

Low capacitance CMOS silicon photodetectors for optical clock injection

S. Latif · S.E. Kocabas · L. Tang · C. Debaes ·
D.A.B. Miller

Received: 27 August 2008 / Accepted: 16 December 2008 / Published online: 18 February 2009
© Springer-Verlag 2009

Abstract We have studied the response of CMOS compatible detectors fabricated in a silicon-on-sapphire (SOS) process, operated under short pulse excitation in the blue. These high speed, low capacitance detectors would be suitable for very precise, surface-normal clock injection with silicon CMOS. We characterize the capacitance of the detector structure through a combination of experimental techniques and circuit-level and electromagnetic simulations. The transit-time-limited response of the detectors is validated through pump–probe experiments. Detector response times of ~ 35 ps have been measured, and devices have capacitance as low as ~ 4 fF.

PACS 42.82.Ds · 85.40.-e · 85.60.Bt · 85.60.Gz

1 Introduction

The clock signal is used to synchronize the sequencing, transfer, and sampling of data in digital systems. In microprocessor chips, a global clock signal has to be distributed across the entire area of the chip to drive the various logic

elements that need synchronous timing. The clock distribution network must keep the clock skew (variation in clock edge arrival between different nodes) and jitter (variation in clock edge arrival time at the same node) within manageable limits. As data rates scale to higher speeds, the task of clock distribution in modern microprocessors has become increasingly complex and daunting. Maintaining jitter and skew below a constant percentage of the clock cycle is progressively harder when the clock period is scaled down. There are many factors behind this increasing complexity. First of all, for reasons described in [1], the maximum bandwidth at which electrical wires on chip can operate does not increase as the physical dimensions of the wires are scaled down. Therefore, while the transistors on chip get faster with every generation of CMOS scaling, the wires on chip do not see an increase in bandwidth, and, in fact, their contribution to the delay as a percentage of the clock period increases. This increase in delay can be alleviated by breaking up long global wires, and inserting buffering stages in between, to regenerate the clock signal. However, the buffering stages add their own delay and increase power dissipation. The buffer stages also inject noise from the power supply and add to the skew and jitter. As a consequence, the power required for clock distribution, expended in driving the clock across the global grid and switching hundreds of thousands of latches at the local level, consumes upwards of 25% of the total chip power dissipation in modern microprocessor designs [2]. Conventional clock distribution schemes use a variety of network topologies to distribute the clock across the chip. For example, a grid topology produces the lowest skew, but dissipates more power, whereas a distribution network in the shape of a balanced H-tree consumes less power, but has the potential of poorer skew performance [3]. In practice, hybrids of these techniques are used, in order to meet both the skew and power constraints, and additional

S. Latif (✉) · S.E. Kocabas · D.A.B. Miller
Department of Electrical Engineering, Ginzton Laboratory,
Stanford University, Stanford, CA 94305, USA
e-mail: slatif@stanford.edu

L. Tang
Department of Applied Physics, Ginzton Laboratory, Stanford
University, Stanford, CA 94305, USA

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

power savings are achieved by incorporating clock gating techniques, in which parts of the clock distribution network can be turned off when not in use [4]. A number of techniques have been proposed to alleviate the difficulties of designing high speed clock distribution networks. These include 3-dimensional architectures that reduce wire delay [5], standing-wave (salphasic) clock networks [6], clock distribution with RF waves coupling to on-chip antennas [7], clocks generated locally by multiple mutually synchronized oscillators on chip [8], and optical clock distribution [9].

The use of optics to generate and distribute clock signals on chip is an approach that offers significant performance advantages. For interconnects in general, the introduction of optics can overcome many fundamental limitations in interconnects, such as frequency dependent cross-talk, loss and dispersion, and impedance matching concerns. Summaries of the advantages offered by optical interconnects are given in [10] and [11]. Optical clock injection has been proposed for clock distribution in microprocessors [12], optical triggering of analog-to-digital converters [13], and for the generation of multi-phasic clocks for multiplexing or demultiplexing circuits [14].

