Analysis of high-bandwidth low-power microring links for off-chip interconnects

Noam Ophir and Keren Bergman

Department of Electrical Engineering, Columbia University, New York, NY USA 10027.

ABSTRACT

Performance scalability of computing systems built upon chip multiprocessors are becoming increasingly constrained by limitations in power dissipation, chip packaging, and the data throughput achievable by the interconnection networks. In particular, today's systems based on electronic interconnects suffer from a growing memory access bottleneck as the speed at which processor-memory data can be communicated out of the chip package is severely bounded. Silicon photonics provide a CMOS-compatible solution for integrating high bandwidth-density off-chip optical I/O which can overcome some of these packaging limitations while adhering to pJ/bit-scale power efficiency requirements. Microrings in particular pose an attractive option for realizing optical communication functionalities due to their low footprint, low power dissipation, and inherent WDM-suitability due to their wavelength-localized operation. We analyze a terabit-per-second scale microring-based optical WDM link composed of current best-of-class devices. Our analysis provides quantitative measures for the maximal achievable bandwidth per link that could be reasonably realized within several years. We account for the full optical power budget to determine the achievable bandwidth as well as to enable a power consumption analysis including transmit and receive circuitry, photonic-device power dissipation, and laser power. The results highlight key device attributes that require significant advancement and point out the need for improvements in laser wall-plug efficiencies to provide sub-pJ/bit scale optical links.

Keywords: Silicon photonics, chip I/O, microrings, wavelength-division multiplexing.

1. INTRODUCTION

Performance scalability of computing systems built upon multicore architectures of chip multiprocessors (CMP) are becoming increasingly communication bound due to power dissipation, chip packaging, and data throughput constraints. Processors-to-memory communication has become a substantial determinant to overall system performance, as current I/O significantly under-provisions communications bandwidth to memory compared to the desired 1-Byte/Flop ratio of memory bandwidth to compute power.

While electrical-communication technology is advancing to provide low-power high-bandwidth I/O solutions, [1-3] these solutions are either limited to within-package multi-chip module interconnects or require the inclusion of relatively large footprint (hundreds of microns pitch) electrical RF connectors[2] to go beyond the package.

While optical solutions may not necessarily be more power efficiency in very short links compared to electronics (more so if restricted to non-monolithic integration), the relative immediate advantage of optics is the bandwidth density it can deliver over a single optical fiber. Furthermore, once fibers are used for transmission, the link performance is relatively link agnostic for < 1 km distances – providing potential for equal bandwidth connectivity between CMPs on a shared board, rack, or multiple separate racks.

Silicon microrings, in particular, are of interest for power-efficient dense-WDM solutions due to their relatively small footprint, low-power operation, and their inherent wavelength-selective operation. In the following sections, an unamplified bandwidth-maximized microring-based point-to-point optical link is analyzed in terms of the optical power budget and achievable aggregate bandwidth. The components are based on current reported best of class devices as well as several assumptions on performance improvements which are key to realizing such links.

Optoelectronic Integrated Circuits XV, edited by Louay A. Eldada, El-Hang Lee, Proc. of SPIE Vol. 8628, 86280N · © 2013 SPIE
CCC code: 0277-786X/13/$18 · doi: 10.1117/12.2007759

2. MICRORING LINK DESIGN

2.1 Integration options in intermediate future

Full monolithic integration of silicon photonics within a computing chip offers a significant potential performance boost thanks to close integration with the electronics,[4,5] but integration of photonics on the same silicon as the processors is likely to become commercially feasible in the long term (greater than 10 years) because it will require significant modifications and the addition of fabrication steps to existing chip-fabrication facilities. In the near term, board-level optical modules offer significant flexibility but don't provide a solution for getting the data in and out of the chip package. Therefore, we limit the discussion of feasible intermediate-future systems to integration of silicon photonic technology as separate optical dies within a package. These dies are connected to the processor either through a shared substrate, a silicon carrier, or die stacked. Such optical dies would include a limited amount of electronics—as much as required to drive the optical link.

2.2 Design objectives and assumptions

This analysis focuses on a link that maximizes aggregate WDM bandwidth per link (single data transmitting fiber) while keeping power efficiency at reasonable values. The design includes a dense WDM implementation with multiple ring modulators coupled to a shared bus waveguide on the transmit (Tx) side and similarly rings used for filtering from a shared bus waveguide for the wavelength demultiplexer (demux) on the receive (Rx) side as shown in Figure 1b.

