

The Benefits of Ultrashort Optical Pulses in Optically Interconnected Systems

Gordon A. Keeler, *Student Member, IEEE*, Bianca E. Nelson, *Student Member, IEEE*,
 Diwakar Agarwal, *Student Member, IEEE*, Christof Debaes, *Student Member, IEEE*,
 Noah C. Helman, *Student Member, IEEE*, Aparna Bhatnagar, *Student Member, IEEE*, and
 David A. B. Miller, *Fellow, IEEE*

Abstract—Many properties of an optically interconnected system can be improved through the use of a modelocked laser. The short pulse duration, high peak power, wide spectral bandwidth, and low timing jitter of such a laser lead to these benefits. Timing advantages include simplified synchronization across large chip areas, receiver latency reduction, and data resynchronization. Lower power dissipation may be achieved through improved receiver sensitivity. Additional applications of short optical pulses include time-division multiplexing, single-source wavelength-division multiplexing, and precise time-domain testing of circuits.

Several of these concepts were investigated using a high-speed chip-to-chip optical interconnect demonstration link. The link employs a modelocked laser and surface-normal optoelectronic modulators that were flip-chip bonded to silicon CMOS circuits. This paper outlines experiments that were performed on or simulated for the link, and discusses the important benefits of ultrashort optical pulses for optical interconnection.

Index Terms—CMOS integrated circuits, modelocked lasers, optoelectronic devices, optical interconnections, optical pulses, optical receivers, synchronization, wavelength division multiplexing.

I. INTRODUCTION

THE DESIGN of electrical interconnects between silicon microelectronic chips is becoming increasingly difficult as CMOS transistor technology improves. At gigahertz frequencies, the propagation of electrical signals on printed circuit board traces and other wiring is severely impacted by frequency- and distance-dependent loss and distortion. This and several other factors have led to the introduction of high-speed optical signaling at all levels of digital communication beyond a few meters in length, and to analyses of the benefits of optical interconnects within CMOS systems [1], [2].

Manuscript received October 14, 2002; revised January 31, 2003. This work was supported by the Defense Advanced Research Projects Agency (DARPA) under the University Optocenters program (Contract Number MDA972-00-1-0024) through the University of New Mexico, and by DARPA/MARCO under the Interconnect Focus Center through the Georgia Institute of Technology. The work of C. Debaes was supported by a fellowship from the FWO Vlaanderen. The work of N. Helman and A. Bhatnagar was supported by a Gerhard Casper Stanford Graduate Fellowship and a Herbert and Jane Dwight Stanford Graduate Fellowship, respectively.

G. A. Keeler, B. E. Nelson, D. Agarwal, N. C. Helman, A. Bhatnagar, and D. A. B. Miller are with the Department of Applied Physics and Electrical Engineering, Ginzton Laboratory, Stanford University, Stanford, CA 94305 USA (e-mail: gkeeler@stanfordalumni.org).

C. Debaes is with the Department of Applied Physics and Photonics, Vrije Universiteit Brussel, Brussels 1050, Belgium.

Digital Object Identifier 10.1109/JSTQE.2003.813317

Because modulation frequency and propagation distance have little effect on optical signal integrity (unlike with electrical interconnects), the use of optics enables a radically new method of signaling within computers: using ultrashort pulses, which cannot be generated or propagated by conventional electrical means. Extremely short optical pulses can be produced by modelocked lasers, which readily provide pulse widths of picoseconds or less, repetition rates from megahertz to hundreds of gigahertz, and reasonably high powers (i.e., up to a few watts) [3]. These short pulses can provide many benefits in a CMOS system when compared to more conventional optical signaling. Precise timing is a particularly important advantage. Several authors have investigated the concept of optical clock distribution (OCD) [4], [5] and some have proposed the use of modelocked lasers for the task [6]. Data transmission is another area where short pulses can yield advantages in an optically interconnected system [1], [7]. To date, the vast majority of optical interconnects that have been demonstrated or proposed have employed continuous-wave (CW) lasers, either driven directly or externally modulated. The use of a modelocked laser in conjunction with a modulator-based optical interconnect enables a new low duty-cycle return-to-zero (RZ) data format, which we refer to as “short pulse signaling.” Short pulse signaling can provide several benefits over the conventional nonreturn-to-zero (NRZ) approach.

This paper discusses several experiments and simulations that were performed to investigate the benefits of short pulses in optical interconnects. Section II describes the details of a complete chip-to-chip optical interconnect demonstrator that employs short pulse signaling. Section III discusses the results of system testing, including measurements of receiver sensitivity enhancement and signal retiming. In addition, other link benefits, such as reduced power dissipation, receiver latency reduction, and time-division multiplexing (TDM), are described. Section IV highlights other applications of ultrafast optical pulse trains for silicon CMOS microelectronics, beginning with optical clock distribution. Because modelocked lasers can provide a very low-jitter clock, they are also useful for precise noncontact chip testing on short time scales; the results of a pump-probe style measurement of circuit delay are described. Finally, a demonstration link that uses a modelocked laser as a broadband source for chip-scale wavelength-division multiplexing (WDM) is briefly discussed. Section V summarizes the various benefits of ultrafast techniques for optical interconnection.




Fig. 1. Schematic of the short pulse interconnect demonstration link. A single laser is used to operate many interconnect channels through the use of a diffractive optical element. The beams are modulated by reflective electro-absorption modulators on the transmitter chip and imaged on the receiver chip using polarizing optics to minimize losses.

II. SHORT PULSE DEMONSTRATION LINK

In order to demonstrate the benefits of short pulses for optical interconnection, an optical test link was first created. Fig. 1 shows a conceptual illustration. Two silicon CMOS circuits with integrated optoelectronics are interconnected in the free-space multichannel link. The interconnect is a simple chip-to-chip imaging design that uses standard bulk optics (i.e., lenses, beam splitters, quarter-wave plates, and Risley prisms). The optical components are centered in cylindrical steel cells and mounted on stainless steel slotted baseplates for stability [8]. A diffractive optical element is used to generate an array of beams from a single input laser beam, which allows several channels to be operated simultaneously with the same source. The interconnect laser can be either a modelocked source or a CW source, allowing a comparison of system performance between the two cases.