In this paper, we will discuss optical detectors for a receiver-less scheme employing clock injection on chip with short optical pulses. A detailed analysis of this receiver-less clock distribution scheme is presented in [12]. The receiver-less approach removes the amplifier traditionally placed in front of the photo-detector in the receiving end of optical links; instead, the photo-detector directly produces enough voltage swing to trigger the logic circuits that need clocking. For example, a rail-to-rail clock signal can be generated by shining short pulse laser beams on a totem pole of photo-detectors [12]. This approach requires high repetition rate (10's of GHz) mode-locked laser sources. Such lasers have now been demonstrated [15] with extremely low amounts of jitter. This receiver-less approach removes the additional power, delay, skew, and jitter contribution that would arise from any electronic receiver amplifier between the detector and electronic latch input. Furthermore, if we were able to generate the clock signal optically at the local level (which would require simultaneous optical clocking of hundreds of thousands of latches on chip), we would be able to remove the need for local clock buffering and amplification. This optical approach can potentially lead to significant reduction in clock power dissipation and skew, since most clock power dissipation and skew occurs at the local level (around 80% in recent microprocessors [3]), where the local clock signal is generated by buffering and amplification.

It is possible to uniformly distribute a light beam onto very large numbers of points on chip using a diffractive optical element. In such a distribution, the required optical energy has to be minimized, however. Since the receiver-less approach removes the voltage amplification provided by the

optical receiver, it must provide a way to generate full rail voltage swing at the optical insertion points without consuming too much optical power. Therefore, photo-detector capacitance becomes a crucial parameter, since detector capacitance determines the amount of optical power or energy required for a voltage swing.

Low capacitance detectors are desirable for several reasons. First, they lead to lower power dissipation (for a given voltage swing). Second, if receiver amplifier circuits are used, then low detector capacitance results in low power dissipation and latency for the receiver, and better noise performance [11]. Third, low capacitance detectors are smaller in size, and hence consume less area on chip. Furthermore, having detectors fabricated entirely in silicon in a CMOS process is particularly attractive because it may allow extremely low capacitances and high speeds in small, integrated devices. Detector speeds in excess of 10 GHz are desirable for clock injection.

In this paper, we present a characterization of the behavior of fully-integrated detectors capable of very low capacitances and high speeds in a CMOS process. We designed a test chip in a commercial silicon-on-sapphire (SOS) CMOS process, developed by Peregrine Semiconductor. Previous work in our group showed that the SOS process was a feasible platform for integrating photodetectors with CMOS electronics, and we measured the large signal rise-time of a 6.2 μm finger-spacing detector [16]. In this paper, we describe a full optical and electrical characterization of various different lateral P-I-N photodetectors on the test chip. To fully understand the time response mechanism, we measured rise times for three different finger spacings using optical pump-probe measurements with integrated modulators. This technique allows high-sensitivity small-signal measurements with excellent time resolution. As we discuss below, the fastest response was 35 ps for the 1.2 μm finger spacing, corresponding to a bandwidth of about 15 GHz. The measured responsivity for a detector with 1.2 μm finger spacing was 0.01 A/W. We predict that the responsivity can be increased by an order of magnitude with a full optimization of the active region layer thickness and contact metal area coverage.


Since detector capacitance is such a crucial parameter here, it is vital to confirm low capacitance operation with high confidence. To this end, we use a ring-oscillator circuit technique to estimate various photodetector capacitances. A combination of experimental data and circuit simulations allows us to extract well calibrated estimates of photodetector capacitances. These estimates are compared with Finite Element Method (FEM) simulations of the detector structure in order to better understand the physical origins of device capacitance for this particular detector design. We find that the capacitance numbers are much different from a simple "parallel plate" type of estimate, and that the capacitances

of our smaller detectors to be of the order of a few femtofarads.


