ORB: An On-chip Optical Ring Bus Communication Architecture for Multi-Processor Systems-on-Chip

Sudeep Pasricha, Nikil Dutt

Center for Embedded Computer Systems
University of California, Irvine, CA
{sudeep, dutt}@cecs.uci.edu

Abstract – **As application complexity continues to increase, multi-processor systems-on-chip (MPSoC) with tens to hundreds of processing cores are becoming the norm. While computational cores have become faster with each successive technology generation, communication between them has become a bottleneck that limits overall chip performance. On-chip optical interconnects can overcome this bottleneck by replacing electrical wires with optical waveguides. In this paper we propose an optical ring bus (ORB) based on-chip communication architecture for next generation MPSoCs. ORB uses an optical ring waveguide to replace global pipelined electrical interconnects while preserving the interface with today's bus protocol standards such as AMBA AXI. We present experiments to show how ORB has the potential to provide superior performance (more than 2×) and significantly lower power consumption (a reduction of more than 10×) compared to traditionally used pipelined, all-electrical bus-based communication architectures, for 65-22 nm technology nodes.**

I. Introduction

With technology scaling down into the ultra deep submicron (UDSM) domain, component integration levels are increasing rapidly to allow billions of transistors and hundreds of cores to be integrated on a single chip. Such multi-processor system-on-chip (MPSoC) designs require efficient on-chip communication architectures to support high data bandwidths and increase parallelism. Unfortunately, deep submicron effects such as capacitive and inductive crosstalk coupling [1] are becoming highly dominant in new technologies, leading to an increase in propagation delay of signals on traditional electrical (i.e. copper) interconnects. Additionally, since in synchronous digital design a signal must propagate from source to destination within a single clock cycle to ensure correct operation, global interconnects that span the chip (and can be several mm in length) have to be clocked at very low frequencies. Such low clock frequencies on global interconnects, coupled with increasing propagation delay, puts serious limits on the achievable bandwidth and overall system performance. According to the International Roadmap for Semiconductors (ITRS), global interconnect delay has already become a major source of performance bottlenecks and is one of the semiconductor industry's topmost challenges [2].

To reduce global interconnect delay, designers today make use of repeater insertion on interconnects [3] to transform the quadratic dependency of propagation delay on interconnect length into a linear one. Another technique that is frequently used in addition to repeater insertion is to pipeline global interconnects by inserting flip-flops, latches or register slices [4-5]. Pipelining allows signals to travel shorter distances (i.e. the segment length from one stage to the next) in a single clock cycle. This enables the global interconnect to be clocked at higher clock frequencies and potentially support larger bandwidths. Fig. 1 shows an MPSoC design with four computation clusters that communicate with other clusters using pipelined global interconnects. Each cluster has bus-based local interconnects that can handle high data bandwidths of the cores in the cluster locally. The global interconnects in such systems can be shared or be point-to-point, and are operated at higher frequencies. Even though the number of global interconnects in a design is typically much less than the number of smaller local interconnects, they are the primary source of bottlenecks [1] and hence critically affect overall performance.

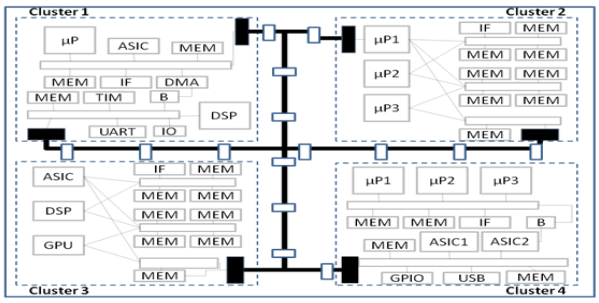

Fig.1. Traditional multi-cycle (pipelined) on-chip global communication in MPSoCs

Pipelined global electrical interconnects, such as the ones seen in Fig. 1 have two serious drawbacks. Firstly, a large number of pipeline stages are inevitably required for MPSoCs with high bandwidth requirements, resulting in high data transfer latencies. Secondly, the large number of latches (and repeaters) required to support multi-GHz frequencies for high performance MPSoCs leads to high power dissipation. These drawbacks are due to the fundamental limitations of using copper as a global interconnect.

