

Receiver-Less Optical Clock Injection for Clock Distribution Networks

Christof Debaes, *Student Member, IEEE*, Aparna Bhatnagar, *Student Member, IEEE*,
 Diwakar Agarwal, *Student Member, IEEE*, Ray Chen, *Student Member, IEEE*,
 Gordon A. Keeler, *Student Member, IEEE*, Noah C. Helman, *Student Member, IEEE*,
 Hugo Thienpont, *Associate Member, IEEE*, and David A. B. Miller, *Fellow, IEEE*

Abstract—We present a new technique of injecting clocks optically onto CMOS chips without the use of a receiver amplifier. We discuss the benefits of such a direct approach and present proof-of-principle experiments of the technique. We analytically compare a receiver-less optical clock distribution and an electrical clock distribution in a fan-out-of-four clock tree to evaluate the timing and power benefits of the optical approach for present microprocessors. We also compare receiver-less direct injection of optical clocks to trans-impedance receiver based injection within the same distribution framework.

Index Terms—CMOS integrated circuits, optical clock distribution, receiver-less clock injection, short optical pulses, trans-impedance receivers.

I. INTRODUCTION

PERIODIC clock signals play a central role in the design of synchronous systems today. Data in these systems must be kept synchronized despite unequal logic path delays. Clocking maintains synchronicity in a straightforward manner by providing timing boundaries. Data only proceeds in the system at clock edges, demanding thus an ever-increasing clock frequency to maximize the computational bandwidth. Transistor scaling can increase computational bandwidth by shrinking clock cycle times. However, the associated rise time of the clock and the allowable variation in its arrival time have to shrink proportionally, making reliable clock distribution of growing importance in synchronous systems.

Shrinking rise time, jitter and skew budgets mandate continual improvement in clock generation and distribution

circuitry. Gigahertz-rate phase-locked-loops (PLLs) operating in a noisy environment must be designed with picosecond jitter. Wires used for distributing clocks require even greater design attention because their bandwidth does not keep pace with scaling. Global wires become more resistive as line-dimensions shrink, resulting in an increased delay. Hence, a sharp edge at the input of such a wire will have a much slower rise time at the destination. Repeater-amplifiers are inserted along these wires to maintain sharp rise times, but they are themselves a source of jitter and skew. A slow rising clock edge is not only more susceptible to noise but will also result in a higher variation in the delay.

Clock distribution in current microprocessors involves delivering a sharp noise-free clock to over a hundred thousand static and dynamic latches, which combined, present a capacitive load of several nanofarads. For example, in the Itanium processor [1] a total of 7 nF of capacitance has to be switched with accurate timing every cycle. The power required to charge and discharge this capacitance alone ($C \cdot V_{dd}^2 \cdot f$) is 23 W.

Unfortunately, most circuit design approaches used for correcting the aforementioned shortcomings of wires add to the chip power consumption. Clocking routinely consumes of the order of 30% of the chip power budget. Consequently, clock design today has a large influence not only on the chip's performance, but also on its power consumption, area, wiring resources and time to market.

A variety of techniques are presently used to alleviate the problems in clock distribution. Global routes are carefully handcrafted with local shielding to reduce jitter and accurately control impedance. Differential signaling is sometimes necessary to reduce signal coupling [1], and extensive modeling is done to minimize skew, jitter and inductive effects [2].

Skew reduction techniques begin with a balanced H-tree configuration at the top levels. Load variations across the chip are compensated by precisely engineering wire-widths and buffer sizes with specialized optimization tools [3]. This can be followed by the use of clock grids [4], i.e., meshes of wires tying various clock tap-points together at the lowest level. Grids at any level allow a dramatic decrease in skew but at the expense of additional power consumption associated with charging the capacitance of the grid and the power consumed via the short-circuit currents of skewed clock buffers. An alternate electrical solution is to actively adjust skew rather than attempting to minimize it. A nonzero skew configuration can improve synchronous performance by providing more time for the critical worst-case

Manuscript received October 31, 2002. This work was supported in part by the Defense Advanced Research Projects Agency under the University Optocenters program under Contract MDA972-00-1-0024 through the University of New Mexico, in part by DARPA/MARCO under the Interconnect Focus Center through the Georgia Institute of Technology, in part by DARPA through the Consortium for Optical and Optoelectronic Technologies in Computing (COOP) at George Mason University, and in part by the COOP-Peregrine-USC workshop and foundry run. C. Debaes was supported by a fellowship from the FWO Vlaanderen. A. Bhatnagar and N. Helman were supported by a Herbert and Jane Dwight Stanford Graduate Fellowship, and a Gerhard Casper Stanford Graduate Fellowship, respectively.

C. Debaes and H. Thienpont are with the Laboratories for Applied Physics and Photonics, Vrije Universiteit Brussels, 1000 Brussels, Belgium (e-mail: christof.debaes@vub.ac.be).

A. Bhatnagar, D. Agarwal, R. Chen, G. A. Keeler, N. C. Helman, and D. A. B. Miller are with the Departments of Electrical Engineering and Applied Physics, Ginzton Laboratory, Stanford University, Stanford, CA 94305-4090 USA (e-mail: treasure@stanford.edu).

Digital Object Identifier 10.1109/JSTQE.2003.813319

data paths [5], [6]. Such active de-skewing is achieved by implementing adjustable delay stages inside sector buffers in the Itanium microprocessor.

While circuit solutions are able to overcome the shortcomings of wires in the near term, more radical approaches for clocking may be needed for future chips. Indeed, radically different clocking approaches—from novel architectures to alternative signaling media—are being researched in academia. Distributed PLL networks across the chip [7], on-chip wireless clock distribution [8], salphasic clocks [9] and optical clocks are among the most promising approaches.

This paper evaluates the benefits of optical clock distribution versus conventional electrical clock distribution. We present a new approach to optical clocking using direct receiver-less injection of short pulses from a mode-locked laser to maximize these benefits. After describing free-space optical clock distribution in Section II, we introduce our receiver-less approach in Section III. Section IV presents experimental results showing the injection of a noise-immune sharp optical clock onto CMOS chips using the receiver-less approach. Section V quantifies the benefits of optical clock distribution versus electrical clock distribution using a delay metric for comparing skew and jitter and a capacitive loading metric for comparing power consumption. The model presented in Section V and in the Appendix evaluates the applicability of optical clocking to large-scale clock distribution using a fan-out-of-four model based on current microprocessors, and attempts to answer the question to what depth optics must be introduced into a conventional electrical clocking tree to significantly reduce clock power, jitter and skew.

