Near Speed-of-Light Signaling Over On-Chip Electrical Interconnects

Richard T. Chang, *Student Member, IEEE*, Niranjan Talwalkar, *Student Member, IEEE*, C. Patrick Yue, *Member, IEEE*, and S. Simon Wong, *Fellow, IEEE*

***Abstract*—The propagation limits of electrical signals for systems built with conventional silicon processing are explored. A design which takes advantage of the inductance-dominated high-frequency regime of on-chip interconnect is shown capable of transmitting data at velocities near the speed of light. In a 0.18-μm six-level aluminum CMOS technology, an overall delay of 283 ps for a 20-mm-long line, corresponding to a propagation velocity of one half the speed of light in silicon dioxide, has been demonstrated. This approach offers a five times improvement in delay over a conventional repeater-insertion strategy.**

***Index Terms*—Integrated circuit interconnections, microstrip, monolithic integrated circuits, on-chip interconnect, phase modulation, transmission lines.**

I. Introduction

INTERCONNECT has long been perceived to be a bottleneck in present and future high-performance digital integrated circuits because of its inability to keep pace with advances in transistor speeds [1]. Consequently, alternative solutions such as on-chip optical interconnects have been proposed in order to avoid the problems associated with global on-chip wires altogether [2]. However, due to technology incompatibility, cost considerations, and nonnegligible delays in converting the signals between optical and electrical domains, this has not yet been shown to be practical. In this brief, we explore the velocity limitations for systems built with conventional silicon processing and show that data transmission at near the speed of light is possible in an all-electrical system.

II. Limitations of Conventional Repeater-Insertion Approach

The most common means of global communication is through the use of wires with appropriately spaced repeaters [3]. Based on a 0.18-μm 1.8-V CMOS six-level aluminum interconnect technology, simulations were carried out in HSPICE for minimum-width (0.44-μm), minimum-spacing (0.46-μm) lines with varying number of repeaters for a fixed overall distance (Fig. 1) [4]. The optimized minimum delay, including both wire and buffer delays, is about 1.35 ns for propagation over a length of 20 mm, translating into an effective velocity of about one tenth of the speed of light in silicon dioxide $c_{\rm ox}$. This optimized delay configuration yields a power dissipation of 30 mW, assuming 1-GHz switching, with an average transition time of 200 ps.

Manuscript received July 27, 2002; revised December 23, 2002. This work was supported by the MARCO Interconnect Focus Center and the Taiwan Semiconductor Manufacturing Corporation (TSMC).

The authors are with the Center for Integrated Systems, Stanford University, Stanford, CA 94305-4070 USA (e-mail: rtchang@stanfordalumni.org).

Digital Object Identifier 10.1109/JSSC.2003.810060

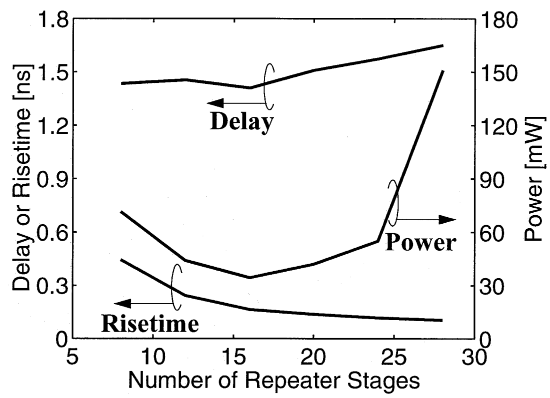

Fig. 1. Simulated performance of minimum-sized interconnect (0.44-μm wide, 0.46-μm spacing for aluminum/silicon dioxide interconnects) with repeaters.

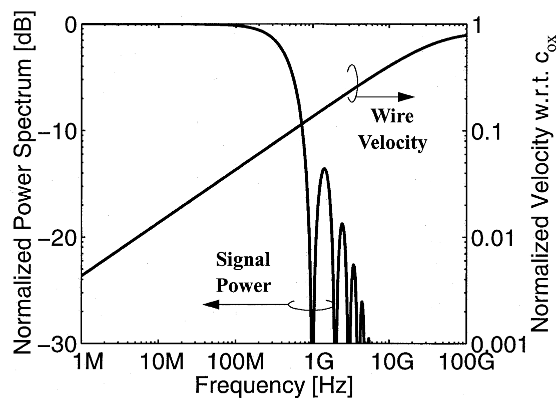

Fig. 2. Frequency characteristics of 500-ps-wide digital pulse and typical minimum-sized global wires. $c_{\rm ox}$ is the speed of light in silicon dioxide.