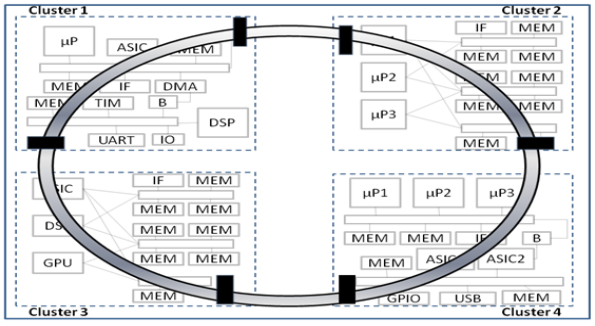

Fig.2. Proposed Optical Ring Bus (ORB) on-chip communication architecture for MPSoCs

Recently, it has been shown that it may be beneficial to replace global on-chip electrical interconnects with optical interconnects [6]. Optical interconnects have many advantages over electrical wires, such as higher signal propagation speed, lower power consumption, high bandwidth and lower electromagnetic interference [7]. While optical interconnects at the chip-to-chip

978-1-4244-1922-7/08/$25.00 ©2008 IEEE

level are already being actively developed [8], on-chip optical interconnects have only lately begun to receive attention. This is due to the recent development of CMOS compatible silicon-based optical components such as light sources [9], waveguides [10], modulators [11-12] and detectors [13-14], which have opened up exciting possibilities for on-chip optical interconnect realization. There is now enough data regarding CMOS compatible on-chip optical components to explore possible architectures and potential trade-offs for design decisions. The latest ITRS even suggests the possibility of on-chip optical interconnects completely replacing global interconnects by 2013 [2]. However, for this to happen, it is crucial to develop on-chip optical communication architectures and study their impact on chip-level power and performance.

In this paper, we describe a novel opto-electric on-chip communication architecture that uses an optical ring bus as a global interconnect between computation clusters in MPSoC designs. Fig. 2 shows how an optical ring bus can replace the global, pipelined electrical interconnects in the MPSoC depicted in Fig. 1. Our proposed optical ring bus (ORB) communication architecture makes use of a laser light source, opto-electric converters, an optical waveguide and wave division multiplexing (WDM) to transfer data between clusters on a chip, while preserving the standard bus protocol interface (e.g. AMBA AXI [5]) for inter- and intra-cluster communication. Our experimental results indicate that compared to a traditional pipelined, all-electrical global interconnect architecture, the ORB architecture dissipates significantly lower power (more than a 10× reduction) and also improves overall performance (more than 2×) for MPSoC designs. In the very likely scenario that bus-based on-chip optical interconnects become a reality in the future, this work takes the first step in developing an optical interconnect based on-chip communication architecture that is compatible with today's standards, and quantifying its benefits over traditionally used all-electrical, pipelined and bus-based communication architectures.

II. Related Work

Early work by Goodman et al. [15] first introduced the concept of on-chip optical interconnects. Comparisons between (non-pipelined) on-chip electrical and optical interconnects have been presented in [6] [16-19]. These studies have shown the promise of on-chip optical interconnects, but have primarily focused on clock networks and non-pipelined point-to-point links. One of the contributions of this paper is to compare on-chip optical interconnects with all-electrical pipelined global bus-based communication architectures that are used by designers to support high bandwidth on-chip data transfers today.

Network-on-chip (NoC) architectures [20-21] have received much attention of late as an alternative to bus-based architectures for future MPSoCs. Similar to pipelined interconnects (shared or point-to-point), NoCs split larger interconnects into smaller segments (links) separated by routers to enable multi-GHz frequencies and high bandwidths. However, electrical NoCs suffer from the same drawbacks as pipelined copper interconnects: high latencies and much higher power dissipation [22] due to buffer overhead in the routers and network interfaces. Some recent work has proposed hybrids of optical interconnects and NoC fabrics [23-24]. A silicon-on-insulator (SOI) 2D mesh-type optical-NoC hybrid architecture was proposed in [23]. However, the high power overhead of electrical routers and opto-electric/electro-optic conversion at the interface of each component, as well as a lack of availability of wideband photonic switching elements makes realizing such an architecture a difficult proposition in the near future. An SOI optical loop-based NoC was proposed in [24], which also suffers from some of the same drawbacks described above. In contrast to these SOI-based optical NoC architectures, we propose a novel polymer-based optical ring bus (ORB) based opto-electric communication architecture that does not require the complexity of network interfaces and packet routers. Our architecture has the significant advantage of seamlessly interfacing with existing bus-based protocols and standards, while providing significant improvements in on-chip power consumption and performance.

III. Optical Ring Bus Architecture: Building Blocks

Optical interconnects offer many advantages over traditional electrical (copper-based) interconnects: *(i)* they can support enormous intrinsic data bandwidths in the order of several Gbps using only simple on-off modulation schemes, *(ii)* they are relatively immune to electrical interference due to crosstalk and parasitic capacitances and inductances, *(iii)* their power dissipation is completely independent of transmission distance at the chip level, and *(iv)* routing and placement is simplified since it is possible to physically intersect light beams with minimal crosstalk. Once a path is acquired, the transmission latency of the optical data is very short, depending only on the group velocity of light in a silicon waveguide: approximately 6.6×10^7 m/s, or 300ps for a 2−cm path crossing a chip [25]. After an optical path is established, data can be transmitted end to end without the need for repeating or buffering, which can lead to significant power savings.

