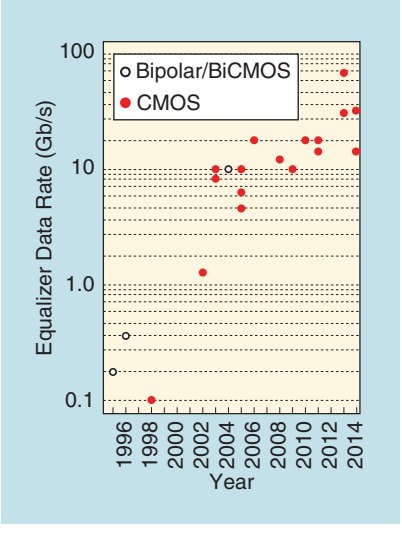

FIGURE 5: The speed of equalizers versus time.

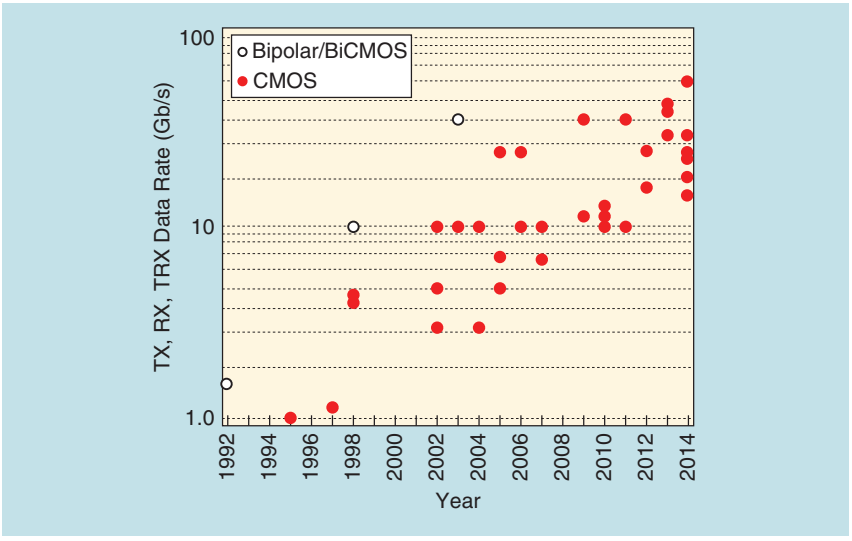

FIGURE 6: TX, RX, and TRX speeds versus time.

The Need for Speed

The need for greater speeds in wireline communications has arisen from both the increasing number of users and the higher data rates that they demand. This trend manifested itself in optical communications in the 1980s and 1990s and continued in copper media thereafter. In the mid-2000s, it was observed that the number of input/output (I/O) pins on large digital chips was increasing rapidly, leading to a greater silicon area and, more importantly, raising the power consumption; if multiplied by hundreds, the value given by (1) becomes prohibitive. The interest in "serialization" thus grew considerably.

Figure 3 plots the speed of MUX and DMUX circuits since the mid-1980s for bipolar and CMOS implementations. The data points represent a MUX or DMUX factor of at least four (rather than two) to account for flip-flop speed limitations. Reflecting the general trend in the semiconductor industry, the plot shows the early dominance of bipolar realizations and the growth of CMOS counterparts from the mid-1990s. The maximum speed is 108 Gb/s. We also note that stand-alone MUX/DMUX circuits are reported primarily up to 2006, indicating the emphasis on TRX integration thereafter.

Depicted in Figure 4 is the speed of CDR circuits as a function of time. Both bipolar and CMOS technologies had humble beginnings, starting at well below 1 Gb/s. Bipolar CDRs continued to rise in speed until the early 2000s, when optical communication fell out of favor and the market focus turned to backplanes and CMOS technology. It is also important to recognize that the notion of multiphase sampling (e.g., half-rate or quarter-rate operation) was proposed in the early 1990s, perhaps for the first time by [1].

Unlike MUX/DMUX and CDR circuits, equalizers were not an integral part of earlier digital links. (Of course, analog equalizers had been used in phone lines decades before.) It was the advent of copper media that imposed the problem of channel imperfections and hence the need for equalization. As plotted in Figure 5, equalizers began to appear in the late 1990s and quickly migrated to CMOS technology. Early designs were in the form of what we today call continuous-time linear equalizers (CTLEs), thus dealing with only the loss and not the impedance discontinuities. Although conceived in the late 1960s [2] and realized in digital