The format of the paper is as follows. In Sect. 2, we describe the details of the structures we use. In Sect. 3, we present the ring-oscillator capacitance measurement technique and results, and compare these with finite element simulations. In Sect. 4, we report on detector responsiveness, and study the transient response of the photodetectors through the use of a pump–probe technique. Section 5 concludes the paper.

2 Device structure

The physical structure of the photodetectors fabricated on this chip is illustrated in Fig. 1. The 70 nm thin active silicon layer forms a lateral P–I–N structure. The width of the intrinsic region between the heavily doped P+ and N+ contact areas is set by the finger spacing parameter, which we varied between 1.2 μm and 6 μm. Multiple fingers are combined in parallel to extend the detector area in the lateral direction. We operate these detectors at blue wavelengths (~425 nm). The absorption depth at that wavelength is approximately 180 nm, so reasonably efficient photodetection is possible. Most importantly, the absence of a semiconductor substrate results in extremely low capacitance devices. As a representative number, we calculate below (after extrapolating from simulations and experimental data), that a 5 μm × 5 μm size detector with a finger spacing (1.2 μm) corresponding to the fastest rise time has a capacitance of roughly 4 fF. Detectors with bigger finger spacings would have even lower capacitance, but with lower speeds.



a



b

Finger Spacing = 2 μm (3 fingers)

Fig. 1 Photodetector structure, (a) schematic of side view, (b) picture of top view

3 Capacitance measurement

We estimate photodetector capacitance using a combination of on-chip circuit techniques and electromagnetic simulations.

3.1 Measurement techniques

We have implemented a technique to measure the capacitance of our photodetector structures on chip. The technique is based on an inverter-based ring-oscillator circuit [17]. An odd number of current-starved digital inverters connected in a ring are known to oscillate and produce a square wave output, given that the individual inverters are fed with enough current to switch between logic levels. The frequency of oscillation is a function of the propagation delay of a transition edge as it propagates once around the loop. This delay, in turn, depends on the capacitive load that each inverter has to charge or discharge. If we were to load each inverter with an unknown capacitance and measure the resulting oscillation frequency, then we could potentially obtain an estimate of the load capacitance. However, this technique requires careful calibration to make sure we are getting a reliable and meaningful estimate. We calibrate the technique by comparing the ring-oscillator output frequency to SPICE simulations of the circuit with different process corners. Once there is good agreement between simulation and measured ring-oscillator data, capacitance estimates are obtained by varying the load capacitance in the simulation domain and comparing against measured data obtained for each ring-oscillator loaded with a detector with different finger spacing. More details on the circuit are given in [18]. For comparison, and to confirm our physical intuition about the capacitance of the detector structure, we performed 2-D static electromagnetic simulations with the help of the COMSOL tool. The detector can be envisioned as a two-conductor “waveguide” formed by the two metal contact lines and the two doped semiconductor regions, wrapped around in a multi-fingered configuration. The p-contact metal line is held at the same electric potential as the P+ doped region, and similarly for the n-contact metal line and the N+ doped region. The tool is used to extract the capacitance per unit length of a 2-D slice composed of a single finger pair of the detector structure. Figure 2 shows a drawing of the top view of a typical detector layout, and also the corresponding cross-sectional view that forms the 2-D “waveguide” simulated in COMSOL.

The total detector capacitance is calculated by multiplying the capacitance per unit length by the total “waveguide” length. The capacitance of the corner regions connecting adjacent metal fingers is also estimated using the same waveguide approximation. The capacitance numbers as a function of finger spacing for both the ring-oscillator and




Fig. 2 (a) Top view of detector layout, and (b) corresponding 2-D slice (cross-section) (the same layer dimensions were used in finite-element simulations). Not to scale

COMSOL method are given in Table 1. The simulated “waveguide” capacitance per unit length for each finger spacing is also given for reference. All detectors are approximately $50 \mu\text{m} \times 50 \mu\text{m}$ in size in these measurements and simulations.