II. CLOCKING FOR PRESENT DAY CHIPS: THE OPTICAL SOLUTION

Optical clocking refers to the use of optical signals to provide timing boundaries and keep data paths synchronous within digital systems. Typically the clock is a full voltage swing, 50% duty cycle square wave. Although the generation of 10 GHz clocks is currently being researched in industry, it is hard to generate them with sharp and steady rise times. Distributing them through a network of low-speed wires without degrading rise times and without adding timing noise (jitter) or skew is an even more challenging task.

In optics, it is relatively straightforward to generate 10 GHz and higher repetition rate short pulse trains using mode-locked lasers. The repetition rate of the pulses from a mode-locked laser is solely determined by the round-trip time in the laser. Since the light in a laser cavity effectively makes several round trips prior to emission, the Q of a mode-locked laser is very high, making the generated pulse stream a very stable clock source. Because losses in optical cavities have very little dependence on frequency, it is relatively straightforward to make high- Q optical resonators, even in the tens or even hundreds of gigahertz range of repetition rates. Essentially the only sources of jitter are spontaneous emission and mechanical fluctuations of the cavity length. Hence, fundamentally, a mode-locked laser producing sub picosecond pulses with gigahertz repetition rates can have jitter on the order of a few hundred femtoseconds or less [10], [11]. Mode-locked lasers have the additional advantage that the

peak optical power of the pulses can be in the kilowatt range. High average power, gigahertz repetition rate lasers suitable for optical clocking have been demonstrated [12].

It is also relatively easy to propagate these broadband, temporally short pulses in free-space systems using conventional off-the-shelf optics. Clock distribution in the optical case may also be greatly simplified by the use of only one diffractive optical element (DOE) to achieve the entire distribution. DOEs can generate any specified pattern of spots from one incoming beam. A two-dimensional array of spots might be ideal for distributing the optical clock to as many as a million points on the chip. DOEs have low optical loss and can distribute the input power fairly uniformly (up to 1% variation can be specified) over the resulting array of beams. Since DOEs are computer generated to custom specifications, the spot array can even be tailored to compensate for nonuniform load distributions (as explained in Section V).

An alternative to the mode-locked laser for optical clock generation is an externally modulated continuous wave laser. However, these lasers require a high-speed modulator driver whose jitter determines the jitter of the clock. The high- Q , low-jitter benefits of mode locking cannot be achieved via this approach, which offers few other benefits. An alternative to free-space clock distribution is to use embedded wave-guides or fibers to divide and distribute the signal to many points [13]. The fabrication technologies for this approach are a subject of research, and this approach can exploit the technology being developed for planar lightwave circuits for other applications. However, as of today, this method suffers from inefficiency because each division point represents a considerable loss in optical power. Furthermore, the nonuniformity in delivered power to various spots is larger than in the DOE case.

In this paper, we will introduce the concept of direct optical injection by allowing the photogenerated carriers to directly create a full digital-level voltage swing in low-capacitance detectors. This method of clock injection is most suitable for a free-space distribution using short optical pulses from mode-locked lasers, and the experimental results presented here adopt this approach to optical clocking.

III. RECEIVER-LESS CLOCK INJECTION

The receiving end of an optical link typically consists of a detector and an electronic receiver circuit. The detector is optimized to convert impinging photons efficiently into electric current. The receiver circuit then amplifies the small photo-generated current into a large signal suitable for standard digital logic.

The input stage of the receiver circuitry generally consists of either a trans-impedance or a clocked regenerative amplifier. Both topologies prefer a detector with minimal capacitance, because the lower this capacitance, the higher the voltage swing induced at the receiver input by a given number of carriers. By connecting the detector to a high impedance capacitive load and by allowing a sufficient amount of optical power on the device, we can directly create large enough voltage swings to trigger small blocks of digital circuits. This scheme does not require any receiver circuit at all, and therefore eliminates the delay, skew and jitter a normal amplifier would introduce. We are thus




Fig. 1. Schematic view of receiver-less direct clock injection.

directly trading the receiver power consumption and latency for optical power.

To create a 50% duty cycle clock from the pulse stream, we implement a totem-pole of two detectors as shown in Fig. 1 (i.e., two detectors in series, one “above” the other, hence the term “totem-pole”). The top diode, connected to the supply, is used to inject a rising edge. The bottom diode, connected to ground, can be used to reset the voltage back to its initial ground state. Both diodes limit the swing of the node in the middle. Since the voltage over a diode cannot rise above the built-in voltage in forward bias, the voltage at the middle node can, rise above V_{dd} or fall below ground by up to the built-in voltage of the diode. By alternating pulses on both detectors we are able to inject a precise square wave clock onto the chip.

The energy (E_{pulse}) needed to create a voltage swing V with capacitance C when using a detector with external quantum efficiency η is:

$$\eta \frac{E_{\text{pulse}}}{h \cdot \nu} = \frac{C \cdot V}{q} \quad (1)$$

where $h \cdot \nu = h \cdot c/\lambda$ is the energy in one photon and q the elementary charge. Since two pulses are necessary every cycle, the minimum averaged optical power will be:

$$\begin{aligned} P_{\text{opt}} &= 2 \frac{hc}{\eta q \lambda} V_{dd} \cdot f (2C_{\text{diode}} + C_{\text{load}}) \\ &= 2 \frac{V_{dd} \cdot f}{R} (2C_{\text{diode}} + C_{\text{load}}) \end{aligned} \quad (2)$$

where R is the DC-responsivity (A/W) of the detector as generally defined in the literature [14]. In order to drive a $C_{\text{load}} = 30$ fF load with $C_{\text{diode}} = 15$ fF detectors (a conceivable number for a well-integrated detector), an optical power of $430 \mu\text{W}$ is needed, provided the detector has 0.5 A/W responsivity and we are using a 1-GHz clock.

The approach thus requires a reasonable amount of optical power, but gives a number of inherent advantages in exchange:

- No amplifier is required and thus we are eliminating the associated *jitter*, *electrical power* and *delay*. The technique is furthermore compatible with the use of multigigahertz repetition rate mode-locked lasers allowing the creation of sharp low-jitter clock edges that are within the specification of current clock requirements.
- The technique only requires locally a small silicon footprint. Furthermore, we eliminate the need to route a high-speed electrical signal across the chip.
- One can straightforwardly shift the delay of the impinging short pulse stream with femtosecond accuracy merely by changing the optical path length (e.g., in the labora-




Fig. 2. Schematic diagram of the optical clock injection demonstration with a four-stage PRBS.

tory, using the combination of a translation stage and a corner-cube reflector in the optical path). Hence, it is possible to make accurate multiphase clocks for high speed multiplexing or demultiplexing circuits.

