Physical Layer Analysis and Modeling of Silicon Photonic WDM Bus Architectures

Robert Hendry, Dessislava Nikolova, Sebastien Rumley, Noam Ophir, Keren Bergman
Columbia University
116th St. and Broadway
New York, NY 10027
{rh2519, dnn2108, sr3061, no2180, kb2028}@columbia.edu

ABSTRACT
Silicon photonic single mode links employing dense wavelength-division multiplexing (WDM) offer a large bandwidth-distance product, high bandwidth density, and potential for close integration with compute logic leveraging CMOS-compatible fabrication processes. These unique characteristics open a new and large design space for compute-systems' interconnects. This work analyzes several architectures in that design space, involving multiple on-chip data-transmitting sites organized around a photonic bus waveguide. We determine the maximum feasible bandwidth of these architectures through a comprehensive analysis of the optical power budget. We also estimate their power consumption and analyze the efficiency of the designs. The results indicate scalability of such links to Tb/s but at the same time illustrate and quantify an inherent trade-off between link complexity and the maximum achievable bandwidth.

Categories and Subject Descriptors
C.4 [**Performance of systems**]: Design studies, B.4.3 [**Input/output and Data Communications**]: Interconnections (Subsystems) - *Bus Topology*.

General Terms
Performance, Design

Keywords
Silicon Photonics

1. INTRODUCTION
As the demand for communication bandwidth in computing systems continues to grow, integrated silicon photonics is seen as a promising solution to meet the increasing demands for high bandwidth density, low power links. Research has shown many potential applications of silicon photonics, including network-on-chips (NoCs)[1],[2], processor-to-memory links [3],[4], and in network switches [5], to name a few. The characteristics of silicon photonic links, in particular those employing dense wavelength-division multiplexing (WDM), open a new and wide design space for interconnect network architectures. This paper focuses on a subset of this diverse design space: providing interconnection between the cores of a chip-multiprocessor (CMP) and an off-chip receiver (.e.g. a memory controller) via a multiple-writer, single-reader WDM Silicon Photonic bus, operated in a time-division multiplexing (TDM) way. Previous work [6] has shown an example application of such a network for delivering efficient, non-local memory access.

In order to achieve high aggregate bandwidth densities per photonic bus waveguide, WDM can be utilized. WDM allows independent data signaling on different wavelength signals (or channels) co-propagating on a single waveguide or fiber. Targeting such WDM designs, a new class of photonic devices has risen in the optics community based on compact optical resonators. Optical resonators formed from microrings, microdisks, and photonic-crystals structures [7-9] are uniquely suitable for WDM as each resonator interacts only with wavelengths which correspond to its resonant modes. These devices have an extremely small footprint (down to several micrometers wide) which results in low power operation as well as allowing integration of thousands of such components on a single die. They can also be leveraged to achieve different functionalities such as very fast optical modulation, switching, and filtering. However, special attention has to be paid in such systems to thermal stabilization of such devices and their unique device performance limitations. Furthermore, the exact physical spectral features of these elements have to be taken into consideration when calculating the power budget of a WDM system.

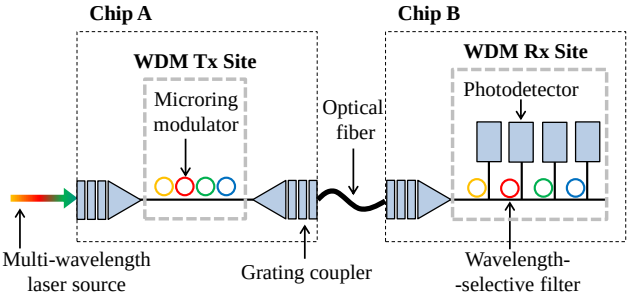

Figure 1. Illustration of a silicon photonic WDM shared bus with the transmitters and receiver on separate chips, connected by an optical fiber coupled to each chip. For simplicity, a link with only four parallel wavelengths is illustrated.

Dense WDM modulation can be accomplished by cascading microring modulators [7], [10], on the same waveguide and tuning each ring so that they resonate on separate wavelengths. A system with four WDM channels has been demonstrated [11], and with further advancement in fabrication techniques, demonstrations involving a much higher number of WDM channels are expected. However, beyond the limitation of the current fabrication capabilities, the maximum number of independent wavelengths on a single waveguide is limited by the optical loss on the link and by the maximal optical power tolerated by the link. In this work, we propose a modeling approach to estimate the losses on such links. Our approach includes the co-optimization of certain silicon photonic device properties, such as the wavelength selective filter

Table I. Summary of Loss and Power Penalty

Loss Contributor	Loss/Penalty
Modulating ring insertion loss	1 dB
Extinction ratio power penalty	2 dB
OOK insertion loss	2.4 dB
Modulator array loss	0.001 – 2.45 dB [1]
Filter insertion loss	0.5 dB
Filter power penalty	0.1 – 2.3 dB [1,2]
Switch through insertion loss	0.03 – 0.26 dB [2]
Switch drop insertion loss	0.52 – 1.8 dB [2]
Switch power penalty	0.04 – 1.5 dB [1,2]
Coupling loss	1 dB / coupler
Waveguide insertion loss	.92 dB / cm
Jitter penalty	2 dB

[1] Calculated as a function of wavelength channel spacing. [2]Optimized for each network size. These calculations and optimizations are described in some detail in the following sections.

quality factor (Q) value, for each WDM and link configuration examined. From these minimal loss estimates, we deduce the maximum feasible number of wavelength channels supported by a given architecture. As the number of wavelength channels directly translates into bandwidth, it is a key comparison metric.