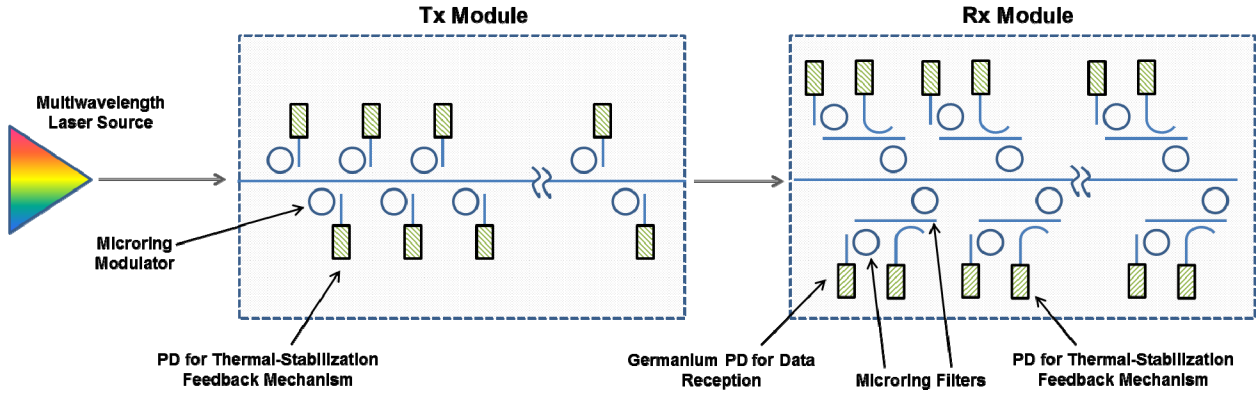

Figure 1. Depiction of a silicon-photonic microring-based WDM link design with Tx and Rx arrays. The modules are connected by fiber.

The design assumes the operational wavelength range is fully populated with equidistant resonances of the microrings which are laid out in a gradually increasing size order. The minimal ring radius for achieving Q values around 10,000 is assumed to be 1.7 μm which dictates a maximal operational wavelength span of ~ 50 nm.[6] For this analysis we assume operation around 1550 nm telecom bands since most reported results are currently for this wavelength range (though 1300 nm is a likely option for systems as well).

Given the assumed non-monolithic integration options assumed, the analysis focuses on two modultion rates per channel which seem likely candidate representative rates -12.5 Gbps and 25 Gbps. Though lower data rates have been shown in a similar study to have the potential for power efficiency optimality,[5] the data rate will likely heavily depend on packaging considerations and the number of pins allocated for communications to the optical module. These data rates seem to be emerging for electrical chip I/O circuitry regardless of optical technologies [12,13] and are therefore good candidates to be compatible with the electronic transceivers which will exist on the chip.

2.3 Link construction

The WDM link presented here (Figure 1) is composed of an off-chip multiwavelength laser source which feeds into a Tx module through polarization maintaining (PM) fiber. The Tx module is composed of a series of microring depletion-mode modulators coupled to a shared waveguide. Each modulator has a drop port with an integrated photodetector (PD) for thermal feedback for stabilization purposes.[9] The output of the Tx module is connected through PM fiber to the Rx module which has a microring-based spectral-demux. The demux is assumed to be composed of two cascaded

microrings to provide sufficient inter-channel-crosstalk suppression. The filtered channels are fed into Germanium avalanche photodiodes (APDs) for detection.

While the modulator drivers can be perhaps located relatively far away from the modulators without significant performance penalties, the receiver front end circuitry must be included in close proximity to the PDs to avoid extra parasitic capacitance. This implies that at least some electronics must be located within the optical module (by limited monolithic integration of dedicated electronics or filp-chip bonding).

2.4 Thermal stabilization

The rings span the full operational wavelength band with frequency-equidistant ring resonances. Assuming the Tx (or Rx respectively) rings laid out tightly on a single die, fabrication variations result in an nearly even shift of the all ring resonances (with up to 100 GHz random variation from the designed frequency comb spacing) [10,11]. Based on the cyclical nature of the ring resonances, this results primarily in a need to reassign the wavelength channels to the specific rings in the Tx or Rx rather than thermally tune each ring to a predetermined wavelength channel. We assume this is done once with a static barrel shifter rather than a dynamic barrel shifter concept [11] and therefore requires thermal tuning of each ring only to "snap" it to the nearest longer-wavelength wavelength channel. The amount of power required to achieve that is therefore proportional to the wavelength spacing of the channels.