While the buffer delays can be improved by using more advanced technology and/or expending more power, the intrinsic wire delay is fundamentally limited. To understand the reasons for this, the power spectral density of a 500-ps digital pulse is compared with the frequency-dependent characteristics of the aforementioned minimum-sized wire without repeaters (Fig. 2). The digital signal is broadband in nature, while the phase velocity for signals propagating along the wire changes dramatically over this frequency range. For transmission line systems, the phase velocity, rather than the group velocity, is meaningful to consider [5]. The telegrapher's equation describes the signal propagation as a function of time:

$$\frac{\partial^2 V}{\partial x^2} = RC\frac{\partial V}{\partial t} + LC\frac{\partial^2 V}{\partial t^2} \quad (1)$$

where x is the distance along the wire, t is time, V is voltage, and R, L, and C are the resistance, inductance, and capacitance

TABLE I
DIFFERENT FREQUENCY REGIMES OF PROPAGATION

	Low Frequency regime (R>>ωL)	High Frequency regime (R<<ωL)
Propagation Equation	$\frac{\partial^2 V}{\partial x^2} = RC\frac{\partial V}{\partial t}$	$\frac{\partial^2 V}{\partial x^2} = LC\frac{\partial^2 V}{\partial t^2}$
Phase Velocity	$v = \sqrt{\frac{\omega}{2RC}}$	$v = \frac{1}{\sqrt{LC}}$

per unit length, respectively. This intuitively suggests that at lower frequencies, the first term on the right side of (1) dominates, and the wire behaves as a distributed RC network. In this RC regime, signals travel slowly by diffusion and undergo frequency dispersion. As the frequency increases, the inductive component of the wire [second term on the right side of (1)] begins to dominate over the resistance, and the wire behaves more as a waveguide. The high-frequency LC regime allows for propagation of an electromagnetic wave; consequently, the peak phase velocity is the speed of light in the dielectric surrounding the interconnect (Table I). This is the key characteristic of the wire that can be exploited in order to achieve high-speed signal propagation. Fig. 3 shows the experimental velocity and attenuation characteristics of an on-chip 1-mm-long coplanar waveguide, with 4.8-μm-wide signal, 20-μm ground wires, and 5-μm spacing, which is obtained through S-parameter measurements [6].

In the example given in Fig. 2, most of the spectral components of the 500-ps pulse are in the RC regime of the minimum-sized wire, traveling at speeds much less than the speed of light. It is not until over 100 GHz that the velocity reaches its maximum and becomes independent of frequency. The effect on digital signals is that different frequency components travel at different speeds and attenuate at different levels, resulting in a significantly distorted output. Repeater insertion does not change the frequency characteristics of the interconnect, but can improve speed by amplifying the faster high-frequency components of the signal, when placed at appropriate intervals. This comes at the cost of increased power dissipation and a delay overhead from the intrinsic delay of the repeater.

III. EXPLOITING THE LC NATURE OF WIRES

The characteristics shown in Fig. 3 suggest that a high-speed system can be built by taking advantage of the wave nature of interconnect. At the same time, it is beneficial to eliminate the low-frequency portion of the signal that lags behind and contributes to inter-symbol interference (ISI). This can be achieved by modulating the digital data with a sufficiently high-frequency carrier, and as a result, concentrating all of the signal power in the faster LC regime.

In order for this system to be realizable, wires must have a low crossover frequency between the RC and LC regimes. By explicitly emphasizing the parasitic inductance and reducing the resistance, this transition can be shifted into the single-gigahertz range. This design strategy of deliberately using the parasitic inductance may be foreign to designers, since the inductance for on-chip wires is often difficult to estimate. In this case, the on-chip transmission lines used have explicit ground returns that provide well-controlled and predictable inductance values.

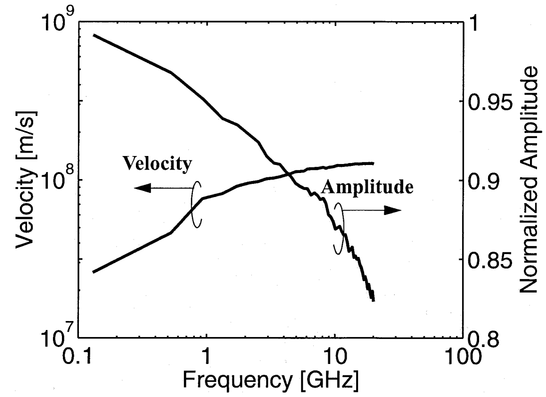

Fig. 3. Experimental results showing velocity and attenuation as a function of frequency for a 1-mm-long coplanar waveguide.

