

Receiverless detection schemes for optical clock distribution

Aparna Bhatnagar^{*a}, Christof Debaes^b, Hugo Thienpont^b and David A. B. Miller^a

^a Department of Electrical Engineering, Ginzton Laboratory, Stanford University, Stanford, CA 94305-4090; ^b Laboratories for Applied Physics and Photonics, Vrije Universiteit Brussels, 1000 Brussels, Belgium

ABSTRACT

This paper summarizes experiments using a modelocked laser to produce a fast-rising, low-jitter electrical signal for clocking CMOS circuits. A simple integrating optical-to-electrical conversion scheme using two photodiodes is used to minimize jitter, area and electrical power consumption. This scheme is called receiver-less because there is no receiver to introduce unwanted skew, jitter and electrical power consumption. Receiver-less optical clock injection to a single CMOS circuit with < 6 ps rms jitter has been demonstrated. To scale the receiver-less concept low capacitance detectors are necessary. Transit-time limited rise times from low capacitance monolithic silicon CMOS photo-detectors have been simulated using 425 nm short pulses. The utility of this 'receiver-less' scheme is examined for larger clock distribution networks and varying photodiode capacitances.

Keywords: Optical clocking, receiver-less, modelocked lasers, flip-chip bonding, CMOS photodetectors.

INTRODUCTION

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 electrical functions in novel ways. For example, modelocked lasers, which emit short optical pulses, have properties which can be exploited to improve the performance of link circuits as described in [1]. Additionally, modelocked lasers with multi-Gb/s repetition rate and extremely good frequency stability may enhance the timing accuracy of electronic circuits at high speeds. Here we explore the potential of modelocked lasers as low jitter timing sources for future electrical circuits.

All synchronous electrical functions require a timing signal or clock. For example, high-speed communication chips need clocks to synchronize their communication channels, while microprocessors and larger systems require clocks to synchronize different functional blocks to each other. Clock generation and distribution are therefore fundamental functions in electronics. The primary figures of merit for a clock are skew, jitter, rise/fall time and power consumption. As future chip designs become more complex the latency of the clock distribution increases, leading to increased skew and jitter. Additionally as chips become faster, greater absolute accuracy is required from the clock. Contrary to this requirement, the ability of electrical distribution networks to distribute accurate clocks decreases with speed. Consequently, many alternative approaches to high-speed clock generation and distribution are being researched [2, 3]. For applications requiring an extremely stable high frequency oscillator, optics offers a solution in the form of the modelocked laser.

Modelocked lasers with sub-picosecond jitter are available in a range of fixed repetition rates from kHz to multi-GHz. As reviewed in this paper, we have used the pulse train from a modelocked laser to deliver a clock to a CMOS chip with picosecond precision using flip-chip integrated photodetectors which directly drive the clock load. Such direct optical clock injection without a receiver amplifier achieves low-jitter timing with minimum area and electrical power consumption. Since there is no receiver amplifier, the characteristics of the photodetectors determine the speed and required optical power for this technique. The clock rise and fall times are given by the carrier transit times in the photodetector, while to minimize the optical power the detector capacitance must be minimized.

Optimal detectors can be designed in III-V semiconductors and integrated in arrays to CMOS chips via flip-chip bonding. Here, we additionally show that by using short pulse blue light, high efficiency and high speed can be obtained using low capacitance detectors fabricated directly in a CMOS process (here a silicon-on-insulator process). Finally, we note that global clock distribution is feasible with modelocked lasers through the use of free space diffractive optics or guided wave optics. We conclude with an analysis of the benefits and limitations of free space optical distribution relative to the dominant electrical distribution approach.

ELECTRICAL CLOCKING BACKGROUND

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 presents challenges as the speed and complexity of chip designs increase [4]. As mentioned above, the figures of merit for a clock are skew, jitter, rise/fall 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. 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, the rise/fall times should also be $< 10\%$ of the clock period. Current microprocessors use grids and adjustable delay clock buffers to reduce clock skew at the expense of power dissipation. Repeater amplifiers sharpen transition times also at the expense of power. Power consumption is presently the most limited resource in clocking.

In practice, it is the combined effects of skew, jitter and power supply fluctuations that determine the final clock edge variation [5]. 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; ~ 1.5 clock cycles in a recent 2 GHz Power 5 microprocessor by IBM [5]. Clock wires are heavily repeated to transmit reasonably precise clock edges across such long path lengths. Repeaters add latency, consume electrical power and are susceptible to power supply noise converting it into clock jitter. Despite these drawbacks however, repeaters are unavoidable for GHz electrical clock distribution. The latency of the clock is proportional to the jitter accumulated in that path. Qualitatively, this is because as clock trees get deeper with scaling, the path from the clock source to the clocked node gets longer, and the factors causing variations increase in proportion.