The chips used in the link were designed in a $0.5\text{-}\mu\text{m}$ silicon CMOS technology. Each chip includes linear arrays of transmitter drivers and receivers, as well as several digital test circuits. The transmitter drivers are simple inverter-based designs that change the bias across the integrated electroabsorption modulators. Two types of receiver were used in the link: an asynchronous transimpedance-amplifier receiver similar to previous designs [9] and a clocked integrating “sense-amplifier” receiver based on a latch design used elsewhere [10]. The sense-amplifier receiver uses positive feedback from a cross-coupled pair of inverters to evaluate an optical input, and its timing diagram illustrated in Fig. 2. While the clock is LOW, input photocurrent is integrated on a low-capacitance input node and the output is held HIGH. When the clock goes HIGH, the integrated charge is evaluated and the output goes to the proper state. This value is latched for the remainder of the clock cycle by a subsequent circuit. Electrical output signals from both receiver types can be sent to on-chip digital test circuits, to electrical output pads, and to additional transmitter drivers that enable high-speed optical readout from the receiver chip.

Digital functions on the chip include a number of convenient bit-error rate (BER) testing circuits. Each transmitter channel in a linear array can be independently driven with data from a $2^{22} - 1$ pseudorandom bit sequence (PRBS) generated on-chip. After it is transmitted through the optical link, the received data can be compared to expected values that are generated by a complementary PRBS circuit on the receiver chip. Hence, the link BER can be measured, allowing a quantitative evaluation of system performance as interconnect parameters are varied.

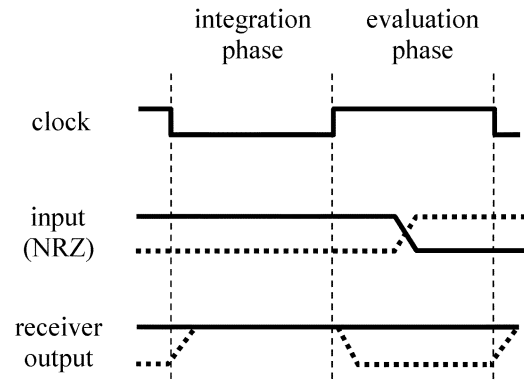


Fig. 2. Timing diagram for the sense-amplifier receiver. Clock LOW corresponds to the integrating phase, while clock HIGH triggers the evaluation phase. The receiver electrical output follows the optical input after evaluation.




Fig. 3. Scanning electron micrograph of an optoelectronic device array integrated to one of the CMOS chips.

The optoelectronic devices used in the demonstration system are molecular beam epitaxy (MBE)-grown, surface-normal GaAs-based multiple-quantum-well $p\text{-}i\text{-}n$ diodes. They are hybridly integrated to the silicon circuits via flip-chip bonding to improve link performance and increase channel density, using an approach that has been described elsewhere [11]. The diodes can be operated as photodetectors by applying a static reverse bias, or as reflective electroabsorption modulators by varying the bias. A single integration step thus adds optical input/output (I/O) functionality to each CMOS chip. Fig. 3 shows a portion of the $2 \times 2\text{-mm}$ silicon die with an integrated array of 200 optoelectronic devices. Each diode has a capacitance of about 250 fF after integration. For 3.3-V operation, the modulator contrast ratio is about 2:1 and centered at 850 nm. Because of the limited contrast ratio, optically differential signaling is employed (i.e., using two beams per channel). This coding technique overcomes the requirement of setting an optical threshold at the receiver and helps to reject amplitude noise from the input laser. The spectral bandwidth of the modulators is wide enough (about 5 nm) that they can efficiently modulate even very short optical pulses, which tend to have a broadened optical spectrum.




Fig. 4. Single-channel eye diagram from the receiver chip of the short pulse interconnect at 82 Mb/s. The signal shown here is the optical output from a modulator on the receiver chip. Hence, this data represents signal transmission through the complete short pulse optical link, followed by optical retransmission to a high-speed oscilloscope.

Two short pulse optical sources were used in the experiments. The first is a commercial modelocked Ti : sapphire femtosecond laser with a pulse width of approximately 100 fs and a repetition rate of 82 MHz. The second is an actively modelocked external cavity diode laser with a spectral linewidth of about 0.7 nm and a pulse width of less than 35 ps. (This measurement was limited by detector bandwidth and is possibly much shorter, since other lasers of this design have been shown to give pulse widths less than 6 ps [12].) The diode laser repetition rate is variable; it can be modelocked at harmonics of the 200 MHz fundamental frequency in order to drive the link at high data rates and has been operated above 3 GHz. The wavelength of each laser was adjusted to achieve optimum performance with the integrated modulators.

During link operation, data was transmitted between the chips at the repetition rate of the modelocked laser. The CMOS circuits were electrically clocked at the same frequency, with the clock phase adjusted such that the short pulses were incident on the modulators while their output was valid. Thus, each pair of modulators in a channel encodes a data bit onto the amplitude of each optical pulse. Fig. 4 shows an eye diagram from a single channel of the short pulse interconnect at 82 Mb/s. In this case, the transimpedance receivers were used with an average received optical power of approximately 100 μ W per beam. The electrical outputs of the receiver circuits were used to drive additional modulators, and high-speed optical readout was performed using a separate CW laser and optical detector. Unlike the sense-amplifier receivers, the transimpedance outputs are neither clocked nor latched; thus, the output can be seen to relax to the off state with a characteristic time constant of ~ 1 ns. (The transimpedance receiver was designed for 1 Gb/s operation, at which speed the receiver output would appear as a typical NRZ signal.)