Silicon photonic microring and microdisk modulators have been demonstrated at modulation rates of up to 30 Gb/s [12]. Before reaching a photodetector, the WDM signal is demultiplexed by another set of cascaded rings that act as wavelength-selective filters, such as those demonstrated by Xiao et al. [13]. The optical signals are then routed to individual photodetectors, which have been demonstrated with a sensitivity of -22 dBm [14]. Silicon waveguides are coupled to optical fibers to provide connectivity between multiple silicon chips using couplers such as [15] or [16]. Lastly, microrings can be used to switch light from one waveguide to another. If the microring is designed such that its free spectral range (FSR) is equal to the spacing between the WDM wavelengths, it acts as a multi-wavelength switch (called a "comb switch") which can switch an entire set of WDM signals simultaneously, as demonstrated in [17]. Provided a multi-wavelength laser source (for example [18]), these components form the building blocks for a diverse design space of very high bandwidth interconnection links. Figure 1 illustrates all of these components in the context of a simple point-to-point silicon photonic link where the transmitter and receiver are on separate chips.

The following sections present our approach and report the power and performance predictions gathered by applying it to various link architectures. These results illustrate a trade-off between the size (determining the number of CPUs and memory sites connected) and complexity (i.e. number of devices and channels used) of a silicon photonic WDM link and the maximum number of parallel wavelength channels that can feasibly be employed on that link. This trade-off will be present in any silicon photonic network architecture and requires careful consideration at design time. In addition, we provide a power consumption analysis for each of the architectures that captures the cost associated with making architectural changes and the tradeoffs toward maximizing performance.

Section 2 describes the models for the various silicon photonic devices and the optimization of some of the microring parameters in the context of the considered network architectures. Section 3 proposes various link designs and shows how architectural changes affect link performance. The final section concludes the methods and results of this work.

Table II. Summary of Power Contributions

Device	Type/Origin	Power/Device (mW)
Modulator	Thermal	0.875 [1]
	Driver circuitry	1.35 [23]
	Dissipation in ring	0.1 [24]
Switch	Thermal	3.5 [1]
Filter	Thermal	0.875 [1]
Detector	Static	3.95 [23]
Laser	Static	1250 [2]

[1] Derived from equations in [22]. [2] 125mW laser delivered to first coupler with 10% wall-plug efficiency.

2. SILICON PHOTONIC LINKS

2.1 Optical Power Budget

In order to ensure that signals propagating through a silicon photonic link reach their destination before attenuating bellow the sensitivity threshold of the receiver, the loss along that link must fall within an acceptable range. This constraint is called the *optical power budget* and can be summarized as:

$$P_{Budget}^{dB} \geq P_{Loss}^{dB} + 10log_{10}(N_{\lambda}) \quad (1)$$

This model is described in detail in [19]. P_{Budget}^{dB} is the maximum power dissipation allowed in the link relative to the input power (in dB), P_{Loss}^{dB} is the total loss (in dB), and N_{λ} is the number of wavelength channels used in the link. There is an upper bound, T_1, on the aggregate amount of optical power that can be injected into a waveguide due to emergence of nonlinear effects (primarily two photon absorption) at high power levels. This upper bound represents the power level above which silicon photonic devices are expected to exhibit non-linear behavior that causes too significant signal distortion. There is also a lower bound, T_2, representing the receiver sensitivity, i.e. the minimal amount of power required for each receiver to successfully detect the transmitted information. P_{Budget}^{dB} is defined simply as $T_1 - T_2$.

Shot-noise limited avalanche photodetectors, as reported in [14] can be projected to have sensitivity of -22dBm with non-return-to-zero modulation at 10 Gb/s, i.e. T_2 = -22dBm. Non-linear effects in silicon waveguides start to become significant around 100mW [20]. Since a coupler incurs a loss of 1dB [16], the maximum optical input power for a silicon waveguide is 125mW or 21dBm, thus T_1 = 21 dBm. This yields a 43dB optical power budget. P_{Loss}^{dB} is calculated using device models detailed in the following subsections. This loss represents the sum of the contributions of all the components present along a link.

