Optical Interconnection and Clocking for Electronic Chips

Aparna Bhatnagar and David A. B. Miller
Department of Electrical Engineering
Stanford University, Stanford CA 94305

ABSTRACT

As the speed of electronic circuits approaches 10 Gb/s and beyond, the idea of bringing high speed optical signals directly to CMOS (complementary-metal-oxide-semiconductor) chips seems increasingly imminent. This convergence of electronic speeds with optical network speeds offers opportunities for using light to aid critical electrical functions in novel ways, especially in interconnects and clocking. Modelocked lasers, which emit short optical pulses, have properties which can be exploited to improve the performance of circuits, including low jitter timing sources for future multi-Gb/s electrical circuits.

Keywords: Clock distribution, CMOS detectors, modelocked lasers, jitter, receiver-less, high-speed clocking.

INTRODUCTION

Light is the dominant communication medium between cities, continents and even galaxies. The properties of light which make it useful for communication include its low loss, and low noise for transmission in free-space and over fiber, as well as its high frequency which enables tremendous data capacity. The fundamental physics which underlies these has been described in [1]. Optical communication between electronic chips however, has not been the norm for many reasons. Metal wires have traditionally provided sufficient performance for this purpose. Therefore, the manufacturability and packaging of wired interconnects for electronics is very well developed. As bandwidth requirements reach 10 Gb/s and beyond, the use of optics to link electronic chips on computer boards may prove more efficient than copper wire even for distances less than 1 meter.

Wires are an imperfect communication medium especially at high data rates. When a fast signal is sent down a wire, it gets heavily attenuated and distorted, and can interfere with subsequent data. Without proper termination it can get reflected multiple times and circulate inside the wire. Wires are also susceptible to crosstalk, or noise from other circuits, and have a limited lifetime due to electromigration. This can lead to skew and timing jitter in wires used for clock distribution. Careful electrical design and increased power consumption allow the continued use of wires in moderately high speed systems. However, these effects become more and more limiting at high speeds and density. Optical interconnection and clocking of electronic chips eliminates or mitigates problems with high frequency loss and distortion, wave reflection, electromigration, crosstalk and power dissipation.

In the interconnect application, the distances range from centimeters to ~ 1 m for chip-to-chip and chip-to-board interconnection. Optical interconnection offers several orders of magnitude data rate improvement compared to wires. It also eliminates the electrical power consumed by circuits used to compensate for high frequency loss and distortion. Electrical crosstalk is mitigated and reflections are avoided without power consuming terminations. Additionally, optical interconnection enables new approaches such as wavelength division multiplexing (WDM) and precise optical clocking for time-division multiplexing (TDM) of data are also possible with optical interconnects. WDM offers flexible bandwidth, while optical clocking enhances interconnect performance.

Optical clocking can also be useful for distances less than 1 cm such as on an electrical chip, where precise, noise immune clocks are needed. Typical clock wires can not distribute GHz clocks without the use of amplifying circuits. These circuits add noise on the clock, consume power and cause skew. Electrical approaches to overcome the bandwidth limitations of wires without introducing skew and jitter inevitably incur power penalties. At high speeds clock distribution becomes increasingly difficult. Optically distributed clocks reduce skew and jitter and can enable multi-GHz clocking.

We have investigated various approaches to the use of optics in interconnects and clocking, and found modelocked lasers to be particularly advantageous. Modelocking is a unique mechanism in optics, whereby a laser can be made to emit light in a train of short pulses. The duration of these pulses can be as short as femtoseconds (10^{-15}s) while the repetition rate can be as high as 100 GHz, determined by the round-trip travel time of light in the laser cavity. This repetition rate is quite stable over time-scales on the order of hours, because it is fixed by the cavity length. Such stable short optical pulse trains have a number of unique properties which can be useful for optical interconnection and clocking of electronic chips.

Because the light from a modelocked laser is emitted in

short pulses, it is composed of a band of wavelengths in the frequency domain. Thus, short pulses can be used for WDM, without the need for multiple lasers running at different wavelengths. Short pulses also deliver high peak powers because the energy is concentrated in a short time interval. This property leads to improved efficiency in the amplification of data in an electrical receiver. Finally, the modelocked pulse train can be used as a precise timing reference or clock. The timing stability of an optical pulse is difficult to match electrically at frequencies in the range of several GHz.