III. BENEFITS OF SHORT PULSE SIGNALING

A. Receiver Sensitivity Improvement

The sensitivity of an optical interconnect receiver (i.e., the minimum optical energy required to achieve a given BER) is typically worse than for a telecommunications receiver, because very small circuit sizes and low power budgets limit the maximum gain available. However, the use of short pulses can poten-

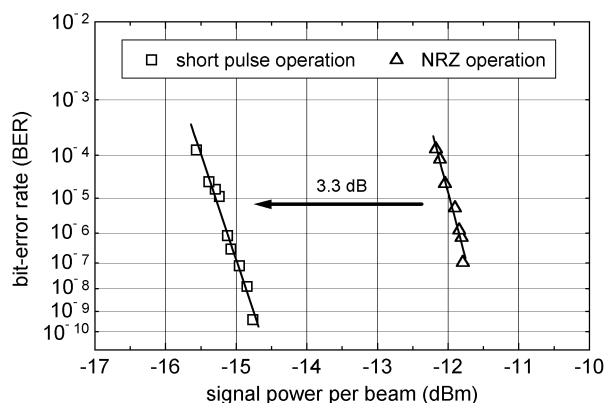


Fig. 5. Measured single-channel BER at 400 Mb/s for NRZ and short pulse link operation. A power savings of 3.3 dB is obtained when the interconnect is operated with the modelocked diode laser.

tially decrease optical power requirements for certain receiver types. This improvement in sensitivity would lower power dissipation, and could decrease the overall circuit size if fewer gain stages are needed.

To experimentally demonstrate this principle, we used the on-chip BER tester to compare link performance between conventional NRZ signaling and short pulse signaling. To do this, the link was operated first with the modelocked diode laser and then with a CW diode laser tuned to the same central wavelength. The sense-amplifier receivers were used for this experiment, and the interconnect was operated at 400 Mb/s in both cases with the operating conditions unchanged between tests. Fig. 5 shows a BER plot for the link as a function of transmitted signal power. The measurements show that an improvement in receiver sensitivity of 3.3 dB was obtained through the use of short pulses. Two separate effects provide this enhancement, which is the greatest improvement attainable for these receivers. The first effect, yielding a gain of 3 dB, is due to more efficient use of energy by the receiver. Because the integration phase (shown in Fig. 2) corresponds to only half of the clock period, any input energy that arrives during the evaluation phase is wasted. Using short pulse RZ signaling, or RZ signaling with a very well-timed pulse, avoids the loss, provided the clocked receiver is properly synchronized to the incoming data. (Note that while increasing the duty-cycle of the integrating phase would help reduce the loss, the evaluation phase must be long enough to ensure sufficient amplification, and the output data must be valid long enough to be sampled by a subsequent latch.) The remaining 0.3-dB gain is attributed to RC speed limitations of the transmitters. The modulator output eye diagram begins to close when using NRZ signaling at these data rates, with the maximum signal amplitude achieved only at the center of the bit. When short pulses are used, they can sample the data at the center of the bit and transmit the maximum signal.

The sensitivity of transimpedance-amplifier receivers can also be substantially enhanced with short pulses, albeit by a different mechanism [13], [14]. It has been shown that the sensitivity of commercial transimpedance receivers can be improved by greater than 5 dB by employing short pulse signaling using optical pulses with a duration of a few picoseconds [13].

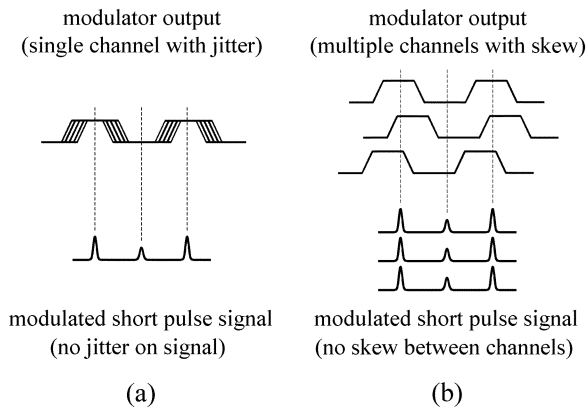


Fig. 6. Illustration of (a) jitter removal and (b) skew removal from the transmitter outputs using optical short pulses from a modelocked laser. The laser output has very low jitter and optical distribution of the beams creates very little interchannel skew.

When the bandwidth of the short pulse input is higher than that of the receiver front-end, the impulse response of the receiver is obtained. Thus, for a given optical energy, the receiver drives the largest possible voltage swing and gives the sharpest rising and falling edges [13].

Receiver performance gain depends on the capacitive load of the photodiodes and the capacitance of the integrated photodiodes used in the link is higher than was expected during design. Therefore, sensitivity enhancement was not observed for our link with the transimpedance-amplifier receivers. However, simulations show that by reducing the diode capacitance to 40 fF, a receiver sensitivity improvement of 5 dB should be possible if optical pulses with a duration of less than 10 ps are used.

B. Signal Retiming

To exploit the bandwidth of a multichannel parallel interconnect fully, one must ensure good signal synchronization between and within channels. Differences in signal timing (either static differences or variable changes) are known as interchannel skew and timing jitter. This timing uncertainty can limit the maximum data rate of parallel interconnects. Such variations are present in the signals sent to the output modulators of a transmitter chip due to process variations, clock skew, gate delays from dissimilar circuits, and changing local conditions on-chip. Minimizing the problem with electrical techniques increases design complexity, but the extremely short pulse duration and low-jitter periodicity of a modelocked laser can be used to remove essentially all transmitter-related skew and jitter [15].