Each wavelength channel must be injected into the link with enough input power such that the injected power minus the loss on the link will still be greater than the sensitivity of the photodetector that ultimately receives the modulated signal. At the same time, there is a finite amount of loss that can be overcome by increasing input power due to the non-linear threshold. Therefore, the maximum number of wavelengths that can feasibly be used in a link can be calculated from Equation (1).

2.2 Modeling Loss

The components included in this study are limited to waveguides, microrings acting as switches, modulators, or passive filters, fiber-waveguide couplers and photodetectors (losses in fibers are assumed negligible for this study for the sub-km reach relevant to

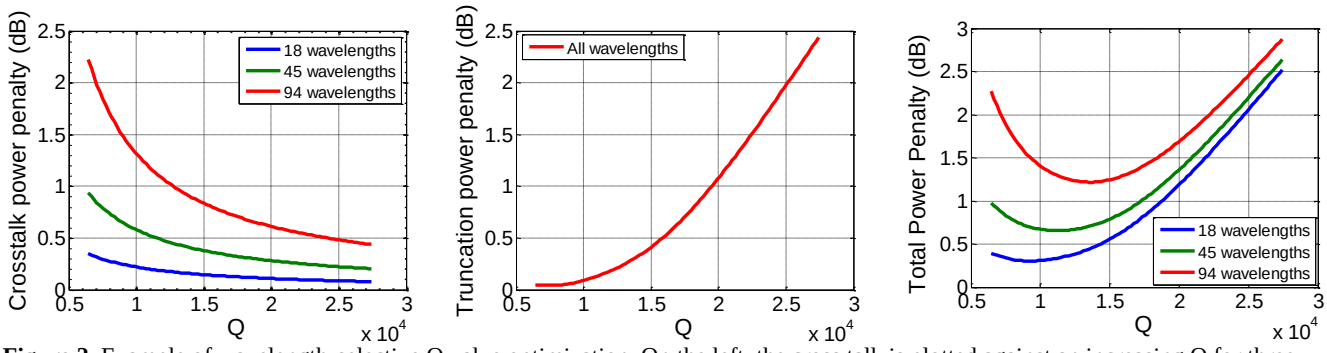

Figure 2. Example of wavelength-selective Q value optimization. On the left, the cross talk is plotted against an increasing Q for three different numbers of wavelengths for a 16-site multiple-writer single-reader bus. The middle plots the truncation caused by the narrowing of the filter as the Q increases. The figure on the right plots the sum of these contributions, illustrating how the optimum Q value differs depending on the number of wavelengths used.

these systems). These losses are summarized in Table I. The values presented as ranges differ depending on both or either the link architecture and/or the number of wavelengths employed. These ranges indicate the minimum and maximum values used.

2.2.1 Insertion Loss

As light propagates through silicon photonic devices, optical power is lost. The ratio of power lost is termed insertion loss. These losses can vary due to the material used the variability and limitations of the fabrication processes, and fundamental properties of the optical structures. Overall, the insertion loss for most silicon photonic devices is expected to decrease as research and technology improvements reveal better methods of fabrication.

Microring modulators can imprint amplitude modulation on light by varying the loss through the device. This is achieved by slightly shifting the device resonance between two states. In the logic "0" state, light is coupled on resonance to the device away from the waveguide. In the "1" state, light passes slightly off resonance with a much smaller portion of light coupled out of the bus waveguide (however, still resulting in some loss). Therefore, even inactive modulator banks (set to the "1" state) contribute to insertion loss if they share the same bus waveguide. We denote as *ring insertion loss* the insertion loss due to coupling to modulator rings to on-resonance wavelengths. *OOK insertion loss* represents the average optical power lost through the act of amplitude modulation (extinguishing light in the "0" state bits), under the assumption of equal probability of 1 and 0 bits transmitted. *Filter insertion loss* stands for the insertion loss that occurs when a signal is filtered through a microring filter before detection. The *waveguide insertion loss* is simply the loss per unit length as light travels down a waveguide. Each waveguide-fiber interface, required to move light on and off chips, incurs a 1dB *coupling loss*.

The *switch through insertion loss* and *switch drop loss* describe the power losses corresponding to the two possible output ports of a microring switch [17]. From the physical operation of a ring switch it comes out that these two loss values are inversely proportional and can in fact be controlled at design time by changing the coupling of the ring to the waveguide. These losses are computed as a function of the ring coupling coefficients according to the ring transmission models in [21]. The optimum coupling coefficients that minimizes the total loss depends, again, on the link architecture and is subject to optimization.

2.2.2 Power penalty

As signals are switched and filtered, impairments are introduced due to inter-wavelength-channel crosstalk and signal sideband truncation. These impairments can be accounted for by increasing the input signal power to maintain an acceptable signal-to-noise ratio. Similar to the switch insertion loss, the *filter power penalty* depends on parameters that can be controlled at design time. Specifically, the Q factor of the microring filters can be controlled by varying the width of the waveguide in the microring. Figure 2 illustrates the tradeoff between a higher Q factor and the total power penalty. The optimal Q value changes for different values of N_λ. A higher Q results in less crosstalk, but more truncation. The best Q value to choose, therefore, depends on the spacing between wavelength channels. P_{Loss}^{dB} varies between link architectures of different size and complexity (which will be detailed in the next section) and consequently so does N_λ and, in effect, the channel spacing. This relationship implies that there is an optimum filter Q value for each architecture that minimizes the combined effect of crosstalk and signal truncation. Finally, a 2dB *jitter penalty* due to imperfect clocking is assumed.

