

Survey on Unified Inter/Intrachip Optical Network for Chip Multiprocessors

Priyanka Rajendran¹ and Dr.Gnana Sheela K²

^{1,2}*Electronics and Communication Engineering, Toc H Institute of Science and Technology, Kerala, India*

Abstract: *As modern computing systems become increasingly complex, communication efficiency among and inside chips has become as important as the computation speeds of individual processing cores. Traditionally, to maximize design flexibility, interchip and intrachip communication architectures are separately designed under different constraints. Jointly designing communication architectures for both interchip and intrachip communication could, however, potentially yield better solutions. This paper presents unified inter/intrachip optical network called UNION, for chip multiprocessors (CMPs). UNION is based on recent progresses in nanophotonic technologies. It connects not only cores on a single CMP, but also multiple CMPs in a system. UNION employs a hierarchical optical network to separate interchip communication traffic from intrachip communication traffic. It fully utilizes a single optical network to transmit both payload and control packets. The network controllers on each CMP not only manage intrachip communications, but also collaborate with each other to facilitate interchip communications. It compared UNION with a matched electrical counterpart in 45-nm process. The results for eight real CMP applications show that on average UNION improves CMP performance by 3× while reducing 88% of network energy consumption.*

Keywords: *Chip multiprocessors, Nano photonics, Inter/intra chip, Network-on-chip, Nano scale technologies*

I. Introduction

Modern computing systems have become increasingly complex to satisfy the growing performance demanded by applications. As the number of transistors available on a single chip increases to billions, chip multiprocessor (CMP) has become an attractive platform delivering high performance with limited power budget. For interchip communications, bus-based and ad-hoc architectures are still popular, and signals are mostly transmitted by electrical interconnects on printed circuit boards (PCB). The limitations of electrical interconnects such as high-delay and high-power consumption, are already shown in high performance systems. For intrachip communication, architectures gradually move from ad-hoc and bus-based architectures to network-on-chip (NoC) to alleviate issues including poor scalability and limited bandwidth [1], [2]. Optical interconnects, with advantages including ultra high throughput, low-delay and low-power consumption, are proposed to replace both inter and intrachip electrical wires. Traditionally, interchip/ intrachip communication architectures are separately designed.

There is a huge performance gap between intra and interchip electrical interconnects. First, the delay of on-chip wires is much smaller than the off-chip ones because of the substantially smaller physical length, resistance, and capacitance. For optical interconnects, unlike electrical wires, inter and intrachip channels can be interconnected seamlessly. Both on-chip and off-chip channels can be implemented with optical waveguides and they can be interconnected with passive couplers. The allowed operating bandwidths of both waveguides are broad enough for real applications. Given the high propagation speed of light, the delay difference for on-chip and off-chip links is minor. In addition, if the small transmission loss is neglected, the transmission power is independent of transmission length. All these make the optical interchip and intrachip interconnects well matched, and a unified design becomes natural. There are different methods available for the inter/ intra chip communication. Apart from the electrical connection optical interconnections are the better method. The following sections deal with literature background of different algorithms and their comparative analysis.

II. Literature Review

Kapur et al (2002) proposed that the optical interconnects from both delay and power perspective. It involved the properly modelling power and delay of cache component in an optical signalling system as well as repeated wires. It shows that favourable to switch to optical interconnects for global signalling, on both power and delay account, for long wires, whereas, shorter wires favour an electrical operation on lower power grounds [2]. The scaling paradigm adopted by semiconductor industry, while conducive to devices, severely degrades interconnect performance. On-chip optical interconnects present a possible solution. While the usefulness of optics in chip to chip communication is becoming clearer, there exists a considerable ambiguity in its utility in high performance on-chip applications. It addresses optical system specifically for global signalling application.

The comprehensively model both optical and electrical links and subsequently compare their performance in the light of the two most important figures of merit: delay and power consumption. The delay of driving the modulator can be minimized by using an exponentially sized buffer chain. The transmitter delay, thus calculated, is dependent on the technology generation and is found to be about 75ps at 50nm node. The waveguide delay is, next, modelled as a function of dimensions using a full 2-dimensional treatment called the crlective index method [13]. This length depends on the receiver conditions and drops with increasing RPD and or IOP. Thus, the optical interconnects are much faster, however, the power expenditure to get this speed advantage is an important consideration and is discussed next in the form of power delay trade-offs for the two systems. The electrical repeated wire exhibits a minimum delay at ascertain spacing and sizing of repeaters, which also expends most power.

The aforementioned optimization, the delay can be efficiently traded for power savings, however, the power asymptotically saturates to a minimum value dictated by the wire capacitance. The global signalling: delay and power for, both best available repeated electrical wires and optical interconnects. For optical interconnects, this included modelling the receiver. Subsequently, it compared the two extensively at 50nm node as well as a function of scaling. They find that the optical interconnects for long links have advantages on both merits over even repeated electrical wires, while, the repeated electrical wires have a power advantage and delay disadvantage at shorter link lengths.