Experiments using modelocked lasers for optical interconnection between CMOS chips have demonstrated wavelength division multiplexing as well as receiver sensitivity improvement [2, 3]. Here we summarize in particular experiments using a modelocked laser for clocking CMOS circuits. A simple integrating optical-to-electrical conversion scheme using two photodiodes is used to produce a fast-rising, low-jitter full voltage swing clock with minimum area and electrical power consumption. Receiver-less optical clock injection to a single CMOS circuit with < 6 ps rms jitter is shown. Low capacitance monolithic silicon CMOS photo-detectors using 425 nm short pulses for optical clocking are discussed briefly.

ELECTRICAL WIRES

The use of short pulse lasers for optical interconnection and clocking is interesting because it exploits some of the fundamental benefits that optics can provide such as wavelength diversity, highly coherent output and timing stability arising from frequency independent loss. In proposing optical interconnects as a replacement for copper wire, it is necessary to use those benefits suitably to address otherwise insurmountable problems. Therefore an understanding of the limitations and capabilities of electrical interconnects is important.

Semiconductor scaling allows CMOS chips to run faster and process a greater quantity of information with each generation. Fig. 1 is a graph of microprocessor clock speeds projected by the international technology roadmap for semiconductors (ITRS) [4]. Copper wires used for clocking and chip-to-chip interconnection must supply the bandwidth required to achieve the projected clock rates. Wires by themselves are not a good transmission medium for GHz signals, primarily due to high frequency resistive loss. Qualitatively, the bandwidth of any wire is directly proportional to its cross sectional area and inversely proportional to the square of it length. Therefore, the bit rate (B) supported by a wire can be written as a function of its area (A) over length squared (l^2), also called aspect ratio, according to the formulae given in [5]:

$$B = 10^{15}\, A/L^2 \quad \text{(LC off-chip wires)}$$
$$B = 10^{17}\, A/L^2 \quad \text{(RC on-chip wires)}$$

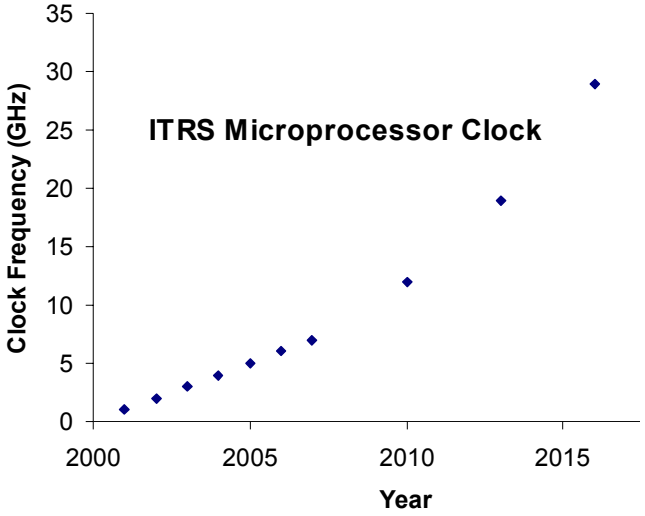

Fig 1. ITRS projections for microprocessor clock frequency

As chips scale, the transistors get faster, therefore the wires must somehow be made faster to keep pace. A simple three-dimensional shrink of a wire does not increase its bandwidth because the aspect ratio remains the same. As the wire interconnection capacity on each metal level fills up, additional wiring levels are added, and wires begin to fill up the third dimension on top of the chip. The aspect ratio constraint implies that once all the available space is filled with wires system capacity can only be increased by switching to a fundamentally different interconnect technology, not merely by scaling the existing approach.

One way to increase the bandwidth of existing copper wires is to buffer them periodically with electrical repeaters. Another way, also known as equalization, is to electronically compensate for wire frequency response. Both consume electrical power and area. Both are limited solutions which can delay the fundamental bandwidth bottleneck but do not overcome it. In contrast, optics allows a few orders of magnitude higher inherent bandwidth. The natural progression of interconnect technology is to add more wires until no more can be added, then electronically enhance the bandwidth until it becomes unprofitable and then switch technologies. Historically, this trend gradually shifts optical interconnects into shorter and shorter length domains.