3.2 Discussion

We can make a few observations about the capacitance estimates obtained above. First, the detector capacitance decreases with increasing finger spacing, as expected from a parallel plate/fringing capacitance model. Second, these calculated and measured capacitances are much higher than those one would obtain by treating the capacitance as just the “plane-parallel” capacitance between the detector fingers, showing the importance of fully including fringing capacitance effects. We therefore conclude that simple plane parallel models are not helpful in estimating capacitance for this structure, and a full “waveguide” view analysis has to be adopted. With this good understanding of the capacitance

Table 1 Capacitance estimates

Finger spacing (μm)	Ring-oscillator method	COMSOL simulations	Capacitance/unit length pF/m
	Capacitance (fF)	Capacitance (fF)	
1.2	135	110	188.2
2.0	85	74	153.7
3.0	80	62	132.5
6.0	40	33	106.2

mechanisms for this kind of detector structure, we can calculate the capacitance for detector sizes used in practical applications. For a $5 \mu\text{m} \times 5 \mu\text{m}$ detector with $1.2 \mu\text{m}$ finger spacing (the finger-spacing with the fastest time response), the capacitance is calculated to be ~ 4 fF. Since this comparable to the gate capacitance of a small size transistor in the $0.25 \mu\text{m}$ process we are using, we can conclude that this detector design is suitable for high speed, low capacitance applications.

4 Optical response measurement

Photodetectors for clock-injection applications are required to be fast (e.g., >10 GHz), and efficient. We therefore need to measure responsivity, model the speed of operation for our detectors, and measure detector speed to compare against theory.

4.1 Responsivity

The responsivity varies with detector finger spacing. The smaller the finger spacing, the greater the fraction of the detector surface area blocked by metal contacts. For the highest speed (and the smallest finger spacing of $1.2 \mu\text{m}$) detectors, we measured a responsivity of 0.01 A/W at a wavelength of 430 nm , which corresponds to an external quantum efficiency of 3%. This number is low, but does not arise from any fundamental design limitations. For example, for the $1.2 \mu\text{m}$ finger spacing detectors, 70% of the detector surface area is covered with metal used for contact, which blocks the same percentage of incoming light from being converted into photocurrent. However, this amount of metal coverage was required because of the layout design rules and minimum feature sizes of the particular CMOS process used, so in principle it should be easy to fabricate these kinds of detectors in a CMOS process with a much smaller portion of detector surface area covered by metal. Similarly, the silicon absorbing region, being 70 nm thick, can only absorb 30% of incoming light, given an absorption depth of 180 nm at a wavelength of 430 nm [19]. Again, this absorption could be brought close to 80% without sacrificing on speed, say for

a silicon layer thickness of 300 nm. Further optimizations, such as the addition of an anti-reflection coating layer and careful engineering of silicon-insulator interface trap density, would ensure much higher responsivity without losing the advantages of easy CMOS integration.

4.2 Modeling detector speed

To model the time response of the detector, we start with some simplifying assumptions. The photogeneration is assumed to occur only in the intrinsic region of thickness d , which is taken to be fully depleted. Any diffusive transport mechanisms are assumed negligible. The photogenerated carriers are assumed to travel with the saturated drift velocity, v_d (taken here for simplicity to be the same for electrons and holes), under the influence of the applied electric field. We start with a short pulse of light injecting total charge Q_{opt} uniformly into the depletion region. We first calculate the charge collected at the edge of the depletion region as a function of time. The charge dQ passing through one edge of the depletion region in time dt can be calculated by using the flux of carriers passing through the depletion region boundary (Fig. 3)

$$dQ = \rho A v_d dt,$$

where ρ is the charge density of one type of carrier (electrons or holes) in the depletion region, and A is the cross-sectional area of the diode.