To perform signal retiming (as shown in Fig. 6), short optical pulses should be placed at the center of the bit in the timing reference frame of the transmitter chip. An actively modelocked laser pulse train typically has very low jitter. Thus, a short pulse interconnect can remove the jitter that occurs on the transmitter chip by *instantaneously* sampling the output state of a modulator with great precision—*theoretically up to $\pm 1/2$ of a bit of jitter can be removed*. Static skew in a multichannel parallel link can be removed in a similar fashion. An array of well-timed short pulse beams can sample every modulator *simultaneously*. After retiming is achieved, the signals can be propagated without jitter and skew.






Fig. 7. Jitter removal from a single interconnect channel at 82 Mb/s. The upper trace shows the electrical input signal, while the lower trace shows the optical readout of the receiver. The unlatched transimpedance-amplifier receiver output returns to the LOW state with its characteristic time constant.



(a)



(b)

Fig. 8. Skew removal for two channels operating at 82 Mb/s, using electrical inputs that are skewed by $3/8$ of a bit period. The interconnect is operated using (a) a CW laser and (b) a modelocked laser. Because the data was measured between the transmitter and receiver chips, the finite contrast ratio of the modulators is evident in (b). The finite width of the short pulses is due to the limited bandwidth of the optical detectors used for testing.

To demonstrate these concepts using the demonstration link, skew and jitter were intentionally added to the data at the transmitter chip. Electrical data from a pattern generator was provided to the transmitters rather than using the on-chip PRBS generator. First, jitter was added, with a magnitude of up to $3/8$ of the bit period. The NRZ optical interconnect transferred all of this jitter to the receiver. However, as shown in Fig. 7, the short pulse interconnect removed the transmitter jitter.

Neighboring transmitter channels in the array can be simultaneously operated with different electrical data. In this case, two square waves were intentionally skewed by $3/8$ of a bit relative to one another. Fig. 8 shows the optical outputs of two

channels measured directly after the transmitter chip (i.e., intercepted halfway through the link). One beam is shown from each differential channel, first for the CW and then for the short pulse interconnect. While the interchannel skew remains for the NRZ interconnect, the short pulse interconnect effectively removes all skew.

C. Additional Link Benefits

Many additional benefits can be obtained by operating an optical link using short pulse signaling. Perhaps one of the most common uses of short pulses in telecommunications is for time-division multiplexing. This technique allows an optical fiber to carry more data than a single transmitter has the bandwidth to provide. Recent experimental electrical links have also employed TDM to attain high aggregate data rates [16]. However, in most electrical interconnects, it is the wiring bandwidth and not the speed of the devices that typically limits the maximum signaling rate. In contrast, optical channels have so much bandwidth that they are very difficult to use efficiently. To achieve higher data rates than a single transmitter can obtain [i.e., 10Gb/s for a state-of-the-art vertical cavity surface emitting laser (VCSEL)], TDM can be performed using modulators and short optical pulses. As in the fastest electrical links, one of many clock phases would drive each transmitter in an array. The modulators could be sampled with a phase-delayed array of optical pulses and the bits interleaved on a single fiber. A simple demultiplexing approach at the receiver chip could use gated receivers and the same multiphase clock. (This simple approach would, however, suffer from optical losses.)

Short pulse signaling with modulator-based interconnects can be more energy efficient than conventional NRZ signaling. With short pulses, optical power is only incident on the modulator when the output is valid, while NRZ signaling places power on the devices during the output transitions as well. This wastes energy at the transmitter chip because the modulators dissipate electrical power proportional to the absorbed optical power. It also wastes energy at the receiver chip, where the slow transmitted edges contribute to poorer receiver sensitivity (as described earlier).

The delay, or signal latency, of an optical link is an important parameter that should be minimized. It can be separated into three components: transmitter latency, propagation delay, and receiver latency. In most cases, the propagation delay is the largest component, but receiver latency can be comparable in magnitude for short links such as on-chip interconnects [17]. By simulating the receiver designs used in this work, we have shown [18] that the use of short pulse signaling can substantially reduce the latency of the receiver types used for optical interconnection.

A final link benefit is that larger synchronous areas can be achieved with short pulse signaling. Particularly in free-space systems, optical path lengths can be fixed with great precision. Optically interconnected systems can, therefore, be designed to have extremely low skew. Even very large systems could be operated synchronously [1] with a centralized optical clock and data signals whose timing is derived from the same modelocked source. The timing phase of such a system would be very well defined throughout, because the optical pulses from a modelocked laser have such sharp edges and low jitter.




Fig. 9. Concept of receiverless optical clock injection. A pair of pulse trains from a modelocked laser is incident on a pair of photodiodes. When the pulses are staggered by half of the bit period, a low-jitter square wave appears at the center node. The use of low-capacitance photodetectors reduces the energy required to produce a clock signal with full logic levels.

IV. OTHER APPLICATIONS OF SHORT PULSES

A. OCD

While many benefits can be achieved by using short pulses to transmit data within a system, OCD could very well be the most alluring application. The clock in a digital system is used to synchronize logic operations; it is typically operated just below the frequency at which the circuits begin to fail. However, due to the limitations of electrical signaling, it is becoming increasingly difficult to provide a high-quality clock to all points on a chip. An extremely large fraction of the power dissipated by microprocessors—greater than 33% in a recent design [19]—is now devoted to the problem of clock distribution. In order to keep a 10-GHz microprocessor operating within acceptable timing margins (i.e., with a clock timing uncertainty of less than 10%), future chips will require a distributed clock with no more than 10 ps of combined skew and jitter. While this presents an enormous challenge for future electrical clock distribution schemes, such timing precision may be achieved reasonably easily with optics. Very good control of signal path lengths can be obtained by optical distribution. Additionally, potentially inexpensive external cavity modelocked lasers have been demonstrated with picosecond pulse widths and subpicosecond timing jitter [20], making them an ideal timing source.