We assume dynamic thermal fluctuations within the chip's package to reasonably be ~ 40 °C rather than the standard 100 °C specification. With a 10 GHz/°C resonance shift sensitivity,[12] this results in 400 GHz of dynamic shift to be thermally stabilized.

Overall power to be spent on fabrication-related thermal tuning as well as dynamic stabilization is related to both factors. If dense WDM (~100 GHz) channel spacings are considered, the dynamic thermal tuning requirements dominate over the fabrication variation compensation considerations. For coarse-WDM designs, this may not be the case.

3. OPTICAL POWER BUDGET ANALYSIS

As both insertion loss (IL) and power penalties increase as the wavelength spacing is reduced, we examine the overall loss and penalties for a range of channel spacings at both modulation rates. The loss and penalties determine how much laser power is required per channel as well as the aggregate laser power present in the bus waveguide.

Two primary nonlinear effects - two photon absorption (TPA) and TPA-induced free carrier absorption (FCA) will cause a sharp increase in propagation loss in the bus waveguide at high powers, we limit the aggregate optical power in the waveguide to be ~ 20 dBm (21 dBm launched into the edge coupler). Other potential limits on optical power lie in the stability bound within the microring modulators which limits the per-channel power reaching the microring modulator to ~ 6 dBm.[13]

In the rest of this section the loss and power penalties are analyzed for each module with a component-level breakdown and a module-level loss and penalties as function of channel spacing. The power penalties are computed assuming a shot-noise limited model as APDs are assumed at the receiver side.

3.1 Tx module

Beyond IL and power penalty arising from the use of a non-ideal physical device with limited extinction ratio and inherent loss, additional loss and penalties are accrued from the operation of these within an array of devices since the spectral features of the rings extend to adjacent wavelength channels.[6,14] Modulator Q values are set to be ~ 12,000 (cavity loss limited at small radii) for 12.5-Gbps modulation rate whereas Q ~ 8,000 (photon-lifetime limited) is used for 25-Gbps modulation rate.

Edge-coupler loss is assumed to be 1 dB/facet and bus-waveguide propagation loss is assumed 1.7 dB/cm.[15] The length of the waveguide is scaled proportionally to the number of modulators assuming 60-μm by 60-μm footprint per modulator and layout of modulators on both sides of the waveguide. Table 1 details the loss and penalties accrued in the Tx module.

Table 1. Loss and power penalties break down by component in the Tx module.

	12.5-Gbps Modulation Rate (channel spacing varied from 200 GHz down to 25 GHz)	**25-Gbps Modulation Rate** (channel spacing varied from 400 GHz down to 50 GHz)
Edge coupler (x2)	2 dB	2 dB
Waveguide IL	0.33 – 1.44 dB	0.25 – 0.8 dB
Modulator IL	4 dB [16]	3.5 dB [17]
Modulation Power Penalty	2 dB [16]	2.8 dB [17]
OOK-Modulation IL	2.4 dB [16]	2.1 dB [17]
Array-induced IL	0.06 – 4.42 dB	0.04 - 2.42 dB
Intermodulation Power Penalty	0 – 2.23 dB	0 – 0.65 dB

3.2 Rx module

As with the modulator array, operation of a WDM demux requires consideration of the channel spacing. In particular, design of the filtering stages as well as Q factors of each ring requires balancing of loss and sideband truncation penalties with inter-channel crosstalk penalties. Both filtering and crosstalk penalties are computed from analysis and simulation of the ring filtering response.

For dense WDM settings the crosstalk-related penalties outweigh array-induced losses and therefore a worst-case crosstalk (i.e, the first ring filter encountered in the array) is examined as the limiting reference for the optical power budget. A 2-stage filtering design achieves minimal overall losses plus power penalties at the channel spacings and modulation rates inspected in this work. The channel spacing crossover point beyond which a single stage of filtering outperforms two stages is ~ 180 GHz. Optical feedback for thermal stabilization is assumed to require a power tap and additional PD after the first stage of filtering. The optical feedback for the second stage of filtering is based on the APD used for receiving the data.