In modern systems increasing delay in the clock distribution combined with an increasing use of repeaters leads to increased jitter at higher frequency, whereas the requirement is for the jitter to decrease with increasing frequency. The number of repeaters required continues to increase with clock frequency, as does the clock jitter and power consumption. These trends make clock distribution and clock integrity a serious challenge in electrical systems today. For 10 GHz and beyond, new clock distribution approaches which minimize jitter and electrical power consumption will be required.

OPTICAL CLOCKING

Light is an ideal medium for high speed communication. Optics replaced electrical wires decades ago for long haul communications because the distance dependent loss and low bandwidth of wires limited their capacity. Similar need for higher capacity has led to the introduction of optics at progressively shorter length scales for communication between systems and possibly between chips. With the possibility of light coming down to CMOS chips, the idea of using light to enhance the timing accuracy of high speed electrical circuits becomes relevant.

The idea of optical clock distribution was first put forth in a seminal paper by J. W. Goodman et. al. in 1984 [6]. The primary motivation for optical clocking then was to minimize global clock skew. It was assumed that optical signals would be distributed at a high level on a chip or board with the lower levels of distribution done electrically. Experimental research efforts since then have concentrated on two different approaches to the distribution of light beams. The ‘guided wave’ distribution approach is so called because the light paths are defined by waveguides, which can be fibers or integrated on-chip waveguides. In contrast, the free-space approach is based on the diffraction of light from an element similar to a grating to obtain an array of beams from a single beam and image these onto the clock nodes.

A fiber based approach is suitable for long distance global clock schemes such as inter-board distribution, where the optical clock source feeds a fiber splitter and the clock is distributed via multimode fibers and focused on appropriately placed detectors. For a reasonably high fan-out the primary concerns with this approach are fiber to detector alignment, and uniformity. In 1991 Delfyett [7] demonstrated the distribution of a 302 MHz optical clock to 1024 ports via multimode optical fiber. The optical source was a hybrid modelocked 830 nm semiconductor laser-amplifier with 10 mW average power, and a measured jitter ~ 400 fs. The 6-sigma jitter between any two receiving ports was measured to be 12 ps over one hour. In 1998 a planar H-tree board level optical clock distribution to 64 points using embedded thin-cladding PMMA polymer fiber bundles was demonstrated [8]. The flexible fibers were bent through via holes in the PCB to illuminate receivers on the other side. The light output intensity at different fibers varied by over 100%, corresponding to 3.5 dB uniformity. There was also an excess loss of 6.2 dB while coupling the light into the bundle in addition to some bending and propagation losses.

For chip-scale clock distribution integrated planar waveguides fabricated on chip have been researched, particularly for CMOS compatible fabrication methods. A 1-to-48 clock distribution targeting a Cray supercomputer board application was constructed using polymer waveguides and TIR (total internal reflection) couplers at 850 nm [9]. The proposed optical clocking source for the waveguide distribution has been an 850 VCSEL modulated by an electrical driver. Coupling losses into and out of the integrated waveguides are the major drawback while propagation loss and bending losses are also significant. The integrated approach is also inflexible once fabricated.

In free-space optical clock distribution light beams propagate in air and through specific optical elements to achieve the distribution. Free space distributions are viable at much shorter length scales ranging from a few mm to ~ 1 m for on-chip or chip-to-chip communication. They utilize a hologram which acts as a grating and lens to generate focused spots at the detectors. Optical signals traveling in free space do not incur propagation loss or distortion. The efficiency of this scheme can therefore be as high as 80 % with less than 5 % spot intensity variation [10].

Finally, though the method of distribution of the clock is an important implementation question, it is also necessary to recall that the major problems in the electrical distribution of clocks today are power consumption and jitter. Thus whatever the method of distribution, it must mitigate these problems considerably.

RECEIVER-LESS SHORT-PULSE DISTRIBUTION

The use of short pulse lasers for optical clocking is interesting because it exploits some of the fundamental benefits that optics can provide such as timing stability and wavelength diversity. 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 (wavelength division multiplexing), 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.