For OCD to become a commercial reality, it must be shown that optics can provide ample performance at a reasonable cost. To satisfy this condition, several key issues need to be addressed. Inexpensive optical packaging and integration techniques that are compatible with CMOS fabrication processes must be developed. Appropriate optical detectors do exist, including the integrated III-V devices described here, III-V integrated MSM detectors [21] and silicon detectors that can be designed in a standard CMOS process [22]. However, detectors with low-enough capacitance have only recently been integrated with CMOS.

The receivers used for OCD are of critical importance, because they can easily increase power dissipation and add skew and jitter to the local clock. One attractive solution is the use of very low-capacitance photodetectors in a “receiverless” fashion [23]. As shown in Fig. 9, this approach uses a totem pole of photodiodes, and places temporally staggered optical pulses on the detectors. (Note that any of several detector types could be

used, but monolithic or hybrid integration would likely be necessary to minimize parasitic capacitance.) The resulting photocurrent is rapidly integrated on the low-capacitance center node, creating a rail-to-rail clock signal with precise timing. With no subsequent gain stages present to add skew or jitter, the receiver-less approach allows a high-quality clock to be derived without degrading the original signal. In this way, optics could provide a method of distributing a high-integrity clock to thousands of points, thereby eliminating many of the complicated electrical schemes required today. More details on the concept of receiver-less optical clock injection can be found elsewhere in this issue [24].

While OCD would reduce the design time and on-chip power dissipation related to clock distribution, some degree of electrical distribution will be necessary. Millions of latches will still be used in future microprocessors, and electrical wires and buffers will drive them. Thus, even an optically distributed clock will have some amount of skew and jitter added before it reaches the latch level. A significant advantage of the optical distribution method, however, is that the very precise clock can be provided *without* buffers to key points on the chip. For example, the high data-rate electrical interconnects that employ time-division multiplexing would benefit from very the precise timing of short pulses.

B. Measurements With High Temporal Precision

The low jitter and short pulse duration of a modelocked pulse train enable another interesting application for optically interconnected systems: precise, noncontact on-chip timing measurements. Modelocked lasers have long been used to measure ultrafast physical phenomena with very high precision. For measurements of silicon microelectronics, the use of optically triggered photoconductors to perform on-chip electrical sampling is an interesting approach [25]. Low-temperature-grown GaAs photoconductors can serve as extremely fast electrical pass-gates [26], and integrating these devices to CMOS would enable high-speed analog sampling without heavily loading the circuit. This approach is currently being investigated for use in photonically assisted analog-to-digital converters [27].

Another approach, the well-known technique of pump-probe testing, could be very useful for testing electronic chips that have both optical input and output capabilities. To demonstrate this technique, a silicon chip with integrated optoelectronics was employed. It is similar to the chips described earlier, but fabricated in a 0.25- μm CMOS technology. The measurement was designed to determine the electrical latency of a circuit that consists of a transmitter driver and transimpedance receiver [28]. The circuit was chosen because its delay, together with the delay of optical signal propagation, accounts for the entire latency of an optical link. The latency is difficult to measure by electrical means because of CMOS speed limitations and the capacitive loading effects of probing. Thus, it was measured with the optical pump-probe approach. A schematic of the circuit under test is shown in Fig. 10.

The measurement was performed in an optical pump-probe setup, using the ~ 100 -fs duration pulses of the modelocked

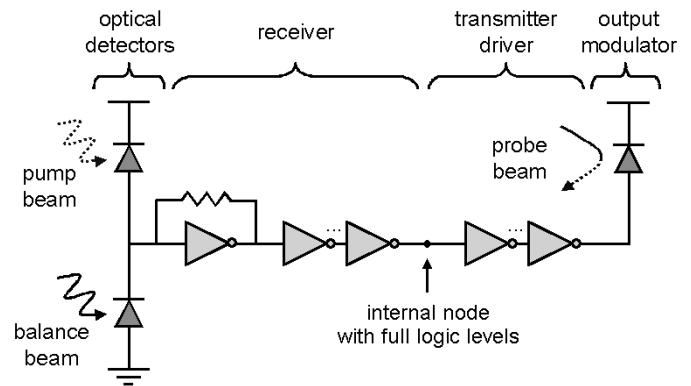


Fig. 10. Simplified schematic diagram of transimpedance-based receiver and transmitter driver circuit used in the latency measurements. Full logic levels are generated by the receiver at the point indicated. The functions of the optical test beams are described in the text.

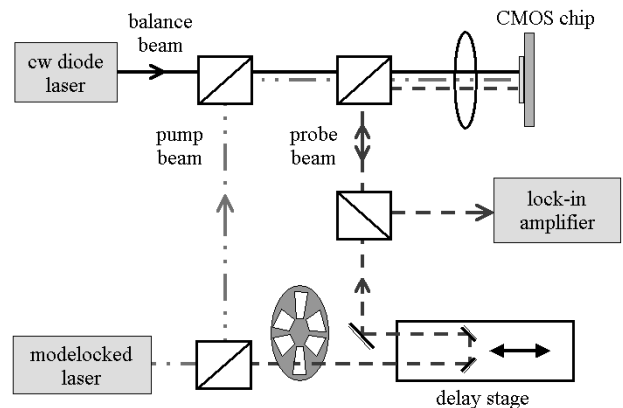


Fig. 11. Schematic of the optical pump-probe setup used for interconnect latency measurements.

Ti:sapphire laser. Fig. 11 shows the experimental setup. Pseudodifferential input data, comprised of a repetitive optical pulse train (pump beam) and a CW diode laser beam (balance beam), was incident on the detector pair of the optically differential receiver. The output voltage of the circuit was another pulse train, whose pulses were broadened by the limited bandwidth of the circuit and time-delayed by the latency of the circuit. The transmitter driver output voltage controlled the modulator reflectivity, which was optically sampled with a short pulse readout beam (probe beam). The intensity of the modulated probe beam could then be measured using a lock-in amplifier. By varying the relative delay between the pump and probe, the temporal response of the optoelectronic circuit was mapped with picosecond resolution.

