Bit-Error-Rate Monitoring for Active Wavelength Control of Resonant Modulators

A new method uses bit-error-rate measurements to acquire and stabilize the wavelength of an optical resonant modulator to an optical carrier wave. This is attractive because it uses the pertinent metric, bit error rate, to optimize the modulator resonance independent of other system variations, meaning it can compensate for system aging and drift even in the heater element itself.

William A. Zortman
Anthony L. Lentine
Douglas C. Trotter
Sandia National Laboratories

Michael R. Watts
Massachusetts Institute of Technology

Optical interconnections can significantly reduce the power dissipation and greatly increase the aggregate connection bandwidth in high-performance systems. These include multiprocessor digital computers and data centers with silicon photonics stacked on CMOS[1] or directly (monolithically) integrated with CMOS on chip[2,3] and intrasatellite communications for supporting ultra-high-bandwidth data links from large focal plane arrays to digital signal processing engines.[4] Silicon photonic microring and microdisk modulators for the transmit side of the links are an active area of research[5-7] because they're compatible with silicon electronics processing and have been demonstrated at data rates higher than 10 Gbits per second (Gbps) with sub-10 fJ/bit switching energies.[6,7]

A key problem for these devices, however, is control of the resonant optical wavelength, which varies as a function of fabrication tolerances (such as thickness and litho/etch bias),[8,9] temperature, and optical incident power.[10]

This article presents a new method of tuning a resonant optical modulator's wavelength to match the incident carrier wavelength (about 1,550 nm). We use independent logic one and logic zero bit errors from a local receiver and simple logic circuitry to drive an integrated microheater that adjusts the device's temperature. As a result, we make the feedback loop independent of heater aging or other link parameters that could change over time or environmental conditions. The feedback loop will inherently compensate for all sources of variations over the system lifetime because it looks at what we care about—the bit errors. Additionally, the system can initialize the link. The resonator can be off resonance and the feedback loop will cause heating of the resonator until the laser line is found (providing of course that the desired resonant wavelength is at a longer wavelength than the unheated resonator's frequency).

We designed the system to compare the data stream coming from the modulator as

Published by the IEEE Computer Society

Related Work in Modulator Wavelength Control

Previous efforts involving silicon photonic microring and microdisk modulators have focused on controlling the resonant optical wavelength. Heaters have been successfully integrated with microring and microdisk filters and modulators.[1-4] Researchers have also demonstrated closed-loop control circuits to stabilize the wavelength.[1,5-7] DeRose et al. integrated a heater and sensor with a modulator to stabilize its wavelength,[1] but the device didn't perform as a modulator. Other drawbacks to this type of method are that the sensor and other link components age over time, the background thermal environment is sampled, and every sensor and modulator must be precalibrated and recalibrated. Other researchers have demonstrated a wavelength stabilization loop, which assesses the received power versus wavelength.[5,6] Qui et al. used a dual resonator to match power levels, using scattered light levels as the detection mechanism for locking.[7] However, locking power levels for a modulator only gives an indirect measure of the optimum resonant wavelength position. Each of these techniques lacks the ability to look at the end result of all the factors that affect the bit error rate (BER) and generate a feedback signal that includes those inherent degradations to the performance and correct them.

Other options to improve the error rate and mitigate the wavelength shift might include feed-forward equalization (FFE) and decision feedback equalization (DFE). These methods rely on changing the drive strength (FFE) or the decision threshold in the receiver (DFE) on the basis of data sent and received data streams. This can improve the error rate while the resonator is locked, but these methods do not maintain a resonant frequency (although they could add some margin together with a resonant locking scheme like the one we describe in the main article). Engineers can certainly use these methods to train a link and preadjust parameters for known distortions arising from strings of one or zero emphasis.

References

1. C.T. DeRose et al., ''Silicon Microring Modulator with Integrated Heater and Temperature Sensor for Thermal Control,'' *Proc. Conf. Lasers and Electro-Optics,* Optical Soc. of America, 2010, paper CThJ3.
2. W.A. Zortman et al., ''Integrated CMOS Compatible Low Power 10Gbps Silicon Photonic Heater Modulator,'' *Proc. Optical Fiber Communication Conf.,* Optical Soc. of America, 2012, no. OW4I.5.
3. M.R. Watts et al., ''Adiabatic Resonant Microrings (ARMs) with Directly Integrated Thermal Microphotonics,'' *Proc. Conf. Lasers and Electro-Optics,* Optical Soc. of America, 2009, paper CPDB10.
4. G. Li et al., ''25Gb/s 1V-Driving CMOS Ring Modulator with Integrated Thermal Tuning,'' *Optics Express,* vol. 19, no. 21, 2011, pp. 20,435-20,443.
5. K. Padmaraju et al., ''Dynamic Stabilization of a Microring Modulator Under Thermal Perturbation,'' *Proc. Optical Fiber Communication Conf.,* Optical Soc. of America, 2012, paper OW4F.2.
6. E. Timurdogan et al., ''Automated Wavelength Recovery for Microring Resonators,'' *CLEO: Science and Innovations,* Optical Soc. of America, 2012, paper CM2M.1.
7. C. Qui et al., ''Wavelength Tracking with Thermally Controlled Silicon Resonators,'' *Optics Express,* vol. 19, no. 6, 2011, pp. 5143-5148.

detected by a local receiver against what the electronic driver sent to the modulator for imprinting on the carrier wave. This is done using two stages of logic. The first is an XOR gate that decides if there was an error, and the second is an AND gate that determines the type of error. We define 0 errors as events in which a 0 was sent to the modulator and a 1 was detected. We define 1 errors as the opposite; a 1 was expected but a 0 was received. By accumulating the errors, we can determine on which side of the resonance the carrier wave resides (we demonstrate this using simulation results). Some other logic-gate implementations provide the same basic logic functionality and might be slightly more power efficient or simpler, substituting inverters and AND gates for the XOR and AND, for example.