- The signal we inject in the chip has potentially a very high slew-rate, which will be larger than anything that can be created by transistors on chip. The creation of these very sharp edges can be exploited on chips to trigger very specific circuits, such as samplers. These optically triggered on-chip samplers can lead to more accurate timing and jitter measurements of time-critical signals on a chip.
- Large swings are directly created on-chip, thus the effect of noise on the supply rail is reduced.

In the next section, we show it is possible with this technique to drive actual digital circuits. We measure the properties of the direct clock injection technique such as the jitter and demonstrate also their immunity toward supply noise.

IV. EXPERIMENTAL VERIFICATION OF RECEIVER-LESS OPTICAL CLOCK INJECTION

A. Proof-of-Principle Demonstration

As a first demonstration we show a proof-of-principle of using the technique of receiver-less direct optical clock injection to drive actual digital circuitry on a chip. We implement, therefore, a small pseudorandom bit sequence generator (PRBS) consisting of 4-D flip-flops and an XOR gate on a CMOS chip as shown in Fig. 2. The chip is fabricated using a commercial $0.5 \mu\text{m}$ Ultrathin Silicon-on-Sapphire (UTSi) CMOS process and integrated with detectors which are GaAs-Al_{0.3}Ga_{0.7}As multiple-quantum-well (MQW) p-i-n diodes. The detectors were flip-chipped to designated pads on the chip and have an active area of $12 \times 12 \mu\text{m}^2$.

To drive the optical clocking experiment, short-pulses from a Ti:Sapphire mode-locked laser with a repetition rate of 80 MHz are carefully aligned onto the detectors of the totem-pole diode pair. We injected short-pulses at the top and bottom diodes temporally shifted by 6.1 ns to create a 50% duty-cycle clock on the chip. The laser is generating sub-picosecond pulses centered at 850 nm, which is much smaller than any time-scale for electrical signals on the chip.

In Fig. 3, a close-up of the eye-diagram at the output of the PRBS circuit is shown when optically clocked with a power of $160 \mu\text{W}$ 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




Fig. 3. Closeup of eye diagram of the output of PRBS drive by the optical clock. The histogram of the jitter on the falling edge is shown.




Fig. 4. Experimental setup.

jitter includes also the jitter introduced by both the flip-flops and the source follower electrical output driver that switches large currents. With this experiment, we have demonstrated a stable, functioning optically clocked digital circuit, and measured an upper bound on the jitter on the output data.

B. On-Chip Samplers as a Technique to Measure Jitter of the Injected Optical Clocks

We used on-chip electrical samplers to measure the injected electrical signals directly inside the chip. Through the technique of repetitive sub-sampling, on-chip samplers offer a unique way to reconstruct signals without large loading, as first proposed by Svensson and Larsson [15]. It allows us to measure a more accurate upper bound on the jitter of the receiver-less technique.

A schematic of the setup is given in Fig. 4. A synchronization signal from the Ti:sapphire laser is used to trigger a pulse-generator, which, in turn, triggers the on-chip samplers. By measuring the output of the sampler at varying delays programmed in the pulse generator, we can reconstruct the repetitive signal at the input of the samplers.

Since the output of the samplers settles in a much smaller time-scale than the repetition period of the laser, it will represent the input voltage of just one sample. Now by using another subsampling device such as a digital oscilloscope, we measure multiple samples of the voltage at the input for one delay. By combining the statistics of these voltage uncertainties at different delays, we are able to deduce the timing uncertainty of the original signal.

Fig. 5 shows the result of such a measurement. The samplers and the receiver-less totem-pole are now implemented in a standard 0.25- μm CMOS process. Similar MQW p-i-n diodes as mentioned above are flip-chipped on this chip but have now an active area of $20 \times 20 \mu\text{m}^2$. The resulting graph is in fact a gray-scale image where, for given coordinates, the intensity of a




Fig. 5. Jitter plot showing the result of the sampled rising edge of the MQW totem pole. Subplot is showing the histogram of the signal crossing point 1.5 V.

point is proportional to the probability of a sample. The subplot shows the histogram of the measured rising edge crossing the 1.5-V boundary. The pass-gate switches of the on-chip samplers consist of PMOS transistors and therefore, we are only able to measure internal voltages above 1.2 V.

With this method we measure an rms-jitter of 4.4 ps. This value matches well with the expected timing accuracy of the system and its triggering set-up. Indeed, directly measuring the optical pulses on a high-speed detector leads up to 3.7 ps rms-jitter when the scope is similarly triggered with the pulse-generator.

C. Supply Noise Immunity: A Comparison With a TIA Receiver

We now compare the jitter of a clock between the receiver-less direct injection technique and a conventional TIA asynchronous amplifier. The implemented receiver consists of an inverter-based trans-impedance amplifier and two additional inverter stages very similar to the receivers implemented in [16]. In order to demonstrate the immunity of the receiver-less method against supply bounce we superimpose an uncorrelated 1-MHz square wave on the supply voltage with peak-to-peak values ranging from 0 to 150 mV. Table I presents the measured rms-jitter. The optical power per detector for the TIA-receiver is $16 \mu\text{W}$ and we are using $150 \mu\text{W}$ of optical power for the receiver-less direct injection technique.

We see that the jitter in the receiver-less technique is smaller than the jitter for the clocks received by the TIA-receivers. Furthermore, we demonstrated the immunity of the technique toward supply bounce. Indeed, the jitter generated with the receiver-less technique does not change significantly for larger supply noise, while the jitter in the TIA receiver in contrast is highly dependent on varying levels of supply noise.

V. MODELING OF AN OPTICAL CLOCK DISTRIBUTION TREE

We have demonstrated the potential of direct receiver-less clock injection to deliver a clock, at a single point on a chip in a manner that is both accurate and tolerant of power-supply

TABLE I

STANDARD DEVIATION ON THE JITTER OF AN OPTICALLY INJECTED CLOCK WHEN NOISE IS INJECTED ONTO THE SUPPLY RAIL. THE INJECTED NOISE IS AN UNCORRELATED SQUARE WAVE AT 1 MHz WITH THE PEAK-TO-PEAK VOLTAGE DEPICTED IN THE TABLE

Injected Noise ($V_{\text{peak-to-peak}}$)	Receiver-less injection	TIA receiver
0mV	4.4 ps	11.0 ps
25 mV	4.3 ps	12.1 ps
50 mV	4.6 ps	11.1 ps
75 mV	5.0 ps	12.0 ps
100 mV	4.9 ps	14.3 ps
125 mV	5.2 ps	15.2 ps
150 mV	5.3 ps	18.9 ps

noise. We will now investigate the usability of this approach in a realistic clock distribution network in modern microprocessors.