The results are shown in Fig. 12 for different supply voltages. The measured latency, defined as the delay between the input and the 50% transition of the output, was found to be about 550 ps. This simple measurement technique can easily be extended to other circuits. By adding circuits between the receiver and transmitter driver shown in Fig. 10 (i.e., at the internal node with CMOS logic levels), one could perform a precise measurement of its electrical delay. This technique would make it possible to measure even very small delays, such as the switching time of a single minimum-size inverter.




Fig. 12. Data from a series of pump-probe circuit delay measurements as a function of supply voltage.




Fig. 13. The broadband output of a modelocked laser can be used for single-source WDM by employing spectral slicing. Here, each wavelength is modulated by a separate integrated device and multiplexed into a single fiber. At the receiver chip, a similar demultiplexing operation is performed.

C. Single-Source WDM Optical Interconnect

A final but intriguing application of ultrashort optical pulses is the concept of single-source wavelength-division multiplexing. In telecommunications, WDM has enabled very high data rates to be transmitted on a single optical fiber. At shorter length scales, four-channel “wide” WDM appears to be a cost-effective approach for 10 Gb Ethernet. There, it achieves a high aggregate bandwidth with cheaper, lower performance devices than a single 10 Gb/s link would require [29]. Applying WDM to waveguide-based optical interconnects may lower the cost and complexity of packaging by reducing fiber counts.

Optoelectronic devices that would allow WDM at the interconnect level include either multiwavelength laser arrays or modulators in conjunction with a broadband source. Multiwavelength VCSEL arrays typically rely on complicated growth techniques, and although some such arrays have been fabricated [30], wavelength control during fabrication has remained difficult. An additional problem lies in the stability of integrated transmitter wavelengths as the CMOS chip temperature changes. This problem can be reduced by using a single broadband source located off-chip and performing spectral slicing to define the wavelength channels [31]. The concept

is shown in Fig. 13. The system uses a modelocked laser as a broadband WDM source, and employs spectral slicing by modulating each wavelength component with a separate reflective electroabsorption modulator. This WDM interconnect also benefits from the short pulse signaling advantages mentioned in earlier sections. A complete, chip-to-chip link based on this principle was recently demonstrated; the experimental details can be found elsewhere in this issue [32].

V. CONCLUSION

The ultrashort low-jitter pulses from a modelocked laser present a number of benefits to optically interconnected systems. In order to investigate these benefits, we demonstrated a chip-to-chip multichannel optical link between CMOS chips and employed short pulse signaling. This new data format was achieved by modulating the optical pulses of a modelocked laser with electroabsorption modulators that were flip-chip bonded to the silicon circuits. By measuring the BER of the demonstration link, we experimentally verified that short pulse signaling improves receiver sensitivity when compared to conventional NRZ signaling. We also demonstrated signal retiming (i.e., skew and jitter removal) with the link. In addition to these advantages, signaling with short pulses can decrease receiver latency and enable very high bandwidth channels using time-division multiplexing.

Combining the low jitter of a modelocked laser with the low skew that accompanies free-space optical distribution creates an almost-ideal optical timing source; significant applications of modelocked lasers for optical interconnection include OCD and precise temporal measurements of circuits. A pump-probe measurement of circuit latency was demonstrated with picosecond resolution. Finally, a novel system that uses the broad spectral bandwidth of ultrashort pulses for single-source WDM was discussed.

Advances in laser technology must first occur if many of the benefits discussed in this work are to become practical. While high repetition rate modelocked lasers suitable for precise laboratory testing are currently available and compact versions have been demonstrated [33], the cost of these solid-state lasers would prohibit their use as CMOS optical interconnect light sources. However, recent work [34], [35] has shown that very compact, monolithically integrated modelocked diode lasers (i.e., with a cavity length less than 1 cm) have the potential to fill that role. With continued advances in this field, ultrafast optical techniques may become practical even for low-cost applications such as silicon CMOS optical interconnects. Given this, the use of short optical pulses could prove extremely beneficial in the design of future optically interconnected systems.

ACKNOWLEDGMENT

The authors wish to thank Dr. P. Atanackovic, Dr. T. Pinguet, and M. Wistey for MBE growth of the GaAs materials, and Dr. R. Piestun for design of the diffractive optical element. National Semiconductor Corporation provided CMOS chip fabrication.