Since,

$$\rho = Q_{opt}/Ad,$$

$$dQ = \frac{Q_{opt}}{Ad} A v_d dt = \frac{Q_{opt}}{d} v_d dt.$$

So

$$\frac{dQ}{dt} = \frac{Q_{opt}}{d} v_d,$$

and hence

$$Q(t) = \frac{Q_{opt}}{d} v_d t.$$




Fig. 3 Modeling voltage time response through a charge flux calculation

Now we compute the voltage induced by the collected charge:

$$V(t) = V(t_0) + \frac{Q(t)}{C_d},$$

where $V(t_0) = 0$, and the depletion region capacitance $C_d = \frac{\epsilon_{Si} A}{d}$. So,

$$V(t) = \frac{\frac{Q_{opt}}{d} v_d t}{\frac{\epsilon_{Si} A}{d}} = \frac{Q_{opt} v_d}{\epsilon_{Si} A} t.$$

Therefore, we find that ideally, based on this simple model, the voltage response of the photodetector is a linear function of time, until the point at which all the photogenerated charge has been collected, after which it stays constant.

4.3 Pump-probe measurements of transient response

The optical pump-probe technique can be used to measure very fast transient signals on chip. It is possible to achieve quite high temporal resolution while maintaining good linearity and sensitivity, provided that the electro-absorption modulator used to sample the signal (in our implementation) is biased optimally. The pump-probe technique used here relies on the presence of Multiple Quantum Well (MQW) GaAs modulators on chip (connected to various nodes on the chip via hybrid integration techniques). The reflectivity of these modulators varies as a function of applied voltage, and is probed by an incident mode-locked probe beam. The voltage transients that we want to measure are generated by a mode-locked pulse beam, incident on an SOS detector. By varying the time delay between the pulse and the probe beam, and measuring the probe power reflected from the MQW modulator, we are able to map out fast voltage transients on chip with picosecond time-scale resolution. A 850 nm short pulse laser is used to generate both the 850 nm probe signal, and the 425 nm pump signal (through second harmonic generation). Figure 4 shows a schematic of the optical setup. Further experimental details on this technique can be found in [20].




Fig. 4 Optical setup for pump-probe experiment




Fig. 5 (a) Photodetector speed test—schematic of the structure under test, showing device connectivity (b) transient response for 1.2 μm , 2.0 μm , and 3.0 μm finger spacing detectors

The detector is biased in series with a MQW modulator as shown in Fig. 5. The detector is hit with a pump pulse, which causes the voltage at the middle node (i.e., between the detector and the modulator) to rise, and this time response is measured by probing the MQW reflectivity as explained earlier. To ensure that the modulator was operating in a linear fashion, we kept the pump power at a low enough level, such that the resulting voltage change across the modulator was around a hundred millivolts or less. For such a small signal excitation, the absorption vs. voltage characteristic of the modulator can be safely assumed to be linear. The middle node voltage is reset periodically to a determined value by an electrical pulse reset signal, which is synchronized with the mode-locked laser repetition rate. Figure 5 displays the biasing arrangement and the pump–probe response curves for detectors with three different finger spacings.

One common feature present in all the curves is the presence of two regimes: a fast-rising approximately linear portion, followed by a much slower (possibly exponential) tail. We believe the linear portion corresponds to the transit-time-limited response of the detector. The character of the slow tail changes as the incident beam is moved in and out of focus on the surface of the detector, which leads us to believe

that the slow part of the response is caused by some of the light leaking into the heavily doped silicon contact areas. We expect that the resulting photogenerated carriers from those contact regions are collected through a slow diffusion process. To minimize this part of the voltage response, the beam should be focused as best as possible so that its spatial extent is limited to the area between the metal fingers; then there should be minimum light absorbed in the heavily doped silicon region. Were the metal fingers placed directly above the silicon contact areas (instead of the metal being $\sim 1 \mu\text{m}$ above as in our case), we expect this problem would be eliminated or substantially reduced.