2.3 Modeling Power

All active photonic components present along the link must be powered and induce power consumption. The power consumption from these devices is summarized in Table II. Detailed discussion of the energy of optical links can be found in [25].

During operation of a modulator, i.e. when data is modulated on the carrier signal, the consumed power is determined by the power consumption of the *driver circuitry* and the power *dissipated in the ring* by the photocurrent. Note that only one modulator bank (i.e. a series of N_λ cascaded modulators) is active at any given time on the network. Due to fabrication variation, the initial wavelength position of each microring's resonance needs to be adjusted to match the operating wavelength of the optical path. This is done by using *thermal* tuning [22], [26]. In addition, these microrings operate in a thermally volatile environment which can shift the resonance. The silicon ring resonance has a 10 GHz/°C sensitivity (resonance shift per temperature change). The microrings of the modulators, switches, and filters all require thermal trimming and tuning power to compensate for fabrication inaccuracies and dynamic thermal shifts in the environment in order to maintain network operation. Therefore, such resonant devices exhibit significant static power dissipation, i.e. relatively fixed regardless of whether the device is actively transmitting at a given moment.

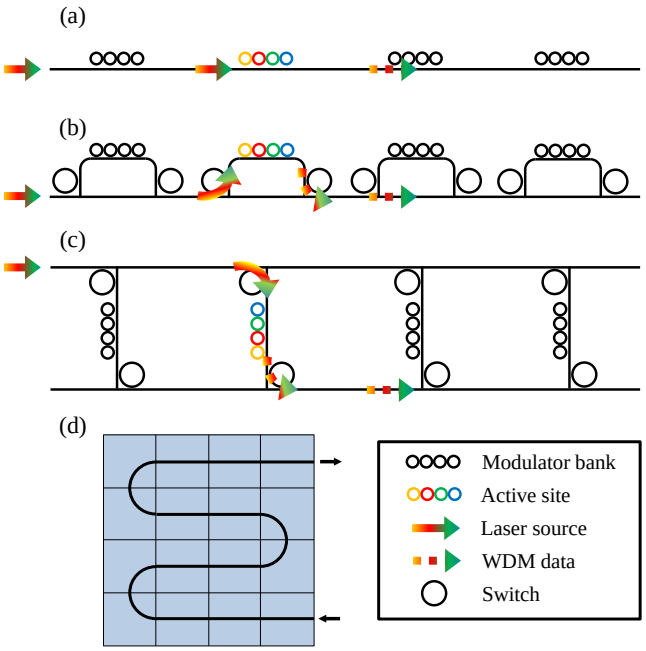

Figure 3. (a) Basic architecture. (b) Switched architecture. (c) Dual-waveguide architecture. (d) Layout of bus on a CMP.

Table III. Architecture Cost Functions

LOSS AND POWER PENALTY			
Contributor	**Occurrences**		
	Basic	**Switched**	**Dual-waveguide**
Modulating ring insertion loss	N	C	C
Extinction ratio power penalty	1	1	1
OOK insertion loss	1	1	1
Modulator array loss	N	C	C
Filter insertion loss	1	1	1
Filter power penalty	1	1	1
Switch through insertion loss	0	$2(\frac{N}{C}-1)$	$\frac{N}{C}-1$
Switch drop insertion loss	0	2	2
Switch power penalty	0	1	1
Coupling loss	3	3	3
Waveguide insertion loss	L	L	L
Jitter penalty	1	1	1
POWER CONSUMPTION			
Modulator	$N \times N_\lambda$	$N \times N_\lambda$	$N \times N_\lambda$
	N_λ	N_λ	N_λ
	N_λ	N_λ	N_λ
Switch	0	$\frac{2N}{C}$	$\frac{2N}{C}$
Detector	N_λ	N_λ	N_λ
Laser	1	1	1

Note: N is the number of network sites, *C* is the number of sites per cluster, *L* is the length of the waveguide, and N_λ is the number of wavelength channels.

There are at most only two switches with their resonance actively being shifted at any single time and the resulting power is negligible (as it is amortized by multiple wavelength channels simultaneously switched by each). The total power required for the laser to supply 125mW optical power to the link is calculated assuming 10% wall-plug efficiency. The total consumption of the receiver accounts for all the power consumed by the receiver circuits during 10Gb/s operation [23].