REFERENCES

- [1] D. A. B. Miller, "Physical reasons for optical interconnection," *Int. J. Optoelectron.*, vol. 11, pp. 155–168, 1997.
- [2] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, pp. 728–749, June 2000.
- [3] P. J. Delfyett, L. T. Florez, N. Stoffel, T. Gmitter, N. C. Andreadakis, Y. Silberberg, J. P. Heritage, and G. A. Alphonse, "High-power ultrafast laser-diodes," *IEEE J. Quantum Electron.*, vol. 28, pp. 2203–2219, Oct. 1992.
- [4] B. D. Clymer and J. W. Goodman, "Optical clock distribution to silicon chips," *Opt. Eng.*, vol. 25, pp. 1103–1108, 1986.
- [5] J. H. Yeh, R. K. Kostuk, and K. Y. Tu, "Board-level H-tree optical clock distribution with substrate mode holograms," *J. Lightwave Technol.*, vol. 13, pp. 1566–1578, 1995.
- [6] P. J. Delfyett, D. H. Hartman, and S. Z. Ahmad, "Optical clock distribution using a mode-locked semiconductor-laser diode system," *J. Lightwave Technol.*, vol. 9, pp. 1646–1649, 1991.
- [7] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proc. IEEE*, vol. 88, pp. 728–749, June 2000.
- [8] F. B. McCormick, F. A. P. Tooley, J. L. Brubaker, J. M. Sasian, T. J. Cloonan, A. L. Lentine, S. J. Hinterlong, and M. J. Heron, "Optomechanics of a free space photonic switch: The components," *Proc. SPIE*, vol. 1533, pp. 88–96, 1991.
- [9] T. K. Woodward, A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, J. E. Cunningham, W. Y. Jan, L. A. D'Asaro, L. M. F. Chirovsky, S. P. Hui, B. Tseng, D. Kossives, D. Dahringer, and R. E. Leibenguth, "1-Gb/s two-beam transimpedance smart-pixel optical receivers made from hybrid GaAs MQW modulators bonded to 0.8- μ m silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 8, pp. 422–424, Mar. 1996.
- [10] J. Montanaro, R. T. Witek, K. Anne, A. J. Black, E. M. Cooper, D. W. Dobberpuhl, P. M. Donahue, J. Eno, G. W. Hoepfner, D. Kruckmyer, T. H. Lee, P. C. M. Lin, L. Madden, D. Murray, M. H. Pearce, S. Santhanam, K. J. Snyder, R. Stephany, and S. C. Theirauf, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, pp. 1703–1714, Nov. 1996.
- [11] K. W. Goossen, J. A. Walker, L. A. D'Asaro, S. P. Hui, B. Tseng, R. E. Leibenguth, D. Kossives, D. D. Bacon, D. Dahringer, L. M. F. Chirovsky, A. L. Lentine, and D. A. B. Miller, "GaAs MQW modulators integrated with silicon CMOS," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 360–362, Apr. 1995.
- [12] J. P. van der Ziel, "Active modelocking of double heterostructure lasers in an external cavity," *J. Appl. Phys.*, vol. 52, pp. 4435–4446, 1981.
- [13] L. Boivin, M. C. Nuss, J. Shah, D. A. B. Miller, and H. A. Haus, "Receiver sensitivity improvement by impulsive coding," *IEEE Photon. Technol. Lett.*, vol. 9, pp. 684–686, May 1997.
- [14] P. J. Winzer and A. Kalmár, "Sensitivity enhancement of optical receivers by impulsive coding," *J. Lightwave Technol.*, vol. 17, pp. 171–177, 1999.
- [15] G. A. Keeler, B. E. Nelson, D. Agarwal, and D. A. B. Miller, "Skew and jitter removal using short optical pulses for optical interconnection," *IEEE Photon. Technol. Lett.*, vol. 12, pp. 714–716, June 2000.
- [16] C.-K. K. Yang, V. Stojanovic, S. Modjtahedi, M. A. Horowitz, and W. F. Ellersick, "A serial-link transceiver based on 8-Gsamples/s A/D and D/A converters in 0.25- μ m CMOS," *IEEE J. Solid State Circuits*, vol. 36, pp. 1684–1692, Nov. 2001.
- [17] C. Debaes, M. Vervaeke, V. Baukens, W. Meeus, P. Tuteleers, M. Brunfaut, J. Van Campenhout, and H. Thienpont, "Multi-channel free-space intra-chip optical interconnections: Combining plastic micro-optical modules and VCSEL based OE-FPGA," presented at the Proc. SPIE, vol. 4652C, 2002.
- [18] D. Agarwal and D. A. B. Miller, "Latency in short pulse based optical interconnects," in *Proc. IEEE/LEOS Annu. Meeting*, San Diego, CA, Nov. 11–15, 2001, pp. 812–813.
- [19] D. W. Bailey and B. J. Benscheider, "Clocking design and analysis for a 600-MHz Alpha microprocessor," *IEEE J. Solid-State Circuits*, vol. 33, pp. 1627–1633, Nov. 1998.
- [20] C. M. DePriest, T. Yilmaz, A. Braun, J. Abeles, and P. J. Delfyett, "High-quality photonic sampling streams from a semiconductor diode ring laser," *IEEE J. Quantum Electron.*, vol. 38, pp. 380–389, Apr. 2002.
- [21] M. K. Hibbs-Brenner, Y. Liu, R. Morgan, and J. Lehman, "VCSEL/MSM detector smart pixel arrays," in *Proc. IEEE/LEOS Summer Topical Meeting*, Monterey, CA, July 20–24, 1998, pp. 3–4.
- [22] T. K. Woodward and A. V. Krishnamoorthy, "1-Gb/s integrated optical detectors and receivers in commercial CMOS technologies," *IEEE J. Select. Topics Quantum Electron.*, vol. 5, pp. 146–156, Mar.–Apr. 1999.
- [23] C. Debaes, D. Agarwal, A. Bhatnagar, H. Thienpont, and D. A. B. Miller, "High-impedance high-frequency silicon detector response for precise receiverless optical clock injection," *Proc. SPIE*, vol. 4654, 2002.
- [24] C. Debaes, A. Bhatnagar, D. Agarwal, R. Chen, G. A. Keeler, N. C. Helman, H. Thienpont, and D. A. B. Miller, "Receiverless optical clock injection for clock distribution networks," *IEEE J. Select. Topics Quantum Electron.*, vol. 9, Mar.–Apr. 2003.
- [25] W. R. Eisenstadt, R. B. Hammond, and R. W. Dutton, "On-chip picosecond time-domain measurements for VLSI and interconnect testing using photoconductors," *IEEE Trans. Electron. Devices*, vol. 32, pp. 364–369, Feb. 1985.
- [26] S. Gupta, J. F. Whitaker, S. L. Williamson, G. A. Mourou, L. Lester, K. C. Hwang, P. Ho, J. Mazurowski, and J. M. Ballingall, "High-speed photodetector applications of GaAs and $\text{In}_x\text{Ga}_{1-x}\text{As}$ /GaAs grown by low-temperature molecular-beam epitaxy," *J. Electron. Mater.*, vol. 22, pp. 1449–1455, 1993.
- [27] R. Urata, R. Takahashi, V. A. Sabnis, D. A. B. Miller, and J. S. Harris, "Ultrafast differential sample and hold using low-temperature-grown GaAs MSM for photonic A/D conversion," *IEEE Photon. Technol. Lett.*, vol. 13, pp. 717–719, July 2001.
- [28] G. A. Keeler, D. Agarwal, C. Debaes, B. E. Nelson, N. C. Helman, H. Thienpont, and D. A. B. Miller, "Optical pump-probe measurements of the latency of silicon CMOS optical interconnects," *IEEE Photon. Technol. Lett.*, vol. 14, pp. 1214–1216, 2002.
- [29] "10 Gigabit Ethernet Technology Overview White Paper," 10 Gigabit Ethernet Alliance, Newport Beach, CA, 2002.
- [30] C. J. Chang-Hasnain, J. P. Harbison, C. E. Zah, M. W. Maeda, L. T. Florez, N. G. Stoffel, and T. P. Lee, "Multiple wavelength tunable surface-emitting laser arrays," *IEEE J. Quantum Electron.*, vol. 27, pp. 1368–1376, June 1991.
- [31] E. A. De Souza, M. C. Nuss, W. H. Knox, and D. A. B. Miller, "Wavelength-division multiplexing with femtosecond pulses," *Opt. Lett.*, vol. 20, pp. 1166–1168, 1995.
- [32] B. E. Nelson, G. A. Keeler, D. Agarwal, N. C. Helman, and D. A. B. Miller, "Wavelength division multiplexed optical interconnects using short pulses," *IEEE J. Select. Topics Quantum Electron.*, vol. 9, Mar.–Apr. 2003.
- [33] J.-M. Hopkins, G. J. Valentine, B. Agate, A. J. Kemp, U. Keller, and W. Sibbett, "Highly compact and efficient femtosecond Cr:LiSAF lasers," *IEEE J. Quantum Electron.*, vol. 38, pp. 360–368, Apr. 2002.
- [34] H. Fan, N. K. Dutta, U. Koren, C. H. Chen, and A. B. Piccirilli, "Modellocked distributed Bragg reflector laser," *Electron. Lett.*, vol. 35, pp. 48–49, 1999.
- [35] S. Arahira, Y. Katoh, and Y. Ogawa, "20 GHz subpicosecond monolithic modelocked laser diode," *Electron. Lett.*, vol. 36, pp. 454–456, 2000.