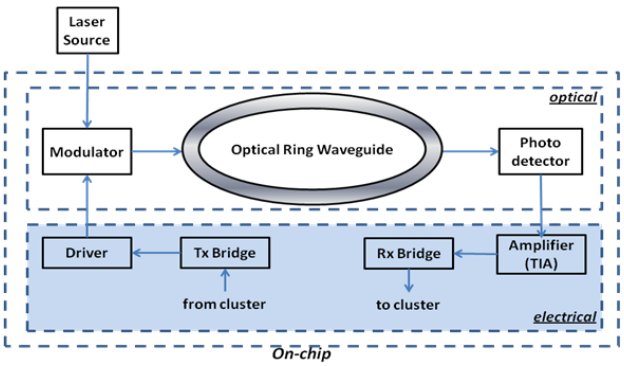

Fig.3. ORB optical interconnect components

Realizing on-chip optical interconnects as part of our proposed ORB communication architecture requires several CMOS compatible optical devices. Although there are various candidate devices that exist for these optical elements, we select specific devices that satisfy on-chip requirements. Fig. 3 shows a high level overview of the various components that make up our ORB optical interconnect architecture. There are four primary optical components: a laser (light source), an opto-electric modulator/ transmitter, an optical ring waveguide and an optical receiver. Integrating such an optical system on a chip requires CMOS compatibility which puts constraints on the types of materials and choices of components to be used. Recent technological advances indicate that it is possible to effectively fabricate various types of optical components on a chip. However, there are still significant challenges in efficiently integrating a silicon based laser on a chip. Using an off-chip laser can actually be beneficial because it leads to lower on-chip area and power consumption. Consequently, in

our optical interconnect system we use an off-chip laser from which light is coupled onto the chip using optical fibers, much like what is done in chip-to-chip optical interconnects today [8] [26].

The transmission part in Fig. 3 consists of a modulator and a driver circuit. The electro-optic modulator converts an input electrical signal into a modulated optical wave signal for transmission through the optical waveguide. The modulators are responsible for altering the refractive index or absorption coefficient of the optical path when an electrical signal arrives at the input. Two types of electrical structures have been proposed for opto-electric modulation: p-i-n diodes [27] and MOS capacitors [11]. Micro-ring resonator based p-i-n diode type modulators [12] [27] are compact in size (10-30 μm) and have low power consumption, but possess low modulation speeds (several MHz). On the other hand, MOS capacitor structures such as the Mach-Zehnder interferometer based silicon modulators [11] [26] have higher modulation speed footprint (several GHz) but a large power consumption and greater silicon footprint (around 10 mm). While these electro-optical modulators today are by themselves not very attractive for on-chip implementation, there is a lot of ongoing research which is attempting to combine the advantages of both these modulator types [12]. Consequently, we make use of a predictive modulator model which combines the advantages of both structures. We assume a modulator capacitance that scales linearly with modulator length at the rate of 1.7pF/mm [17].

The modulator is driven by a series of tapered inverters (i.e. driver). The first stage consists of a minimum sized inverter. The total number of stages N is given as

$$N = log \frac{C_m}{Cg} / log\ 3.6$$

where C_m is the modulator capacitance and C_g is the capacitance of a minimum sized inverter. These drivers receive their input signal from a transmission bridge (*Tx Bridge*) belonging to a cluster. The *Tx Bridge* component is similar to a bridge in a traditional hierarchical shared bus architecture, and logically treats the optical ring bus waveguide as any another shared bus. Any communication request meant for a core in another cluster is sent to the optical ring bus through the transmission bridge.

The optical waveguide is responsible for transporting data via light signals from the source modulator to the destination receiver. The choice of the optical material and wavelength of utilized light are the two main factors affecting waveguide performance. For on-chip optical interconnects, there are two popular candidates for waveguide material: high refractive index silicon on insulator (SOI) and low refractive index polymer waveguides. SOI waveguides have lower pitch (i.e. width) and lower area footprint compared to polymer waveguides. This leads to better bandwidth density (i.e. transmitted bits per unit area). However polymer waveguides have lower propagation delay than SOI waveguides. The area overhead for polymer waveguides is mitigated if they are fabricated on a separate, dedicated layer. Additionally, if wave division multiplexing (WDM) is used, polymer waveguides provide superior performance and bandwidth density compared to SOI waveguides [28]. Consequently, in our optical ring bus, we make use of a low refractive index polymer waveguide with an effective index of 1.4. We chose a ring shaped optical waveguide to avoid sharp turns in the waveguide which can lead to significant signal loss. The optical ring is implemented on a dedicated layer and covers a large portion of the chip so that it can effectively replace global electrical pipelined interconnects.

