Low-power, 2×2 silicon electro-optic switch with 110-nm bandwidth for broadband reconfigurable optical networks

Joris Van Campenhout, William M. J. Green, Solomon Assefa, and Yurii A. Vlasov

IBM T.J. Watson Research Center, Yorktown Heights, NY 10598

jvancam@us.ibm.com

Abstract: We present an ultra-broadband Mach-Zehnder based optical switch in silicon, electrically driven through carrier injection. Crosstalk levels lower than -17 dB are obtained for both the 'on' and 'off' switching states over an optical bandwidth of 110 nm, owing to the implementation of broadband 50 % couplers. Full 2×2 switching functionality is demonstrated, with low power consumption (~ 3 mW) and a fast switching time (< 4 ns). The utilization of standard CMOS metallization results in a low drive voltage (~ 1 V) and a record-low $V_\pi L$ (~ 0.06 V·mm). The wide optical bandwidth is maintained for temperature variations up to 30 K.

OCIS codes: (130.4815) Optical switching devices; (250.6715) Switching

References and links

1. A. F. Benner, M. Ignatowski, J. A. Kash, D. M. Kuchta, and M. B. Ritter, "Exploitation of optical interconnects in future server architectures," IBM J. Res. Dev. **49**(4-5), 755–775 (2005).
2. T. Barwicz, H. Byun, F. Gan, C. W. Holzwarth, M. A. Popovic, P. T. Rakich, M. R. Watts, E. P. Ippen, F. X. Kartner, H. I. Smith, J. S. Orcutt, R. J. Ram, V. Stojanovic, O. O. Olubuyide, J. L. Hoyt, S. Spector, M. Geis, M. Grein, T. Lyszczarz, and J. U. Yoon, "Silicon photonics for compact, energy-efficient interconnects [Invited]," J. Opt. Netw. **6**(1), 63–73 (2007).
3. C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. W. Holzwarth, M. A. Popovic, H. Q. Li, H. I. Smith, J. L. Hoyt, F. X. Kartner, R. J. Ram, V. Stojanovic, and K. Asanovic, "Building Many-Core Processor-to-DRAM Networks with Monolithic CMOS Silicon Photonics," IEEE Micro **29**(4), 8–21 (2009).
4. J. Ahn, M. Fiorentino, R. G. Beausoleil, N. Binkert, A. Davis, D. Fattal, N. P. Jouppi, M. McLaren, C. M. Santori, R. S. Schreiber, S. M. Spillane, D. Vantrease, and Q. Xu, "Devices and architectures for photonic chip-scale integration," Appl. Phys. A **95**(4), 989–997 (2009).
5. A. V. Krishnamoorthy, R. Ho, X. Z. Zheng, H. Schwetman, J. Lexau, P. Koka, G. L. Li, I. Shubin, and J. E. Cunningham, "Computer Systems Based on Silicon Photonic Interconnects," Proc. IEEE **97**(7), 1337–1361 (2009).
6. A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," IEEE Trans. Comput. **57**(9), 1246–1260 (2008).
7. H. F. Hamann, A. Weger, J. A. Lacey, Z. G. Hu, E. Cohen, and J. Wakil, "Hotspot-limited microprocessors: Direct temperature and power distribution measurements," IEEE J. Solid-State Circuits **42**(1), 56–65 (2007).
8. Y. Vlasov, W. M. J. Green, and F. Xia, "High-throughput silicon nanophotonic wavelength-insensitive switch for on-chip optical networks," Nature Photon. **2**(4), 242–246 (2008).
9. B. G. Lee, A. Biberman, P. Dong, M. Lipson, and K. Bergman, "All-optical comb switch for multiwavelength message routing in silicon photonic networks," IEEE Photon. Technol. Lett. **20**(9-12), 767–769 (2008).
10. W. M. J. Green, M. J. Rooks, L. Sekaric, and Y. A. Vlasov, "Ultra-compact, low RF power, 10 Gb/s silicon Mach-Zehnder modulator," Opt. Express **15**(25), 17,106–17,113 (2007).