3. TDM BUS ARCHITECTURES

3.1 Basic and Switched Architectures

In this section, we present how the elements presented so far can be assembled to provide connectivity between different computer components – in our case, between a multi-core chip and an off-chip receiver. Reference [27] provides a thorough analysis of a point-to-point silicon photonic link that could be used to deliver high bandwidth off a chip. In the present case, to service multiple cores in many-core processors, we propose a multiple-writer, single-reader bus that spans the area of the chip. Figure 3 illustrates three possible ways to accomplish such connectivity that are analyzed in this work. Each architecture assumes TDM operation, so that at most one modulation site is active at any point in time.

The architectures shown on Figures 3(a) and 3(b) are termed the *basic* and *switched* architectures, respectively. The light is coupled to a waveguide which transverses the chip passing by all cores, illustrated in Figure 3(d). Each core has a modulator bank allowing it to modulate data onto the incoming light and transmit it to the detector chip. In the basic architecture, in the worst case, the modulated data will pass by the modulator banks of all cores. The switched architecture removes the modulator banks from the common waveguide so that data does not pass by other modulator banks, but must pass by two switches per bank instead.

The cascading of modulation sites in the *basic* architecture becomes problematic as the number of modulation site increases. Each modulator ring induces some loss in the optical signal passing down the waveguide, causing the link optical loss to rise. This relationship sets a limit to the scalability of such a link.

The *switched* architecture utilizes comb switches to improve the scalability of the link. Here, when it is time for a site to begin modulating (i.e. its TDM slot starts), the nearby comb switches are shifted on resonance and light is switched onto the waveguide on which the modulators are coupled. The data is modulated onto the career signal which is then switched back onto the waveguide bus. In this design, adding a modulation site corresponds to adding only two comb switches (i.e. two rings) to the main waveguide bus (as opposed to one ring for every wavelength channel). When large numbers of parallel wavelengths are used, this simple alteration is a dramatic improvement over the initial design, as it will be shown in the next sub-section.

Figure 3(c) illustrates an alternative way to accomplish the switched architecture using two parallel waveguides. In this configuration, input optical power travels down one waveguide, it is switched onto the active modulation site, and then the modulated data is switched onto the second parallel waveguide. The advantage here is that there is only one *switch through insertion loss* incurred for every modulation site.

3.2 Initial Performance Comparison

We applied our approach to estimate the maximum number of wavelengths that can be employed in each architecture. All the losses listed in Table I have to be computed for each device on the bus. The number of times each loss is accounted for is given in Table III. The sum of all losses in the network gives the total loss P_{Loss}^{dB}. The maximum number of wavelengths is the maximum value of N_λ that still satisfies Equation 1.

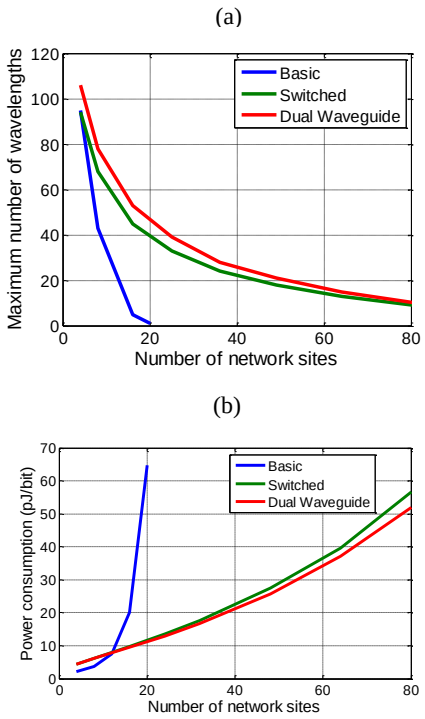

Figure 4. Analysis of basic, switched, and dual-waveguide architectures when C=1. (a) The maximum number of wavelengths that can be used in parallel on each link. (b) The power efficiency of each link.

The length of the waveguide L is calculated assuming a 2cm by 2cm chip where the modulation sites are laid out in a square grid and the waveguide "snakes" back and forth in order to pass by all of the sites, as illustrated in Figure 3(d). In this layout, each core requires approximately a $2cm/\sqrt{N}$ length of waveguide, either running horizontally or both horizontally and vertically to turn the waveguide. Therefore, the total length of the waveguide is approximately $L = (2cm \times \sqrt{N})$. In this analysis, 10 Gb/s modulation per wavelength is assumed. The total available spectral bandwidth is 50 nm, centered on 1550 nm, which is about the maximum FSR possible under realistic design parameters [20].

Figure 4(a) plots the maximum number of wavelengths that can be employed in each architecture as a function of the number of modulation sites sharing the waveguide. The maximum number of wavelengths employed on each link falls off quickly as sites are added to the shared bus. This reduction is due to both the loss incurred by increasing the number of devices coupled to the waveguide and the increased length of waveguide required to reach all the sites. The basic link architecture quickly loses performance and is infeasible with more than 20 sites. The switched architectures generally yield a higher number of wavelengths and allow over 80 modulation sites to share the same waveguide.