Since real chips need a clean clock at over a hundred thousand nodes, a realistic network must distribute the clock to all the points with the same precision and efficiency as we have demonstrated for a single point. The total capacitive clock load, which is represented by the sum of all capacitance of the latches on the chip is large: presently, every cycle a total capacitance of up to 5 to 10 nF needs to be switched in the newest microprocessors.

Clearly, it will not yet be feasible to insert clocks at the lowest level of the clock distribution tree. It would be hard to implement over one hundred thousand detectors on a chip; furthermore, the additional capacitance of all the detectors could significantly increase the total capacitance that has to be switched and hence increase the electrical power consumption unacceptably. There are also limits in the optical domain: a mode-locked laser can only produce a fixed amount of averaged optical power of up to a 5 to 10 W with current research solid state mode-locked lasers [12].

Therefore, in this section, we investigate at what depth in the clock distribution structure would the introduction of optics be more beneficial than conventional distribution networks. We investigate both the case of using receiver-less direct optical clock injection and the implementation of asynchronous TIA-receivers.

In particular, we compare the total electrical consumed *power* in the network as well as the *total delay* from the input of the clock distribution network to the arrival of the clocks at the latches. The total amount of jitter and skew that will be picked up by the distribution network is typically proportional to the total delay in the distribution network, and therefore, the total delay is a very important performance parameter of the configuration.

The particular clock distribution network we are modeling is a symmetric distribution network consisting of inverter buffers which each drive four inverters in the next level of a balanced H-tree configuration as shown schematically in Fig. 6. We require additionally that each inverter drive four times its input

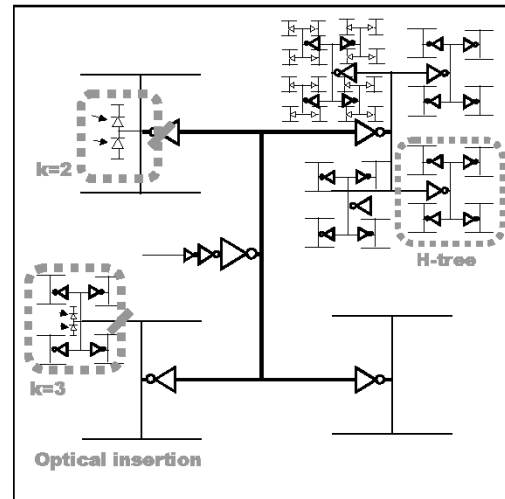


Fig. 6. Schematic view of the modeled clock distribution network.

capacitance. Such a fan-out-of-four (FO-4) configuration generally results in a close-to-minimum delay [17]. As a consequence, at every level, the driving inverter is slightly larger than that at the next level since it has to drive not only four inverters but also the wire capacitances.

The top of the distribution tree is usually an active time alignment circuit such as a PLL or DLL (delay-locked loop) but has only a limited driving capability. Therefore we force the input capacitance of the top of the distribution tree to be equal to the inverter input capacitance at the leaf nodes. This means we have to add an additional number of inverter gain stages at the top of the H-tree for extra capacitance drive. These additional inverters are again chosen to be a FO-4 chain.

This clock distribution network, which is taken as a starting point, is compared to an optical equivalent where the top of the network has been replaced by an optical clock distribution. For a clock tree with n H-tree branches, we define the optical insertion level (parameter k) as the level above which the distribution is done optically as shown in Fig. 6. The electrical power consumption, the delay of the clock distribution network and the required optical power have been modeled for various optical insertion levels in the Appendix. The modeling was done for both a receiver-less direct optical injection technique as well as a more traditional approach with TIA receivers.

The parameters used in this modeling are summarized in Table II. The FO-4 delay, the wire resistance and the wire capacitance characteristic of a typical 0.18 μm technology are used in this model [17]. We underestimate the wire delays in the distribution network by taking the resistance of all the wires in the model to be that of the global wires with the lowest resistance. The wire capacitance per unit length, in contrast, is not significantly different between local, semi-global or global wires and thus the type of wire in the distribution has no influence on the power consumption. The parameter C_{clock} in the table is defined as the total capacitive load of static latches and dynamic latches and the local wires that are required to carry the clock signals on the lowest level of the tree. With the values given in the table we found we would need a nine-level ($n = 9$) clock distribution tree. The extra capacitance of the higher level

TABLE II
NOMINAL VALUES OF THE CLOCK DISTRIBUTION MODEL

	Symbol	Typical value
Leaf node inverter capacitance input capacitance	C_{in}	30fF
Wire capacitance/unit length	C_w	430fF/mm
Resistance of global wires	R_w	20Ω/mm
Fan-out of 4 delay	t_{FO4}	90pS
Supply voltage	V_{dd}	1.8V
Clock frequency	f	1Ghz
Chip total dimension	L	20mm
Total capacitance of the latches	C_{clock}	9nF
Detector transition time	t_{trans}	10pS
Direct injection detector capacitance	C_{rec}	30fF
DC-Responsivity	R	0.5 A/W
TIA-Receiver latency	t_{TIA}	300 pS
TIA-Receiver Analog Power	$PTIA$	500 μW




Fig. 7. Power ratio between optical clock distribution tree and an optical equivalent with different levels of insertion.

wires and inverters is estimated to be 7 nF corresponding to an extra power consumption of 22.7 W. The anticipated delay of this balanced fan-out-of-four electrical network will be 1.7 ns.