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

11. J. Van Campenhout, W. M. J. Green, and Y. A. Vlasov, "Design of a digital, utra-broadband optical switch for reconfigurable optical networks-on-chip," Opt. Express (to be published).
12. R. A. Soref and B. R. Bennett, "Electrooptical Effects in Silicon," IEEE J. Quantum Electron. **23**(1), 123–129 (1987).
13. S. L. Tsao, H. C. Guo, and Y. J. Chen, "Design of a 2 x 2 MMI MZI SOI electro-optic switch covering C band and L band," Microw. Opt. Technol. Lett. **33**(4), 262–265 (2002).
14. K. Jinguji, N. Takato, A. Sugita, and M. Kawachi, "Mach-Zehnder Interferometer Type Optical Wave-Guide Coupler with Wavelength-Flattened Coupling Ratio," Electron. Lett. **26**(17), 1326–1327 (1990).
15. T. Kitoh, N. Takato, K. Jinguji, M. Yasu, and M. Kawachi, "Novel Broad-Band Optical Switch Using Silica-Based Planar Lightwave Circuit," IEEE Photon. Technol. Lett. **4**(7), 735–737 (1992).
16. C. Lavoie, F. M. d'Heurle, C. Detavernier, and C. Cabral, "Towards implementation of a nickel silicide process for CMOS technologies," Microelectron. Eng. **70**(2-4), 144–157 (2003).
17. S. Assefa, C. Jahnes, and Y. Vlasov, "CMOS compatible integrated dielectric optical waveguide coupler and fabrication," US patent application, 12/164580 (2008).
18. P. Dumon, G. Priem, L. R. Nunes, W. Bogaerts, D. Van Thourhout, P. Bienstman, T. K. Liang, M. Tsuchiya, P. Jaenen, S. Beckx, J. Wouters, and R. Baets, "Linear and nonlinear nanophotonic devices based on silicon-on-insulator wire waveguides," Jpn. J. Appl. Phys. 1 **45**(8B), 6589–6602 (2006).
19. J. Van Campenhout, P. Rojo-Romeo, D. Van Thourhout, C. Seassal, P. Regreny, L. Di Cioccio, J. M. Fedeli, and R. Baets, "Thermal characterization of electrically injected thin-film InGaAsP microdisk lasers on Si," J. Lightw. Technol. **25**(6), 1543–1548 (2007).
20. F. N. Xia, M. Rooks, L. Sekaric, and Y. Vlasov, "Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects," Opt. Express **15**(19), 11,934–11,941 (2007).

1. Introduction

The processing power of highly parallel computing systems is increasingly being constrained by the limitations of the electrical interconnects used in such systems. As a result, high-bandwidth, low-power optical interconnects are replacing copper-based electrical interconnects in these systems at ever shorter interconnect lengths [1]. At the chip level, silicon-based optical interconnects have been proposed to enable high-bandwidth and power-efficient communication, both for chip-to-chip as well as for global on-chip interconnects [2]. Optical network-on-chip (ONoC) architectures capable of providing reconfigurable communication paths between the processor cores and memory systems on a chip multiprocessor (CMP) are currently being studied. Both wavelength-selective [3, 4, 5] as well as non-wavelength-selective routing schemes [6] are being considered. In order to obtain an ONoC with a sufficiently high bandwidth capacity, the use of aggressive wavelength-division multiplexing (WDM) is deemed inevitable.

A key device for non-wavelength-selective, reconfigurable ONoCs is a broadband, 2×2 silicon optical switch, capable of simultaneously routing all WDM channels comprising a data stream in a waveguide of the ONoC. The optical bandwidth of these switches directly defines the number and spacing of WDM channels that can be simultaneously switched. As such, switches with a very wide (aggregate) optical bandwidth with low crosstalk are highly desired, as they enhance the throughput capacity of the network. In order to garantuee fast reconfiguration as well as scalability of the ONoC, the switches should have a fast switching speed ($\sim$ns), low power consumption ($\sim$mW), a relatively small footprint ($\sim 0.01\,\mathrm{mm}^2$), and a low drive voltage ($\sim 1\,\mathrm{V}$). Furthermore, the switches should be insensitive to on-chip temperature fluctuations [7, 8].

Previously, resonant ring-based switches have been proposed as broadband switches for routing messages in on-chip interconnection networks [8, 9]. Alternatively, non-resonant switches based on a Mach-Zehnder interferometer could be used for this purpose. In this paper, we present the design, fabrication and characterization of a silicon 2×2 Mach-Zehnder electro-optic switch with an optical bandwidth of 110 nm. This wide optical bandwidth is realized by implementing broadband 50% couplers. Switching is obtained by injecting free carriers into a p-i-n diode phase shifter, with power consumption as low as 3.1 mW and a switching time of

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

4 ns. Furthermore, the wide optical bandwidth is maintained for temperature variations of up to 30 K peak to peak.

2. Design of an ultra-broadband electro-optic switch in silicon

2.1. The conventional 2×2 Mach-Zehnder (MZ) switch

An appropiate starting point for the design of a wavelength-insensitive switch is the conventional, balanced 2×2 Mach-Zehnder (MZ) interferometer, as depicted in Fig. 1a. This basic interferometer consists of two 50 % directional couplers and an intermediate phase-tuning section. The spectral response of this four-port structure is given by the four transmittance spectra $T_{ij}(\lambda) = |S_{ij}(\lambda)|^2$, where $S_{ij}(\lambda)$ are the wavelength-dependent complex transfer functions of the optical field from input port a_i to output port b_j, with $i, j = 1, 2$. In the switch 'off' state, the optical signal arriving at the input port a_1 (a_2) will be fully transferred to the b_2 (b_1) output port. In the switch 'on' state, a π phase shift is induced in the active phase shifter, causing the optical input signal from a_1 (a_2) to destructively interfere at the b_2 (b_1) output port, and therefore switching it to the b_1 (b_2) output port.

