

Near Speed-of-Light On-Chip Electrical Interconnect

Richard T. Chang, C. Patrick Yue, S. Simon Wong

Center for Integrated Systems
Stanford University, Stanford, CA 94305-4070 USA

Abstract

The propagation limits of electrical signals for systems built with conventional silicon processing are explored. Data transmission near the speed of light with an all-electrical system can be achieved by taking advantage of the inductance-dominated high-frequency regime of on-chip interconnect. In a 0.18 μm , 6-level Aluminum CMOS technology, an overall delay of 278ps for a 20mm long line corresponding to a propagation velocity of one half the speed of light in silicon dioxide has been demonstrated.

Introduction

Interconnect has long been recognized as a bottleneck to high-performance digital systems for its inability to keep pace with advances in transistor speeds [1]. Solutions such as the use of optical interconnect have been suggested in order to achieve speed of light propagation for across-chip signaling [2]. However, due to technology incompatibility, cost considerations, and non-negligible delays in converting the signals between optical and electrical domains, this has not yet been shown to be practical. In this paper, 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 with reasonable system overhead.

Limitations of 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, 6-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 stages. Figure 1 shows that the optimized minimum delay, including both wire and buffer delays, for such a system is near 1.35 ns for propagation over a length of 20mm, translating into a velocity of about one-tenth of c_{ox} , the speed of light in silicon dioxide. At the same time that delay is minimized, power consumption and rise time, which is inversely proportional to the total bandwidth that the wire can support, must also be minimized. From Fig. 1, we see that there is a trade-off between power and rise time.

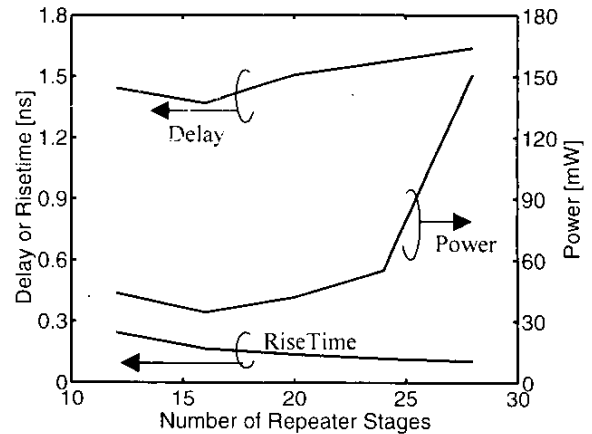


Fig. 1. Performance of minimum-sized interconnect (0.44 μm wide, 0.46 μm spacing in Aluminum) with repeaters.

While the buffer delays can be improved by using more advanced technology, expending more power, and utilizing more silicon area, the wire delay is fundamentally limited. Figure 2 shows the reason behind this by overlaying the power spectral density of a 500ps digital pulse and the intrinsic frequency characteristics of the wire without repeaters. At low frequencies, the wire behaves as a distributed resistance-capacitance (RC) network. In this RC regime, signals travel very slowly and undergo frequency dispersion. As the frequency increases, the impedance due to the inductive component of the wire begins to dominate over the resistance, and the wire behaves more as a waveguide. In this high-frequency inductance-capacitance (LC) regime, the signal group velocity peaks at the speed of light. This is the key characteristic of the wire that can be exploited in order to achieve high-speed signal propagation. However, the power of the digital pulse is concentrated in the low frequency regime of the wire. In the example given in Figure 2, most of the power of a 500-ps pulse is in the RC-regime of the wire, travelling at speeds less than a tenth of the speed of light. Although the high frequency spectral components propagate at extremely high speeds, the low-frequency components lag behind and limit the overall bandwidth and latency of the system. Insertion of repeaters does not change the frequency characteristics of the interconnect, but does help to amplify the high frequency components of the signal. On the other hand, the repeaters contribute delay and power dissipation.




Fig. 2. Frequency characteristics of normal 500-ps wide digital pulse and typical minimum-sized global wires.

Low Loss Wires

The interconnect frequency characteristics shown in Fig. 2 suggest that a high-speed system can be built by taking advantage of the high frequency nature of the interconnect and not expending any power on the low-frequency portion of the signal. This can be accomplished by designing better wires and by changing the transmitted signals to operate primarily in the LC-regime of the wires.