Pande et al (2003) proposed that most of past evaluations of fat-trees for on-chip interconnection networks rely on oversimplifying or even unrealistic architecture and traffic pattern assumptions, and very few layout analyses are available to relieve practical feasibility concerns in nanoscale technologies [2]. This work aims at providing an in-depth assessment of physical synthesis efficiency of fat-trees and at extrapolating silicon aware performance figures to back-annotate in the system-level performance analysis. A 2D mesh is used as a reference architecture for comparison, and a 65 nm technology is targeted by our study. Finally, in an attempt to mitigate the implementation cost of k -array n -tree topologies, we also review an alternative unidirectional multi-stage interconnection network which is able to simplify the fat-tree architecture and to minimally impact performance. Networks-on-chip (NoCs) closely resemble the interconnect architecture of high performance parallel computing systems

The interconnection topologies used in the early NoC prototypes can be traced back to the field of parallel computing. In particular, NoC architectures aiming at low latency communication, performance scalability and flexible routing selected selected fat-trees as their reference topology. The switch for the butterfly fat-tree network or SPIN micro network is examples thereof. This paper focuses on on a specific implementation of Fat-trees(FT): the k -ary n -trees a parametric family of regular multistage topologies. k -ary n -trees are implemented by using stages of identical switches, where k is the number of links of a switch that connect to the previous or to the next stage (i.e., the switch radix is $2k$). All the switches have the same number of ascending and descending links. For fat-tree routing, uses the deterministic routing algorithm. .

During the ascending phase, consecutive destinations are shuffled among the different ascending links of the switches; each ascending port is labelled in italic switch the destination cores that are reachable through it. Also, shows in bold how destinations are distributed in the descending phase. As can be seen, each descending link is only used by a single destination. To analyze 2-ary 4-tree for a 16-core system, and a 2-ary 6-tree for 64 cores. It proves that fat-trees are feasible for on-chip networks from a physical design viewpoint. Unfortunately, for small scale systems, they are not able to capitalize on their better performance scalability yet. As the system size scales up, 2D meshes however suffer from poor performance scalability. Hence, the need for alternative topologies becomes more stringent. k -ary n -trees can provide that performance scalability, but at an impractical power and area cost. .

Nevin et al (2006) developed that optical interconnections are better than thee electronic connections. Current research and technology trends indicate that future chip multiprocessors (CMPs) may comprise tens or even hundreds of processing elements. An important hurdle towards attaining this scale, however, is the need to feed data to such large numbers of on-chip cores [3]. This can only be achieved if architecture and technology developments provide sufficient chip-to-chip and on-chip communication performance to these future generations of CMPs. Optical technology and 3D integration are two potential solutions to current and projected limitations in chip-to-chip communication performance. Still, on chip communication faces considerable technological and architectural challenges of its own. On the one hand, global on chip interconnects do not scale well with technology. Although delay-optimized repeater insertion and proper wire sizing can keep the delay nearly constant, this comes at the expense of power and active area, as well as a reduction in wire count .

The techniques for optimizing the power-delay product have been developed, but unfortunately their most obvious shortcoming is that neither power nor latency is optimal. This, combined with various other technological issues such as manufacturability, conductivity, crosstalk, etc. constitute important roadblocks for future electrical interconnects. On the other hand, to date there is no clear consensus on the architecture of on-chip core interconnects. Although well-understood solutions exist for off-chip interconnects, the on-chip power, area, connectivity, and latency constraints make it challenging to port those solutions to the context of CMPs.

Optical transmission requires a laser source, a modulator, and a modulator driver (electrical) circuit. The laser source provides light to the modulator, which transducers electrical information (supplied by the modulator driver) into a modulated optical signal.

The modulator translates the modulator driver's electrical information into a modulated optical signal. High-speed electro-optic modulators are designed such that injection of an electrical signal changes the refractive index or the absorption coefficient of an optical path. Among different types of proposed modulators, the most recent optical resonator-based implementations are preferable for integrated circuit design, due to their low operating voltage and compact size. They assume this type of modulator in our work Waveguides are the paths through which light is routed. The refractive index of the waveguide material has a significant impact on optical interconnect bandwidth, latency, and area. An optical receiver performs the optical-to-electrical conversion of the light signal. It comprises a photo detector and a trans-impedance amplifier (TIA) stage.

Optical waveguides do not lend themselves gracefully to tree or highly angled structures that may be more common in electrical topologies, for turns and waveguide crossings may result in significant signal degradation. Before delivering into the details of a design space exploration, give a high-level description of the bus protocol. The specifics of the cache coherence protocol are not relevant here; the focus on the handling of coherence requests by the split transaction, fully pipelined hierarchical bus. It categorize the power consumption of the interconnect system into two: the power consumed in the electrical sublevels (switches and wiring), and the power consumed in the optical bus. Power calculations for each component assuming full switching activity, but report total power consumption at full, as well as 50% activity. By carefully modelling the speed, area, and power characteristics of electrical and recently developed optical components, and projecting to 32nm technology, determining a hierarchical bus consisting of both optical and electrical levels yields significant performance within reasonable power and area constraints. The approach exploits wave division multiplexing technology (WDM) to provide each node on the optical bus with unique wavelength, which are used to build a high-bandwidth multi-way bus. This speeds up several protocol operations, especially data transfer and arbitration.