The optical bandwidth of this switch type is limited by the wavelength dependence of its 50 % couplers. Indeed, maximum extinction at the output ports can only be obtained when the two interfering light beams have equal amplitude, which assumes perfect 50 % power splitting in the directional couplers. This condition is typically hard to achieve over a wide wavelength range. To illustrate this bandwidth limitation for MZ switches implemented on the SOI platform, we calculated the spectral response of a MZ switch built from silicon rib waveguides with cross-sectional dimensions of 500×220 nm^2, including a 50-nm-thick slab at the bottom of the waveguide, as shown in Fig. 2b. The thin bottom slab serves the purpose of accomodating the lateral, 200-μm-long p-i-n diode in the phase-tuning section of the switch, which is used to inject carriers into the waveguide core in the switch 'on' state [10]. More details on the simulation method can be found in the appendix of [11].

The $T_{ij}(\lambda)$ transmittance spectra were calculated for a MZ switch built from such waveguides using the transfer-matrix method, for the lowest-order transverse-electric (TE) polarized guided mode and assuming a gap width in the directional coupler of 300 nm and a phase-shifter length $L = 200\,\mu$m. For the switch 'on' state, these calculations included the effect of an injected-carrier density N on the real part as well as on the imaginary part of the refractive index of the silicon waveguide core [12], such that not only the phase shift $\Delta\phi(N)$ but also the associated free-carrier absorption (FCA) loss $\alpha(N)$ was taken into account. No optical losses due to bending or scattering were included. As such, the obtained switching response can be considered the intrinsic response of the device.

The results of these simulations are shown in Fig. 1c. It can be seen that in the switch 'off' state, the optical bandwidth with -20 dB crosstalk is limited to 35 nm. The 'on'-state $T_{ij}(\lambda)$ transmittance spectra are shown as well. This state is reached for an injected-carrier density $N \simeq 1.44 \times 10^{18}\text{cm}^{-3}$, which is equivalent with a π phase shift. The optical absorption loss resulting from these free carriers limits the 'on'-state crosstalk to about -20 dB, and produces an 'on'-state insertion loss of about 0.8 dB.

2.2. Design of a wavelength-insensitive Mach-Zehnder (WIMZ) switch

The optical bandwidth of the MZ switch can be improved by reducing the wavelength sensitivity of the 50 % couplers. This can be done by reducing the gap width of the directional coupler, or by using 2×2 multi-mode interference-based power splitters [13]. However, such approaches will most likely result in tighter tolerances during fabrication or increased optical insertion loss. Alternatively, wavelength-insensitive couplers can be built from two directional couplers with an intermediate phase delay, as was shown by Jinguji et al. [14]. Following this

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

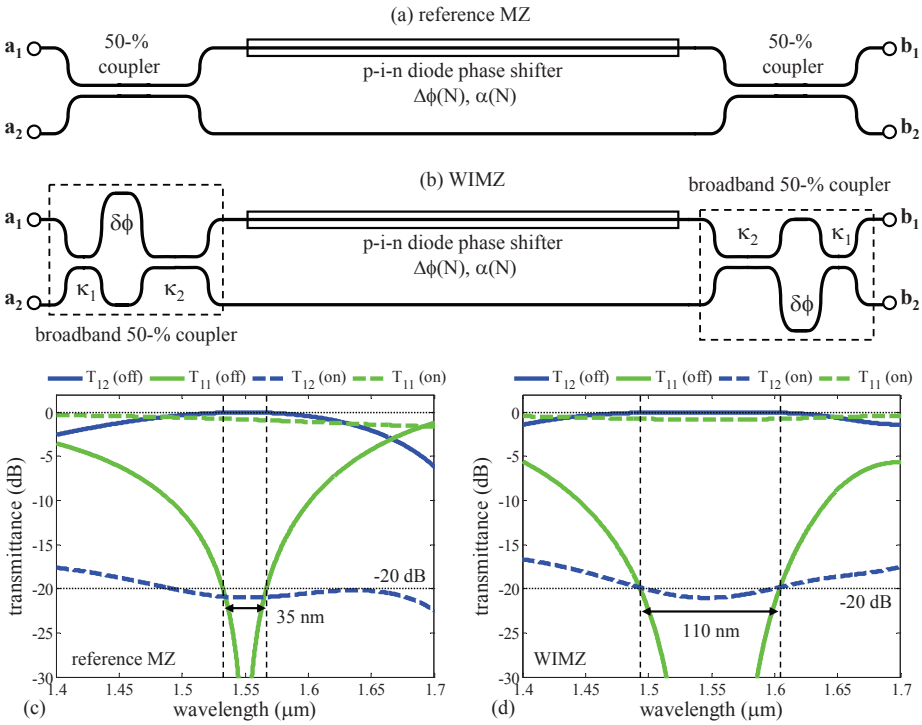

Fig. 1. (a) Schematic of the reference MZ switch. (b) Schematic of the wavelength-insensitive MZ switch, showing the broadband 50 % couplers consisting of two directional couplers with an intermediate phase delay. (c) Simulated transmittance spectra for the reference MZ switch, illustrating the limited optical bandwidth of 35 nm. (d) Transmittance spectra of the wavelength-insensitive MZ switch, showing 110-nm optical bandwidth.