Assuming an applied field of 10^4 V/cm , the drift velocities for electrons and holes in silicon are nominally $6 \times 10^6 \text{ cm/s}$, and $3 \times 10^6 \text{ cm/s}$, respectively [21]. If the voltage rise time is limited by the time it takes for holes, being the slower carriers, to transit across the intrinsic region, then we can estimate transit times of 40 ps, 67 ps, and 100 ps for the 1.2 μm , 2 μm , and 3 μm finger spacing detectors, respectively. From Fig. 5, the duration of the initial approximately linear portion of the curves is approximately 35 ps, 70 ps, and 90 ps for the corresponding detectors. Therefore, we find that the experimentally determined rise-times correspond well to the transit time duration expected for the carriers (moving with the drift velocity of the slower carrier) to traverse the extent of the intrinsic region.

5 Conclusions

To conclude, in this paper we have presented a full characterization of silicon-on-sapphire low capacitance photodetectors aimed at clock injection applications. We have shown that through the use of a calibration procedure that takes into account process corner variation, we can now obtain a reliable estimate for the capacitance of detectors with varying finger spacings. From 2-D static electromagnetic simulations, we have shown that capacitance of this structure corresponds to that between the combined metal and doped semiconductor regions treated as the plates of the capacitor, both arranged in a two-conductor waveguide geometry. The capacitance decreases with increased finger spacing, as expected. We estimate a capacitance of $\sim 4 \text{ fF}$ for our smallest detectors. The measured responsivity was 0.01 A/W for the fastest detector, which could be increased significantly with modifications to the current CMOS process. We have also obtained an estimate of the temporal response of the SOS detectors through the pump–probe technique, and we observe response as fast as $\sim 35 \text{ ps}$, close to the theoretically predicted time response from carrier transit times. Therefore, we have verified the operation of such low capacitance detectors and shown that such devices should be suitable for very precise optical clock injection with silicon

CMOS. Extrapolating from the results obtained here, we can be more confident in predicting the capacitance and bandwidth performance of this kind of detector structure, even for future nanometallic structures with physical size scaled down to dimensions of an order of 100 nm or so. Such sub-wavelength scale photodetectors [22, 23] offer the promise of optoelectronic integration at the scale of transistor dimensions, and coupled with resonantly enhanced detection techniques, would result in significant power, speed, and area gains.

Acknowledgements The authors would like to thank Peregrine Semiconductor for fabricating the test chip in their silicon-on-sapphire process through the USC–DARPA–Peregrine COOP. The authors acknowledge the support of the Interconnect Focus Center, one of five research centers funded under the Focus Center Research Program, a DARPA and Semiconductor Research Corporation program, and of the AFOSR “Plasmon Enabled Nanophotonic Circuits” MURI Program.