To test our approach, we used an integrated resonant heater modulator device[11] operating at 3.125 Gbps with a field-programmable gate array (FPGA) and external receiver driving the control loop. Our results show a tuning range of at least 15°C to 32°C (~1.4 nm). No dithering or calibration is required; the technique isn't susceptible to sensor aging, and it can compensate for long-term drift in the modulator and link characteristics. (For information on other approaches, see the ''Related Work in Modulator Wavelength Control'' sidebar.)

System overview

Figure 1 shows the transmitter side of the link, including the local receiver used to monitor the bit errors generated due to misalignment of the modulator's resonant wavelength with the incoming wavelength. The output optical signal is also routed to the ''link'' receiver shown in the upper right. Importantly, the local receiver must have

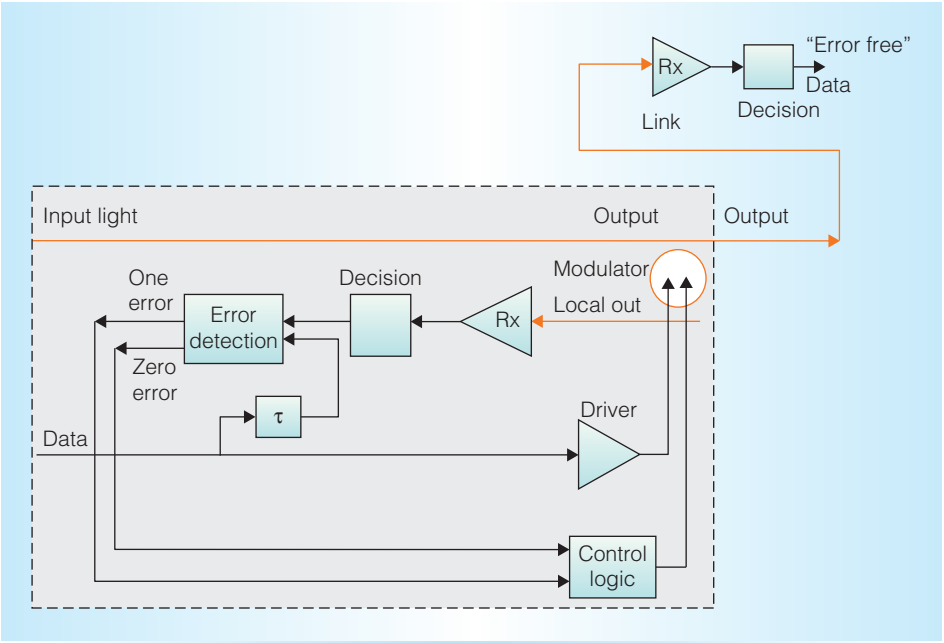

Figure 1. Block diagram of the control technique. The heart of the error detection is an XOR gate that compares logical bits to the sent pattern. In the upper right, a remote receiver is shown. The local receiver in the dashed box ensures that the remote receiver sees no errors by using the control loop to keep the modulator aligned to the carrier wave. Ideally, the modulator will have a drop port that will be used to send data to the local receiver.

higher noise characteristics than the link receiver so that it is generating errors, which the control loop uses, while the link receiver is error free. The electrical data is routed to the modulator and to the logic gates that detect the bit errors. A continuous wave (CW) laser source is connected to the modulator's optical input. The signal from either the drop port or an output tap is routed to an integrated photo receiver. We believe that a drop port on the modulator is the ideal solution because no local filter alignment is necessary. The output from the local receiver is routed to a circuit that performs a threshold (comparator) operation to determine whether the received bit is a logic zero or one. The circuit must be sampled with the clock signal to generate the digital signals, but does not need a clock recovery circuit because the clock is local to the circuit; the transmitter and logic use the same clock.

After the decision (thresholding) circuit, the signal is routed to an error detector that compares the received data against the data that was transmitted and determines when errors are made. The circuits that detect logic one and zero errors are XOR and AND gates. The modulator output is also routed to the interconnect receiver with better noise characteristics, ensuring that the interconnection link will be error free when the algorithm is functioning.

Simulation

Figures 2a to 2e provide an overview from a simulation of the modulator's operation as a function of both wavelength and temperature in the upper set of plots and the paired eye diagram in the middle set of plots. The upper plot is a spectral representation of the modulator Lorentzian resonance and a 193-THz (1,550 nm) CW laser carrier wave at a relative frequency on the plot of zero. We show two resonances for each. The higher frequency is the modulator biased with a 0-V bias, and the lower frequency is the modulator biased with a 2-V reverse bias. In our experiment, the bias levels were approximately 0.5-V forward biased and 0.5-V reverse biased so that we could drive the modulator directly with the LVDS output from the FPGA, but the basic operation