Table 1. Overview of the simulated optical bandwidth (BW) with maximum crosstalk XT^{max} and FCA-induced insertion loss (IL^{FCA}), for both the 'off' and 'on' state of the MZ and WIMZ switch.

	BW (nm)	XT^{max}_{off} (dB)	IL^{FCA}_{off} (dB)	XT^{max}_{on} (dB)	IL^{FCA}_{on} (dB)
MZ	35	-20	0	-20	0.8
WIMZ	110	-20	0	-19	0.8

approach, we designed a wavelength-insensitive 50 % coupler implemented in the sub-micron silicon rib waveguides, as shown in Fig. 2b. It consists of two directional couplers with power-coupling coefficients $\kappa_1 = 0.4$ and $\kappa_2 = 0.8$ respectively, and a phase delay $\delta\phi = 0.54\pi$, as illustrated in Fig. 1b. This is equivalent with directional-coupler lengths of 15 μm and 24.5 μm respectively, and a length imbalance of 160 nm in the phase-delay section of the coupler.

When these broadband couplers are implemented in the MZ interferometer, following a point-symmetric configuration [15], a wavelength-insensitive MZ (WIMZ) switch is obtained. Indeed, the calculated -20-dB bandwidth in the 'off' state is improved by more than a factor of three to 110 nm, as is illustrated by the 'off'-state transmittance spectra T_{11} and T_{12} shown in Fig. 1d. In the 'on' state, a crosstalk level lower than -19 dB is obtained in the same spectral window, as well as an insertion loss of 0.8 dB. It should be noted that further improvement of the optical bandwidth of the WIMZ switch may be obtained by further optimization of the coupling coefficients and phase delay of the broadband couplers. The simulated performance

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

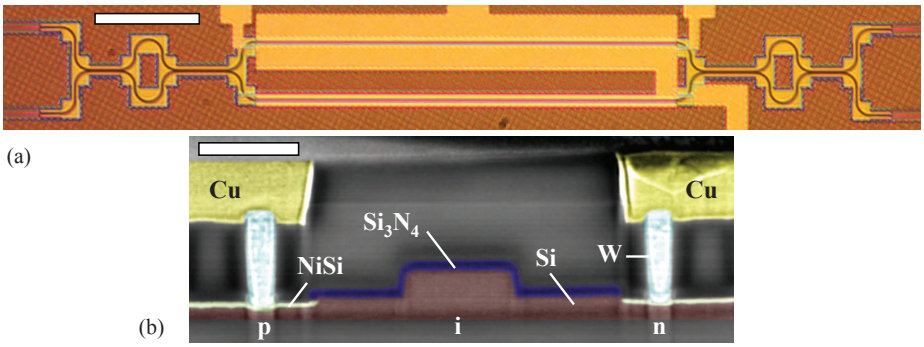

Fig. 2. (a) Microscope image of the fabricated WIMZ switch. The scale bar is 50-μm long. (b) Scanning-electron-microscope image of the cross section of the employed p-i-n diode phase shifter. The scale bar is 500-nm long.

of both the MZ and WIMZ switch is summarized in table 1. The footprint of the WIMZ device is $50\times400\,\mu\text{m}^2$ ($0.02\,\text{mm}^2$), and only slightly larger than that of the reference MZ switch ($50\times330\,\mu\text{m}^2$).

3. Device fabrication

The proposed WIMZ switches were fabricated using 10-Ωcm p-type, 200-mm silicon-on-insulator (SOI) wafers with a 2-μm-thick buried-oxide layer and a 220-nm-thick top silicon layer. The fabrication was performed by utilizing a subset of processing modules from a standard IBM front-end CMOS process flow. The integration of the WIMZ switches required eight lithography levels, with mask designs using 90-nm CMOS ground rules.

The silicon access waveguides were defined utilizing the shallow trench isolation (STI) module, which includes 193-nm deep-UV lithography and dry etch of the full SOI layer with a hybrid Si_3N_4/SiO_2 hard mask. Following this etch, a thick oxide layer was deposited and chemically-mechanically polished, leaving a planarized top surface. Subsequently, the silicon rib waveguides were defined using 193-nm deep-UV lithography and a 170-nm-deep dry etch of the top silicon layer using a hybrid organic-oxide hard mask.

Next, typical CMOS source/drain ion implantation and rapid thermal anneal (RTA) conditions were applied to the rib waveguide to form a lateral p-i-n diode. Highly doped (10^{21}cm^{-3}) p-type and n-type regions were formed in the 50-nm-thick silicon slab, extending to approximately 500 nm away from the rib waveguide core. After the implant activation anneal, standard NiSi ohmic contacts were formed over the implanted regions [16], while NiSi formation on the waveguide core was prevented by depositing and patterning a Si_3N_4 silicide-blocking layer.