Off-Chip Interconnect

Chip size is projected to stay fairly constant despite speed scaling. As a result the density of off-chip wiring can not increase much. Thus, either existing off-chip wires must be engineered to support greater data rates or optical chip-to-chip interconnects can be employed. Electrical wires can support slightly greater data rates via equalization and crosstalk management, at the cost of increased power consumption and complexity. In the long run when an inherently high bandwidth interconnect is required optics can be used. An

optical interconnection consists of a laser or modulator driven by a CMOS circuit on the transmitting chip, and a photodetector integrated with a CMOS receiver on the receiving chip. The optical alternative has many advantages as discussed above and also enables new approaches such as WDM and TDM. Here we elaborate on one advantage, which is the ability to add a precise optical clock as an extra channel for time division (de)multiplexing of the data.

In both the electrical and the optical case, the off-chip data rate per line is higher than the on-chip clock frequency for the critical off-chip interconnects. Therefore, time division (de)multiplexing is used for (de)serializing data from the chip. For example, 10 Gb/s off-chip data rate can be achieved with a 2.5 GHz internal clock by multiplexing four bit streams at the lower rate onto one high-speed line using a four phase shifted clock as shown in Fig. 2. An identical (phase locked) clock is needed at both the transmitting and receiving chips. The timing accuracy of the shifted clocks directly affects the maximum bit rate of the link. The generation and distribution of precise multiphase GHz clocks is a challenging task electrically for three reasons. First, the design of CMOS oscillators which produce multi-GHz clocks with the required frequency stability (typically < 10 % of clock frequency) is difficult. Second, the wires in the distribution introduce skew making precise phase shifts challenging. Finally buffering and distribution of clock phases significantly increases the total electrical power consumption.

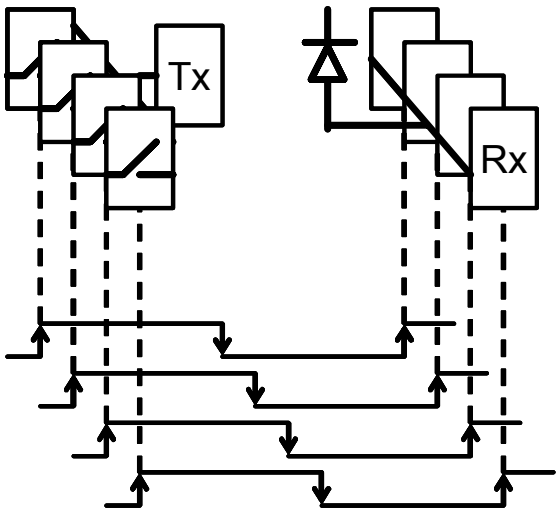

Fig. 2. Four phase clocks for (de)serializing an optical link

Modelocked lasers could supply multi-GHz clocks with excellent frequency stability, and precisely controllable phase is quite readily achieved with optics. Typical semiconductor modelocked lasers have repetition rates ~ 20 – 40 GHz with sub-picosecond jitter. This provides a stable high-speed clock that can be distributed and delivered with precisely controlled phase shifts to multiple points on a CMOS chip with minimal electrical power consumption. Therefore, this is an attractive application for the demonstration of optical clocking with modelocked lasers.

On-Chip Clocking

All synchronous electrical chips require a timing signal or clock. Microprocessors require clocks to synchronize different functional blocks to each other. Electrical clocks are generated by electrical phase-locked-loop (PLL) circuits and distributed by symmetric trees and/or grids of metal interconnect. This generation and distribution method is challenging as the speed and complexity of chip designs increases.