Figure 4(b) similarly plots the power consumption of each architecture, calculated by determining the total power requirement of each configuration and dividing by the achievable aggregate peak bandwidth on the link, which is simply the product of the number of wavelengths and 10 Gb/s. This metric ignores the time for the switches to change state, and so it is a measure of power efficiency at peak bandwidth with 100% link utilization. Although these assumptions may not reflect typical real-life

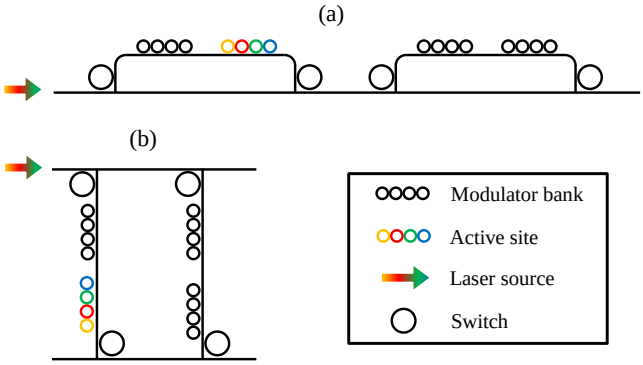

Figure 5. (a) Clustered switched architecture. (b) Cluster dual-waveguide architecture.

operation of such a link, they provide a medium for comparing the link architectures.

The addition of comb switches incurs additional thermal tuning power. This extra power is unfavorable when the number of sites is small and the energy efficiency of the three architectures is roughly the same. However, by 16 sites the bandwidth improvement in the switched architectures yields a better power efficiency than the basic architecture. The switched architectures' power efficiency continues to worsen with the number of sites because, although there are many modulation sites, only one site is active at any given time and yet all of the modulators and corresponding comb switches require thermal tuning.

The addition of comb switches enables link architectures with many modulation sites sharing a single waveguide. Due to the inherent TDM nature of the link, however, the power efficiency worsens with the number of sites. From these results one can conclude that when some target system requires 8 sites or less, the dual-waveguide can support the most wavelengths, but the basic architecture is the most power-efficient. When the system has 16 sites or more, the switched architectures are better in both metrics. Although the power efficiency here appears to compare poorly to typical link transmission efficiencies (generally targeting < 1pJ/bit), this design is not necessarily power inefficient. In fact, in the link with four modulation sites, the transmission power cost is kept below 2 pJ/bit transmission. Additionally, the TDM bus architecture provides networking functionality. Increases in the power consumption are therefore also to be expected with a regular electronic based equivalent system.

3.3 Effective Bandwidth

In the switched architectures, the comb switches take some non-negligible amount of time to change state. A guard time must be introduced between TDM slots in order to ensure that the switches have fully changed state before data begins to propagate through them. This required guard time reduces the bandwidth of the links. Therefore, in the switched architectures, the maximum effective bandwidth of the link is not simply the product of the number of wavelength channels and the data rate. In order to compare the architectures in terms of maximum effective bandwidth, the link efficiency η must be included. It is defined as:

$$\eta = \frac{t_{data}}{t_{data}+t_{guard}} \quad (2)$$

where t_{data} is the time period for which data is being serialized onto the wavelength channels during a TDM slot, and t_{guard} is the guard time required in between time slots to guarantee the switch has fully changed state before signals propagate through the

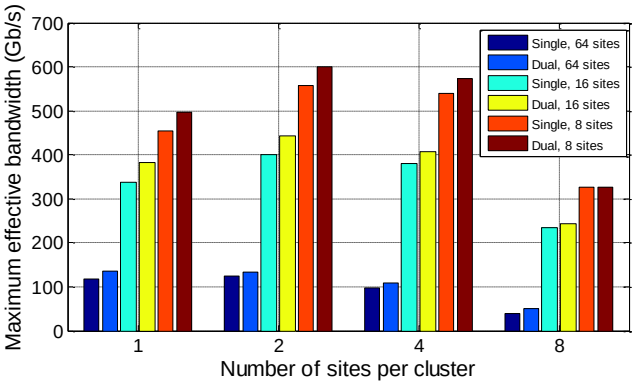

Figure 6. The maximum effective bandwidth as a function of the number of modulation sites per cluster, for 8-site links, 16-site links, and 64-site links, assuming a 4096bit message size and 3ns switching time.

device. Giving this efficiency, the effective bandwidth of each architecture is:

$$BW_{eff} = \eta N_{\lambda} D \quad (3)$$

where D is the data rate per wavelength channel. The basic architecture has an efficiency of 1, because the activation or deactivation of modulator banks happens at the same rate as the data rate; deactivating a modulator bank is simply turning all the rings to the same state that is used to modulate a "one" bit.