An additional 2-3 dB estimated power penalty is added for overall link jitter to account for temporal closing of the eye. Table 2 details the loss and penalties accrued in the Rx module filtering stages up to the APDs.

Table 2. Loss and power penalties break down by component in the Rx module.

	12.5-Gbps Modulation Rate (channel spacing varied from 200 GHz down to 25 GHz)	**25-Gbps Modulation Rate** (channel spacing varied from 400 GHz down to 50 GHz)
Edge coupler (x1)	1 dB	1 dB
Waveguide IL	0.09 dB	0.09 dB
Demux IL	1.5 dB	1.5 dB
Demux power penalty (crosstalk + sideband truncation)	0.17 – 3.02 dB	0.17 – 3.02 dB

3.3 Fiber and connector loss

Fiber and connector loss are assumed to be minimal 0.5 dB overall. Optical propagation to distances < 1 km are negligible in the examined system but single mode (SM) connectors might be a much bigger obstacle to realizing such systems because of the high cost, cleanliness required, and connection accuracy required to achieve low losses.

4. AGGREGATE ACHIEVABLE LINK BANDWIDTH

4.1 Germanium detector sensitivity

We interpolate a sensitivity of -16 dBm at 12.5 Gbps and assume a sensitivity of -12 dBm at 25 Gbps based on recent results.[18-19] Though some of these reported values are relevant for detection at short wavelengths, we assume that sufficient advancements can be made to enable these sensitivities to apply to 1550 nm. Closer integration with electronics (with lower parasitic capacitance) and improved APD designs provide potential for sensitivity improvements.[20] Therefore, we assume a uniform 4-dB improvement to be possible at all bit rates within a few years, which implies a sensitivity of -20 dBm at 12.5 Gbps and -16 dBm at 25 dBm.

High sensitivity receivers are one of the critical components that require further work. Theoretical bounds of closer integration of CMOS electronics with germanium structures give potential for much improved sensitivities – mostly due to low parasitic capacitances.[21] The integration of the trans-impedance amplifier (TIA) and limiting amplifier (LA) or sense amplifier (SA) in very close proximity to the germanium within the optical die is important to avoid excessive parasitics.

4.2 Optical power system limits

With the assumed receiver sensitivities and computed loss and penalty budget the required laser power per channel can be ascertained per channel spacing (Figure 2). Given an aggregate optical power limit in the silicon waveguides this implies maximal aggregate achievable of such a link design.

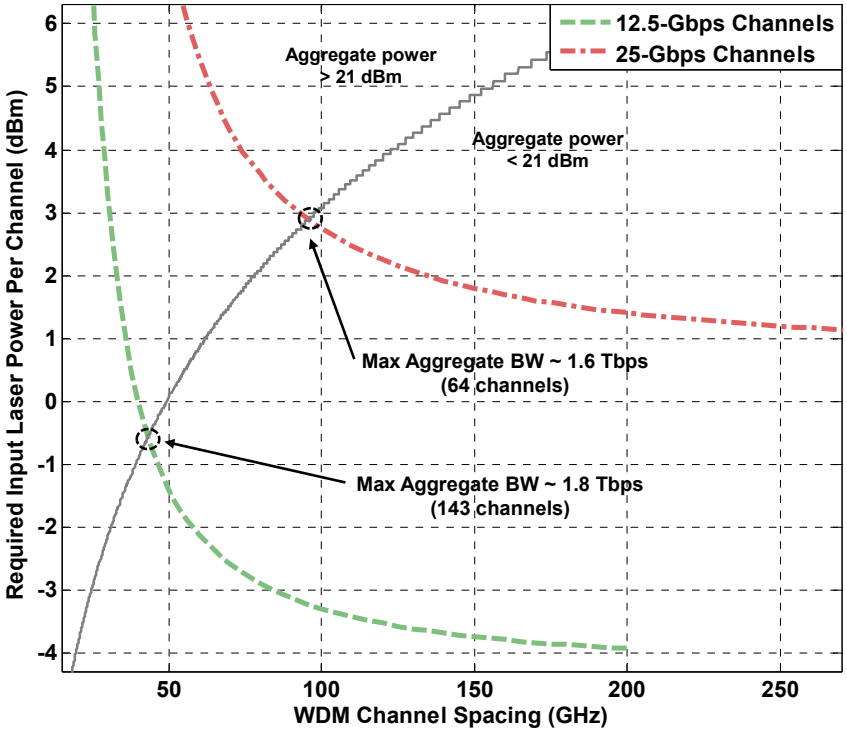

Figure 2. Required optical power per channel for both modulation rate cases. The grey line indicates the border of the area in which the aggregate launched power is less than 21 dBm. The intersection points with the grey line mark the maximal bandwidth operational point for both cases, corresponding to ~ 1.6 Tbps for 25 Gbps modulation rate and ~ 1.8 Tbps for 12.5 modulation rate.