The figures of merit for a clock are skew, jitter, transition time and power consumption. Skew is the unintended relative delay between clocks which are supposed to arrive at the same time at different locations. Clock jitter at any given location is the fluctuation in clock frequency over time. Finally, power is determined by the amount of capacitance driven by the clock distribution. The skew and jitter of the clock must remain within certain budgets to ensure error-free function of an electrical system. Higher clock rates require proportionally greater absolute timing accuracy. Specifically, the combined skew and jitter must remain less than 10% of the clock period for most applications. Thus for example, at 10 Gb/s the total skew and jitter must be below 10 ps. Additionally, transition times should also be ~ 10% of the clock period. Finally, power consumption is presently the most limited resource in clocking.

The use of grids and adjustable delay clock buffers reduces clock skew at the expense of power dissipation. Repeater amplifiers sharpen transition times also at the expense of power. However, jitter is a random noise effect and difficult to compensate electrically. Clock jitter is primarily the effect of power supply noise on transition edges, though it is also composed of PLL jitter. High frequency loss in electronics ultimately limits the Q-factor or frequency stability of an electrical oscillator. Thus although picosecond jitter PLLs have been reported operating ~ 5 GHz, realizing high Q / low jitter PLLs in the high GHz regime is difficult.

Currently, however, clock jitter is dominated by the distribution. On-chip the longest wires are the ones used for clock distribution. The latency from the top of the clock tree to the transistors can be quite high. It is ~ 1.5 clock cycles in the recent 2 GHz Power 5 microcprocessor by IBM [6]. The latency of the clock is proportional to the jitter accumulated in that path. To transmit reasonably precise clock edges across such long path lengths clock wires are heavily repeatered. Increased supply noise with scaling combined with heavy repeatering is causing higher jitter.

OPTICAL CLOCKING EXPERIMENTS

Optical clocking was proposed as an alternative clocking approach by Goodman et. al [7]. An early experiment demonstrated the use of a semiconductor modelocked laser for low jitter clock distribution at a board to board level using optical fiber [8]. The majority of chip-level optical clocking

research has focused on directly modulator laser sources and 'guided wave' distributions, so called because the light paths are defined by waveguides patterned in the geometry of H-trees. Guided wave distribution aimed to minimize global clock skew. However, practical waveguides imposed optical power losses from the input coupling, propagation and bending of light. Also, waveguide process compatibility with CMOS was not achieved. Previous approaches assumed the use of photo-receiver amplifiers to generate the electrical clock from the optical signal, thereby sacrificing optical power and introducing receiver jitter.

Receiver-less Clocking

Jitter is a potentially difficult challenge for electrical clocking at high speeds, as described above. To benefit from the low timing jitter and fast rising edges of modelocked laser pulses we believe it is best to introduce as little circuitry between the photodetector and the clocked node as possible. Therefore we propose the use of a receiver-less detection scheme. The receiver-less ideal is to drive the input capacitance of a clocked element directly with the photocurrent from the detector, without an intervening receiver circuit. This eliminates the power, jitter and latency of the clock receiver, thereby addressing key clocking challenges.

We showed the operation of a CMOS digital logic block clocked by receiver-less injection of pulses from an 82 MHz TiSapphire mode-locked laser. The CMOS chip was fabricated in a 0.5 um Ultra-thin Silicon-on-Sapphire (UTSi) process. The laser generates pulses of about 100 fs, much smaller than any time-scale on the chip. As shown in Fig. 3 two beams are generated using a beam splitter and used to drive a totem-pole of detectors for direct clock injection. Pulses from the two beams are temporally offset by half a bit period or 6.1 ns to generate a 50% duty-cycle clock on chip.