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 μm Ultra-thin Silicon-on-Sapphire (UTSi) process. The laser generates pulses of about 100 fs duration, 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. Importantly, the phase of each distributed clock can be set very precisely using optical path delays. Since the speed of light in air is $\sim 300 \mu\text{m}/\text{ps}$, and μm length adjustments are easily implemented with a motor, sub-picosecond phase adjustment can be achieved if phase variation is desired.

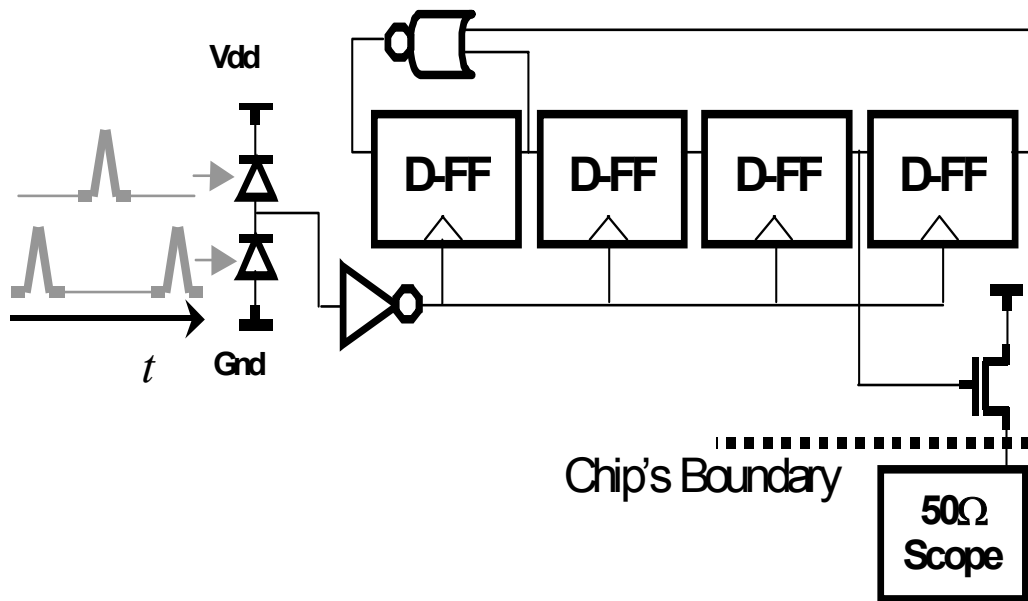


Fig 1. Receiver-less clocking of a digital circuit

As shown in Fig. 1 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 $12 \times 12 \mu\text{m}^2$ GaAs/ $\text{Al}_{0.3}\text{Ga}_{0.7}\text{As}$ 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 2. The output of one of the flip-flops is routed to a wire bond pad and observed on a 20 GHz oscilloscope. A buffer chain and source follower is used to provide sufficient current to drive the oscilloscope.

Fig. 3 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.