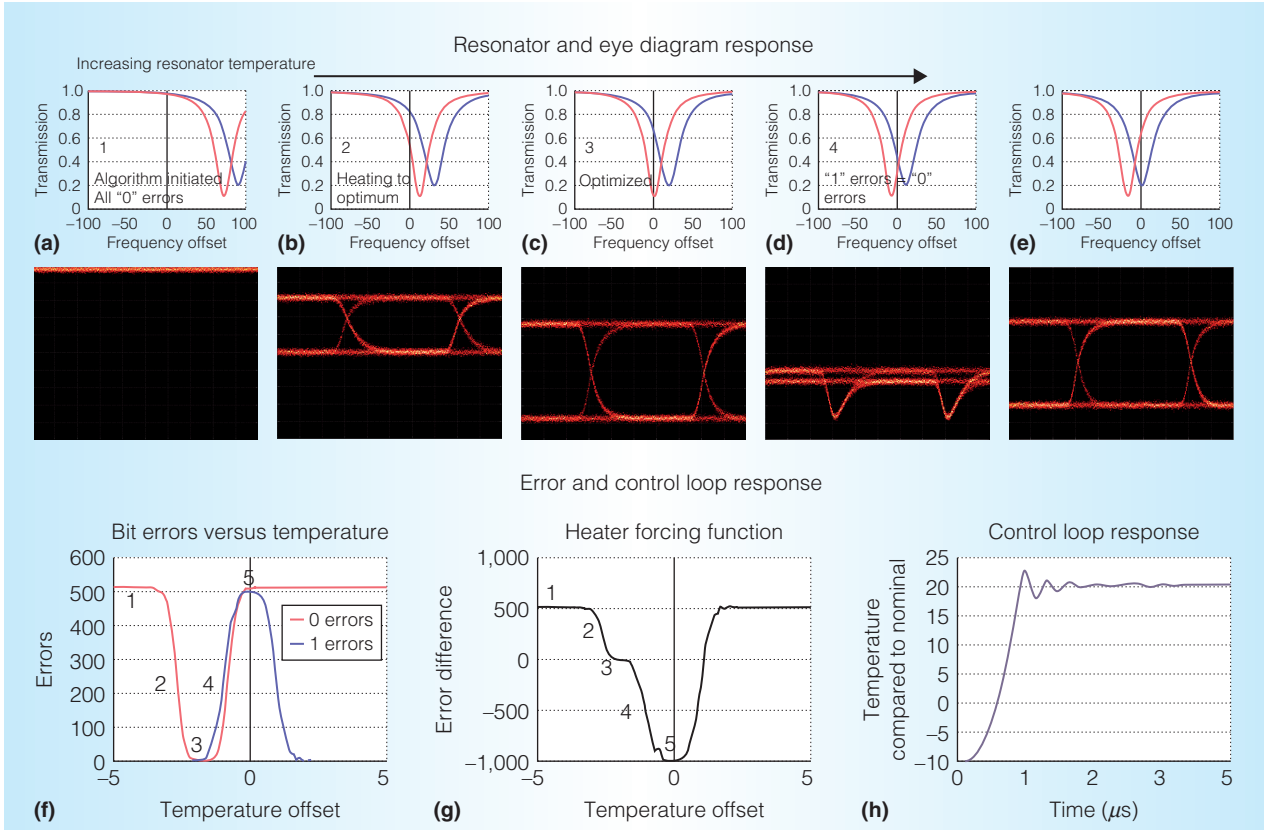

Figure 2. Simulations of our method. A simulated demonstration of the resonance's progression during tuning to higher temperature and lower frequency (a–e). The frequency offset is in gigahertz. An error map shows what occurs in the receiver during tuning using the number assignments from the resonance plots in Figures 2a to 2e (f–g). The total 1 and 0 errors are shown along with their difference (three times logic 1 errors minus the logic 0 errors) illuminating the data-stability points along the progression with labeling, which relate the points back to the resonance positions in Figures 2a to 2e. Finally, we show the stable locking response (h).

of the control method is independent of exactly how the modulator is biased. The two resonances represent the wavelength shift in the modulator resonant frequency as a function of applied voltage that occurs because of the carrier-induced index of refraction changes described, for example, by Watts et al.[12] This voltage is the *modulating voltage*. The second action is the resonance's movement toward the laser line with increased temperature and is represented with the progression of figures with temperature in five steps.

At the beginning of the sequence, the resonance is far from the laser line, so a series of 1's is imprinted on the carrier wave that travels via waveguide to the receiver. The input signal to the modulator has nearly equal numbers of 1's and 0's. As the resonator is tuned so that the carrier wave aligns with the resonance at reverse bias, an eye begins to appear (Figure 2b) until it is fully open with noninverted data at point 3 (Figure 2c). If the resonator is tuned beyond the carrier, a corrupted eye is sent until finally the data is inverted on the other side of the resonance with a well-shaped, but symbolically incorrect eye. Further tuning will result in off-resonance operation and a series of 1 symbols. Silicon resonators' sensitivity to temperature is approximately 10 GHz/°C. The tuning operation naturally favors the red side of the resonance (lower frequency side), which has a steeper slope and therefore greater extinction ratio.[13]

Figures 2f and 2g are both error maps showing what is occurring in the receiver

during tuning using the number assignments from the resonance plots in Figures 2a to 2e. Again, a 0 error is what happens when a 0 is sent and a 1 is received, and vice versa for 1 errors. The progression shows that initially all errors are 0 errors that begin to lessen at point 2, followed by optimization at point 3, an increase in 0 and 1 errors at point 4, and finally, all data are in error at point 5 because the data is inverted. At point 4, there are a similar or equal number of 1 and 0 errors, so the logic implementation requires proper weighting of the errors to allow tuning of the resonator in the cooling sense toward the optimum. We found that weighting 1 errors three times provides good performance.

Figure 2g shows the heater driving function, the 0 errors minus three times the 1 errors as a function of detuning. Values above zero correspond to an increase in heat being applied, and values below zero correspond to a decrease in heat being applied. Importantly, there is an operating range of several degrees above the optimum point where the control loop will continue to operate. However, if the device is heated beyond the point where the forcing function becomes positive again, the circuit must automatically recycle the process.