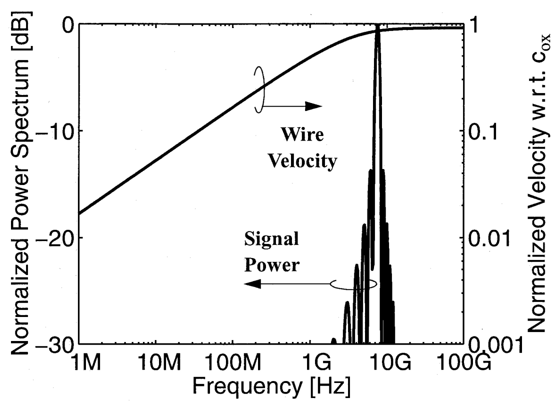

Fig. 4. Frequency characteristics of modulated pulse and low-loss on-chip interconnect. c_{ox} is the speed of light in silicon dioxide.

While the high-frequency LC regime offers high-speed frequency-dispersionless propagation, the interconnect is substantially more lossy at higher frequencies, as demonstrated by Fig. 3. Therefore, the wire must be optimized to minimize loss, while occupying a reasonable amount of area. Thicker top layers of metal and dielectric facilitate the realization of on-chip transmission lines [7], [8]. For example, a microstrip structure with a 6-μm-wide 2-μm-thick copper signal wire on a 2-μm-thick dielectric provides a loss of about 0.5 dB/mm [6]. This totals 10–15-dB loss for a cross-chip global wire. This amount of attenuation is tolerable for a system to reliably recover the signal. In this case, no gain elements, such as repeaters or amplifiers, are present or necessary along the length of interconnect.

Fig. 4 shows the impact of using modulated signaling in combination with the use of low-loss interconnect. Using optimized interconnect lowers the crossover frequency between the RC and LC regimes to a few gigahertz. Using a high-frequency carrier pushes the signal spectral components to lie predominately in the high-speed inductance-dominated region.

IV. SYSTEM IMPLEMENTATION

A simple implementation of this system uses direct conversion from baseband to RF, as illustrated in Fig. 5. The carrier

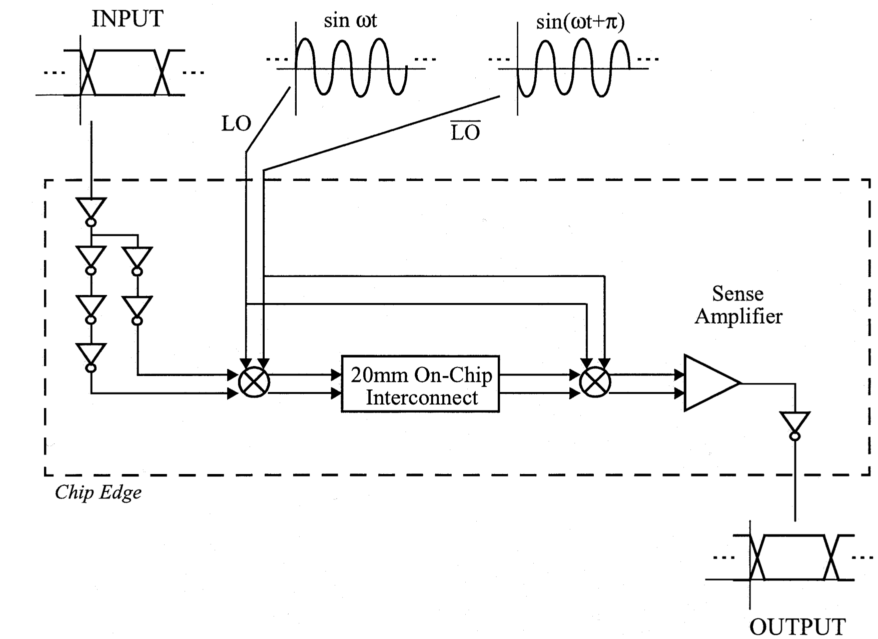

Fig. 5. System block diagram.