Under the assumptions in this analysis, these limits come out to be 1.8 Tbps for 12.5-Gbps modulation rate and 1.6 Tbps for 25 Gbps. The inferior scalability of the 25-Gbps rate is mostly because of the fundamentally lower detector sensitivity at the higher data rate. As well, in both cases the per-channel optical power does not approach the microring-modulator instability threshold.

5. LINK POWER EFFICIENCY

As the required laser power is dependent on the channel spacing, we fix the aggregate at 1.55 Tbps for both modulation rate scenarios. For 12.5-Gbps modulation, this implies 124 wavelength channels whereas for 25-Gbps modulation it requires 62 wavelength channels. Power consumption of different components is based on reported values from literature as well as some extrapolation and assumptions.

For non-monolithic integration with a computing chip, electronic data communications will be required. Based on recent advances in low-power transceivers and assuming further scaling of power with smaller CMOS nodes, it seems realistic to achieve this with 1-2 pJ/bit power cost.[7,8]

Table 3 details the component power dissipation of the system. With current DFB technology of individually packaged lasers a 1% wall-plug laser efficiency would make the laser power overshadow all other link components. 10% wall-plug efficiencies which might be possible with multiwavelength lasers [23,24] (by amortizing the temperature control and cooling of a single cavity over many wavelength channels) make approaching a 1 pJ/bit target objective more realistic.

Table 3 – Link power efficiency with component / functionality / operation breakdown at 1.55 Tbps aggregate bandwidth.

	12.5-Gbps Modulation Rate (50-GHz spacing, 124 channels)	**25-Gbps Modulation Rate** (100-GHz spacing, 62 channels)
Microring modulation	0.01 pJ/bit [18]	0.01 pJ/bit [17]
Modulation driver	0.1 pJ/bit [18]	0.3 pJ/bit
Modulator thermal stabilization	0.15 pJ/bit	0.08 pJ/bit
Demux thermal stabilization	0.22 pJ/bit	0.11 pJ/bit
PD and receiver circuitry	0.4 pJ/bit [18]	1 pJ/bit
Laser source	5.8 pJ/bit @1% efficiency 0.58 pJ/bit @10% efficiency	7.5 pJ/bit @1% efficiency 0.75 pJ/bit @10% efficiency
Electronic data transmission to and from optical module	1 pJ/bit	2 pJ/bit

As seen from Table 3, only with a projected 10% laser wall-plug efficiency, and at 12.5-Gbps modulation rate does the link achieve an overall power efficiency of 2.5 pJ/bit which is within the ballpark for reasonable link operation. Closer integration on the chip with flip-chip bonding can reduce the excessive electronic data transmission costs. Along with further reduction of laser power (by reducing the channel density and therefore sacrificing some of the aggregate bandwidth) should allow the link to reach 1-pJ/bit operation.

6. CONCLUSIONS

As this case study showed, microring-based silicon photonics should be able to provide bandwidth densities as high as 1.8 Tbps on a single optical link requiring only three optical connectors to the optical die. The analysis also provides an estimate of the power efficiency of such a link approaching only a few pJ/bit.

Our analysis shows a high dependence of the overall power efficiency on laser wall-plug efficiency, driving the need for further development of higher-sensitivity receivers as well as more efficient multiwavelength laser sources. Furthermore, because laser power plays a key role in the overall efficiency, the efficiency is inversely proportional to the channel spacing and hence the link aggregate bandwidth.

As shown in the analysis, given some technology improvements silicon resonator-based devices could provide the potential for high-bandwidth, high-density, low-power I/O that could viably be inserted within the chip packaging solution and overcome the I/O bottleneck.

A more in depth analysis is published in reference 24.

REFERENCES

[1] Knickerbocker, J. U. et al., "3-D Silicon Integration and Silicon Packaging Technology Using Silicon Through-Vias," J. Solid State Circuits 41(8), 1718-1725 (2006).