Sung et al (2006) proposed that to demonstrate the system, three components were prepared: a fibre-embedded optical printed-circuit board (OPCB), optical transmitter/receiver modules, and 90°-bent fibre connectors. All components were assembled using precise guide pins and holes so that complete passive alignment was achieved in the OPCB. An optical link of 5-Gb/s/ch signals with a total link loss of 1.5 dB has been successfully demonstrated from the assembled system [3]. The processing chips in present computer systems operate above gigahertz clock speed. However, severe bottleneck of data transmission in the electrical interconnection between chips restricts high-speed processing of the systems [1]. Optical interconnection is a promising solution for this problem [2].

To achieve optical interconnection between chips in computer systems, appropriate optical components and packaging strategies for these components are required. Recent studies [3-10] for chip-to-chip optical interconnection have focused on the development of interconnection schemes and components based on optical printed-circuit boards (OPCBs) which have optical layers and electrical layers in one board. One of the key issues is to devise the optic elements to link the light between the light source (or detector) and the optical layer in the OPCB. The architecture of the optical interconnection consists of three main parts: an OPCB, optical transmitter/receiver (Tx/Rx) modules, and connectors to link light paths between the OPCB and the modules. OPCB is prepared using a fibre array as the optical layer.

A 90°-bent fibre array mounted in a small box is used as a connector. The key feature of this system is the packaging structure to intend passive assembling of the connectors and the modules on the board. The optical Tx/Rx modules were electrically contacted using ball grid array (BGA). The BGA were soldered by a hot air reflowing process after optical alignment of the connection blocks and modules with guide pins. It is important to note that the inaccuracy of the position of electrodes formed on the module and the board is much larger, in a range of a hundred micrometres, than that of the position of the guide pins and holes. Such a large misalignment of electrodes can be absorbed by distortion of the melt balls. This misalignment-absorbing function of melt balls contrasts with the self-alignment of the melt bumps [11]. The optical Tx/Rx modules were prepared to test the data link for four channels with 5 Gb/s/ch. The Tx module has a 1 × 4 array of 850-nm VCSELs and a SiGe BiCMOS driver IC chip.

The Rx module has a 14 PIN-PD array and a SiGe receiver IC chip. Other elements packaged in the modules are an FR4-PCB substrate with differential micro strip line, heat sink, solder balls, and guide pin/hole part. The guide pin/hole part was prepared by cutting the MT ferrule and attaching it on the substrate. The data link performances were tested from the optical interconnection platform where the connection blocks and Tx/Rx modules are assembled in the OPCB. The measured first the optical loss in the light path through the connection blocks and OPCB. Note that after the Tx/Rx modules are once electrically contacted on the OPCB, the optical loss could not be measured. Thus, before attaching the Tx/Rx modules, a pig-tailed Tx module with fiber ribbon and MT ferrule were assembled on the input side of the 90° connection block installed in the OPCB. On the output side of 90° connection block, optical output power was measured using a power meter. Here, index

matching oil was filled in the gaps between the connection block and the cross section of the OPCB to reduce scattering loss.

The total optical loss measured from the Tx module to the power meter through the connection blocks and the OPCB was about 1.5 dB. This loss is 0.2 dB higher than the best data in our previous work [6], obtained from an active alignment of each component under a similar interconnection scheme. Considering that in this work all components are passively assembled, the loss of 1.5 Db near the best alignment data is a remarkable result. All optical paths in the OPCB and the 90 bent fiber connectors were constituted with silica fibers so that very low link loss of about 1.5 dB was achieved. Employing the guide pins and conventional ferrule parts, complete passive alignment was achieved with accuracy within several micrometers in the assembling of the connectors and Tx/Rx modules in the OPCB. A data link of 5 Gb/s/ch 4 channels was successfully demonstrated in the assembled interconnection system. The graft of the existing precise optic components into PCB technology could open commercialization of PCB-based optical interconnection systems.

Asaaf et al (2008) formulated the design and performance of next-generation chip multiprocessors (CMPs) will be bound by the limited amount of power that can be dissipated on a single die [4]. Thus presenting photonic networks-on-chip (NoC) as a solution to reduce the impact of intrachip and off-chip communication on the overall power budget. The low loss properties of optical waveguides, combined with bit-rate transparency, allow for a photonic interconnection network that can deliver considerably higher bandwidth and lower latencies with significantly lower power dissipation than an interconnection network based only on electronic signalling. They explain why on-chip photonic communication has recently become a feasible opportunity and explore the challenges that need to be addressed to realize its implementation. It introduces a novel hybrid micro architecture for NoCs that combines a broadband photonic circuit-switched network with an electronic overlay packet-switched control network. In the continual drive toward improved microprocessor performance, power efficiency has emerged as a prime design consideration.

The limitations on power dissipation imposed by packaging constraints have become so paramount that performance metrics are now typically measured per unit power [1]. At the chip scale, the trend toward multi core architectures and chip multiprocessors (CMPs) for driving performance-per-watt by increasing the number of parallel computational cores is dominating new commercial releases [2-4]. The insertion of photonics in the on-chip global interconnect structures for CMP can potentially leverage the unique advantages of optical communication and capitalize on the capacity, transparency, and fundamentally low energy consumption that have made photonics ubiquitous in long-haul transmission systems. The construction of photonic NoC could deliver performance-per-watt scaling that is simply not possible to reach with all-electronic interconnects.