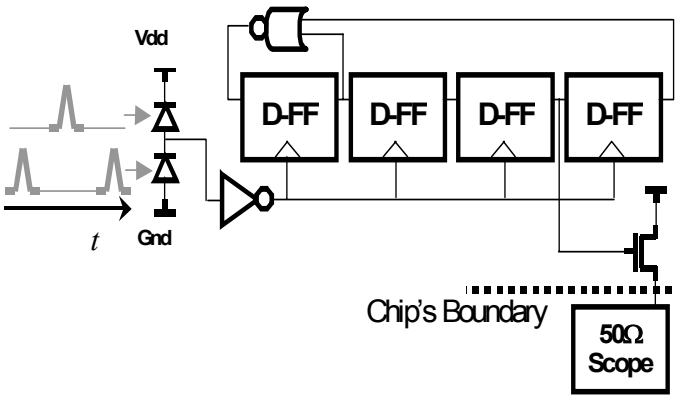

Fig 3. Receiver-less clocking of a digital circuit

As shown in Fig. 3 the digital circuit being clocked is a pseudo random bit sequence generator (PRBS) constructed from four static D-flip-flops and an XOR gate. This closed loop circuit does not require an input, and runs by itself while it receives a good clock. The detectors in the totem-pole are 12 X 12 um^2 $GaAs/Al_{0.3}Ga_{0.7}As$ multiple-quantum-well (MQW) p-i-n diodes integrated in arrays of 200 devices to the CMOS chip via flip-chip bonding as shown in Fig 4. The output of one of the flip-flops is routed to a wire bond pad and observed on a 20GHz oscilloscope. A buffer chain and source follower is used to provide sufficient current to drive the oscilloscope.

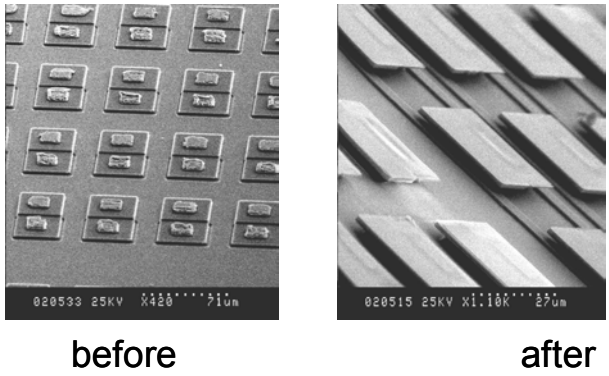

Fig 4. SEMs of flip-chip bonding of an array of detectors to CMOS chip. "Before" shows the bonding bumps. "After", shows the bonded devices.

Fig. 5 is a close-up of the eye-diagram at the oscilloscope when 160 μW is shone on each detector. The histogram on the falling edge shows an rms-jitter of 6 ps. The jitter of the clock at the input of the flip-flop is likely to be significantly less. The measured jitter includes also the jitter introduced by both the flip-flops and the source follower electrical output driver that switches large currents, as well as the jitter from our measurement scope.

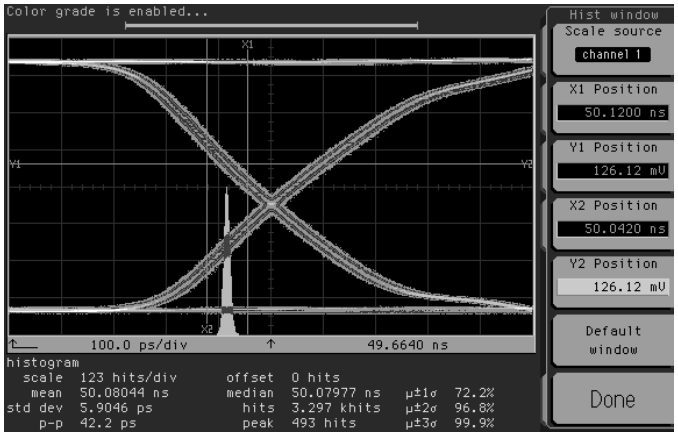

Fig 5. Close-up of eye diagram of the output of PRBS driven by the optical clock. The histogram of the jitter on the falling edge is shown.

Proposed clock distribution

The optical clock can be distributed with low loss and better

than 10 % uniformity in optical power to a small number of points using a diffractive optical element (DOE). The DOE is essentially a grating which, with the help of a lens converts one incoming light beam into an array of focused spots. As opposed to the guided wave distributions mentioned above, a DOE does not incur power losses from coupling, propagration or bending, because the light propagates though free-space, and is effectively 'bent' through constructive interference. The DOE is manufactured in a small piece of glass or plastic separate from the CMOS processing of the chip. The chip to DOE alignment can be lithographically defined in an appropriate packaging scheme.