Subsequently, metal contacts were formed by utilizing tungsten (W) plugs and copper (Cu) interconnects. A dielectric stack was deposited and planarized, and contact apertures overlaying the NiSi contacts were etched and filled with W plugs. Next, similar deposition, etch, and metal fill process steps were performed to form copper (Cu) metal interconnects. Finally, SiO_xN_y-based optical couplers were formed overlaying the fully etched, inversely tapered access waveguides [17]. These optical couplers provide efficient optical coupling from lensed and tapered fibers to the access wire waveguides.

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

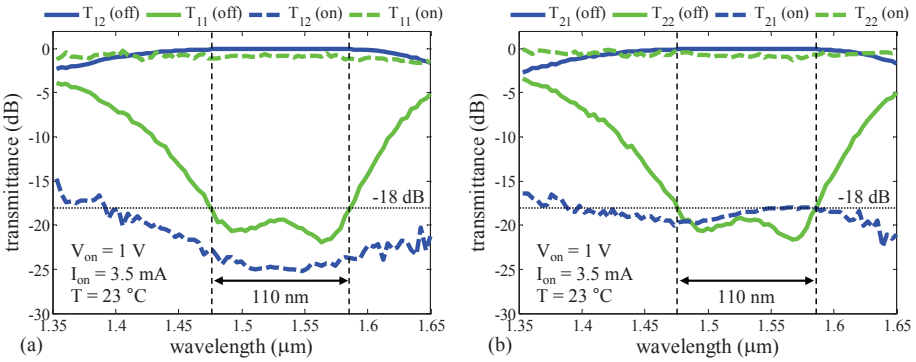

Fig. 3. Steady-state transmittance spectra of the wavelength-insensitive MZ switch. The switching performance for using input port a_1 is shown in (a), and switching performance for using input port a_2 in (b). Crosstalk levels of lower than -17 dB are obtained for both input ports and for both the switch 'off' and switch 'on' state.

Table 2. Overview of the measured switching performance of the WIMZ switch.

input port	BW (nm)	XT_{off}^{max} (dB)	IL_{off}^{total} (dB)	XT_{on}^{max} (dB)	IL_{on}^{total} (dB)
a_1	110	-18	1.1 to 2.0	-23	2.0 to 2.9
a_2	110	-18	1.1 to 2.0	-17	2.0 to 2.9

4. Measurement results

4.1. Optical bandwidth

The optical bandwidth of the fabricated WIMZ devices was characterized by coupling TE-polarized light from a broadband LED source to one of the input ports of the switch and analyzing the intensity spectrum of the transmitted light signal at both output ports. First, this analysis was performed to obtain all four transmission spectra $T_{ij}(\lambda)$ for the 'off' state. These measurements were performed at room temperature (23 °C). In order to allow straightforward comparison with the simulated transmittance spectra in Fig. 1, the intensity spectra measured at the respective output ports were normalized against the sum of the intensity spectra of both output ports, with the input signal at the same input port. The resulting $T_{1j}(\lambda)$ spectra using input port a_1 are shown in Fig. 3a and the $T_{2j}(\lambda)$ spectra using input port a_2 are shown in Fig. 3b. The crosstalk for transmission from a_1 to b_2 is lower than -18 dB over the designed spectral range of 110 nm centered around a wavelength of 1530 nm. The same crosstalk is obtained for transmission from a_2 to b_1, illustrating the 2×2 functionality of the switch.

The less than optimum crosstalk levels within the switch bandwidth result from undesired phase errors in the phase-tuning section, and – to a lesser extent – in the broadband coupling sections. These phase errors are most likely caused by random, fabrication-related variations in waveguide width and etch depth along the phase-sensitive sections of the switch [18]. They are believed to result both in a deviation from the desired 50 % coupling in the coupling section, as well as in a built-in relative phase delay in the phase-tuning section, causing the switch to be detuned from its nominal 'off' state. The deviation of the center wavelength of 1530 nm with respect to the design wavelength of 1550 nm is caused by coupling in the waveguide bends leading to the directional couplers, which was not accounted for in the device design.

Subsequently, a forward-bias voltage V_D was applied to the p-i-n diode in the phase-tuning section, and this voltage was fine-tuned to obtain maximum extinction of both the T_{12} and T_{21} transmittance within the optical bandwidth of the WIMZ switch, which we refer to as the 'on'

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

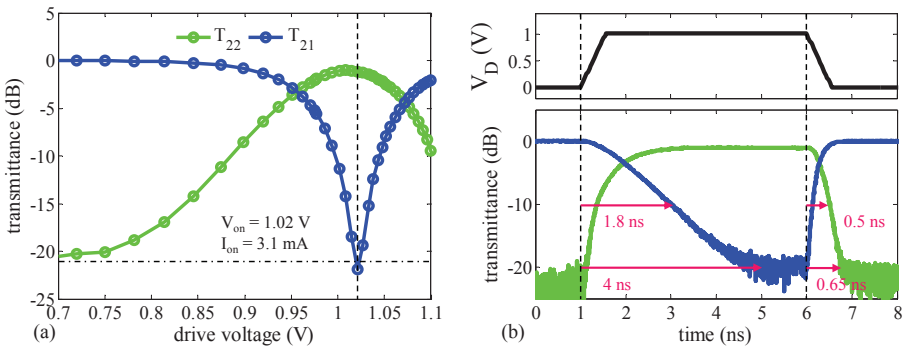

Fig. 4. (a) Switching response measured as a function of peak drive voltage V_D of 100-ns-long drive pulses, at a wavelength of 1518 nm. The 'on' state is reached at $V_{on} = 1.02$ V (± 0.02 V), with -20-dB 'on'-state crosstalk. The horizontal line shows the noise floor. (b) Time-resolved switching response for a 5-ns-long drive pulse with 0.45-ns-long rise/fall times (shown above). The obtained switching times are shorter than 4 ns.