The photonics opportunity is made possible now by recent advances in nanoscale silicon photonics and considerably improved photonic integration with commercial CMOS chip manufacturing. They developed POINTS, an event-driven network traffic simulator, to quantitatively evaluate critical design aspects such as deadlock avoidance/recovery, optimal message size, path multiplicity (PM), and alternative flow control mechanisms. The compelling conclusion of the study is that the power expended on intrachip communications can be reduced by nearly two orders of magnitude when high-bandwidth communications is required among a large number of cores. To illustrate the operation of the proposed NoC, we describe the typical chain of events in the transmission of a message between two ports placed on different cores in the CMP, for example, a write operation that takes place from a processing unit in a core to a memory that is located in another core.

As soon as the write address is known, possibly even before the contents of the message are ready, a path-setup packet is sent on the electronic control network. The packet includes destination address information and, perhaps, additional control information such as priority or flow ID. After the message transmission is completed, a path tear down packet is sent to free the path resources for usage by other messages. Once the photonic message has been received and checked for errors, a small acknowledgment packet may be sent on the electronic control network to support guaranteed-delivery protocols. In the case where a path-setup packet is dropped in the router due to congestion, a path-blocked packet is transmitted by the dropping router to the source, back tracking the path travelled by the path-setup packet.

The topology of choice in design reflects the characteristics of the entire system CMP, where a number of homogeneous processing cores are integrated as tiles on a single die. The communication requirements of a CMP are best served by a 2D regular topology such as a mesh or a torus. These topologies well match the planar, regular layout of the CMP and the application-based nature of the traffic—any program running on the CMP may generate a different traffic pattern. As explained above, a regular 2D topology requires switches which are overly complex to implement using photonic technology. Using a folded-torus topology as a base and augment it with access points for the gateways. It developed a new event-driven simulator that is specifically tailored to provide support for the design exploration of the proposed photonic NoC.

The results of several studies performed using the simulator. The first study is about avoiding deadlock. Deadlock in torus networks has been studied extensively. When dimension-order routing is used, no channel depend on cycles are formed between dimensions, so deadlock involving messages travelling in different dimensions cannot occur. Virtual channel flow control has been proven successful in eliminating intra

dimension deadlocks and make dimension-order routing deadlock free. The power analysis of the electronic control network is based on the fact that this is essentially an electronic packet switched NoC, dimensions (12 _ 12 compared to 6). Assume that each photonic message is accompanied by two 32-bit control packets and the typical size of a message is 2 K bytes. Then the total power consumed by the electronic control network can be approximated.

The electronic control network is utilized lightly, and then the impact of static power becomes more dominant in the overall NoC power budget. However, recent technological breakthroughs in semiconductor processes, namely, Intel's 45 nm process leveraging high-K dielectrics, have been shown to reduce the gate leakage more than 10-fold. Having dramatically reduced the gate leakage, channel leakage remains the major challenge. To generate the 960 Gbps peak bandwidth. The modulated data streams are grouped using passive WDM multiplexers, so power is dissipated mainly in the 24 modulators and 24 receiver circuits in each gateway. Since there is presently no equivalent to the International Technology Roadmap for Semiconductors (ITRS) [1] for the photonic technology, predictions on the power consumption of photonic elements vary greatly. A reasonable estimate for the energy dissipated by a modulator/detector pair, at 10 Gbps, today is about 2 pJ/bit, based on recent results reported by IBM. It estimate that, using silicon ring-resonator modulators and SiGe photo detectors, the energy will decrease to about 0.2 PJ/bit in the next 8-10 years.

The power analysis used here is rather simplistic and uses many assumptions to ease the calculation and work around missing data, its broader conclusion is clear. The potential power difference between photonics-based NoCs and their purely electronic counterparts is significant. Importantly, once generated, the power consumed by propagating the optical signals off-chip across the system is essentially negligible and can enable true scaling for off-chip CMP high-bandwidth communications. Even when one accounts for inaccuracies in our analysis and considers predicted future trends, the advantages offered by photonics represents clear leap in terms of bandwidth-per-watt performance. First, multi core processors step into an era where high band width communications between large numbers of cores is a key driver of computing performance. Second, power dissipation has clearly become the limiting factor in the design of high-performance microprocessors.

The power dissipated on intrachip communication is a large and growing fraction of the total power budget. Third, recent breakthroughs in the field of silicon photonics suggest that the integration of optical elements in the CMOS electronics is likely to become viable in the near future. The intersection of these three factors suggests that silicon photonic technology can be used to construct NoCs, offering a promising low-power solution for high performance on-chip communication. The design of photonic NoC presents interesting and challenging problems. To address these problems, proposed a hybrid NoC micro architecture that combines a photonic circuit switched network with an electronic packet-switched network so that each technology is used advantageously: photonics for bulk-data transmission and electronics for network control.

Electronic packets are used to establish transparent light paths that carry high-bandwidth optical messages across a network of broadband optical switches. The technology required implementing the photonic devices (PSEs and 4 _ 4 switches) and their integration in large-scale NoCs is still immature. From the recent progress it is understood that within a small number of years, the enabling technologies will gradually become available to the designers of silicon-integrated circuits.