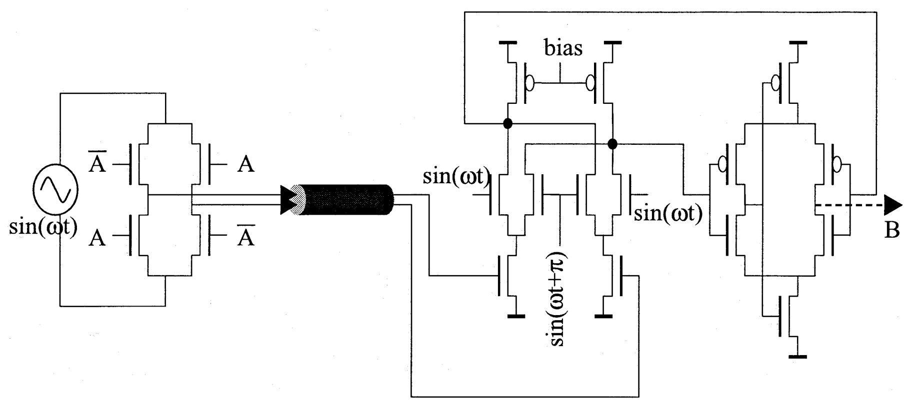

Fig. 6. Simplified circuit schematic of system. A and $\overline{A}$ represent the differential digital data to be transmitted, and $\sin(\omega t)$ is the high frequency carrier. B is the recovered digital signal.

frequency must be in the LC regime and sufficiently above the desired bandwidth of the signal. Because the overall latency of the system not only includes the time of flight across the interconnect, but also includes the transmitter and receiver delays, minimizing the delay of these components is an important criterion in designing the circuits so that the speed gains are not negated in the overall system.

As a demonstration for transmission of 1-GHz data across chip, a 7.5-GHz local oscillator (LO) carrier has been chosen. A passive ring mixer is used to perform this upconversion while minimizing power consumption [9]. This ring mixer then drives an interconnect wire that extends uninterrupted across a chip-edge length of tens of millimeters. This modulation system requires no repeaters, has low power consumption, and simplifies floorplanning since the interconnect does not need a dedicated channel of silicon to accommodate any repeaters. The transmission line is a differential microstrip topology, each with a 1-μm-thick aluminum 16-μm-wide signal line over 2.1-μm intervening SiO_2 layer. This microstrip topology confines the electromagnetic fields to minimize crosstalk. Additionally, this microstrip configuration suffers from less current crowding because the current flows across the entire width of the wire, even at high frequencies. On the receiver end, a double-balanced mixer downconverts the signal back to baseband with a voltage gain of 2. A sense amplifier follows to restore the signal to digital logic levels [10]. A simplified circuit schematic is shown in Fig. 6. Additional buffers, as depicted in Fig. 5, are used exclusively for testing and measurement purposes.

V. Experimental Results

To demonstrate the feasibility of across-chip speed-of-light data transmission, the system has been fabricated in a TSMC 0.18-μm standard logic CMOS technology with six levels of aluminum/silicon dioxide interconnect. The differential LO carrier signals and digital input pulse are generated off-chip, while the transmitter, receiver, and all other components shown in

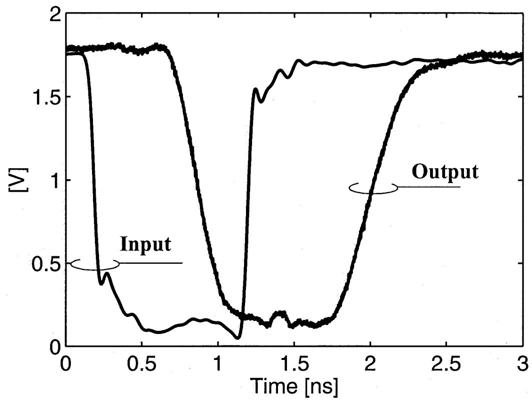

Fig. 7. Sample as-measured input and output waveforms for system described in Fig. 6.

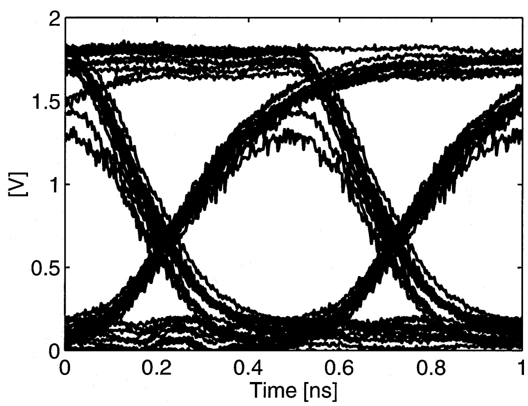

Fig. 8. Eye diagram of system output operating at 1 GHz.