The receiver part in Fig. 3 consists of a photo-detector to convert the light signal into an electrical signal, and a circuit to amplify the resulting analog electrical signal to a digital voltage level. In order to support WDM, where transmission occurs on multiple wavelengths, the receiver includes a wave-selective filter for each wavelength that is received. An important consideration in the selection of a photo detector is the trade-off between detection speed and sensitivity (quantum efficiency) of the detector. Interdigitated metal-semiconductor-metal (MSM) Ge and SiGe photo-detectors have been proposed [13] [29] that have fast response, excellent quantum efficiency and low power consumption. These attributes makes the MSM detector a suitable candidate as a photo-detector in our optical interconnect architecture. We assume a detector capacitance of 100 fF based on a realistic detector realization [18].

A Trans-impedance amplifier (TIA) is used to amplify the current from the photo-detector [17]. The TIA consists of an inverter and a feedback resistor, implemented as a PMOS transistor. Additional minimum sized inverters are used to amplify the signal to a digital level. The size of the inverter and feedback transistor in the TIA is determined by bandwidth and noise constraints. The amplified digital signal is subsequently sent to the receiving bridge (*Rx Bridge*) component, which decodes the destination address, and passes the received data to a specific core in the cluster.

TABLE I
Delay (in ps) of optical components for 1 cm optical data path

Tech Node	65 nm	45 nm	32 nm	22 nm
Modulator Driver	45.8	25.8	16.3	9.5
Modulator	52.1	30.4	20.0	14.3
Polymer Waveguide	46.7	46.7	46.7	46.7
Photo Detector	0.5	0.3	0.3	0.2
Receiver Amplifier	16.9	10.4	6.9	4.0
Total Optical Delay	**162.0**	**113.6**	**90.2**	**74.7**

Table I shows delays of the various optical interconnect components used in ORB and described above, for a 1 cm optical data path, calculated based on estimates from [2] [17]. It can be seen that the optical interconnect delay remains constant for the waveguide, while the delay due to other components reduces with each technology generation. This is in contrast to the minimum electrical interconnect delay which is expected to remain almost constant (or even increase slightly) despite optimal wire sizing (i.e. increasing wire width) and repeater insertion to reduce delay.

IV. ORB Communication Architecture

The previous section gave an overview of the various components that are part of our on-chip optical interconnect solution. In this section we elaborate on the operation of our optical ring bus (ORB) architecture.

In the ORB communication architecture, the various cores within each cluster are locally interconnected using high speed and low complexity, low power and low area footprint electrical bus-based communication architectures (such as hierarchical buses or crossbar buses). When a core in a cluster must communicate with a core in another cluster, the request is routed to an ORB interface, which interfaces with the optical ring waveguide. The interface consists of a transmitting bridge, which is similar to a standard bridge used in hierarchical bus-based architectures. The bridge sends the request to the modulator drivers, which pass it on to the modulator. The electrical signal is

converted into a light signal, transmitted through the optical ring waveguide, and finally reaches the receiver interface. The receiver interface consists of a photo-detector and TIAs which convert the light signal back into a digital voltage signal, and a receiving bridge, which accepts the request and is responsible for sending it to the appropriate core in its corresponding cluster. A cluster can have more than one transmitting and receiving interfaces, depending on its communication needs.

For a global interconnect with an address bus width a, a data bus width d and c control bits, there are $a+d+c$ optical ring waveguides. These optical waveguides must be spaced 0.5–3 μm to avoid significant crosstalk. It was shown in [6] that a single wavelength optical link is inferior to a delay optimized electrical interconnect in terms of bandwidth density. To improve bandwidth density of the optical interconnect, we make use of wavelength division multiplexing (WDM) [17] [30]. This involves using multiple wavelength channels to transmit data on the same waveguide. WDM can significantly improve optical interconnect bandwidth density over electrical interconnects. We assume that each of the waveguides has λ available wavelengths for WDM. This creates a λ-way bus, and necessitates a mechanism for determining how the λ wavelengths are distributed among various data streams. The value λ has significant implications for performance, cost and power since using a larger number of wavelengths improves bandwidth but requires more processing, area and power overhead at the transmitters and receivers. Based on predictions in [6] which indicate that the number of wavelengths will increase with every technology node, reaching around 10 for the 22 nm node, we limit the maximum number of wavelengths λ in our ring waveguides, and consequently the number of allowed transmitter/receiver interface pairs to 10.