3.4 Clustered Architectures

The effects of the guard time on the maximum bandwidth can be mitigated by combining the switched architecture approach with the basic architecture's cascaded modulation sites, as illustrated in Figure 5(a) and 5(b). We call these link architectures *clustered* architectures. In Table III, the term *C* represents the number of modulation sites per cluster. Figure 6 plots the maximum achievable bandwidth for both the switched and the dual waveguide architectures using 1, 2, 4, or 8 sites per cluster, assuming a 3ns guard time and a 4096-bit message size (i.e., the TDM slots are exactly long enough to serialize 4096 bits divided amongst all the wavelength channels). The 3ns value corresponds to the switching time of a comb ring, as reported in [5]. The maximum effective bandwidth is calculated using equations (2) and (3). As it appears in Figure 6, there is an optimal design configuration for each network size.

Guard times are required when the light is being switched from one cluster to another. During this time, no modulation can occur. Increasing the number of modulation sites per cluster diminishes the number of guard times required, and consequently improves the efficiency. However, increasing the number of modulation sites induces more loss, and therefore reduces the number of available wavelength channels. Assuming 4096 bits are modulated during each TDM slot and a 3ns switching time, the link architectures with two modulation site per cluster maximizes the effective bandwidth. If shorter messages are considered, the guard time impact increases, leading to optimal architectures with more sites per cluster.

Finally, Figure 7(a) plots the maximum effective bandwidth for the basic, switched, and dual waveguide switched architectures using the best clustering configurations for each network size. This plot illustrates how, despite the guard times required for switching, the performance and scalability of the silicon photonic multiple-writer, single-reader bus is drastically improved by introducing ring switching elements in appropriate places along

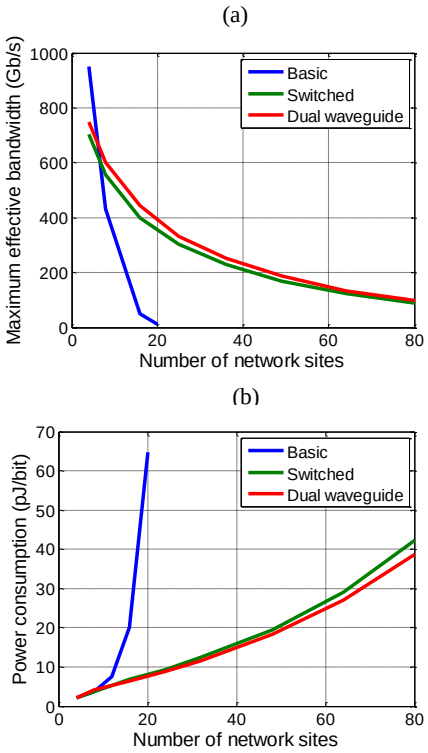

Figure 7. (a) Maximum effective bandwidth for each type of bus architecture, with only the best clustering configurations shown. (b) Power consumption of architectures with the same clustering as above (optimized for bandwidth, not power).

the waveguide. Note that even though the maximum number of wavelengths for the switched architectures is higher than for the basic one as shown on Figure 4(a), due to the guard time they have lower maximum effective bandwidth for less than 8 network sites.

Figure 7(b) shows the energy efficiency of the proposed network for the optimal clustering in the case of the switched architectures. For less than 8 sites the basic bus has the lowest power consumption. For increased number of interconnected sites switched architectures have significantly lower consumption per bit. The reported values for larger networks (i.e. bigger than 16 sites) are comparable with the efficiency of nowadays electronic interconnects with the same connectivity (10-50 pJ/bit) [25].

4. CONCLUSIONS

The paper proposes an approach to analyze the maximum performance achievable in various silicon photonic TDM bus architectures. First, it was shown that silicon photonic device models can, and should, be optimized to suit particular link configurations. Without this optimization, the performance of the links would be underestimated. Then, the device models were applied to show that introducing comb switches can drastically improve the scalability of such a bus. A performance and size/complexity tradeoff inherent in silicon photonic link design was illustrated. As buses become longer and more complex, less wavelength parallelism should be expected. Lastly, it was demonstrated how clustering can further improve the performance and scalability of these links. As parallelism in CMPs continues to grow, such links can be utilized to provide high bandwidth chip I/O, providing connectivity to many injection sites spread throughout a chip.

ACKNOWLEDGMENTS

The authors would like to thank Takashi Shiraishi and Kishore Padmaraju for their advice regarding Silicon Photonic device modeling. This research is supported in part by the Department of Energy under grant DE-SC0005114.

REFERENCES

[1] Bergman, K., 2007. Silicon Photonic On-Chip Optical Interconnection Networks. *LEOS* (October 2007).

[2] Petracca, M., et al., 2008. Design Exploration of Optical Interconnection Networks for Chip Multiprocessors, *16th IEEE Symposium on High Performance Interconnects* (August 2008).

[3] Stojanovic V., et al., 2010. Design-Space Exploration for CMOS Photonic Processor Networks. *OFC* (March 2010).