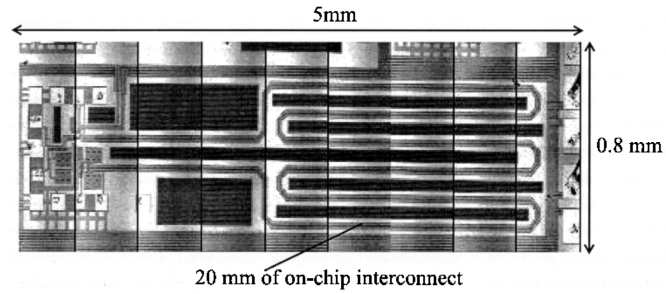

Fig. 9. Die photo of test chip.

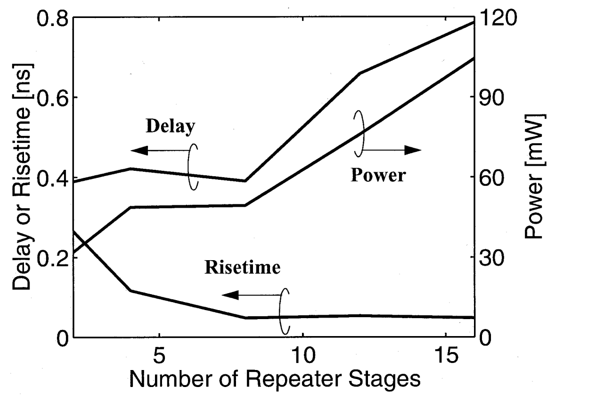

Fig. 10. Simulated performance of wide-sized interconnect (16-μm-wide microstrip in aluminum/silicon dioxide) with repeaters.

Fig. 5 are integrated on-chip. Fig. 7 shows sample as-measured input and output waveforms, and Fig. 8 shows a sample eye diagram of the output operating at 1 GHz. The overall delay from the test chip input to the output is 322 ps for the falling edge and 490 ps for the rising edge. This includes a delay of 123 ps due to the inverters needed for driving signals on and off of the chip for testing and measurement. Subtracting these inverter delays yields an average delay of 283 ps with a power consumption of 16.1 mW (including the power delivered by the external LO, but excluding the testing inverters). This total delay over the length of 20 mm corresponds to an effective signal propagation speed of nearly one half of c_{ox}. Fig. 9 shows a die photo of the 0.8-mm×5-mm test system. The interconnect is laid out in a serpentine manner around dummy fill metal.

VI. Scalability of the System

One of the main features of this design is the low-loss transmission line that uses wide wires to reduce the loss seen at high frequencies. Alternatively, the wire width may be reduced at the expense of increased power dissipation. As the wire width is reduced, the resistance increases, but so does the inductance, at a slightly slower rate. This indicates the crossover frequency between the RC and LC regimes is relatively insensitive to the wire width, and there is little or no delay penalty to scaling down the wire. The higher inductance translates into a higher impedance, which reduces the amount of power required to drive the interconnect. On the other hand, the loss along the length of the interconnect, which, at high frequencies, is approximately proportional to the ratio of the resistance to the characteristic impedance, has increased. Therefore, these effects compete to determine the required power. Simulations show that reducing the wire width to 4 μm requires an additional 2 mW of power to produce the same signal level at the receiver input [6].

In addition, this system design will benefit from impending technology enhancements, such as low-K dielectrics, copper, more metal layers, and faster devices.

VII. Comparisons With Other Techniques

Some of the performance improvement in delay for modulated signaling can be attributed to the use of optimized low-loss on-chip transmission lines. Fig. 10 shows that using the same wide microstrip lines in conjunction with repeaters improves the optimal delay from 1.4 ns using minimum-sized wires to 0.4 ns, over a three times improvement. On the other hand, this comes at the cost of higher power dissipation; the power has increased from 30 to 50 mW, assuming 100% activity.