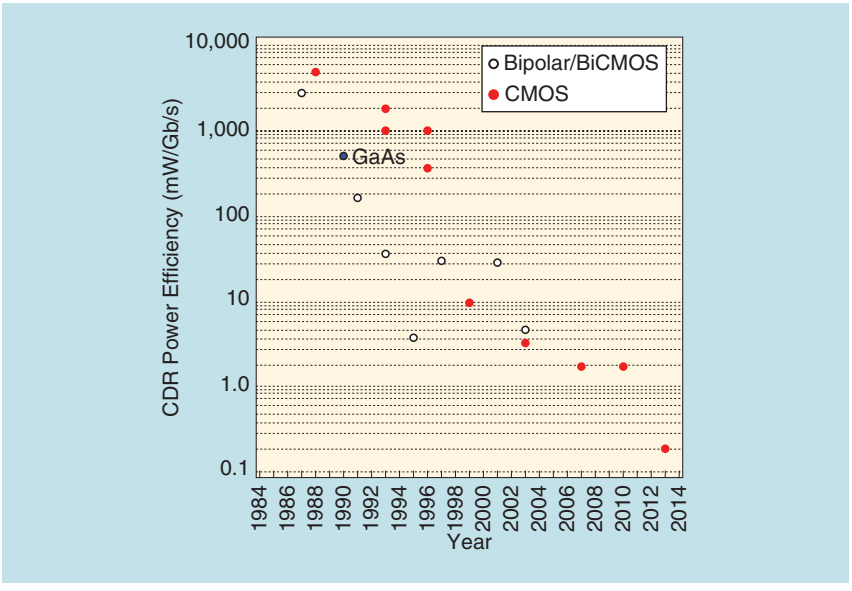

FIGURE 7: CDR PE versus time.

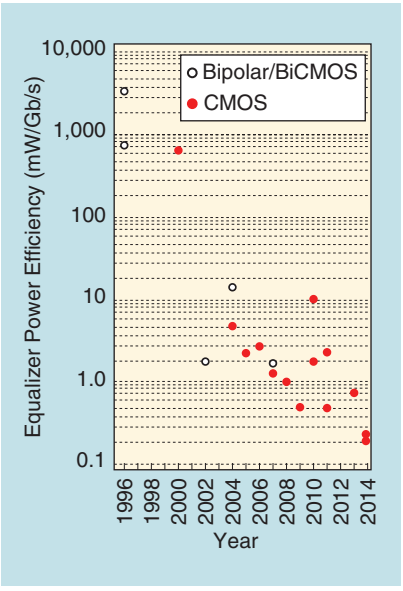

FIGURE 8: Equalizer PE versus time.

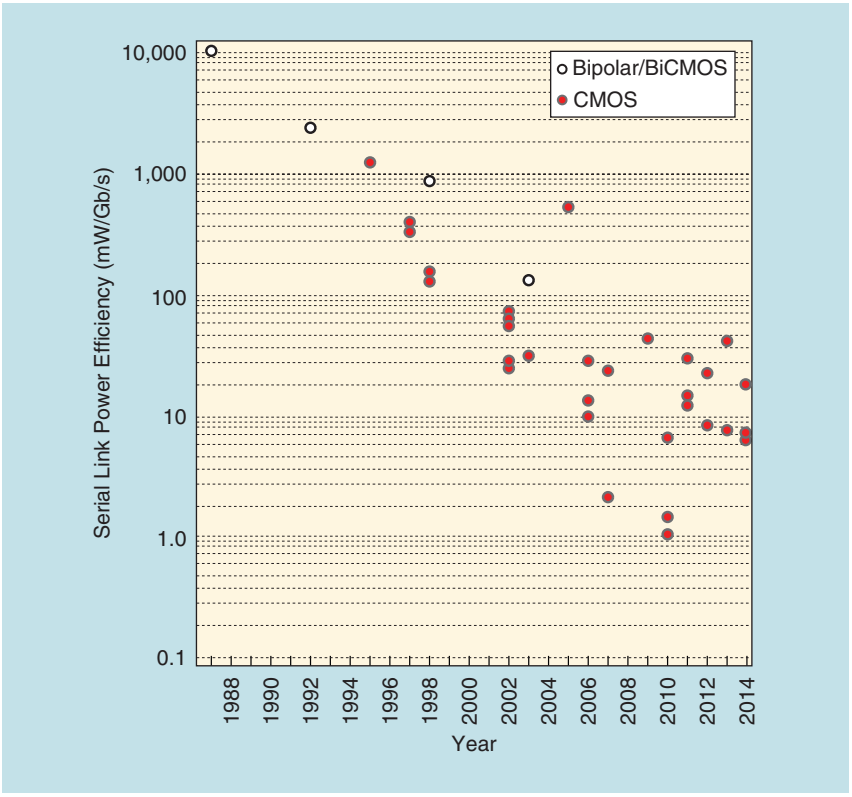

FIGURE 9: Serial link PE versus time.

signal processors later, decision-feedback equalizers (DFEs) did not appear in the analog domain until the 2000s. This is because feedback timing constraints limited the speed of DFEs but not the CTLEs.