In Fig. 7, we depict the ratio between the electrical power consumption of a distribution network with different optical insertion levels and the conventional FO-4 electrical distribution tree. We present this ratio for various different capacitances for the optical detectors that are implementing the receiver-less injection method. In the case of an optical clock distribution with trans-impedance receivers the same plot can be used pro-

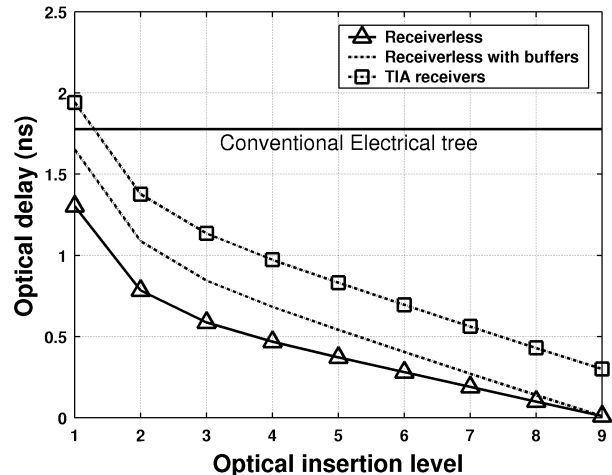


Fig. 8. Absolute delay comparing the optical clock distribution networks with direct injection and TIA receivers vs. the electrical distribution delay.

vided one calculates the equivalent capacitance of the receiver's power consumption (as explained in the Appendix). A 500 μW TIA-receiver will thus correspond to an equivalent capacitance ($C_{rec,equiv}$) of 154 fF.

It is clear that by introducing optics in the distribution tree only an insignificant amount of power can be gained both for the receiver-less and the TIA receiver except for the very lowest levels of detector capacitance in the receiver-less case. In fact, when the capacitance of the detector or the effective capacitance of the receiver is large compared to inverter capacitance at the leaf nodes, the power can even grow radically for optical insertion levels close to the latches. We can only be slightly more optimistic when opting for the receiver-less technique with very small capacitance detectors integrated on the chip.

In Fig. 8, the total distribution delay is compared between a conventional clock distribution tree, and the optical equivalents using either direct injection or transimpedance receivers. One can see that, in contrast with the power gain, the introduction of optics in the clock distribution tree does make a large difference. The graph compares the receiver-less approach, in which we first assume that the totem-pole diodes are capable of driving capacitive loads no matter how large. The second curve shows again the delay using the receiver-less direct injection technique in a tree in which additional buffers are placed such that the load on the receiver-less diodes is always four times the final leaf node capacitance C_{in} . Finally, the delay in a distribution with TIA-receivers is presented.

The reduction in delay, and the associated reduction in jitter and skew, is thus much more significant than the savings in electrical power. For example, with an optical insertion at level 7 of this nine-level distribution tree, we would reduce the delay from 1.7 ns to only 190 ps in the buffered receiver-less case, corresponding to a 90% reduction.

Since the receiver-less direct injection technique directly trades optical power for reduced electrical power consumption, delay and jitter, the total required optical power for a clock distribution using the receiver-less technique will be crucial. Fig. 9 shows the result of the modeling of the averaged total

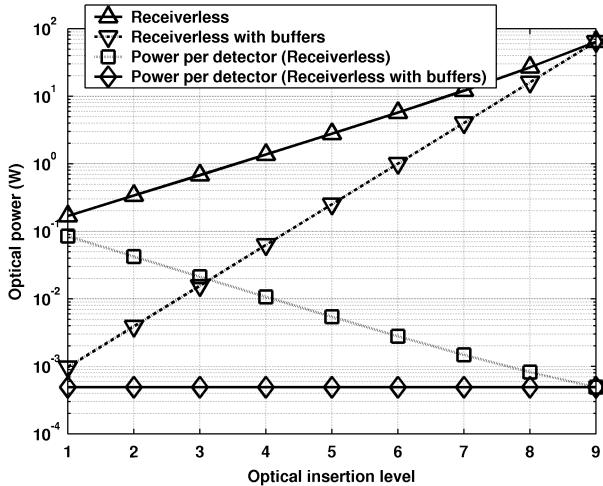


Fig. 9. The total required optical power and the optical power per device for direct clock injection.

required optical power. We present as well the optical power required on each detector. The modeling was done for both the simple receiver-less case, and the receiver-less case in which we add additional capacitance-drive buffers.

If we assume that the mode-locked laser can generate up to 5 W of average optical power, we can inject an optical clock up to the fifth level for the receiver-less distribution and up to the seventh level for the receiver-less technique with extra capacitance-drive buffers. This corresponds respectively to an optical power of 5.4 mW and 490 μ W per detector.

As a conclusion to the modeling of the electrical power consumption, the delay, and the required optical power, we found that choosing optical clock distribution networks will likely not create significant savings in the power consumption, even if the higher wire capacitance of local wires is taken into account. There can however be a large decrease in the total delay through the distribution tree, making the optical solution potentially superior in terms of skew and jitter performance. By using the receiver-less direct injection approach, the delay in the tree can be substantially lower than with the TIA approach. In addition, the direct injection technique is shown to be very immune to supply noise and it can easily be driven by very-high-repetition-rate precise mode-locked laser sources, suggesting a large gain in absolute jitter and skew performance.

In comparing our model to clock distribution in present microprocessors [1], [4], we note that although H-tree configurations are increasingly more widely adopted, they are generally not as uniformly constructed as in our model. Firstly, the inverters from sets of levels in our model are in practice grouped together, to give roughly three total stages of larger buffers (for the Itanium microprocessor the three buffer stages are called the Primary Driver, Second Level Clock Buffers (SLCB) and Gates). Often the first two stages have fanout smaller than four. However, lumping sets of levels to produce larger buffers would not significantly alter the results of the modeling presented here.¹ In fact disregarding wire effects,

¹The power consumption is found by summing all capacitance in wires and inverters, a result that will not change by grouping the inverters differently. The delay in a multi-level H-tree without inverters is more difficult to model, but is unlikely to change drastically.

the three buffer stages approach is exactly the same as our constant FO-4 approach, because each large buffer is composed of smaller buffers in a FO-4 chain.

Secondly, wire dimensions and inverter strengths are carefully tuned and individually designed to compensate for the nonuniform clock loads over a chip's core. By opting for a free-space optical clock distribution, the drive strength at optical insertion points can be accurately tuned by changing the design of the DOE. Note that changing the power distribution will not affect the delay at which photons are impinging on the chip, and thus we can keep the clock of different strengths synchronized.

VI. CONCLUSION

In this paper, we have shown the potential of direct receiver-less optical injection onto p-i-n GaAs detectors. When a totem-pole of the detectors is optically driven by the stable pulse stream of a mode-locked laser, well-defined sharp electrical rising and falling edges can be directly created on-chip, capable of driving small digital circuits. We experimentally demonstrate low-jitter receiver-less injection of clocks that can drive a PRBS on a chip. We characterized an upper bound on the jitter of the injected clock with on-chip samplers and showed that the receiver-less technique is more noise-immune than a TIA-receiver implementation.