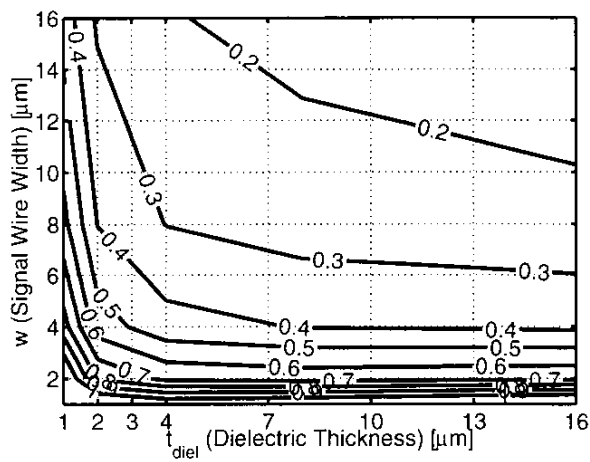


Fig. 3. Design of low-loss on-chip transmission lines. Contours of constant loss in dB/mm for a 2- μ m thick Aluminum microstrip.

While the high-frequency LC-regime offers high-speed, frequency-dispersionless propagation, the interconnect is substantially more lossy at higher frequencies. Therefore, the wire must be optimized to minimize loss, while being restricted to a reasonable amount of area. Thicker top layers of metal and dielectric facilitate the realization of on-chip

transmission lines (Fig. 3). For example, a microstrip structure with a 6 μ m wide, 2 μ m thick signal wire on a 2 μ m thick dielectric provides a loss of about 0.5 dB/mm. This wire would be able to span a cross-chip distance of 20 mm with a total loss of 8 dB. Repeaters will no longer be required because an amplifier at the end of the wire can compensate for this loss.




Fig. 4. Frequency characteristics of modulated pulse and low-loss on-chip interconnect.

System Implementation

In order to take advantage of these low loss wires, we must adopt a signaling scheme that concentrates power in the high-frequency regime. One way to achieve this is to modulate a digital signal up to a sufficiently high frequency such that the line operates in the inductance-dominated regime. Figure 4 shows the impact of such a signaling scheme in combination with the use of low-loss interconnect. Using optimized interconnect, as described in the previous section, allows for near speed of light transmission above 5 GHz. Furthermore, modulating the digital pulse with a high-frequency carrier concentrates the signal power in this high-speed region.

A simple implementation of this system uses direct conversion from baseband to RF. This can be accomplished by mixing the digital data with a high-frequency carrier and directly driving the long length of interconnect. This carrier frequency must be in the LC-regime and sufficiently above the desired bandwidth of the signal. As a demonstration for transmission of 1 GHz data across chip, we have chosen a 7.5 GHz local oscillator (LO) carrier. Furthermore, since the impedance of on-chip interconnect is, in practice, limited to a few tens of ohms, it is difficult to achieve a gain greater than one when driving this load, especially at these multi-GHz frequencies. Therefore, a passive ring mixer is used to perform this up-conversion while minimizing power consumption.

This ring mixer then drives an interconnect wire that extends uninterrupted across a length of tens of millimeters. In the conventional approach, a single, continuous wire is not possible because of the frequency dispersion which leads to slow rise times and limited bandwidth. Repeaters are therefore needed, which expend more power, use silicon area, and contribute delay. In contrast, this modulation system requires no repeaters, saves power and simplifies floorplanning, since the interconnect does not need a dedicated channel of silicon to accommodate the repeaters.

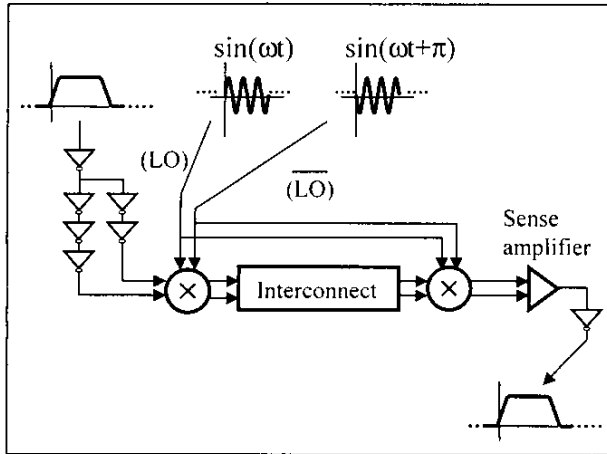


Fig. 5. System block diagram.