There are two ways to allocate the wavelengths (i.e. multiplex the optical bus): by address space and by cluster. In the address space based scheme, wavelengths are allotted to different address spaces, whereas in the cluster based scheme each cluster has exclusive use of a subset of the λ wavelengths. It was shown in [24] that even though the cluster based allocation scheme allows only λ cluster interfaces to the optical bus, it is more beneficial in terms of power consumption compared to the address space allocation approach. Consequently, we use a cluster based wavelength allocation approach in ORB.

If simplicity in design is a key concern, each of the N clusters in an MPSoC application can be allocated an equal number of wavelengths λ/N. However, this does not take into account the specific performance requirements of the application. It is very possible that certain clusters have greater communication bandwidth needs than others. Consequently, the fraction of total wavelengths λ_i allocated to a cluster i is calculated as

$$\lambda_i = \lambda.\left(\frac{BW_i}{\sum_{j=1}^{N} BW_j}\right)$$

where BW_i is the bandwidth requirements of a cluster i, and the number of allocated wavelengths λ_i being rounded to the nearest integer. The total number of transmitters for cluster i on the optical ring bus is

$$T_{i\_total} = \lambda_i.(a_i+d_i+c_i)$$

and the total number of receivers for cluster i is

$$R_{i\_total} = (\lambda - \lambda_i).(a_i+d_i+c_i)$$

One final important design consideration is to ensure that light does not circulate around the optical ring for more than one cycle, because that could lead to undesirable interference from older data. This is resolved by using attenuators with each modulator, to act as a sink for the transmitted wavelength(s), once the signal has completely traversed the optical ring.

V. Experiments

In this section we present comparison studies between ORB and traditional all-electrical on-chip communication architectures. The ORB communication architecture uses an optical ring bus as a global interconnect, whereas the traditional all-electrical communication architecture uses electrical pipelined global interconnects. Both configurations use electrical buses as local interconnects within clusters.

A. Applications

We select several MPSoC applications for the comparison between our ORB architecture and the traditional pipelined electrical architecture. These applications are selected from the well known SPLASH-2 benchmark suite [31] (*Barnes*, *Ocean*, *FFT*, *Radix*). We also select applications from the networking domains (proprietary benchmarks *Netfilter* and *Datahub* [32]). These applications are parallelized and implemented on multiple cores. Table II shows the characteristics of the implementations of these applications, such as the number of cores (e.g. memories, peripherals, processors), programmable processors and clusters on the MPSoC chip. The die size is assumed to be 2cm×2cm.

TABLE II
MPSoC applications and their characteristics

MPSoC Application	Description	# of cores	# of proc.	# of clusters
Radix	Integer radix sort	18	4	3
Barnes	Evolution of galaxies	26	6	4
FFT	FFT kernel	28	6	4
Ocean	Ocean movements	35	10	5
Netfilter	Packet processing and forwarding	49	22	6
Datahub		68	26	8

B. Experimental Setup

The applications described above are modeled in SystemC [33] and simulated at the transaction level bus cycle accurate abstraction [34] to quickly and accurately estimate performance and power consumption of the applications. The various cores are interconnected using the AMBA AXI [5] standard bus protocol. A high level simulated annealing floorplanner based on sequence pair representation PARQUET [35] is used to create an early layout of the MPSoC on the die, and Manhattan distance based wire routing estimates are used to determine wire lengths.

For the global electrical interconnect, wire delay and optimal delay repeater insertion points are calculated using an RLC transmission line wire model described in [36]. Latches are inserted based on wire length (obtained from routing estimates), wire delay and clock frequency of the bus, to pipeline the bus and ensure correct operation [4]. For instance, a corner to corner wire of length 4 cm for a 2cm×2cm die size has a projected delay of 1.6 ns in 65 nm technology, for a minimum width wire size W_{min} [36]. To support a frequency of 2.5 GHz (corresponding to a clock period of 0.4 ns), 4 latches need to be inserted to ensure correct (multi-cycle) operation. It has been shown that increasing wire

width can reduce propagation delay at the cost of area. For our global interconnects, we therefore consider wider interconnects with a width $3W_{min}$ which results in a near optimal power delay product at the cost of only a slight area overhead. The delay of such a wide, repeater-inserted wire is found to be approximately 26 ps/mm, varying only slightly (±1 ps/mm) between the 65-22 nm nodes.