Dana et al (2008) presented a new methodology that many-core microprocessors will push performance per chip from the 10 giga flop to the 10 teraflop range in the coming decade. To support this increased performance, memory and inter-core bandwidths will also have to scale by orders of magnitude. Pin limitations, the energy cost of electrical signalling, and the non-scalability of chip-length global wires are significant bandwidth impediments [5]. Corona is a 3D many-core architecture that uses nanophotonic communication for both inter-core communication and off-stack communication to memory or I/O devices. Its peak floating-point performance is 10 teraflops. Dense wavelength division multiplexed optically connected memory modules provide 10 terabyte per second memory bandwidth.

A photonic crossbar fully interconnects its 256 low-power multithreaded cores at 20 terabyte per second bandwidth. Nano photonics offers an opportunity to reduce the power and area of off- and on-stack interconnects while meeting future system bandwidth demands. Optics is ideal for global communication because the energy cost is incurred only at the endpoints and is largely independent of length. Dense wavelength division multiplexing (DWDM) enables multiple single wavelength communication channels to share a waveguide, providing a significant increase in bandwidth density. Recent nanophotonic developments demonstrate that waveguides and modulation/demodulation circuit dimensions are approaching electrical buffer and wire circuit dimensions. Several benefits accrue when nano photonics is coupled to emerging 3D packaging [1]. The 3D approach allows multiple die; each fabricated using a process well-suited to it, to be stacked and to communicate with through silicon via (TSVs). Optics, logic, DRAM, non-volatile memory (e.g. FLASH), and analog circuitry may all occupy separate die and co-exist in the same 3D package.

Utilizing the third dimension eases layout and helps decrease worst case wire lengths. Advances in silicon nano photonics have made complete photonic on-stack communication networks a serious alternative to

electrical networks. Photonic interconnects are interesting because they can be much more energy efficient than their electrical counterparts, especially at high speeds and long distances. A complete nanophotonic network requires waveguides to carry signals, light sources that provide the optical carrier, modulators that encode the data onto the carrier, photodiodes to detect the data, and injection switches that route signals through the network. It is imperative that the optical components be built in a single CMOS-compatible process to reduce the cost of introducing this new technology. Waveguides confine and guide light and need two optical materials: a high refraction index material to form the core of the waveguide and a low index material that forms the cladding.

Crystalline silicon (index $n \approx 3.5$) and silicon oxide (index $n \approx 1.45$) both are commonly used in CMOS processes. A silicon oxide waveguide has typical cross-sectional dimensions of ≈ 500 nm with a wall thickness of at least $1 \mu\text{m}$. These waveguides have been shown to be able to carry light with losses on the order of 2–3 dB/cm and can be curved with bend radii on the order of $10 \mu\text{m}$. Corona is a tightly coupled, highly parallel NUMA system. As NUMA systems and applications scale, it becomes more difficult for the programmer, compiler, and runtime system to manage the placement and migration of programs and data. The architecture is made up of 256 multithreaded in order cores and is capable of supporting up to 1024 threads simultaneously, providing up to 10 teraflops of computation, up to 20 terabytes per second (TB/s) of on-stack bandwidth, and up to 10 TB/s of off-stack memory bandwidth. Synthetic workloads stress particular features and aspects of the interconnects.

The SPLASH-2 benchmark suite indicates their realistic performance. The SPLASH-2 applications are not modified in their essentials. For the SPLASH-2 applications, In four cases (Barnes, Radio city, Volrend, and Water-Sp) the L Mesh/ECM system is fully adequate. These applications perform well due to their low cache-miss rates and consequently low main memory bandwidth demands. FMM is quite similar to these. The remaining applications are memory bandwidth limited on ECM-based systems. For Cholesky, FFT, Ocean, and Radix, fast memory provides considerable benefits, which are realized only with the fast crossbar. LU and Ray trace are like Hot Spot: while OCM gives most of the significant speedup, some additional benefit derives from the use of the fast crossbar. They posit below possible reason for the difference between Cholesky, FFT, Ocean, and Radix on the one hand, and LU and Ray trace on the other, when examining the bandwidth and latency data.

To investigate the potential benefits of nano photonics on computer systems developed an architectural design called Corona. Corona uses optically connected memories (OCMs) that have been architected for low power and high bandwidth. A set of 64 OCMs can provide 10 TB/s of memory bandwidth through 128 fibers using dense wavelength division multiplexing. Once this memory bandwidth comes on chip, the next challenge is getting each byte to the right core out of the hundreds on chip. Corona uses a photonic crossbar with optical arbitration to fully interconnect its cores, providing near uniform latency and 20 TB/s on-stack bandwidth. The systems using optically connected memories and an optical cross bar between cores coupled perform 2 to 6 times better on memory-intensive workloads than systems using only electrical interconnects, while dissipating much less interconnect power. Thus nanophotonics can be a compelling solution to both the memory and network-on-chip bandwidth walls, while simultaneously ameliorating the power wall.

Keith et al (2011) formulated that the Optical links have successfully displaced electrical links when their aggregated bandwidth–distance product exceeds ~ 100 Gb/s-m because their link energy per bit per unit distance is lower [6]. Optical links will continue to be adopted at distances of 1 m and below if link power falls below 1 pJ/bit/m. The optical links directly to a switching/routing chip can significantly improve the switched energy/bit. The early experimental switched CMOS-vertical-cavity surface-emitting laser (VCSEL) system operating at Gigabit Ethernet line rates that achieves a switched interconnect energy of less than 19 pJ/bit for a fully non blocking network with 16 ports and an aggregate capacity of 20 Gb/s/port. The CMOS-VCSEL switch achieves an optical bandwidth density of 37 Gb/s/mm² even when operating at a modest line rate of 1.25 Gb/s and is capable of scaling to much higher peak bandwidth densities (~ 350 Gb/s/mm²) with 5–10 pJ/switched bit.