On the receiver end, another mixer down-converts the signal back to baseband. Since the received signal has been slightly attenuated through the interconnect, an active double-balanced mixer provides gain to recover the signal. A sense amplifier follows to provide further gain and ensures compatibility with digital logic levels [4]. The actual system implementation is shown in Fig. 5, along with the additional buffers needed for testing purposes. Although the second harmonic of the LO is present after demodulation, simulations show that a filter is not necessary to remove this higher order harmonic.

While the signal travels along the interconnect at the speed of light, additional delay is incurred through the mixers and sense amplifier. This delay due to the active circuitry must be minimized so that the speed gains are not negated in the overall system.

Experimental Results

To demonstrate the system feasibility, the system shown in Figure 5 has been fabricated in a TSMC 0.18- μm logic CMOS technology with 6 levels of aluminum/silicon dioxide interconnect, to demonstrate near speed-of-light data

transmission across chip. The differential LO carrier signals and digital input pulse are generated off-chip, while the transmitter, receiver, and all other components shown in Fig. 5 are integrated on-chip. The interconnect uses a differential microstrip topology, each with a $16\mu\text{m}$ wide line and either $1.9\mu\text{m}$ or $4.9\mu\text{m}$ of oxide between the signal line and the underlying ground line. This differential configuration with a ground shield improves the noise immunity to underlying and neighboring signals. Furthermore, it eliminates loss due to the substrate and provides a well-defined return path.




Fig. 6. Sample as-measured input and output waveforms for system described in Fig. 5.

The system designed with $1.9\mu\text{m}$ dielectric microstrip line is more easily realized in current fabrication processes. The system is optimized with respect to this technology option. The overall delay from the test chip input to the output is 322ps for the falling edge and 490ps for the rising edge. Figure 6 shows sample as-measured input and output waveforms. The inverters that are present to facilitate testing and measurement incur 123ps of this delay. Subtracting these inverter delays yields an average delay of 278ps with a power consumption of 16.1 mW (including the external LO, but excluding the inverters). This total delay corresponds to an effective signal propagation speed of nearly one-half of c_{ox} .

The system designed with $4.9\mu\text{m}$ dielectric yields an average delay of 322.5ps, after subtracting the inverter delays. However, because of the higher impedance of the interconnect, the power consumption is reduced by 25% to only 12.0 mW, including the external LO, but excluding the inverters. Figure 7 shows an eye diagram of the output operating at 1 GHz. This eye may be further improved by including circuitry to reset the internal nodes.




Fig. 7. Eye diagram of system output operating at 1 GHz.

Figure 8 shows a die photo of the test system, 0.8mm in width by 5mm in length. The interconnect occupies the majority of the die area and is laid out in a serpentine manner around dummy metal. In comparison, the transmitter and receiver, located at the top of the die, consume relatively little area. Though the metal area usage of the interconnect is relatively high, no active circuitry is required along the length of the line, allowing for routing over other active circuitry.




Fig. 8. Die photo of test chip.

Summary

In summary, this paper demonstrates, for the first time, the feasibility of all-electrical near speed-of-light cross-chip communication by designing signaling systems around optimized low-loss on-chip transmission lines. The measured overall propagation speed including all circuit delays is roughly one-half the speed of light in oxide, while consuming less power than the conventional repeater approach. In addition, this design will benefit from impending technology enhancements, such as low-K dielectrics, copper, and more metal layers. At the same time, faster devices will decrease the delay due to the active circuitry of the system.

Acknowledgements

The authors wish to acknowledge MARCO Interconnect Focus Center and TSMC for supporting this research.

References

- [1] M. Bohr. "Interconnect Scaling – The Real Limiter to High Performance ULSI." *International Electron Devices Meeting Digest*, 1995, 241-244.
- [2] D. Miller. "Rationale and Challenges for Optical Interconnects to Electronic Chips." *Proceedings of the IEEE*, vol. 88, Issue 6, June 2000, p. 728-749.
- [3] H. B. Bakoglu. *Circuits, Interconnections, and Packaging for VLSI*. Addison-Wesley, 1990.
- [4] M. Bazes. "Two Novel Fully Complementary Self-Biased CMOS Differential Amplifiers." *IEEE Journal of Solid-State Circuits*, vol. 26, No. 2, Feb. 1991, p. 165-168.