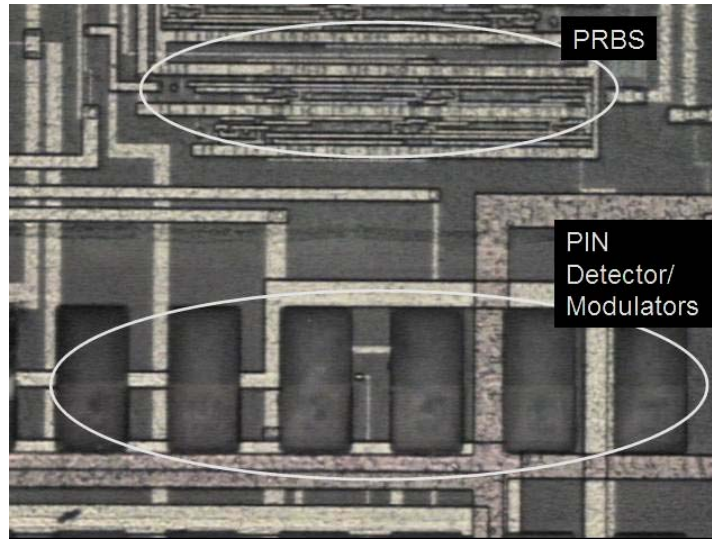


Fig 2. Flip-chip bonded array of detector/modulators to CMOS chip

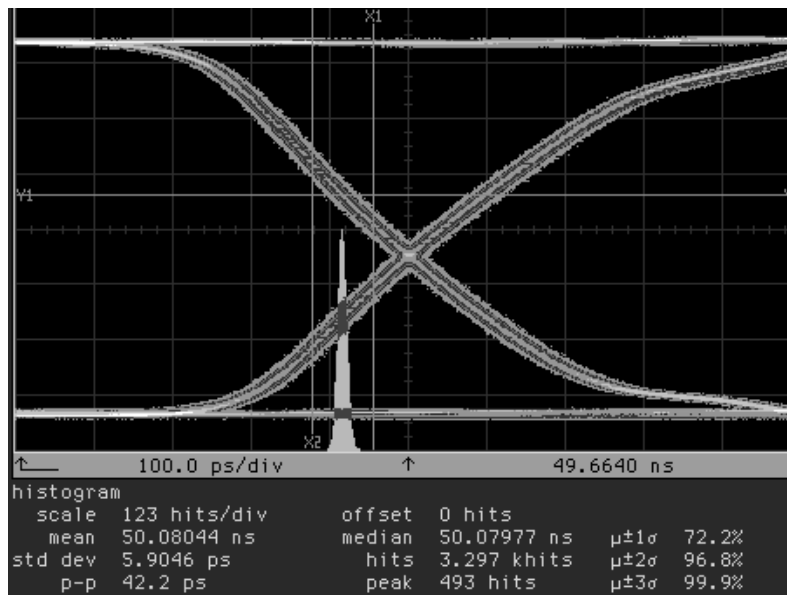


Fig 3. Close-up of eye diagram of the output of the optical clocked PRBS.
The histogram of the jitter on the falling edge is shown.

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, propagation or bending, because the light propagates through 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 $30 \mu\text{m} \times 30 \mu\text{m}$. A model for the scaling of the receiver-less scheme for larger clock distribution networks and varying photodiode capacitances is discussed in [11].

In general, synchronization involves distributing the clock signal to multiple locations. This becomes progressively more difficult with electrical wires as clock speeds increase. Receiver-less optical clock distribution can reduce the jitter, skew and electrical power of traditional repeated-wire approaches, but there is a limit to the number of receiver-less distribution points set by the optical power budget and detector capacitance. To determine this limit we note that the amount of optical power required to swing each receiver-less node to the supply is:

$$P_n = CV/T_b R$$

where V is the supply voltage, C is the node capacitance, R is detector responsivity and T_b is the bit period. The node capacitance is comprised of the detector capacitance plus input capacitance of a buffer to drive the final load. The buffer is not an amplifier, since its input will be a full-swing signal, but is included to provide capacitive gain. In practice the optical power budget is fixed. We note that 10 GHz mode-locked lasers with 2W optical power are in research [12] and assume that total laser power budget is unlikely to exceed 500 mW at higher bit rates. The total number of receiver-less nodes possible is therefore

$$N = P_{tot}/P_n = 10^5/C(\text{fF})$$

The buffer capacitance is ~ 2 fF in $0.13 \mu\text{m}$ CMOS and scales down with technology; thus we C will be dominated by the detector capacitance. A reasonable assumption for present flip-chip technology might be 20 fF, which would allow ~ 2000 to 5000 receiver-less points. We can also conclude that the number of possible receiver-less nodes remains constant unless detector capacitance scales down. Hence, the receiver-less approach can be useful for moderate density clock distribution only unless low capacitance detectors are available.

CMOS DETECTORS

One possibility for implementing low capacitance photodetectors is to fabricate them in the CMOS process alongside the circuits. Silicon is a good material for photodetection especially at visible wavelengths. However, at 850 nm detectors in bulk CMOS processes have a responsivity-speed tradeoff [13]. To overcome this tradeoff silicon-on-insulator (SOI) CMOS technology has been used to fabricate high speed, low efficiency detectors [14, 15]. One interesting effect of an insulating substrate is that the capacitance of detectors in SOI can be very low for shallow silicon layers. Planar PIN detectors have been fabricated in 100 nm thick silicon on sapphire having an estimated capacitance less than 5 fF for detectors as large as $30 \mu\text{m} \times 30 \mu\text{m}$.

Because its bandgap is ~ 1.1 eV, silicon has a fairly long absorption depth of $\sim 14 \mu\text{m}$ at 850 nm. Hence the low capacitance 100 nm thin silicon detectors have very poor efficiency at 850 nm. For free space clock distribution however, there is flexibility in the wavelength of light that can be used. Since low capacitance is essential it is ideal not to use 850 nm light in favor of shorter wavelengths. A particularly attractive shorter wavelength for clocking may be 425 nm, obtained by frequency doubling traditional 850 nm sources. Silicon has an absorption depth of ~ 135 nm in the blue ($\lambda = 400$ nm). Therefore, at this wavelength most of the photogeneration would occur within the depletion region, making CMOS detectors fast, efficient and low capacitance in the blue.