Figure 2h shows a simulation of the modulator's temperature as a function of time under an initial temperature detuning of 30°C. The low amplitude ringing that happens as lock is achieved further supports the thesis that the system is stable under locking.

We performed these simulations in Matlab Simulink. In the simulation, the modulator's frequency response was approximately 10 GHz/°C. We set the heater's frequency response to 1 MHz, and the local receiver bit error rate (BER) was approximately 10^{-4}. The tuning range of 30°C was limited by the current through the heater. We performed simulations with a 2^7-1 pseudorandom pattern.

Thus, if the modulator temperature is initially less than the desired temperature, the control loop will apply an increasingly positive signal to the heater and heat up the device. When the number of errors is about equal, just to the left of the resonance on a frequency plot, the positive and negative stimulus to the control loop will cause the device to settle near its optimum wavelength. Extra weighting on the logic 1 errors keeps the operating point from climbing up the curve to the right of the resonance where the 1 and 0 errors are comparable, but rapidly increasing.

FPGA-implemented logic block

We performed our tests using a Virtex 6 FPGA-mounted on a commercially available ML605 board. We designed the logic around a 3.125-Gbps signal primarily because of the board performance and availability of low-jitter crystal oscillators for clocking at that speed. The 3.125-Gbps signal is received from the modulator and is demultiplexed down to a 32-bit-wide signal running at 97.6 MHz, and it follows three paths in the FPGA.

We split the received signal into three paths (P1, P2, and P3) in the FPGA (see Figure 3). The first two are primarily to match the pattern. On start-up, the modulator device initially imprints only logic 1's. The device is initially heated until there is roughly an equal number of 1's and 0's (P2) so that there is a discernible pattern that can be matched to the desired pattern by shifting one bit at a time (P1). Once the pattern is matched, P3 measures the bit errors and does the fine-temperature tuning described earlier, required to optimally adjust the resonant wavelength.

P1 takes the incoming data to a pattern comparator. This is a bank of look-up tables (LUTs) that compare the incoming pattern to the known sent pattern in memory and goes to a logic 0 when a match is found. This signal, S_{pl}, determines the state of the logic (which we describe later). As long as S_{pl} is high, a counter increments the output, D_{ml}, of a RAM block that scans through 16 different starting points of the test pattern, which in this case is a 16-bit repeating pattern (0100011110111000). When a match is found, the counter ceases, S_{pl} goes low, and D_{ml} is frozen to the matched sequence. A pattern match is needed to properly implement the bit error control logic, L_{BER}, in the second logic path. D_{ml} is routed into the L_{BER} block, although it isn't useful until