state. Again, the reported transmittance spectra were normalized against the sum of the 'off'-state intensity spectra of both output ports. Maximum extinction was found to occur at the 'on'-state voltage $V_{on} = 1$ V (± 0.01 V), and an 'on'-state current $I_{on} = 3.5 \pm 0.1$ mA. The resulting T_{ij} spectra for the 'on' state are also shown in Figs. 3a and 3b. It can be seen that the 'on'-state crosstalk levels between T_{12} and T_{11} are lower than -23 dB over the wavelength window of interest, whereas the 'on'-state crosstalk levels between T_{21} and T_{22} are lower than -17 dB. This difference in crosstalk is most likely caused by a slight deviation from 50 % coupling in the coupling section. The T_{11} 'on'-state insertion loss was measured to be 0.9 ± 0.2 dB, whereas the T_{22} 'on'-state transmission was 0.9 ± 0.4 dB. It should be noted that these values for 'on'-state insertion loss only contain the contribution from FCA as a result of the normalization. The obtained values agree well with the simulated value of 0.8 dB.

In order to estimate the passive optical losses in the device due to bending, scattering and optical absorption at the electrical contacts, the sum of the intensity spectra recorded at both output ports in the 'off' state was normalized against the intensity spectrum recorded for transmission through a reference waveguide on the same chip. The passive insertion loss of the WIMZ switch was found to vary from 1.1 ± 0.2 dB at 1480 nm to 2.0 ± 0.2 dB at 1590 nm. As such, the total 'on'-state insertion loss is estimated to vary from 2.0 ± 0.2 dB at 1480 nm to 2.9 ± 0.2 dB at 1590 nm. The measured performance of WIMZ switch is summarized for both input ports in table 2.

The series resistance R_s of the p-i-n diode phase shifter was measured to be $8 \pm 2\Omega$. The low 'on'-state voltage $V_{on} \sim 1$ V originates in part from the low series resistivity $r_s = R_s \times L$ of the p-i-n diode phase shifter, which was measured to be $r_s = 1.6 \pm 0.4\,\Omega$·mm. This low resistivity is a direct result of using NiSi-based electrical contacts, which are commonly used in standard CMOS technology. Voltages as low as 1 V are compatible with advanced CMOS drive circuits.

4.2. Intrinsic switching response

While the transmittance spectra measured under steady-state drive conditions demonstrate the wavelength insensitivity of the WIMZ switch, the obtained 'on'-state switching response is affected to some degree by parasitic self heating of the phase-tuning section. Indeed, the electrical power dissipated in the p-i-n diode causes a temperature increase of the active phase shifter. As a result, part of the free-carrier induced phase delay is offset by a counteracting,

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

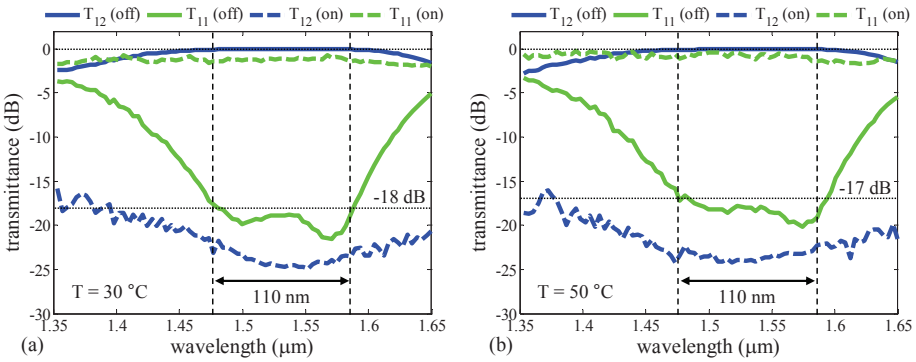

Fig. 5. Steady-state transmittance spectra of the WIMZ switch, measured at 30 °C (a) and 50 °C (b). Only a minor crosstalk degradation (-17 dB) is observed at 50 °C, indicating a tolerance against temperature variations of up to 30 K peak to peak.

thermally induced phase delay, owing to the large thermo-optic coefficient of silicon and the poor heat-sinking ability of SOI devices with a thick buried-oxide layer [19].