[4] Wang H., et al., 2008. Nanophotonic Optical Interconnection Network Architecture for On-Chip and Off-Chip Communications. *OFC* (2008).

[5] Biberman A., et al.,2011. Broadband Silicon Photonic Electrooptic Switch for Photonic Interconnection Networks, *IEEE Photonic Technology Letters*, vol. 23, no. 8 (April 2011).

[6] Whelihan, D., et al. 2013. P-sync: A Photonically Enabled Architecture for Efficient Non-local Data Access. *International Symposium on Parallel and Distributed Processing* (2013), pp.189-200.

[7] Bogaerts et al., 2012.Silicon microring resonators, *Laser Photonics Rev.*, vol. 6, no. 1 (2012), pp. 47-73.

[8] Soltani, M., Yegnanarayana, S., Adibi, A. 2007. Ultra-high Q planar silicon microdisk resonators for chip-scale silicon photonics. *Optics Express*, vol. 15, no. 8 (2007).

[9] Altug, H., Vuckovic, J. 2003. Two-dimensional coupled photonic crystal resonator arrays. *Applied Physics Letters*, vol. 84, no. 2 (2003).

[10] Li, G. et al, 2013. Ring Resonator Modulators in Silicon for Interchip photonic links. *IEEE Journal on Selected Topics in Quantum Electronics*, vol. 19, no. 6 (2013).

[11] Manipatruni, S., Chen, L., and Lipson, M. 2010. Ultra high bandwidth WDM using silicon microring modulators. Optics Express, vol. 18, no. 16, 2010.

[12] Rosenberg, J. C., et al. 2011. High-speed and low-power microring modulators for silicon photonics. *IEEE Photonics Conference* (Oct, 2011).

[13] Xiao, S., Shen, H., Khan, M. H., Qi, M.. 2008. Silicon microring filters. *Conference on Lasers and Electro-Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies*, OSA Technical Digest (CD) (Optical Society of America, 2008), paper JWA84.

[14] Masini, G., et al. 2012. CMOS photonics for optical engines and interconnects. *OFC/NFOEC Technical Digest* (2012).

[15] Selvaraja, S. K., et al. 2009. Highly efficient grating coupler between optical fiber and silicon photonic circuit. *Conference on Lasers and Electro-Optics/International Quantum Electronics Conference*, OSA Technical Digest, (2009).

[16] Pu, M., Liu, L., Ou, H., Yvind, K., Hvam, J. 2010. Ultra-low-loss inverted taper coupler for silicon-on-insulator ridge waveguide. *Optics Communications*, vol. 283, no. 19 (October 2010).

[17] Lee, B., Biberman, A., Dong, P., Lipson, M., Bergman, K. 2008. All-optical comb switch for multiwavelength message routing in silicon photonic networks. *IEEE Photonics Technology Letters*, vol. 20, no. 10 (2008).

[18] Zheng, X., et al. 2013. Efficient WDM Laser Sources Towards Terabyte/s Silicon Photonic Interconnects. *Journal of Lightwave Technology*, vol. 31, no. 15 (December 2013).

[19] Biberman, A., Preston, K., Hendry, G., Sherwood-Droz, N., Chan, J., Levy, J. S., Lipson, M., Bergman, K.. 2011. Photonic Network-on-Chip Architectures Using Multilayer Deposited Silicon Materials for High-Performance Chip Multiprocessors. *ACM Journal on Emerging Technologies in Computing Systems*, (June 2011).

[20] Preston, K., et al. 2011. Performance guidelines for WDM interconnects based on silicon microring resonators. *CLEO* (2011).

[21] Yariv, A., Yeh, P. 2007. *Photonics: Optical electronics in modern communications*. Oxford University Press (2007).

[22] Padmaraju, K., Logan, D.F., Zhu, X., Ackert, J. J., Knights, A. P., Bergman, K. 2013. Integrated thermal stabilization of a microring modulator. *Optics Express*, vol. 21, no. 12, pp. 14342–14350 (2013).

[23] Zheng, X., et al. 2011. Ultra-efficiency hybrid integrated silicon photonic transmitter and receiver. *Optics Express*, vol. 19, no. 6, pp. 5172-5186 (2011).

[24] Li, G., et al. 2011. 25 Gb/s 1V-driving CMOS ring modulator with integrated thermal tuning. *Optics Express*, vol. 19, no. 21, pp. 20435-20443 (2011).

[25] Krishnamoorthy, A. V., et al., 2009. Computer systems based on silicon photonic interconnects. *Proceedings of the IEEE*, vol. 97, no. 7. (July 2009).

[26] Krishnamoorthy, A.V., et al., 2011. Exploiting CMOS Manufacturing to Reduce Tuning Requirements for Resonant Optical Devices. *IEEE Photonics Journal*, vol 3, no. 3 (June 2011).

[27] Ophir, N., Bergman, K. 2013. Analysis of high-bandwidth low-power microring links for off-chip interconnects. *Proceedings of SPIE*, vol. 8628 (2013).