After demonstrating the potential of injecting the clock into one point on the chip, we modeled the usability of optical clocks for large distribution networks. Since most of the power is consumed in the last stage of a distribution network, we found that the use of an optical clock will not make significant savings in the power consumption even if wire capacitance effects are included. However, the delay of the distribution can be significantly lowered. Hence, large reductions in skew and jitter are possible through the use of receiver-less optical clock injection.

APPENDIX CLOCK DISTRIBUTION MODEL

The fan-out-of-four clock distribution tree we are modeling has four times more inverters in each successive level or stage of the tree. If we assume that the buffers or latches at the leaf nodes are also driving four times their input capacitance, we can calculate the total number of levels that are required in the tree:

$$\frac{C_{\text{clock}}}{4 \cdot C_{\text{in}}} = 4^{n-1} \Rightarrow n = \log_4 \frac{C_{\text{clock}}}{C_{\text{in}}}. \quad (3)$$

Next we calculate how much the inverter buffers are increased in size to compensate for the wire resistance from bottom to top. When all inverter sizes are known we will model the following three important performance metrics: the total electrical power consumption, the delay in the distribution tree, and the required optical power to drive the optical distribution tree.

A. Inverter Size at Various Levels of the Distribution Network

At every stage the driving inverter ($C_{\text{in},i}$ input capacitance) switches the wire capacitance as well as the input capacitance of the 4 inverters at the next level ($C_{\text{in},i+1}$). If we additionally impose that every inverter exactly drives four times its input capac-

itance, we can write down a recursive formula that determines the input capacitance at every stage,

$$C_{\text{in},i} = \frac{1}{4} \left[C_W \cdot \frac{3L}{2^i} + 4 \cdot C_{\text{in},i+1} \right]; \quad C_{\text{in},n} = C_{\text{in}} \quad (4)$$

where C_W is the wire capacitance per unit length and $3L/2^i$ is the wire length of one H-tree configuration at level i for a chip with dimension L . At the leaf inverter nodes, the inverter size is fixed to a minimum value C_{in} that is taken as a reference. Expanding the recursive formula results in:

$$C_{\text{in},i} = C_{\text{in}} + \frac{3L}{2} C_W \left(\frac{1}{2^i} - \frac{1}{2^n} \right). \quad (5)$$

Since we require the input capacitance of the tree to be the same as the reference input capacitance of the leaf nodes, we need to add an additional number of inverter stages. The buffer length (BL) or the number of stages required for such a capacitance-drive amplification is:

$$\begin{aligned} \text{BL} &= \log_4 \left(\frac{C_{\text{in},1}}{C_{\text{in},n}} \right) = \log_4 \left[1 + \frac{3 C_W L}{2 C_{\text{in}}} \left(\frac{1}{2} - \frac{1}{2^n} \right) \right] \\ &\approx \log_4 \left(\frac{3 C_W L}{4 C_{\text{in}}} \right) \quad (\text{for } 2^n \gg 1 \wedge C_W L \gg C_{\text{in}}). \end{aligned} \quad (6)$$

Last equation assumed a large distribution network ($2^n > 1$) and a global wire capacitance, $C_W \cdot L$ that is significantly larger than the inverter capacitance at the leaf node C_{in} .

B. Modeling of the Required Electrical Switching Power

The electrical power consumption is primarily determined by the energy needed to charge and discharge all the capacitances at the tree for every clock cycle. This power is linearly proportional with the capacitance, i.e., $P = V^2 \cdot f \cdot C_{\text{Tot}}$. It suffices thus to account for all capacitances in the distribution network. We calculate first the total capacitance of all the inverters ($C_{\text{totInv},j}$) of a distribution network with j -levels.

$$\begin{aligned} C_{\text{totInv},j} &= \sum_{i=1}^j 4^{i-1} C_{\text{in},i} \\ &= \left(\frac{C_{\text{in}}}{3} - \frac{C_W L}{2^{n+1}} \right) (4^j - 1) + \frac{3 C_W L}{4} (2^j - 1) \\ &\approx \frac{C_{\text{in}}}{3} 4^j + \frac{C_W L}{4} 2^j \quad (\text{if } 2^j \text{ is large } \wedge 2^n \gg 1). \end{aligned} \quad (7)$$

We calculate as well the sum of all the wire capacitance $C_{\text{totW},j}$ of the H-tree configuration with j -levels. Note that we do not include the wires after the inverter at level j .

$$\begin{aligned} C_{\text{totW},j} &= \sum_{i=1}^{j-1} \frac{3L}{4} 2^i C_W = \frac{3}{2} C_W L (2^{j-1} - 1) \\ &\approx \frac{3 C_W L}{4} 2^j \quad (\text{for } 2^j \text{ large}). \end{aligned} \quad (8)$$

Finally, we have to account as well for the capacitance in the buffer at the top of the distribution tree:

$$\begin{aligned} C_{\text{Buf}} &= \sum_{i=0}^{\text{BL}-1} C_{\text{inv}} 4^i \\ &= C_{\text{inv}} \frac{4^{\log_4 \left(\frac{3 C_W L}{4 C_{\text{in}}} \right)} - 1}{3} \approx \frac{C_{\text{inv}} C_W L}{C_{\text{in}} 4} - \frac{C_{\text{inv}}}{3}. \end{aligned} \quad (9)$$

Note that the buffer capacitance is small and nearly independent of the size of the distribution tree n . We can therefore neglect the power consumption of the top buffer.

The total capacitance in an electrical distribution network is:

$$\begin{aligned} C_{\text{tot}E} &= C_{\text{totInv},n} + C_{\text{totW},n} + C_{\text{Buf}} \\ &\approx C_{\text{inv}} \frac{4^n}{3} + C_W L \cdot 2^n. \end{aligned} \quad (10)$$

If we now choose to inject the clock optically at level k of the distribution tree, we eliminate $C_{\text{totInv},k}$ capacitance in the top inverters and $C_{\text{totW},k}$ of capacitance of top level wires. However, the receiver-less technique requires a fixed capacitance of the detector totem-pole for every optical entry point. The total capacitance of a distribution network with optical insertion at level k becomes:

$$C_{\text{totInv},n} + C_{\text{totW},n} - C_{\text{totInv},k} - C_{\text{totW},k} + 4^{k-1} C_{\text{rec}}. \quad (11)$$