The term "serial link" was used, perhaps for the first time, in 1991 to refer to a 1.5-Gb/s TRX [3]. Figure 6 plots TX, RX, or TRX data rates as a function of time. Except for a few bipolar designs, the space has been governed by CMOS technology as integration proves essential here. It is remarkable that the speed in CMOS has climbed by a factor of 60 in 20 years. If this trend continues, the serial data rate will exceed 400 Gb/s by 2025!

The problem of channel loss at high speeds has motivated work on nonbinary signaling schemes that consume less bandwidth for a given data rate. For example, PAM-4, previously investigated in the late 1990s [4], has reemerged as a plausible candidate [5], [6]. Affording a twofold bandwidth reduction, PAM-4 nonetheless requires more complex circuit design and a higher amplitude resolution in the receiver.

Power Efficiency

With a large number of I/Os, the power consumed by each link between chips becomes critical. In this section, we examine the historical trends in the power efficiency (PE) of wireline TRXs and some of their building blocks.

Figure 7 plots the PE of CDR circuits as a function of time. Since the CDR speeds in the late 1980s were quite low (Figure 4), their normalized performance to 1 Gb/s translates to several watts of power. Bipolar implementations continued to improve their PE until early 2000s, after which CMOS took over. We observe that the CDR normalized power dissipation has fallen by a factor of 25,000 over 25 years, i.e., by about 50% per year. If this trend continues, CDR circuits will consume about 4 μW/Gb/s in 2025.

The PE of equalizers is depicted in Figure 8. As with CDR circuits, equalizers have seen a dramatic reduction in their normalized power, by about a factor of 3,000 over 15 years. This 70% per year improvement would yield an equalizer PE of 1 μW/Gb/s in 2025.

Of greatest interest to us is the PE of the overall serial ink. As shown in Figure 9, the normalized power of CMOS designs has dropped from 1.2 W/Gb/s to about 6 mW/Gb/s, with a few exceptions around 1–2 mW/Gb/s. This factor of 200 over 20 years is less impressive than those observed for CDR and equalizer circuits and can be attributed to three issues: the weakly scalable power consumed by the TX output drivers, various functions that have been added to TRXs, and, most notably, the push for higher speeds and hence the need to operate with nonlinear power-speed tradeoffs. Nevertheless, if sustained, this 30% per year reduction would lead to a link PE of about 0.4 mW/Gb/s by 2025.

References

[1] T. H. Hu and P. R. Gray, "A monolithic 480 Mb/s parallel AGC/decision/clock-recovery circuit in 1.2-um CMOS" *IEEE J. Solid-State Circuits*, vol;. 28, pp. 1314–1320, Dec. 1993.

[2] M. E. Austin, "Decision-feedback equalization for digital communication over dispersive channels," Tech. Rep. 347, Lincoln Lab., Aug. 1967.

[3] R. C. Walker, T. Hornak, C. S. Yun, and J. Doernberg, "A 1.5 Gb/s link interface chipset for computer data transmission," *IEEE J. Selected Areas Commun.*, vol. 9, pp. 698–703, June 1991.

[4] R. Farjad-Rad and M. Horowitz, "An equalization scheme for 10Gb/s 4-PAM signaling over long cables," in *Proc. Mixed Signal Conf.* 1997.

[5] A. Nazemi, K. Hu, B. Catli, D. Cui, U. Singh, T. He, A. Momtaz, and J. Cao, "A 36 Gb/s PAM4 transmitter using an 8b 18 GS/s DAC in 28 nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 58–59.

[6] J. Kim, A. Balankutty, A. Elshazly, Y. Huang, H. Song, K. Yu, and F. O. Mahony, "A 16-to-40 Gb/s quarter-rate NRZ/PAM4 dual-mode transmitter in 14nm CMOS," in *ISSCC Dig. Tech. Papers*, Feb. 2015, pp. 60–61.

About the Author

Behzad Razavi (razavi@ee.ucla.edu) is a professor of electrical engineering at the University of California, Los Angeles, where he conducts research on analog and high-speed circuits. He has received seven awards at ISSCC, CICC, ESCCIRC, and VLSI Circuits Symposium for his research and four awards for his teaching. His books have been published in seven languages. He received the 2012 Donald Pederson Award in Solid-State Circuits.

SSC