In order to decouple the self-heating effect from the intrinsic, free-carrier induced switching response, we performed a series of time-resolved transmittance measurements using a TE-polarized, coherent light source at a fixed wavelength of 1518 nm. A drive signal consisting of 100-ns long pulses with variable peak-to-peak voltage, zero bias voltage, and a 10-% duty cycle was applied to the switch. The time-domain transmittance waveform was recorded on a high-speed oscilloscope and the switch transmittance was evaluated 60 ns after arrival of each pulse. The resulting T_{22} and T_{21} transmittance values are shown as a function of applied peak voltage of the pulses in Fig. 4a. The 'on' state is reached at a peak voltage $V_{\mathrm{on}} = 1.02\,\mathrm{V}$ ($\pm 0.02\,\mathrm{V}$), with an 'on'-state crosstalk of lower than -20 dB. As expected, the 'on'-state drive current $I_{\mathrm{on}} = 3.1 \pm 0.1\,\mathrm{mA}$ obtained under pulsed drive conditions is (slightly) lower as compared to that obtained under steady-state drive conditions. From these numbers, the power consumption of the WIMZ switch in the 'on' state is estimated to be 3.1 mW.

From Fig. 4a, it can be seen that 'off'-'on' switching can be obtained by applying a peak-to-peak drive voltage of only 0.3 V in addition to a bias voltage $V_b = 0.87\,\mathrm{V}$. A such, the $V_\pi \cdot L$ figure of merit can be estimated to be about 0.06 V·mm. Once again, this record-low $V_\pi \cdot L$ is obtained in part by using low-resistance NiSi-based ohmic contacts, used in standard CMOS technology.

4.3. Switching speed

The switching speed was evaluated by applying 5-ns-long pulses with a peak voltage equal to the 'on'-state voltage V_{on}, and 10%-90% rise and fall times of 0.45 ns. The time-resolved T_{21} and T_{22} transmittance waveforms measured for these drive conditions are shown in Fig. 4b. For the 'off'-'on' transition, 20-dB extinction of the T_{21} transmittance is obtained after a switching time $t_{\mathrm{on}}^{20\,\mathrm{dB}} = 4\,\mathrm{ns}$, whereas 10-dB extinction is reached after $t_{\mathrm{on}}^{10\,\mathrm{dB}} = 1.8\,\mathrm{ns}$. For the 'on'-'off' transition, these switching times were measured to be $t_{\mathrm{off}}^{20\,\mathrm{dB}} = 0.65\,\mathrm{ns}$ and $t_{\mathrm{off}}^{10\,\mathrm{dB}} = 0.5\,\mathrm{ns}$ respectively.

4.4. Tolerance to variations of the ambient temperature

In order to assess the influence of temperature variations on the WIMZ switching performance, we performed steady-state transmittance measurements at elevated temperatures of 30 °C and 50 °C, both for the 'off' and 'on' state. The results are shown in Fig. 5. For both temperatures,

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

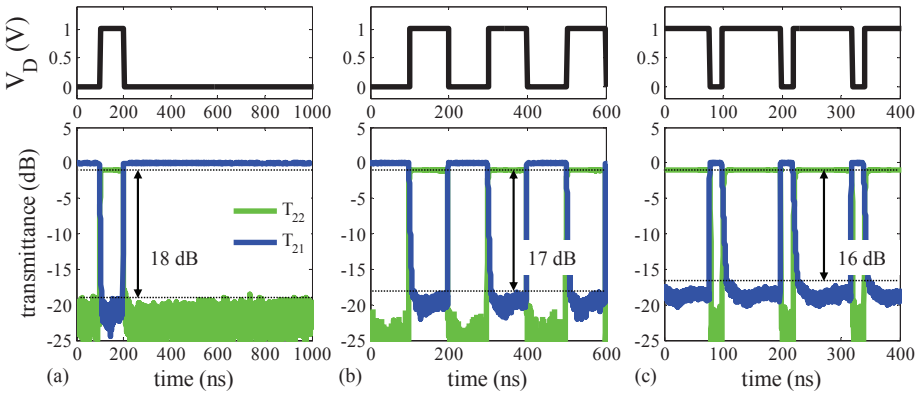

Fig. 6. Switching response measured for 100-ns-long 'on'-state duration with a fixed 'on'-state drive voltage ($V_{on} = 1$ V) and variable duty cycles: (a) 10 %, (b) 50 % and (c) 83 %. A minor degradation in crosstalk levels of up to -16 dB can be observed for high duty cycles. The input wavelength was 1518 nm.

the obtained switching response is fairly similar to the one obtained at 23 °C, which could be expected as the WIMZ switch is a non-resonant optical device, featuring only very short (< 160 nm) built-in phase-delay sections in the broadband 50 % couplers. A slight increase in crosstalk levels is observed for the 'off' state at 50 °C (-17 dB). This degradation is most likely the result of a slight magnification of the phase errors in the structure with increasing temperature. In the 'on' state, similar crosstalk levels can be obtained at 30 °C and 50 °C as compared to the response at room temperature.