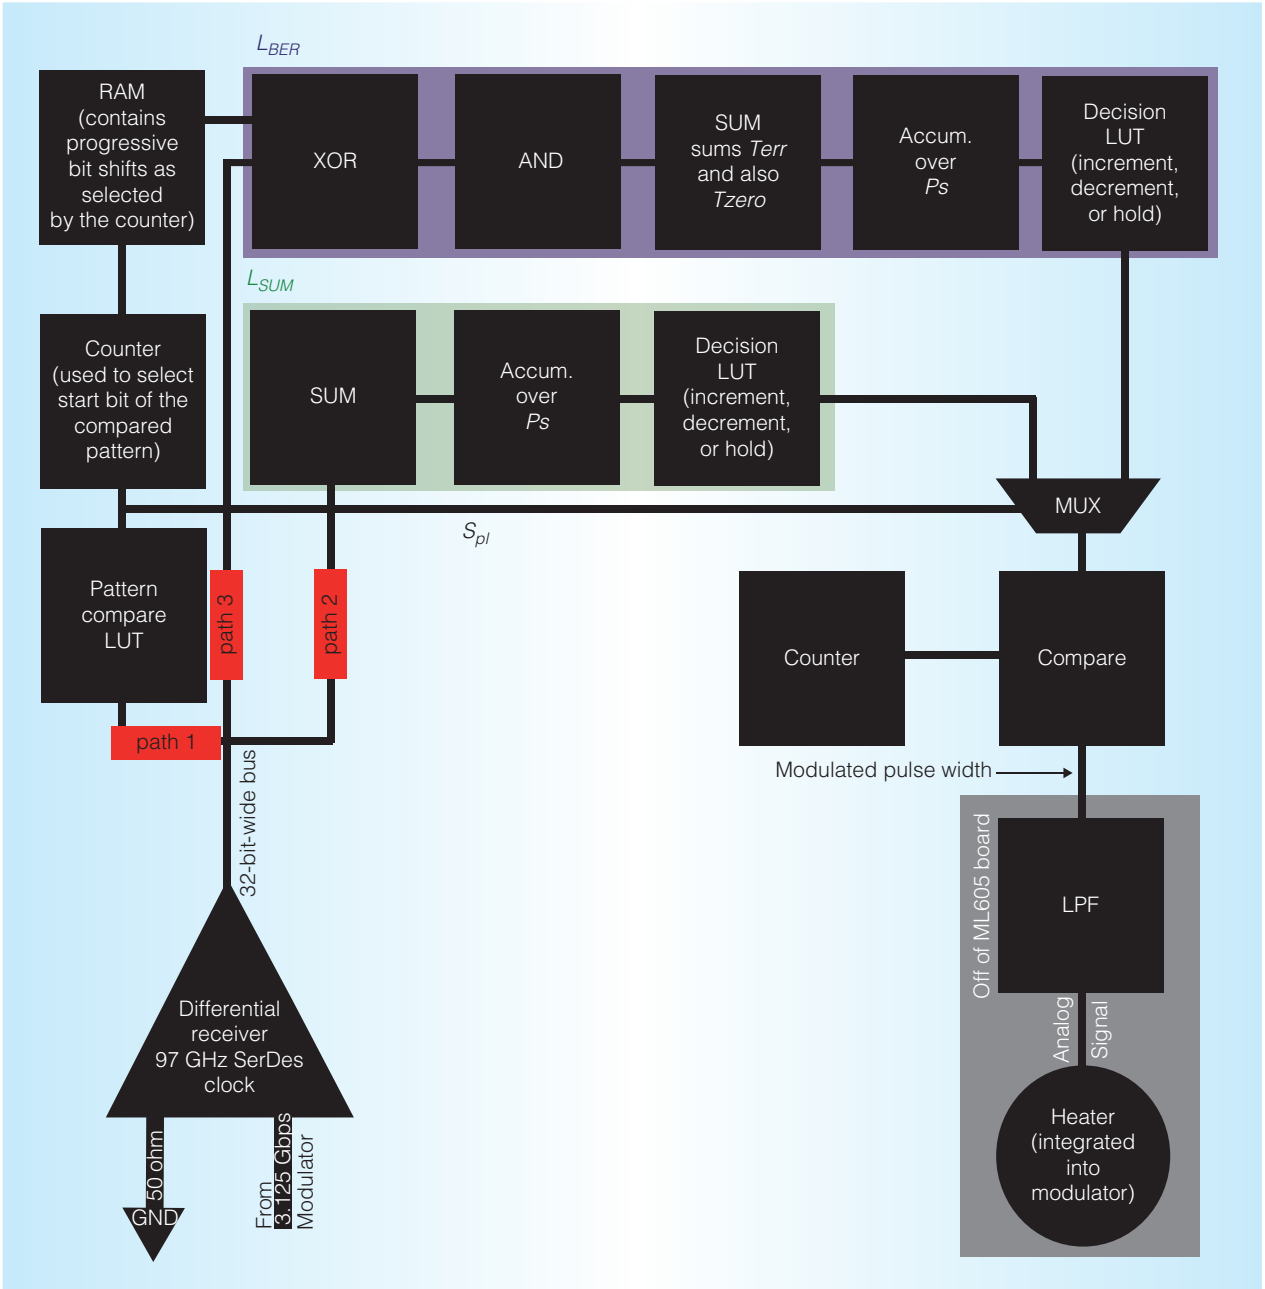

Figure 3. The internal logic blocks of the experimentally implemented locking loop. Path 1 leads to a pattern-matching block that, when locked, allows the feedback loop to operate using the bit error control logic, L_{BER}, which allows for precise locking of the resonance. Until S_{pl} goes high, with the correct pattern, we use a simple summing block to approximate the optimum locking point.

lock is achieved. This method of pattern matching would also work with a buffer for random data such as exists in a real link.

P2 is routed into logic called L_{sum}, which simply counts the total number of 1 values received. L_{sum} is a coarse locking mechanism that operates on the assumption that with any data stream there are approximately equal numbers of 1's and 0's. L_{sum} counts the number of 1 values using an accumulator for 32,000 bits, 10.24 μs, or 1,000 clock cycles, P_s, referred to as the sampling period.

At the end of P_s, the total number of 1's is assessed using the LUTs. If the total is above 17,000 or below 15,000, the value that drives the heater is incremented or decremented, respectively. For values in between, the heater retains the previous value. L_{sum} is the logic that drives the heater until S_{pl} goes low, at which point L_{BER} takes over. The logic in L_{sum} is necessary for initial coarse locking so that the pattern searching done in P1 can occur. This allows a pattern to be found and subsequently enables the next path, P3.

The third path, L_{BER}, contains the XOR and AND gates we described earlier and outputs the total 1 errors, T_{one}, and the total 0 errors, T_{zero}, to two different accumulators over P_s. At the end of P_s, a comparator weights T_{one} three times heavier than T_{zero} and decides which direction the heater should increment, or if it should stay the same. For sampling periods that generate less than four errors (BER about 10^{-4}), incrementing the counter is disabled. This logic could also be instantiated with a state machine, but a simple LUT implementation uses fewer logic resources.

The output of either L_{sum} or L_{BER} feeds into a comparator depending on the value of S_{pl}, which selects either L_{sum} or L_{BER} through a multiplexer (MUX). This comparator goes to logic 0 when an adjacent counter hits the value determined from the logic we described earlier, yielding a varying digital pulse length. This output connects to a low pass filter that converts the digital pulse length to an analog voltage level for the on-device heater. The full pulse length is 330 ns.

The block L_{sum} can be eliminated because the L_{BER} block could be fed a series of 0's until lock is achieved and a pattern is available to match against. Also, the transmitter in our experiment (not shown in the logic block) gets its data from a static 16-bit pattern. In a field implementation, truly random data is used, but locking to random data would be similar because the sent data can be stored in memory for a determined latency while locking is achieved.

In an integrated solution, we could potentially eliminate the pattern-matching circuitry. The delay and delay variations from the local modulator output, through the local receiver and decision circuit, could be relatively small for a low-capacitance simple receiver. If the delay and delay variations over time and temperature are small enough, the delay difference between the received data and the copy of the signal that feeds the modulator can be bounded to within a bit time. In this way, the bits between these two signals can always remain aligned.