The limiting parameter in a receiver-less distribution is the total power available from the modelocked laser. To minimize power it is important to have a low loss distribution and also low capacitance photodetectors. The capacitance of flip-chip bonded photodetectors can be reduced to ~ 10 fF by reducing their area. This would slightly necessitate tighter focusing and better alignment. Low capacitance detectors are also feasible within the CMOS process if Silicon-on-insulator (SOI) CMOS is used. The thin silicon films in these technologies allow detectors with capacitances on the order of a few fF for detectors as large as 20 x 20 μm. A model for the scaling of the receiver-less scheme for larger clock distribution networks and varying photodiode capacitances is discussed in [9].

Importantly, the phase of each distributed clock can be set very precisely using optical path delays. This allows the use of optical clocking for serial links. Since the speed of light in air is ~ 300 μm/ps, and μm length adjustments are easily implemented with a motor, sub-picosecond phase adjustment can be achieved. Moreover, since air is not lossy or dispersive like wire, the same clock can be delivered to the transmitter and receiver without relative skew between the two. This is particularly easy to implement in optical links where an optical clock can be added simply as an extra channel.

CONCLUSIONS

Optical interconnection is the backbone of long distance communication. As electronics speeds approach network speeds a possible convergence in high speed computing and communication devices seems possible. Therefore, the prospect of interfacing silicon electronics with light beams is appealing. At high speeds, wires are a progressively poorer transmission medium subject to high frequency loss, reflections, and crosstalk. Therefore achieving data rates above 10Gb/s requires high levels of design complexity, power consumption and engineering effort, when it is at all possible. For such data rates an inherently high bandwidth channel such as light in free-space or fiber is conceptually simpler and more power efficient for both off-chip communication and on-chip clocking.

Jitter and power consumption are likely to be the primary problems in electrical clocking. The low timing noise property of modelocked lasers can be used to produce a fast-rising, low jitter electrical clock using densely integrated photodetectors on CMOS chips. Experiments showing an optically clocked digital circuit have been shown. The use of low capacitance detectors and low loss diffractive optical elements will enable a distribution on a larger scale. The accurate phase control of optical clocks also makes them interesting in serial link applications.

Both electrical and optical links require multiphase clocking to send data at high speeds. More than 50% of the link power is often consumed in the generation and distribution of the electrical clocks needed for synchronizing and multiplexing the I/O. Optical clocks are an easy addition to optical links and remove much of the power consumption while providing enhanced performance at high speed. The impact of optics will likely be evident well above 10 Gb/s as communication over copper wire becomes intractable, and optics enables reduced electrical power consumption while increasing the performance of silicon.

1. Miller, D., *Physical reasons for optical interconnection.* International Journal of Optoelectronics, 1997. **11**(3): p. 155-68.
2. Nelson, B.E., et al., *Wavelength division multiplexed optical interconnect using short pulses.* Selected Topics in Quantum Electronics, IEEE Journal of, 2003. **9**(2): p. 486-491.
3. Keeler, G.A., et al., *The benefits of ultrashort optical pulses in optically interconnected systems.* Selected Topics in Quantum Electronics, IEEE Journal of, 2003. **9**(2): p. 477-485.
4. *ITRS 2002 Update*.
5. Miller, D. and H. Ozaktas, *Limit to the bit-rate capacity of electrical interconnects from the aspect ratio of the system architecture.* Journal of Parallel and Distributed Computing, 1997. **41**(1): p. 42-52.
6. Restle, P.J., et al. *Timing Uncertainty Measurements on the Power5 Microprocessor*. in *Solid-State Circuits Conference, 2004. Digest of Technical Papers. ISSCC. 2004 IEEE International*. 2004.
7. Goodman, J., et al., *Optical interconnections for VLSI systems.* Proceedings of the IEEE, 1984. **72**(7): p. 850-66.
8. Delfyett, P.J., D.H. Hartman, and S.Z. Ahmad, *Optical clock distribution using a mode-locked semiconductor laser diode system.* Lightwave Technology, Journal of, 1991. **9**(12): p. 1646-1649.
9. Debaes, C., et al., *Receiver-less optical clock injection for clock distribution networks.* Selected Topics in Quantum Electronics, IEEE Journal of, 2003. **9**(2): p. 400-409.