The ratio of the total capacitance (and therefore also the consumed power) between the conventional electrical distribution network and the optical approach at level k is

$$\begin{aligned} \frac{P_{\text{opt},k}}{P_{\text{el}}} &= 1 - \frac{C_{\text{in}}}{3 C_{\text{tot}E}} \cdot 4^k - \frac{C_W L}{2 C_{\text{tot}E}} \\ &\quad \times \left(3 - \frac{1}{2^{n-k}} \right) \cdot 2^k + \frac{C_{\text{rec}}}{4 C_{\text{tot}E}} \cdot 4^k. \end{aligned} \quad (12)$$

We can conclude that by inserting optics at level k , we remove the power related to the eliminated inverters but have to add the power consumed by the added detector capacitances. Both terms are proportional to the number of optical insertion points $N_o = 4^k$ we implement on the chip. If $C_{\text{in}} = 3/4 \cdot C_{\text{rec}}$ holds there would be no power gain or loss when opting for optical clock distribution if wire effects are neglected. The introduction of wire effects gives a slight preference for optical distribution networks; a power proportional to the number of detectors in one dimension (the square root of N_o) will be cancelled in first order.

When a TIA-receiver implementation is chosen instead, we need to introduce an equivalent capacitance that will account for the analog power P_{TIA} consumed in the receiver:

$$C_{\text{req,Eq}} = \frac{P_{\text{TIA}}}{V^2 \cdot f}. \quad (13)$$

All results we previously found will still apply provided we replace the C_{rec} by the equivalent $C_{\text{rec,Eq}}$.

C. Modeling of the Delay in the Optical Clock Distribution

Next to power consumption, delay is a very important performance metric since it is an indicator of the jitter and the skew that will be picked up by the tree. For the fan-out-of-four clock distribution we are modeling we have chosen the following inverter-to-inverter delay [16]:

$$t_{d,i} = t_{\text{FO4}} + \frac{1}{2}C_W R_W l_i^2 \quad (14)$$

where the inverter-to-inverter length l_i on level i is equal to $L/2^i$. To calculate the total delay of the conventional electrical distribution tree we have to sum the various inverter-to-inverter delays $t_{d,j}$ and an additional delay due to the initial capacitance drive buffer with length BL.

$$\begin{aligned} t_{d,\text{el}} &= t_{\text{FO4}} \cdot n + \frac{C_W R_W L^2}{6} \left(1 - \frac{1}{4^{n-1}}\right) + t_{\text{FO4}} \cdot \text{BL} \\ &\approx t_{\text{FO4}} \left(n + \log_4 \frac{3C_W L}{4C_{\text{in}}}\right) + \frac{C_W R_W L^2}{6}. \end{aligned} \quad (15)$$

We assumed that at the lowest level the tree is loaded with our initial total clock load C_{clock} in which the capacitance of the local wires is already included.

If we now chose to inject the clock optically at level k the delay will become

$$\begin{aligned} t_{d,\text{opt},k} &= (n - k)t_{\text{FO4}} + t_{\text{trans}} \\ &\quad + \frac{C_W R_W L^2}{6} \left(\frac{1}{4^{k-1}} - \frac{1}{4^{n-1}}\right). \end{aligned} \quad (16)$$

We added to the difference in delay of an n - and k -stage tree the transition time in the receiver-less detector as well as the additional wire delay from the receiver-less totem-pole to one of the four inverters of the next level. Note that the model is completely valid for an optical injection via a TIA receiver as well provided we replace the transition time by the delay in the receiver $t_{d,\text{TIA}}$.

In this derivation, we assumed that for an optical injection at level k , the detector or the receiver-less technique or TIA-receiver could drive the load at that level, no matter how large its capacitance is. Now, if we assume instead that the injected clock can maximally drive four times the referenced minimal inverter capacitance C_{in} we need to add an additional amount of inverters to the optical delay $t_{\text{del,opt},k}$:

$$\begin{aligned} t_{\text{FO4}} \cdot \log_4 \left(\frac{4C_{\text{in},k}}{4C_{\text{in}}}\right) \\ = t_{\text{FO4}} \cdot \log_4 \left[1 + \frac{3C_W L}{2C_{\text{in}}} \left(\frac{1}{2^k} - \frac{1}{2^n}\right)\right]. \end{aligned} \quad (17)$$

For optical injection close to the leaf nodes the extra wire capacitance is insignificant since the wires are short and hence we do not need an extra buffer. The delay we derived above is still a good approximation to the total delay in the distribution tree. On the other hand, for optical injection closer to the top of the tree we can simplify the extra buffer delay to

$$t_{\text{FO4}} \cdot \log_4 \left(\frac{C_{\text{in},k}}{C_{\text{in}}}\right) \approx t_{\text{FO4}} \left(\log_4 \frac{3C_W L}{4C_{\text{in}}} - \frac{k+1}{2}\right). \quad (18)$$

D. Modeling of the Required Optical Power

The receiver-less injection technique is directly trading consumed electrical power and delay with power in the optical domain. Hence, the amount of required averaged optical power for an optical clock distribution tree will be significant, and will eventually set the ultimate limits on the optical insertion level.

As explained in Section III, the optical power required to drive a certain node is in first order proportional to the amount capacitance that has to be switched. The total switch capacitance at optical insertion level k is

$$\begin{aligned} C_{\text{opt,tot}} &= 4^{k-1}C_{\text{rec}} + 4^k C_{\text{in},k+1} + 3/2 \cdot 2^{k-1}LC_W \\ &= 4^{k-1}C_{\text{rec}} + 4^k C_{\text{in}} + \frac{3}{2}C_W L \cdot 2^k \left(1 - \frac{1}{2^{n-k}}\right). \end{aligned} \quad (19)$$

where $3/2 \cdot 2^{k-1} \cdot L$ is the total length of the wires from the detector to the inverters on level k . However, the required optical power for an optical clock distribution with the buffered receiver-less or TIA-receivers is different. Such an implementation requires a fixed amount of power per detector, thus requiring $P_{\text{opt,tot,TIA}} = 4^{k-1} \cdot P_{\text{opt,TIA}}$.