Experimental demonstration

To test the circuit, we used an integrated resonant silicon heater-modulator shown in Figure 4.[11] Figure 4a is a schematic that shows the modulation section connected to heavily doped *P* and *N* contacts on the left side of the device. This side of the device imprints the microwave signal on the carrier. The right side of the device is a simple resistive heater fabricated using our standard *N*-type doping. Figure 3b is a scanning electron micrograph (SEM) of the device showing the tungsten contacts to the modulator. Two of the metal lines are missing because they delaminated during sample preparation. The oxide strip in the center of the device is present for electrical isolation of the heater circuit from the modulation circuit. We designed the modulator to operate using a differential drive, as shown in Zortman et al.[14]

Figure 5 shows the test setup. The device was mounted on a thermoelectric cooler (TEC). The local receiver was a Nortel transimpedance amplifier (TIA) with a transimpedance of approximately 500 ohms and a 3-dB bandwidth of 11 GHz. We amplified the optical signal coming off the chip using an erbium-doped fiber amplifier (EDFA). This was necessary because the fiber-to-chip coupling losses are a total of about 20 dBm. We then divided the signal using a 3-dB splitter with half the signal going to an Agilent DCA-X sampling scope and the other half routed into the receiver. We attenuated the receiver signal electrically by 6 dB (50 percent of the incoming voltage) to set the logic threshold and then routed this into the FPGA receiver. The FPGA receiver input is differential, so the unused side was terminated with 50 ohms.

The external laser was initially set to a wavelength—0.5 nm, 1 nm, or 2 nm—below the optimum (approximately 1,550 nm).

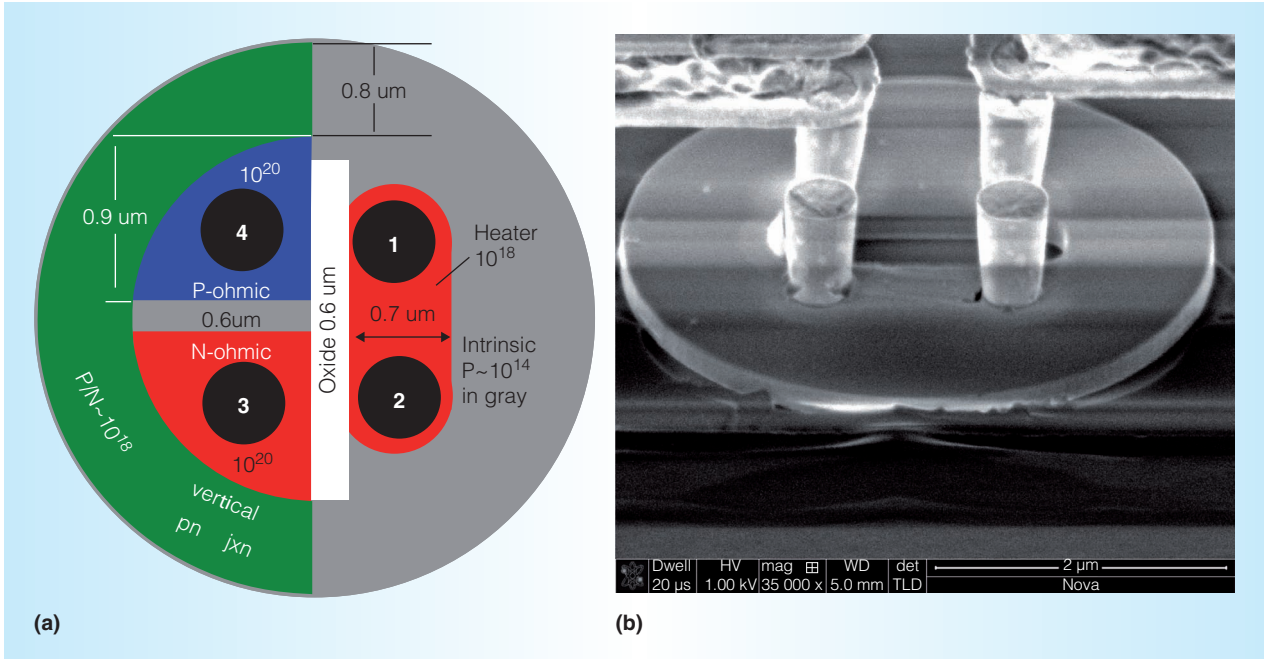

Figure 4. Modulator schematic and SEM. A schematic of the heater modulator used in this demonstration (a). The modulation diode and ohmic contacts are on the left, electrically isolated from the resistive heater element on the right. A scanning electron micrograph (SEM) (b). Although the metal lines of the modulation side are delaminated, the four tungsten plugs and the oxide strip are clearly visible.

Before we enabled the control loop, we observed no visible pattern from the device output. When the control loop was enabled, the device was heated and we observed the correct bit stream on the data output, as shown in Figures 5c to 5e. The eye diagrams shown are for locking when the laser is 500 pm, 1 nm, and 2 nm from optimal, respectively. The closure and amplitude change in the eye diagrams have been observed,[11] and this results from imperfect electrical isolation of the heater in the heater modulator. We do not believe this is inherent in the locking mechanism.

We performed further experiments to test the feedback loop for changes in a hypothetical system during data transmission. We tuned the laser in 10-pm steps from 1,551 to 1,553 nm and back again, and the device resonance tracked the input laser's wavelength shift by tuning the device temperature about 25°C. This would apply to changes in the carrier-wave frequency. Although mechanisms for controlling carrier-wave lasers are robust, such a capability could provide flexibility for channel reconfigurability in future systems.

In a third experiment, we varied the device temperature using the TEC. We increased the temperature 16°C over two minutes (about 0.13°C/s, the TEC's tuning speed limit), followed by a 16°C drop from its initial temperature. In both cases, lock was maintained with no observed instability in the eye diagram. This demonstration shows the ability to compensate for what are likely environmental challenges for silicon photonics intimately integrated with the varying thermal environment of silicon microprocessors. The thermal range was limited by the TEC.