Computing and communications performance must continue to scale to meet the growing proliferation of capability and reductions in price expected from systems manufacturers. Indeed, the past three decades have witnessed over five orders of magnitude improvement in data processing performance per-dollar, in large part due to the investments and improvements in very large scale integration (VLSI) technologies. However, for next-generation systems, designer focus has turned to also reducing the power of computation and communication. Photonics can help. Optical links directly integrated into switch chips would remove the need for electrical chip I/O. This, of course, is beneficial only if dense photonic integration beyond what is expected with scaled VLSI technology is possible and if the integrated photonic links consume less power than electrical ones.

Today's VCSEL device has broad industry acceptance, in part due to its ability to be directly connected to electrical circuits at the chip and wafer levels and directly modulated at speeds in excess of 15 Gb/s [5-9] in the 25 Gb/s devices being actively pursued. Arrays of VCSELs have been bonded CMOS VLSI chips, with each VCSEL capable of up to 15 Gb/s modulation by the CMOS circuits it wide-pen eyes. The chip's power dissipation had three major contributors: $6 \text{ mW} \times 256 = 1.54 \text{ W}$ optical receivers; $12 \text{ mW} \times 256 = 3.1 \text{ W}$ VCSELs

plus drivers; and a 1.4 W switching fabric. A total maximum throughput of 320 Gb/s was possible with this chip. Several chips and system prototypes were assembled. Although the detectors routinely had 100% yield, the best OE VLSI chip had one non-functional VCSEL so that actual total throughput was slightly less than 320 Gb/s. Giga bit Ethernet data traffic at a relatively low bit rate of 1.25 Gb/s was transmitted through the switch. Even so, the system provided a switched energy of 19 pJ/bit including the optical input and VCSEL outputs.

The photonic network on a macrochip provides low power, high bandwidth, and high-density communication between sites. Every site in a microchip is interconnected to every other site via WDM links that run in orthogonal directions on two routing layers. The optical signals from the sites are coupled into, and between, the routing layers through OPxC face-to-face optical couplers. Anticipate a bisection bandwidth on the order of 10 TB/s traversing the microchip. A low-power transmitter requires three critical elements: an energy-efficient modulator, a low-power driver circuit, and intimate integration of these two components with low parasitic. It is possible to further optimize the ring modulator for better modulation quality.

The device maximized the E/O modulation by using an optimized waveguide design and an asymmetric p-n junction design, with a 5×10^{17} cm⁻³ p-doping concentration, a 1×10^{18} cm⁻³ n-doping concentration, and a junction offset of 50 nm to maximize the mode overlap with the depletion region for maximized index change [12]. Ge-p-i-n waveguide diodes were developed for high speed signal detection. They were fabricated using Luxtera's Ge-enabled OE process integrated in Freescale's HIP7_SOI130 nm CMOS. A 200–300-nm-thick Ge film is grown on top of the Si waveguide. Highly efficient optical transmitters employ a compact modulator either a resonant ring or an electro absorptive device in reverse bias, so the modulator appears to the driving circuit as a simple capacitance with no shunt current.

Creating efficient receivers is harder than creating efficient transmitters, because receivers convert small photocurrents into CMOS-level voltages and must tradeoffs energy consumption, performance, and SNR against each other. The SNRs correspond to certain BERs, the SNR becomes a primary design target for optical receivers. Even greater power savings from photonic links could be had, if processing, interconnect, and memory functions could be optically enabled and brought in close proximity to each other to create a multi-vehicle-sized "macro chip" that supports switching, processing, routing, and stacked memory functions. Optical links have successfully displaced electrical links for application domains when the aggregated bandwidth–distance product of the domain has exceeded ~100 Gb/s-m. It is also evident that commercial optical links provide lower bit energy per unit distance than their electrical counterparts.

The aforementioned is true today only for the longer lengths (>10 m) with VCSEL links, it is interesting to note that this advantage of optical links can be exploited in the future even for much shorter links because the communication bit energy per unit distance of electrical links does not scale down as one approaches chip-to-chip and even on-chip length scales. Hence, there is an opportunity to reduce total interconnect energy with photonics for a given system configuration. For switched or routed links, the integration of the photonics and electronics is vital to power savings. Simply providing optical links directly to a chip can not only alleviate constraints on electrical chip I/O but can significantly improve the switched energy per bit. Several experiments in integrating VCSEL-based links at the system level have been undertaken in an effort to increase bandwidth density and reduce system complexity and power. Even greater power savings from photonic links could be had, if processing, interconnect, and memory functions could be optically enabled and brought in close proximity to each other to create a multi-vehicle-sized "macrochip" that supports switching, processing, routing, and stacked memory functions.