ACKNOWLEDGMENT

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

REFERENCES

- [1] F. E. Anderson, J. S. Wells, and E. Z. Berta, "The core clock system on the next-generation Itanium microprocessor," in *Proc. ISSCC Conf.*, vol. 2, Feb. 2002, pp. 110–112.
- [2] Y. I. Ismail, E. G. Friedman, and J. L. Neves, "Exploiting the on-chip inductance in high-speed clock distribution networks," *IEEE Trans. VLSI Syst.*, vol. 9, pp. 963–973, Dec. 2001.
- [3] P. J. Restle, T. G. McNamara, D. A. Webber, P. J. Camporese, K. F. Eng, K. A. Jenkins, D. H. Allen, M. J. Rohn, M. P. Quaranta, D. W. Boerstler, C. J. Alpert, C. A. Carter, R. N. Bailey, J. G. Petrovic, B. L. Krauter, and B. D. McCredie, "A clock distribution network for microprocessors," in *Symp. VLSI Circuits, 2000, Dig. Tech. Papers*, 2000, pp. 184–187.
- [4] P. J. Restle, C. A. Carter, J. P. Eckhardt, B. L. Krauter, B. D. McCredie, K. A. Jenkins, A. J. Weger, and A. V. Mule, "The clock distribution of the power4 microprocessor," in *IEEE ISSCC, 2002, Dig. Tech. Papers*, vol. 1, Feb. 2002, pp. 144–145.
- [5] E. G. Friedman, "Clock distribution networks in synchronous digital integrated circuits," *Proc. IEEE*, vol. 89, pp. 665–692, May 2001.
- [6] L. Xun, M. C. Papaefthymiou, and E. G. Friedman, "Retiming and clock scheduling for digital circuit optimization," *IEEE Trans. Computer-Aided Design*, vol. 21, pp. 184–203, Feb. 2002.
- [7] V. Gutnik and A. P. Chandrakasan, "Active GHz clock network using distributed PLLs," *IEEE J. Solid-State Circuits*, vol. 35, pp. 1553–1560, Nov. 2000.
- [8] B. A. Floyd and O. K. S. C.-M. Hung, "Intra-chip wireless interconnect for clock distribution implemented with integrated antennas, receivers, and transmitters," *IEEE Solid-State Circuits*, vol. 37, pp. 543–552, May 2002.
- [9] V. L. Chi, "Salphasic distribution of clock signals for synchronous systems," *IEEE Trans. Comput.*, vol. 43, pp. 597–602, May 1994.
- [10] A. M. Braun, V. B. Khalifin, M. H. Kwakernaak, W. F. Reichert, L. A. DiMarco, Z. A. Shellenbarger, C. M. DePriest, T. Yilmaz, P. J. Delfyett Jr., and J. H. Abeles, "Universality of mode-locked jitter performance," *IEEE Photon. Technol. Lett.*, vol. 14, pp. 1058–1060, Aug. 2002.
- [11] L. A. Jiang, M. E. Grein, H. A. Haus, and E. P. Ippen, "Noise of mode-locked semiconductor lasers," *IEEE J. Select. Topics Quantum Electron.*, vol. 7, pp. 159–167, Mar. 2001.

- [12] L. Krainer, R. Paschotta, S. Lecomte, M. Moser, K. J. Weingarten, and U. Keller, "Compact Nd:YVO₄ lasers with pulse repetition rates up to 160 GHz," *IEEE J. Select. Topics Quantum Electron.*, vol. 38, pp. 1331–1338, Oct. 2002.
- [13] Y. Liu, L. Lin, C. Choi, B. Bihari, and R. T. Chen, "Optoelectronic integration of polymer waveguide array and metal-semiconductor-metal photodetector through micromirror couplers," *IEEE Photon. Technol. Lett.*, vol. 13, pp. 355–357, Apr. 2001.
- [14] S. M. Sze, *Physics of Semiconductor Devices*, 2nd ed. New York: Wiley, 1981, ch. 13, p. 749.
- [15] P. Larsson and C. Svensson, "Measuring high-bandwidth signals in CMOS circuits," *Electron. Lett.*, vol. 29, no. 20, pp. 1761–1762, Aug. 1993.
- [16] A. V. Krishnamoorthy, A. L. Lentine, K. W. Goossen, J. A. Walker, T.K. Woodward, J. E. Ford, G. F. Aplin, L. A. D'Asaro, S. P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, M. F. Chirovsky, and D. A. B. Miller, "3-D integration of MQW modulators over active submicron CMOS circuits: 375 Mb/s transimpedance receiver-transmitter," *IEEE Photon. Technol. Lett.*, vol. 7, pp. 1288–1290, Nov. 1995.
- [17] R. Ho, K. W. Mai, and M. A. Horowitz, "The future of wires," *Proc. IEEE*, vol. 89, pp. 490–504, Apr. 2001.



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.



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.



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.



Hugo Thienpont (M'97–A'97) was born in Belgium 1961. He graduated from the Vrije Universiteit Brussels (VUB) in 1984 as an electrical engineer with majors in applied physics and applied optics.

In 1994 he became Professor in the Faculty of Applied Sciences with teaching responsibilities in 6 compulsory photonics courses. Today, he is director of research at the Laboratory for Photonics and is promoter of different photonics related research and industrial projects. His research activities comprise materials, modeling, components and devices, packaging and demonstrators for photonic interconnects in computing. He was associate editor of *Optical Engineering*, and guest editor of the special issues on "Optics in Computing" for the *Journal of the EOS* in 1998 and the OSA journal, *Applied Optics* in 2000, and for the *International Journal of Optoelectronics* on "Tutorials on Optics in Computing." He serves in technical program committees of several EOS, IEEE, OSA, and SPIE topical and annual meetings and was Technical Program Chair of the ICO/EOS topical meeting "Optics in Computing'98."

Dr. Thienpont received the ICO international prize 1999 for his contributions in the field of Optical and Photonic Information Processing and served as IEEE LEOS Distinguished Lecturer from 2001 to 2003.



Aparna Bhatnagar (S'99) received the B.S. degree in physics and the M.S. degree in electrical engineering from Stanford University in 1999 and 2001, respectively. She is currently working toward the Ph.D. degree in electrical engineering under Prof. David A. B. Miller at Stanford University, where her research involves optical interconnects and clock distribution to CMOS chips.



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.



David A. B. Miller (M'84–SM'89–F'95) received the B.Sc. degree from St. Andrews University, 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 fundamentals and applications of optics in information, sensing, switching, and processing. He has published more than 200 scientific papers, holds over 40 patents, and has received several awards for his work.

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.



Ray Chen (S'01) received the B.Sc. degree in electronics engineering from National Chiao Tung University, Hsin-Chu, Taiwan in 1997 and the M.S. degree in electrical engineering from Stanford University in 2001. From 1997 to 1999, he served in the Chinese Army as a Communication Technical Officer. He is currently working toward the Ph.D. degree in electrical engineering under Prof. David A. B. Miller at Stanford University, where his research involves optical interconnects and O/E interface module design.