For the optical ring bus (ORB) architecture, we model all the components described in Section III, annotated with appropriate delays. The global optical bus length is calculated using simple geometric calculations and found to be approximately 43 mm. Based on this estimate, as well as optical component delay values from Table I, we determine the maximum operating frequencies for ORB as 1.4 GHz (65 nm), 2 GHz (45 nm), 2.6 GHz (32 nm) and 3.1 GHz (22 nm). To ensure a fair comparison, we clock the traditional all-electrical global pipelined interconnect architecture at the same frequencies as the optical ring bus architecture in our experiments. The cores in the clusters are assumed to operate at twice the interconnect frequencies. We set the width of the address bus as 32 bits and that of the separate read and write data buses as 64 bits. The bus also uses 68 control bits, based on the AXI protocol. These translate into a total of 228 (address + read data + write data + control) optical waveguides. Finally, WDM is used, with a maximum of $\lambda = 10$ wavelengths allocated based on cluster bandwidth requirements, as explained in Section IV.

C. Performance Comparison

Optical waveguides provide faster signal propagation compared to electrical interconnects because they do not suffer from RLC impedances. But in order to exploit the propagation speed advantage of optical interconnects, electrical signals must be converted into light and then back into an electrical signal. This process requires an overhead that must be taken into account while comparing optical interconnects with electrical interconnects.

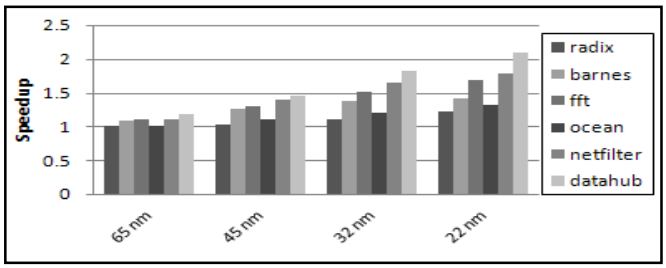

Fig.4. Performance speedup for ORB over traditional all-electrical bus-based communication architectures

To compare the performance of the ORB and traditional pipelined global interconnect based communication architectures, we simulate the MPSoC applications for the two configurations across 65, 45, 32 and 22 nm technology nodes. Fig. 4 shows the results of this comparison. It can be seen that the ORB architecture provides a performance speedup over traditional all-electrical bus-based communication architectures for UDSM technology nodes. The speedup is small for 65 nm because of the relatively lower global clock frequency (1.4 GHz) which does not require as much pipelining. However, from the 45 nm down to the 22 nm nodes, the speedup increases steadily because of rising clock frequencies which introduce more pipeline stages in the electrical global interconnect, increasing its latency and reducing bandwidth, compared to the ORB architecture. The speedup for *radix* is lower than other applications due to the smaller length of global interconnect wires, which reduces the advantage of having an optical link for global data transfers. On the other hand, the lower speedup for *ocean* is due to the smaller number of global inter-cluster data transfers, despite having long global interconnects. Overall, the ORB architecture speeds up global data transfers due to the faster optical waveguide. Despite the costs associated with converting the electrical signal into an optical signal and back, ORB can provide more than a 2× performance speedup. With improvements in optical component fabrication over the next few years, the opto-electrical conversion delay is expected to decrease leading to even greater performance benefits.

TABLE III
Power Consumption of Optical Data Path (in mW)

Tech Node	65 nm	45 nm	32 nm	22 nm
Transmitter	18.4	8.6	6.0	5.0
Receiver	0.3	0.2	0.3	0.3
Total Optical Power	**18.7**	**8.8**	**6.3**	**5.3**

D. Power Comparison

To estimate the power consumption for the electrical interconnects, we must account for power consumed in the wires, repeaters and bus logic components (latches, bridges, arbiters and decoders). For bus wire power estimation, we determine wire lengths using our high level floorplan and routing estimates as described earlier. We then make use of bus wire power consumption estimates from [38], and extend them to account for repeaters [39]. Static repeater power and capacitive load is obtained from data sheets. Capacitive loads for components connected to the bus are obtained after logic synthesis. Other capacitances (e.g. ground, coupling) are obtained from the Berkeley Predictive Technology Model (PTM) [40], and ITRS estimates [2]. The power consumed in the bus logic components is calculated by first creating power models for the components, based on our previous work on high-level power estimation of communication architectures [39]. These power models are then plugged into the SystemC simulation models. Power numbers are obtained for the components after simulation and are within 5% accuracy of gate-level estimates [39]. Simulation is also used to obtain accurate values for switching activity, which is used for bus wire power estimation.