Xiaowen et al (2014) proposed as modern computing systems become increasingly complex, communication efficiency among and inside chips has become as important as the computation speeds of individual processing cores. Traditionally, to maximize design flexibility, interchip and intrachip communication architectures are separately designed under different constraints. UNION is based on recent progresses in nanophotonic technologies [7]. It connects not only cores on a single CMP, but also multiple CMPs in a system. UNION employs a hierarchical optical network to separate interchip communication traffic from intrachip communication traffic. It fully utilizes a single optical network to transmit both payload and control packets. The network controller on each CMP not only manages intrachip communications, but also collaborates with each other to facilitate interchip communications. They compared UNION with a matched electrical counterpart in 45-nm process. Simulation results for eight real CMP applications show that on average UNION improves CMP performance by 3× while reducing 88% of network energy consumption. Modern computing systems have become increasingly complex to satisfy the growing performance demanded by applications.

The number of transistors available on a single chip increases to billions, chip multiprocessor (CMP) has become an attractive platform delivering high performance with limited power budget. For interchip communications, bus-based and ad-hoc architectures are still popular, and signals are mostly transmitted by electrical interconnects on printed circuit boards (PCB). The limitations of electrical interconnects such as high-delay and high-power consumption, are already shown in high performance systems. For intrachip communication, move from ad-hoc and bus-based architectures to network-on-chip (NoC) to alleviate issues

including poor scalability and limited bandwidth [1-2]. Optical interconnects, with advantages including ultrahigh throughput, low-delay and low-power consumption, are proposed to replace both inter- and intrachip electrical wires. For interchip communication, optical interconnects are studied for more than a decade and many promising research results are proposed [3-4].

For intrachip communication, with silicon photonics being mature, optical links are suggested to replace long metal wires on chip [5-8]. For optical interconnects, unlike electrical wires, the inter and intrachip channels can be interconnected seamlessly. Both on-chip and off-chip channels can be implemented with optical waveguides and they can be interconnected with passive couplers. The allowed operating bandwidths of both waveguides are broad enough for real applications. UNION uses a hierarchical optical NoC for each chip. The optical routers are connected in fat tree topology, and both payload and control data is transmitted in this single fat tree network. Fat tree is widely adopted in high performance systems and also NoC designs. In UNION, it is required that an on-chip network can be integrated with inter network in a hierarchical manner, and fat tree especially satisfies this requirement with inherent hierarchical property.

To extend the top-level routers on the tree to interconnect the off-chip network, forming a larger system. Fat tree is also suitable for central-control protocol proposed that helps setup the interchip network and also boosts the performance of intrachip network. Besides the data sub network, the control sub network can also be implemented as fat tree topology, which provides the opportunity of integrating two subnet works together as we proposed here. To highlight these merits of the fat tree topology, they compare it with mesh topology. The mesh topology does not own the same hierarchical property as fat tree. In UNION, if both sides of a transaction are within the same concentrator, packets are transmitted through a local electrical crossbar. On the other hand, if a core needs to send a packet out of the concentrator, it first tries to reserve an optical path to the destination concentrator.

The path is reserved successfully, the packet would be transmitted optically to the destination concentrator that would send it to the destined core through the local electrical crossbar. The interchip network connects all intrachip networks. for interchip communications. though bus-based communication architecture has limited scalability, it is still a viable low-cost choice for systems with a moderate number of chips. The low-cost design can improve the feasibility of the whole system. Another advantage of using bus is that we need not fix the system size at early stages of design time. In UNION's intertie network, the number of data bus channels is proportional to the number of top-level routers in the intrachip network. Specifically, each upward port of the top-level router of fat tree would access a separate bus channel. For 64-core CMPs, only 16 data bus channels are required. Each data bus channel is composed of on-chip silicon waveguides, polymer waveguides embedded on the PCB, and optical connectors connecting on chip waveguides with on-board waveguides.

Performance is measured in terms of the average number of iterations that an application can be finished in a given time. For most applications, CMPs using UNION achieve more than $3\times$ improvement compared with the CMPs using its electrical counterpart. The satellite receiver application is an exception that only 10% improvement is achieved by UNION. This is because that under this application, most traffics are on-chip transmissions. When most of the data flows are confined on an individual chip, the contribution of the unified design would not be well illustrated [11]. They propose an adaptive power control mechanism to improve the power efficiency of the network, and it is applicable to other optical NoC architectures as well. The mechanism is based on the observation that a large portion of power is consumed by lasers in the network. For example, in an 80-nm design, while the power consumed by a simple data link to transfer one bit is about 2.5 pJ, the laser source consumes about 1.68 pJ which considers a large proportion.

The control scheme is to reduce the power consumption of laser dynamically. In UNION, the average packet delay and energy consumption increase very slightly respecting to the number of chips, showing very good scalability. But in the matched electrical network, the packet delay and energy consumption increase quickly. The performance gap between UNION and the electrical network is widened with larger number of chips. In addition, for the electrical network, there is a giant leap between the one- and two-chip system, showing that there is a huge performance gap [10] between on-chip and off-chip interconnects. This phenomenon does not exist in UNION. For optical network, interchip propagation delay is very low and the arbitration delay is independent of the hops. For electrical network, interchip traversing involves much larger propagation delay and the serialization delay. The network controller on each CMP not only managed intrachip communications, but also collaborated with each other to facilitate interchip communications. They compared CMPs using UNION with those using a matched electrical counterpart in 45-nm process. Simulation results of eight applications showed that on average UNION improved CMP performance by $3\times$ while reducing 88% of network energy consumption.