However, it should be noted that at these higher temperatures, the 'on'-state is reached at a slightly different electrical bias level: $V_{on} = 0.99 \pm 0.01$ V and $I_{on} = 3.3 \pm 0.1$ mA at 30 °C, and $V_{on} = 0.97 \pm 0.01$ V and $I_{on} = 3.0 \pm 0.1$ mA at 50 °C respectively. This is caused in part by the temperature dependence of the electrical properties of the p-i-n diode itself. Given the strong sensivity of the 'on'-state crosstalk of the WIMZ switch to fluctuations on the drive voltage, as illustrated in Fig. 4a, 'on'-state crosstalk levels below -20 dB will only be obtained at these higher temperatures if the drive voltage is adjusted accordingly. This issue could be resolved by designing switches with an improved tolerance to 'on'-state phase noise, which can be obtained in a MZ-lattice switch configuration [11].

4.5. Influence of self heating on switching performance

As already mentioned in section 4.2, the 'on'-state current under steady-state drive conditions (3.5±0.1 mA) is slightly higher than that for a drive signal consisting of short pulses (3.1±0.1 mA). This is caused by self heating of the p-i-n diode under forward bias, which detunes the switch from its optimum extinction state. As such, the drive voltage yielding the lowest 'on'-state crosstalk levels will depend on the average 'on'-state duration and duty cycle of the drive signal. A degradation in crosstalk levels can be expected if a fixed 'on'-state drive voltage is used both for long and short 'on'-state durations or for drive signals with varying duty cycles.

In order to quantify this degradation in the present WIMZ device, we evaluated the switching performance for 100-ns-long 'on'-state duration and varying duty cycles. The T_{22} and T_{21} transmittance waveforms were measured at fixed drive voltage, for three different duty cycles (10 %, 50 % and 83 %), and are shown in Fig. 6. It can be seen from these curves that the worst-case crosstalk levels degrade from -18 dB to -16 dB with increasing duty cycle. For a fixed drive

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009

voltage $V_{on} = 1\,V$, a crosstalk level lower than -14 dB can be maintained for arbitrarily long message durations with arbitrarily high duty cycles (not shown). The switching performance for such long message durations could be improved by reducing the 'on'-state power dissipation, or by improving the heat sinking and increasing the length of the p-i-n diode phase shifter to reduce self heating.

5. Discussion

The WIMZ switching performance described in section 4, including low crosstalk over a wide optical bandwidth, low power consumption and fast switching speed, is essential for optical switches in short-range applications such as chip-to-chip and intra-chip optical interconnects. The wide optical bandwidth of the WIMZ switch enables data streams consisting of a large number of WDM channels with coarse channel spacing to be carried and routed using a single waveguide. In addition, the optical bandwidth of the WIMZ switch is essentially temperature insensitive. When combined with silicon-based multiplexing and demultiplexing circuits with wide flat-top pass bands [20], the WIMZ switches are expected to enable high-throughput, temperature-insensitive, reconfigurable optical interconnects integrated in silicon.

In order to illustrate the power efficiency of the present WIMZ switch, for the specific application of a non-wavelength-selective, circuit-switched ONoC as proposed in [6], we can assume a conservative WDM channel spacing of 2 nm and a single-channel data rate of 10 Gbps. Under these conditions, a 500-Gbps WDM data stream could be simultaneously switched by the WIMZ device. Given the power consumption of only 3 mW, the switching energy per bit would be only 6 fJ/bit. Furthermore, as shown in section 4.5, self-heating effects are relatively weak, owing to the low power consumption. As a result, crosstalk levels lower than -16 dB are obtained in the present device, for switching conditions anticipated in circuit-switched ONoC architectures (100-ns-long messages and high duty cycles, [6]).

6. Conclusion

In conclusion, we have demonstrated an ultra-broadband Mach-Zehnder based optical switch in silicon, operated through carrier injection in a p-i-n diode. The optical bandwidth of the switch is 110 nm, which is three times larger than that of a conventional MZ switch. Crosstalk levels lower than -17 dB were demonstrated for both the 'off' and 'on' switching states in fabricated devices over the designed optical bandwidth, with full 2×2 switching functionality. The optical bandwidth of the switch was shown to be largely temperature insensitive. The power consumption was measured to be as low as 3.1 mW and switching times were shorter than 4 ns. A low drive voltage of 1 V as well as a record-low $V_\pi \cdot L$ of 0.06 V·mm were obtained, owing in part to the implementation of low-resistance NiSi-based ohmic contacts. The 'off'-state insertion loss was measured to be in the range 1.1 dB to 2.0 dB, whereas the 'on'-state insertion was in the range 2.0 dB to 2.9 dB. The demonstrated switching characteristics are essential for realizing high-throughput, low-power, reconfigurable, short-range optical interconnects.

Acknowledgments

This work was supported in part by the DARPA APS Program, under contract HR0011-08-C-0102. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of DARPA or the Department of Defense. The authors would like to thank D. M. Kuchta for helping with measurement-automation software. The authors also gratefully acknowledge the efforts of the staff of the Microelectronics Research Laboratory (MRL) at the IBM T. J. Watson Research Center, where the devices were fabricated.

#119837 - $15.00 USD Received 10 Nov 2009; revised 10 Dec 2009; accepted 13 Dec 2009; published 16 Dec 2009