Future work on this system includes the need for a more detailed understanding of the power dissipated in the photonics and electronics. The power consumption and dissipation consists of the power in the local receiver, logic, microheater driver, and the microheater itself. We believe that the dissipation can be dominated by the microheater if we can sufficiently reduce the power of integrated receivers. Estimating some of the parameters gives an idea of the

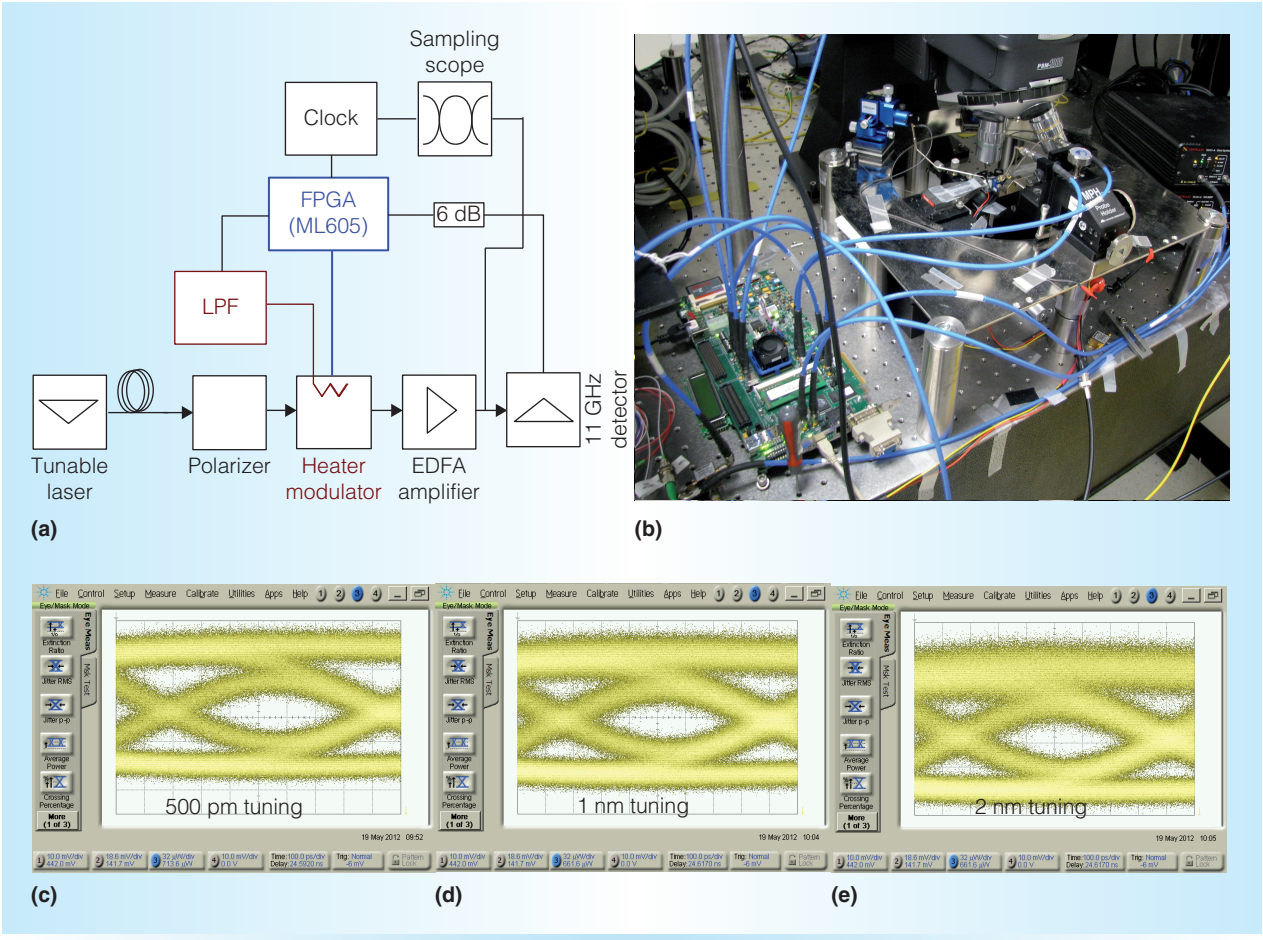

Figure 5. Experimental setup and results. A carrier wave from a tunable laser is imprinted with a microwave signal at the heater-modulator. The signal is then sent through an erbium-doped fiber amplifier (EDFA) to a receiver and subsequently split between the sampling scope and the FPGA board. The signal into the board is attenuated by 6 dB to decrease the signal-to-noise ratio and provide finer tuning (a). A picture of the setup with the FPGA board on the left connected to the modulation device through the SMA microwave cables. The heaters are driven with single-point probes. Other SMA cables connect to clock monitors and the scope trigger (b). Data output from heating the device. The laser is set at 500 pm (c), 1 nm (d), and 2 nm (e) from resonance.

power consumption we can expect to be added to a link by using this control loop. We expect that the electronics (register, decision, and accumulation) will consume 35 fJ/bit based on simulations of the 30 gates needed for these functions. The conversion to the analog domain should add another 10 fJ/bit. For the receiver, our group has a goal of 10 fJ/bit, although no demonstration of this is yet available. The heater itself requires 70 fJ/bit if we assume that the maximum amount of tuning required will be to the next point on a 100-GHz wavelength division multiplexing (WDM) grid. According to these estimates at a 10 Gbps data rate, we believe that an additional 125 fJ/bit will be needed for this system. Implementing this system in intimate integration with the modulator and reducing the logic will allow a full understanding of the power required and enable better comparisons for system designers as they consider alternatives.