For the optical interconnect, power consumption estimates for a transmitter and receiver in a single optical data path are derived from [17] and shown in Table III. It can be seen that the power consumed by the transmitter dominates power consumed by the receiver. The size as well as the capacitance of the modulator is large, requiring a large driving circuit. The power consumed by the optical waveguide is almost independent of interconnect length, since the length is relatively short and consequently the optical power loss in the waveguide is negligible.

Having described how power is estimated for the electrical and optical interconnects, we now show the results of comparing power consumption for the all-electrical pipelined interconnect architecture with our ORB architecture. Fig. 5 shows the results of this comparison. It can be seen that the ORB architecture consumes more power for the 65 nm node, compared to the all-electrical pipelined interconnect architecture. However, for technology nodes from 45 nm onwards, there is a significant reduction in ORB power consumption, due to expected improvements in opto-electrical modulator structure fabrication as well as an increase in electrical power consumption due to higher

operating frequencies and greater leakage. ORB can provide more than a 10× power reduction compared to all-electrical pipelined global interconnect architectures, which is a strong motivation for adopting it in the near future.

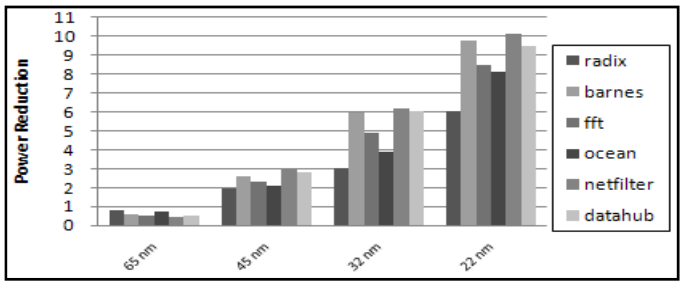

Fig.5. Power reduction for ORB over traditional all-electrical bus-based communication architectures

VI. Conclusion and Future Work

In this paper, we presented preliminary work on a novel on-chip opto-electrical bus-based communication architecture. Our optical ring bus (ORB) communication architecture replaces global pipelined electrical (copper) interconnects with an optical ring waveguide and opto-electric modulators and receivers. While there is a definite performance and power overhead associated with converting electrical signals into optical signals and back today, we showed that ORB can be beneficial for ultra-deep submicron (UDSM) technology nodes below 65 nm. Our experimental results, based on emerging technology trends and recently published studies, have shown that the ORB architecture can provide as much as 2× performance gain, along with a 10× power reduction, compared to the traditional all-electrical interconnect architecture. It is clear that ORB can provide a performance-per-watt that is far superior to electrical alternatives. Furthermore, ORB is scalable to accommodate an increasing number of computational clusters and cores on a chip in the future, and provides a clean separation of concerns as the optical waveguide and components are fabricated in a separate, dedicated layer. Our ongoing work is looking at characterizing bandwidth density and analyzing the implications of emerging optical components for the ORB architecture. Future challenges in this area include the need for active or passive control methods to reduce optical interconnect susceptibility to temperature variations, and better opto-electric modulator designs to reduce delay and power consumption.