On-chip optics promise speed-of-light transmission without the frequency dispersion and loss associated with electrical wires [2]. For on-chip optical interconnects to be viable, an integrated method of transforming signals from electrical to optical domains and back again is necessary. Detailed simulations have been performed in [11] that quantify the delays due to this transformation overhead. We observe that the modulated signaling approach on electrical wires, in terms of delay, performs just as well or even better than the on-chip

TABLE II
PERFORMANCE OF DIFFERENT APPROACHES TO ON-CHIP SIGNALING

Signaling	Propagation Medium	Time of Flight for 20mm [ps]	Delay for 20 mm [ps]	Power [mW]	Comments
Modulation (This work)	Wide metal wires ($e_r = 4$)	133	300	16	Large metal area
Repeaters	Min. sized metal wires ($e_r = 4$)	133	1400	30	Slow
	Wide metal wires ($e_r = 4$)	133	400	50	Large metal area, High power
Optics (Edge-Emitting) [11]	Air ($e_r = 4$)	66	300	80	Packaging, Integration issues
	On-Chip Waveguide ($e_r = 11.7$)	228	500	80	Integration issues
Optics (VCSEL) [11]	Air ($e_r = 1$)	66	400	60	Packaging, Integration issues
	On-Chip Waveguide ($e_r = 11.7$)	228	600	60	Integration issues

optical links (Table II). Furthermore, the metal waveguide in a silicon dioxide environment has a faster time of flight than an integrated optical polysilicon waveguide.

VIII. SUMMARY

This brief demonstrates, for the first time, the feasibility of all-electrical near speed-of-light across-chip communication by designing modulated signaling systems around optimized low-loss on-chip transmission lines. The measured effective speed is roughly one half the speed of light in oxide, for a 20-mm-long interconnect. The system consumes 16.1 mW of power and supports a bandwidth of 2 Gb/s. Furthermore, the microstrip transmission line size may be reduced in exchange for increased power consumption. When compared to conventional repeater techniques, the modulated signaling approach offers a five times improvement in delay.

REFERENCES

[1] M. Bohr, “Interconnect scaling–The real limiter to high performance ULSI,” in *Int. Electron Devices Meeting Tech. Dig.*, 1995, pp. 241–244.
[2] D. Miller, “Rationale and challenges for optical interconnects to electronic chips,” *Proc. IEEE*, vol. 88, pp. 728–749, June 2000.
[3] H. B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*. Reading, MA: Addison-Wesley, 1990.
[4] R. Chang, C. P. Yue, and S. Wong, “Near speed-of-light on-chip electrical interconnect,” in *Symp. VLSI Circuits Dig. Tech. Papers*, 2002, pp. 18–21.
[5] S. Ramo, J. Whinnery, and T. Van Duzer, *Fields and Waves in Communication Electronics*, 3rd ed. New York: Wiley, 1994, p. 263.
[6] R. Chang, “Near speed-of-light on-chip electrical interconnect,” Ph.D. dissertation, Stanford Univ., Stanford, CA, 2002.
[7] B. Kleveland, T. Lee, and S. Wong, “50-GHz interconnect design in standard silicon technology,” in *IEEE Microwave Theory and Techniques Symp. Tech. Dig.*, 1998, pp. 1913–1916.
[8] A. Deutsch, G. Kopcsay, P. Restle, G. Katopis, W. Becker, H. Smith, P. Coteus, C. Surovic, B. Rubin, R. Dunne, T. Gallo, K. Jenkins, L. Terman, R. Dennard, G. Sai-Halasz, and D. Knebel, “When are transmission-line effects important for on-chip interconnections?,” *IEEE Trans. Microwave Theory Tech.*, vol. 45, pp. 1836–1846, Oct. 1997.
[9] A. Shahini, D. Shaeffer, and T. Lee, “A 12-mW wide dynamic range CMOS front-end for a portable GPS receiver,” *IEEE J. Solid-State Circuits*, vol. 32, pp. 2061–2070, Feb. 1997.
[10] M. Bazes, “Two novel fully complementary self-biased CMOS differential amplifiers,” *IEEE J.Solid-State Circuits*, vol. 26, pp. 165–168, Feb. 1991.
[11] E. Kyriaskis-Bitzaros, N. Haralabidis, M. Lagadas, A. Georgakilas, Y. Moisiadis, and G. Halkias, “Realistic end-to-end simulation of the optoelectronic links and comparison with the electrical interconnections for system-on-chip applications,” *J. Lightwave Technol.*, vol. 19, pp. 1532–1541, Oct. 2001.