[2] Jaussi, J. E. et al., "Multi-Gbit I/O and Interconnect Co-design for Power Efficient Links," Proc. IEEE Conf. on Electrical Performance of Electronic Packaging and Systems (EPEPS), pp. 1-4 (2010).

[3] Kim, B. et al., "A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS," J. Solid State Circuits 44(12), 3526-3538 (2009).

[4] Georgas, M. et al., "Addressing link-level design tradeoffs for integrated photonic interconnects," Proc. Custom Integrated Circuits Conference (CICC), 1-8 (2011).

[5] Orcutt, J. et al., "Open foundry platform for high-performance electronic-photonic integration," Opt. Express 20(11), 12222-12232 (2012).

[6] Sherwood-Droz, N. et al., "Device guidelines for WDM interconnects using silicon microring resonators," Proc. Workshop on the Interaction between Nanophotonic Devices and Systems (WINDS) (2010).

[7] Fukuda, K. et al., "A 12.3mW 12.5Gb/s Complete Transceiver in 65nm CMOS," Proc. IEEE Int'l Solid-State Circuits Conf., IEEE CS, no. 20.5 (2011).

[8] Proesel, J. E. and Dickson, T. O., "A 20-Gb/s, 0.66-pJ/bit Serial Receiver with 2-Stage Continuous-Time Linear Equalizer and 1-Tap Decision Feedback Equalizer in 45nm SOI CMOS," Proc. IEEE Symp. VLSI Circuits, IEEE CS, no. 20.1 (2011).

[9] Padmaraju, K. et al., "Thermal stabilization of a microring modulator using feedback control," Opt. Express 20(27), 27999-28008 (2012).

[10] Zortman, W. A. et al., "Silicon Photonics Manufacturing," Opt. Express 18(23) 23598-23607 (2010).

[11] Krishnamoorthy, A. V. et al., "Exploiting CMOS manufacturing to reduce tuning requirements for resonant optical devices," IEEE Photonics Journal 3(3) 567-579 (2011).

[12] DeRose, C. T. et al., "Silicon Microring Modulator with Integrated Heater and Temperature Sensor for Thermal Control," Proc. Conference on Lasers and Electro Optics (CLEO), CthJ3 (2010).

[13] Li, Q. et al., "Experimental characterization of the optical-power upper bound in a silicon microring modulator," Proc. IEEE Optical Interconnects, TuB5 (2012).

[14] Padmaraju, K. et al., "Intermodulation crosstalk from silicon microring modulators in wavelength-parallel photonic networks-on-chip," Proc. Photonic Society Annual meeting (LEOS), ThB5 (2010).

[15] Xia, F., Sekaric, L., and Vlasov, Y. A., "Ultracompact optical buffers on a silicon chip," Nature Photonics 1, 65-71 (2007).

[16] Dong, P. et al., "High-speed and compact silicon modulator based on a racetrack resonator with a 1 V drive voltage," Opt. Letters 35(19), 3246-3248 (2010).

[17] Li, G. et al., "25Gb/s 1V-driving CMOS ring modulator with integrated thermal tuning," Opt. Express 19(21), 20435-20443 (2011).

[18] Zheng, X. et al., "Ultra-efficient 10Gb/s hybrid integrated silicon photonic transmitter and receiver," Opt. Express 19(6), 5172-5186, (2011).

[19] Assefa, S. et al., "20Gbps receiver based on germanium photodetector hybrid-integrated with 90nm CMOS amplifier," Proc. Conference on Lasers and Electro Optics (CLEO), PDPB11 (2011).

[20] Assefa, S., Xia, F., and Vlasov, Y. A.,"Reinventing germanium avalanche photodetector for nanophotonic on-chip optical interconnects," Nature 464, 80-84 (2010).

[21] Miller, D. A. B., "Optical interconnects to electronic chips," J. Applied Optics 49(25), 59-70 (2010).

[22] Wojcik, G. L. et al., "A single comb laser source for short reach WDM interconnects," Proc. SPIE Photonics West, 7230-21 (2009).

[23] Lu, Z. G. et al., "Low noise InAs/InP quantum dot C-band monolithic multiwavelength lasers for WDM-PONs," Proc. Optical Fiber Communications (OFC), JWA27 (2009).

[24] Ophir, N. et al., "Silicon photonic microring links for high-bandwidth-density, low-power chip I/O," IEEE Micro, in press.