References

1. D.A.B. Miller, H.M. Ozaktas, Special issue on parallel computing with optical interconnects. *J. Parallel Distrib. Comput.* **41**, 42–52 (1997)
2. P. Mahoney, E. Fetzer, B. Doyle, S. Naffziger, *Int. Solid State Circuits Conf.*, 2005. *Dig. Of Tech. Papers. ISSCC. 2005 IEEE Int.*, San Francisco, CA, pp. 292–599 (2005)
3. N. Ranganathan, N.P. Jouppi, HP Labs Technical Reports, Palo Alto, CA HPL-2007-163 (2007)
4. C. Yeh, G. Wilke, H. Chen, S. Reddy, H. Nguyen, T. Miyoshi, W. Walker, R. Murgai, *7th International Symposium on Quality Electronic Design (ISQED)* (2006)
5. K.B. Banerjee, S.J. Souri, P. Kapur, K.C. Saraswat, *Proc. IEEE* **89**, 602–633 (2001)
6. F. O’Mahony, C.P. Yue, M.A. Horowitz, S.S. Wong, *IEEE J. Solid-State Circuits* **38**, 1813–1820 (2003)
7. K.K. O, K. Kim, B.A. Floyd, J.L. Mehta, H. Yoon, C.-M. Hung, D. Bravo, T.O. Dickson, X. Guo, R. Li, N. Trichy, J. Caserta, W.R. Bomstad, J. Branch, D.-J. Yang, J. Bohorquez, E. Seok, L. Gao, A. Sugavanam, J.-J. Lin, J. Chen, J. Brewer, *IEEE Trans. Electron. Devices* **52**, 1312–1323 (2005)
8. V. Gutnik, A.P. Chandrakasan, *IEEE J. Solid-State Circuits* **35**, 1553–1560 (2000)
9. J.W. Goodman, F. Leonberger, S.-Y. Kung, R.A. Athale, *Proc. IEEE* **72**, 850–866 (1984)
10. D.A.B. Miller, Physical reasons for optical interconnection. *Int. J. Optoelectron.* **11**, 155–168 (1997)
11. D.A.B. Miller, *Proc. IEEE* **88**, 728–749 (2000)
12. C. Debaes, A. Bhatnagar, D. Agarwal, R. Chen, G.A. Keeler, N.C. Helman, H. Thienpont, D.A.B. Miller, *IEEE J. Sel. Top. Quantum Electron.* **9**, 400–409 (2003)
13. R. Urata, L.Y. Nathawad, R. Takahashi, K. Ma, D.A.B. Miller, B.A. Wooley, J.S. Harris Jr., *IEEE J. Lightwave Technol.* **21**, 3104–3115 (2003)
14. D.A.B. Miller, A. Bhatnagar, S. Palermo, A. Emami-Neyestanak, M.A. Horowitz, *Int. Solid State Circuits Conf.*, 2005. *Dig. Of Tech. Papers. ISSCC. 2005 IEEE Int.*, San Francisco, CA, pp. 86–87 (2005)
15. F. Quinlan, S. Gee, S. Ozharar, P.J. Delfyett, *Opt. Lett.* **31**, 2870–2872 (2006)
16. A. Bhatnagar, S. Latif, C. Debaes, D.A.B. Miller, *J. Lightwave Technol.* **22**, 2213–2217 (2004)
17. A.V. Krishnamoorthy, T.K. Woodward, R.A. Novotny, K.W. Goossen, J.A. Walker, A.L. Lentine, L.A. D’Asaro, S.P. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L.M.F. Chirovsky, G.F. Aplin, R.G. Rozier, F.E. Kiamilev, D.A.B. Miller, *Electron. Lett.* **31**, 1917–1918 (1995)
18. C. Debaes, *Intra Multi-Chip Module Interconnects*. Ph.D. dissertation, Dept. of Appl. Phys. and Photonics, Vrije Universiteit Brussel, Brussels, Belgium (2003)
19. E.D. Palik, *Handbook of Optical Constants of Solids*, vol. 1 (Academic Press, San Diego, 1985)
20. G.A. Keeler, D. Agarwal, C. Debaes, B.E. Nelson, N.C. Helman, H. Thienpont, D.A.B. Miller, *IEEE Photonics Technol. Lett.* **14**, 1214–1216 (2002)
21. R.F. Pierret, *Semiconductor Device Fundamentals* (Addison-Wesley, Reading, 1996). Chap. 3, Fig. 3.4
22. L. Tang, D.A.B. Miller, A.K. Okyay, J.A. Matteo, Y. Yuen, K.C. Saraswat, L. Hesselink, *Opt. Lett.* **31**, 1519–1521 (2006)
23. L. Tang, S.E. Kocabas, S. Latif, A.K. Okyay, D.-S. Ly-Gagnon, K.C. Saraswat, D.A.B. Miller, *Nature Photonics* **2**, 226–229 (2008)