Gordon A. Keeler (S'97) was born in Thunder Bay, Canada, in 1973. He received the H.B.Sc. degree in physics from Lakehead University, Thunder Bay, Canada, in 1996 and the M.S. and Ph.D. degrees in applied physics from Stanford University, Stanford, CA, in 1998 and 2003, respectively. His Ph.D. research was done under Prof. David A. B. Miller.

His research involved optical interconnects and optoelectronic devices.



Bianca E. Nelson (S'97) was born in Milwaukee, WI, in 1973. She received the B.S. degree in physics and astronomy from the University of Iowa, Iowa City, in 1996 and the M.S. and Ph.D. degrees in applied physics from Stanford University, Stanford, CA, in 1998 and 2003, respectively. Her Ph.D. research was done under Prof. David A. B. Miller.

Her research involved wavelength division multiplexed optical interconnect systems and passive optical devices.



Diwakar Agarwal (S'95) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Bombay, India, in 1994 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1996 and 2002, respectively. His Ph.D. research was done under Prof. David A. B. Miller.

His research interests involve optical–electrical electrical–optical interfaces, and optical interconnect systems.



Aparna Bhatnagar (S'99) received the B.S. degree in physics and the M.S. degree in electrical engineering, in 1999 and 2001, respectively, from Stanford University, Stanford, CA, where she is currently pursuing the Ph.D. degree in electrical engineering.

Her research involves optical interconnects and clock distribution to CMOS chips.



Christof Debaes (S'99) was born in Geraardsbergen, Belgium, in 1975. He graduated as an electrotechnical engineer from the Vrije Universiteit Brussel (VUB) in 1998. He received the Ph.D. degree from the Applied Physics and Photonics Department, VUB, in collaboration with the Ginzton Laboratory, Stanford University, Stanford, CA, directed by D. A. B. Miller.

He has been granted a postdoctoral fellowship from the Flemish Fund for Scientific Research.

His research interests involve short-distance optical interconnects and clock distribution networks to CMOS-chips.



David A. B. Miller (M'84–SM'89–F'95) received the B.Sc. degree from the University of St. Andrews, St. Andrews, U.K., and the Ph.D. degree from Heriot-Watt University, Edinburgh, U.K., in 1979.

He was with Bell Laboratories, from 1981 to 1996, as a department head from 1987, latterly of the Advanced Photonics Research Department. He is currently the W. M. Keck Professor of Electrical Engineering at Stanford University, Stanford, CA, and the Director of the Ginzton and Solid State and Photonics Laboratories, Stanford, CA. His research interests include quantum-well optoelectronic physics and devices, and fundamental and applications of optics in information, sensing, switching, and processing. He has published more than 200 scientific papers, and holds over 40 patents.

Dr. Miller has served as a Board Member for both the Optical Society of America (OSA) and IEEE Lasers and Electro-Optics Society (LEOS), and in various other society and conference committees. He was President of the IEEE Lasers and Electro-Optics Society in 1995. He was awarded the Adolph Lomb Medal and the R. W. Wood Prize from the OSA, the International Prize in Optics from the International Commission for Optics, and the IEEE Third Millennium Medal. He is a Fellow of the Royal Societies of London and Edinburgh, OSA, and APS, and holds an honorary degree from the Vrije Universiteit Brussels, Brussels, Belgium.



Noah C. Helman (S'02) received the A.B. degree in physics from Harvard University, Cambridge, MA, in 1998 and the M.S. degree in applied physics from Stanford University, Stanford, CA, in 2002. He is currently working toward the Ph.D. degree in applied physics under Prof. David A. B. Miller at Stanford University.

His research involves semiconductor modulators in optical interconnect systems.