Simulations of SOI CMOS photodetectors with a structure as shown in Fig. 4 were conducted using MEDICI software to determine the efficiency and speed at 425 nm. Corresponding detectors were also fabricated by Peregrine semiconductor in a silicon-on-sapphire process. The silicon layer was assumed 100 nm thick, and an intrinsic silicon with spacing S of intrinsic silicon between heavily doped N and P regions was used. Detectors with S varying from $1.2 \mu\text{m}$ to $6 \mu\text{m}$ were fabricated.




Fig. 4. Schematic cross-section of a two finger lateral p-i-n SOI photodetector.

Fig. 5 shows device simulation results for the rise time of SOI detectors with 5V bias. The device structure was simulated in MEDICI with a 425 nm wavelength short pulse optical input impinging on the detector at time zero. The integrated photocurrent or, equivalently, the total charge collected at the detector terminal is plotted as a function of time for 6 μm and 1.2 μm finger spacing p-i-n SOI detectors. The simulation shows a 10-90% rise time of 120 ps for the 6 μm detector and a rise time of 15 ps for the 1.2 μm detector. These simulations show ~ 20 ps rise times are possible in this technology making these detectors suitable for 10 Gb/s clocking applications. Experimental results corresponding to these simulations will be presented. The feasibility of low capacitance monolithic CMOS detectors with multi-GHz speed and good efficiency would allow receiverless clocking to scale to a greater number of distribution points.




Fig.5. MEDICI simulations of the integrated photocurrent vs. time for planar p-i-n SOI detectors with 5 V bias for 6 μm and, in the inset, 1.2 μm finger spacing.

CONCLUSIONS

As electronic 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 is conceptually simpler and more power efficient.

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 impact of optics will likely be evident well above 10 Gb/s as copper wire becomes intractable, and optics enables reduced electrical power consumption while increasing the performance of silicon.

REFERENCES

1. 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.
2. O'Mahony, F., et al., *A 10-GHz global clock distribution using coupled standing-wave oscillators*. Solid-State Circuits, IEEE Journal of, 2003. **38**(11): p. 1813-1820.
3. Mule, A.V., et al., *Electrical and optical clock distribution networks for gigascale microprocessors*. Very Large Scale Integration (VLSI) Systems, IEEE Transactions on, 2002. **10**(5): p. 582-594.
4. Restle, P.J., et al., *A clock distribution network for microprocessors*. Solid-State Circuits, IEEE Journal of, 2001. **36**(5): p. 792-799.
5. 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.
6. Goodman, J., et al., *Optical interconnections for VLSI systems*. Proceedings of the IEEE, 1984. **72**(7): p. 850-66.
7. 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.
8. Li, Y., et al., *Multigigabits per second board-level clock distribution schemes using laminated end-tapered fiber bundles*. Photonics Technology Letters, IEEE, 1998. **10**(6): p. 884-886.
9. Chen, R.T., et al., *Fully embedded board-level guided-wave optoelectronic interconnects*. Proceedings of the IEEE, 2000. **88**(6): p. 780-793.
10. Walker, S. and J. Jahns, *Array generation with multilevel phase gratings*. Journal of the Optical Society of America A (Optics and Image Science), 1990. **7**(8): p. 1509-13.
11. 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.
12. Krainer, L., et al., *Compact Nd:YVO/sub 4/ lasers with pulse repetition rates up to 160 GHz*. Quantum Electronics, IEEE Journal of, 2002. **38**(10): p. 1331-1338.
13. Woodward, T.K. and A.V. Krishnamoorthy, *1-Gb/s integrated optical detectors and receivers in commercial CMOS technologies*. Selected Topics in Quantum Electronics, IEEE Journal of, 1999. **5**(2): p. 146-156.
14. Csutak, S.M., et al., *CMOS-compatible high-speed planar silicon photodiodes fabricated on SOI substrates*. Quantum Electronics, IEEE Journal of, 2002. **38**(2): p. 193-196.
15. Emsley, M.K., O. Dosunmu, and M.S. Unlu, *High-speed resonant-cavity-enhanced silicon photodetectors on reflecting silicon-on-insulator substrates*. Photonics Technology Letters, IEEE, 2002. **14**(4): p. 519-521.