References

[1] R. Ho, W. Mai, and M. A. Horowitz, “The future of wires”, Proceedings of the IEEE, 89(4):490–504, April 2001.
[2] International Technology Roadmap for Semiconductors, 2006, http://www.itrs.net/
[3] V. Adler, E. Friedman, “Repeater Design to Reduce Delay and Power in Resistive Interconnect”, IEEE TCAS, May 1998.
[4] V. Nookala, S. S. Sapatnekar, "Designing optimized pipelined global interconnects: Algorithms and methodology impact," IEEE ISCAS, Vol.1, pp.608-611, May 2005.
[5] AMBA AXI Specification www.arm.com/armtech/AXI
[6] M. Haurylau et al., "On-chip Optical Interconnect Roadmap: Challenges and Critical Directions," IEEE Journal of Selected Topics in Quantum Electronics, Vol. 12, No. 6, Nov/Dec 2006.
[7] D. A. Miller, “Rationale and challenges for optical interconnects to electronic chips” Proc. IEEE, Jun 2000.
[8] I. Young. “Intel introduces chip-to-chip optical I/O interconnect prototype”, Technology@Intel Magazine, Apr 2004.
[9] H. Rong, et al. “A continuous-wave Raman silicon laser,” Nature, vol. 433, pp. 725–728, Feb. 2005.
[10] S. J. McNab, N. Moll, Yu. A. Vlasov, “Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides,” Opt. Express, vol. 11, no. 22, Nov. 2003.
[11] A. Liu et al., “A High-Speed Silicon Optical Modulator Based on a Metal-Oxide-Semiconductor Capacitor," Nature, Vol. 427, pp. 615-618, Feb 2004.
[12] Q. Xu et al., “12.5 Gbit/s carrier-injection-based silicon microring silicon modulators”, Optics Express, 15(2), Jan. 2007.
[13] M. R. Reshotko, D. L. Kencke, B. Block, “High-speed CMOS compatible photodetectors for optical interconnects,” Proc. SPIE, Oct. 2004, vol. 5564, pp. 146–155.
[14] S. J. Koester, et al. “High-fffficiency, Ge-on-SOI lateral PIN photodiodes with 29 GHz bandwidth,” Proc. Device Research Conf, Notre Dame, IN, 2004, pp. 175–176.
[15] J. W. Goodman et al., “Optical Interconnects for VLSI Systems," Proceedings of the IEEE, Vol. 72, No. 7, July 1984.
[16] M. J. Kobrinsky et al., “On-Chip Optical Interconnects," Intel Technology Journal, Vol. 8, No. 2, pp. 129-141, May 2004.
[17] G. Chen, H. Chen, M. Haurylau, N. Nelson, D. Albonesi, P. M. Fauchet, E. G. Friedman, "Predictions of CMOS Compatible On-Chip Optical Interconnect," Proc. of SLIP, 2005, pp. 13-20.
[18] Ian O’Connor, “Optical solutions for system-level interconnect”, Proc. SLIP, Feb 2004.
[19] A. M. Pappu A. B. Apsel, “Analysis of intrachip electrical and optical fanout”, Applied Optics, 44(30):6361–6372, Oct 2005
[20] L. Benini and G. D. Micheli, “Networks on chip: A new SoC paradigm”, IEEE Computer, 49(2/3):70–71, Jan. 2002.
[21] W. J. Dally and B. Towles, “Route packets, not wires: On-chip interconnection networks”, In Design Automation Conf., pages 684–689, June 2001.
[22] S. Vangal et al, “An 80-tile 1.28 TFLOPS network-on-chip in 65 nm CMOS”, In Proc. ISSCC Feb. 2007.
[23] A. Shacham, K. Bergman, L. Carloni, “The Case for Low-Power Photonic Networks on Chip”, In Proc. DAC 2007.
[24] N. Kirman et. al, “Leveraging Optical Technology in Future Bus-based Chip Multiprocessors”, In Proc. MICRO, 2006.
[25] I.-W. Hsieh et al, “Ultrafast-pulse self-phase modulation and third-order dispersion in si photonic wire-waveguides”, Optics Express, 14(25):12380–12387, Dec. 2006.
[26] C. Gunn, “CMOS photonics for high-speed interconnects”, IEEE Micro, 26(2):58–66, Mar./Apr. 2006.
[27] C. A. Barrios et al., “Low-Power-Consumption Short-Length and High-Modulation-Depth Silicon Electro-optic Modulator" Journal of Lightwave Technology, Vol. 21, No. 4, April 2003.
[28] L. Eldada, L. W. Shacklette, “Advances in Polymer Integrated Optics," IEEE JQE, Vol. 6, No. 1, Jan/Feb 2000.
[29] A. Gupta et al., “High-speed optoelectronics receivers in SiGe”, In Proc. VLSI Design, Jan. 2004, pp. 957–960.
[30] B. G. Lee et al., “Demonstrated 4×4 Gbps silicon photonic integrated parallel electronic to WDM interface”, OFC 2007.
[31] S. C. Woo et al. “The SPLASH-2 programs: Characterization and methodological considerations”, Proc. ISCAS, 1995.
[32] S. Pasricha, N. Dutt, “The Optical Ring Bus (ORB) On-Chip Communication Architecture”, CECS Technical Report, Feb 2008
[33] SystemC initiative. www.systemc.org
[34] W. Müller, J. Ruf,, W. Rosenstiel, “SystemC Methodologies and Applications”, Norwell, MA: Kluwer, 2003
[35] S. N. Adya, I. L. Markov, "Fixed-outline Floorplanning: Enabling Hierarchical Design", In IEEE Trans TVLSI, Dec. 2003
[36] Y. I. Ismail and E. G. Friedman, “Effects of Inductance on the Propagation Delay and Repeater Insertion in VLSI Circuits," IEEE TVLSI, Vol. 8, No. 2, pp. 195-206, April 2000.
[38] C. Kretzschmar, et al., “Why transition coding for power minimization of on-chip buses does not work”, DATE 2004
[39] S. Pasricha, Y. Park, F. Kurdahi, N. Dutt, “System-Level Power-Performance Trade-Offs in Bus Matrix Communication Architecture Synthesis”, CODES+ISSS 2006.
[40] Berkeley Predictive Technology Model, U.C. Berkeley, http://www-devices.eecs.berkeley.edu/~ptm/