Because our demonstration showed the ability of our method to support system initialization, it should be possible to do this for many channels and to reconfigure resonators to different channels with small changes to the logic. Additionally, BER measurement and comparisons of the local and remote

receiver using $2^{31}-1$ pseudorandom bit stream data, or real data such as a video stream, at 10 Gbps will further establish this method as a viable solution for future systems. The heater and modulator themselves require further investigation as to their operability and reliability across the range of temperatures in CMOS systems.

Finally, we plan to develop and improve the link's control coefficients. We can do this in a few ways. One is to investigate link stability and locking time by attenuating the signal into the feedback electronics. Further investigation will also include the logic itself. Because we are using pulse-width modulation, we can achieve finer resolution timing simply by expanding the output driver bus to include more bits. We expect that as the control coefficients are developed they will vary with data rate and between periods where the link is active and inactive. MICRO

Acknowledgments

Sandia National Laboratories is a multiprogram laboratory managed and operated by Sandia, a wholly owned subsidiary of Lockheed Martin, for the US Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

1. R.G. Beausoleil, ''Large-Scale Integrated Photonics for High Performance Interconnects,'' *ACM J. Emerging Technologies in Computing Systems,* May 2011, article no. 6; doi:10.1145/1970406.1970408.
2. J.S. Orcutt et al., ''Open Foundry Platform for High-Performance Electronic-Photonic Integration,'' *Optics Express,* vol. 20, no. 11, 2012, pp. 12,222-12,232.
3. W.A. Zortman et al., ''Monolithic and Two-Dimensional Integration of Silicon Photonic Microdisks with Microelectronics,'' *Photonics J.,* vol. 4, no. 1, 2012, pp. 242-249.
4. M.R. Watts et al., ''Low Power Silicon Microphotonic Communications for Embedded Systems,'' *Proc. 13th Ann. Workshop High-Performance Embedded Computing* (HPEC 09), 2009; http://www.ll.mit.edu/HPEC/agendas/proc09/Day1/F2_1445_Watts_presentation.pdf.
5. Q. Xu et al., ''Micrometre-Scale Silicon Electro-Optic Modulator,'' *Nature,* vol. 435, 19 May 2005, pp. 325-327.
6. M. Watts et al., ''Vertical Junction Silicon Microdisk Modulators and Switches,'' *Optics Express,* vol. 19, no. 22, 2011, pp. 21,989-22,003.
7. P. Dong et al., ''Low V_{pp}, Ultralow-Energy, Compact, High-Speed Silicon Electro-Optic Modulator,'' *Optics Express,* vol. 17, no. 25, 2009, pp. 22,484-22,490.
8. W.A. Zortman, D. Trotter, and M. Watts, ''Silicon Photonics Manufacturing,'' *Optics Express,* vol. 18, no. 23, 2010, pp. 23,598-23,607.
9. A. V. Krishnamoorthy et al., ''Exploiting CMOS Manufacturing to Reduce Tuning Requirements for Resonant Optical Devices,'' *IEEE Photonics J.,* June 2011, pp. 567-579.
10. Q. Li et al., ''Experimental Characterization of the Optical-Power Upper Bound in a Silicon Microring Modulator,'' *Proc. IEEE Optical Interconnects Conf.,* 2012, IEEE CS, pp. 38-39.
11. W.A. Zortman et al., ''Integrated CMOS Compatible Low Power 10Gbps Silicon Photonic Heater Modulator,'' *Proc. Optical Fiber Communication Conf.,* Optical Soc. of America, 2012, no. OW4I.5.
12. M. Watts et al., ''Vertical Junction Silicon Microdisk Modulators and Switches,'' *Optics Express,* vol. 19, no. 22, 2011, pp. 21,989-22,003.
13. P. Dong et al., ''Low V_{pp}, Ultralow-Energy, Compact, High-Speed Silicon Electro-Optic Modulator,'' *Optics Express,* vol. 17, no. 25, 2009, pp. 22,484-22,490.
14. W.A. Zortman et al., ''Low-Voltage Differentially-Signaled Modulators,'' *Optics Express,* vol. 19, no. 27, 2011, pp. 26,017-26,026.

William A. Zortman is a research and development engineer with the Applied Photonics Microsystems group at Sandia National Laboratories. His research interests include silicon photonic and electronic applications, high-speed interconnect, high-volume manufacturing, embedded systems, system integration, and productization. Zortman has a PhD in electrical engineering from the University of New Mexico. He is a member of IEEE and the Optical Society of America (OSA).

Anthony L. Lentine is a principal member of the technical staff at Sandia National Laboratories. His research interests include photovoltaics, renewable energy integration, and silicon photonics for low-energy communications in high-performance computing and other national-security applications. Lentine has a PhD in physics from Heriot Watt University in Edinburgh, UK.

Douglas C. Trotter is a research and development engineer at Sandia National Laboratories, where he led the development of the integrated Silicon Photonic Platform (SPP1). His research interests include diagnostics for pulse power investigations, specifically the development of camera and sensor arrays for x-ray diagnostics, and the development of silicon photonic systems. Trotter has a master's degree in economics from the Monterey Institute of International Studies.

Michael R. Watts is an associate professor in the Electrical Engineering and Computer Science Department at the Massachusetts Institute of Technology. His research interests include photonic microsystems for low-power communications, sensing, and microwave-photonics applications. Watts has a PhD in electrical engineering from the Massachusetts Institute of Technology.

Direct questions and comments about this article to William A. Zortman, Sandia National Laboratories, Applied Photonic Microsystems, 1515 Eubank Blvd., MS-1082, Albuquerque, NM 87123; wzortm@sandia.gov.

Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.