III. Comparative Analysis

Table 1. Comparative analysis of different techniques on UNION

Author	Year	Algorithm	Advantages	Disadvantages	Results
Xiaowen et al	2014	Nanophotonic Technologies	1. Ultrahigh throughput 2. Low-delay 3. Low –power consumption 4. Inter and Intra channels can be interconnected seamlessly 5. Low path setup time	1. Fabrication of waveguide is difficult	1. Improved CMP Performance 2. 88% of reduction in network energy 3. Low power Consumption
Keith W et al	2011	VCSEL-CMOS Vertical- cavity surface emitting Laser Silicon	1. High speed 2. chip area savings	1. system complexity	1. bandwidth–Distance product exceeds 100 Gb/s-m 1. It provides 10
Dana al	2008	Silicon Nanophotonic	1. It provide 2 to 6 times more performance on many memoryintensive workloads, while simultaneously reducing power.	LASER source Selection	TB/s of through 128 fibers using Dense Wavelength Division multiplexing.
Assaf et al	2008	1. photonic networks-on-chip (NoC) 2. novel hybrid microarchitecture for NoCs	1. higher bandwidth and lower latencies with significantly lower power dissipation	1. Difficulty in the case of short message communication	Low power Dissipation
Sung Hwan et al	2006	Fiber-embedded optical printed-circuit board (OPCB)	1. Optical Interconnection increases the speed of processing 2. Low power dissipation	1. fabrication is complex	Increased Processing Speed
Nevin et al	2006	Integration of CMOS-compatible Optical technology to on-chip cache-coherent Buses	1. Reduced power, Area 2. Improved performance	1. components requirement is high	Increased data Communication Speed
Wolf et al	2005	H.264 HDTV Decoder	1. less power, 2. less silicon area 3. less metal area, 4. 63% less switch capacity, 5. 69% less link capacity to achieve 2X performance	1. many topologies are required	2 times Increased Performance
Ludovici et al	2003	2D mesh Architecture	High speed	Complex circuits	Time Management
Pawan et al	2002	optical signaling System	1. Less delay 2. Less power consumption	Selection of optical source is difficult	Less delay
Dally et al	2001	Network on chip	Low power	High implementing cost	Improved Performance

Table shows the comparison between the different techniques used in chip microprocessor to improve its performance. There are different methods are used ,among this the interconnections using optical method is the best because The network controller on each CMP not only manages intrachip communications, but also

collaborate with each other to facilitate interchip communications. It compared UNION with a matched electrical counterpart in 45-nm process. Simulation results for eight real CMP applications show that on average UNION improves CMP performance by 3x while reducing 88% of network energy consumption.

IV. Conclusion

The disadvantages of using electrical network can be overcome by using optical interconnections. It employed hierarchical optical network to separate interchip communication traffic from intrachip communication traffic. It fully utilized a single optical network to transmit both payload and control packets. By using this method the performance of the CMP is improved by 3x while reducing 88% of network energy consumption.

References

- [1]. Xiaowen Wu, Yaoyao Ye, Jiang Xu, Wei Zhang, Weichen Liu (2014). "UNION: A Unified inter/intra
- [2]. Optical Network for Chip Multiprocessors" *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 22, no. 5,
- [3]. Nevin Kirman, Meyrem Kirman, Rajeev K. Dokania, Jos'e F. Martinez, Alyssa B. Apsel
- [4]. Matthew A. Watkins, David H. Albonese. Leveraging Optical Technology in Future Bus-based Chip Multiprocessors. The 39th Annual IEEE/ACM International Symposium on Microarchitecture.
- [5]. Sung Hwan Hwang, Mu Hee Cho, Sae-Kyoung Kang, Hyo-Hoon Park, Han Seo Cho (2006)
- [6]. Passively Assembled Optical Interconnection System Based on an Optical Printed-Circuit Board, DOI: 10.1109/PS.2006.4350170 Conference: Photonics in Switching,
- [7]. Assaf Shacham, Keren Bergman, Luca P. Carloni (2008). Photonic Networks-on-Chip for chip microcontroller. *IEEE Transactions on Computers*, vol. 57, no. 9.
- [8]. Dana Vantreasey, Robert Schreiberz, Matteo Monchieroz, Moray McLarenz, Norman P. Jouppizmarco, Fiorentinoz, Al Davis, Nathan Binkertz, Raymond G. Beausoleilz (2008) Corona: System Implications of Emerging Nanophotonic Technology
- [9]. I. A. Young, E. Mohammed, J. T. S. Liao, A. M. Kern, S. Palermo, B. A. Block, M. R. Reshotko, and P. L. D. Chang (2010). Optical Technology for Tera-scale Computing. *IEEE J. Solid-State Circuits*
- [10]. J. Xu, W. Wolf, J. Henkel, and S. Chakradhar (2005). H.264 HDTV Decoder Using Application-Specific Networks-on-Chip, in Proc. IEEE Int. Conf. Multimedia Expo, Jul. 2005, pp. 1508–1511.
- [11]. P. P. Pande, C. Grecu, A. Ivanov, and R. Saleh (2003). Design of a Switch for Network on Chip Applications. In Proc. Int. Symp. Circuits Syst., vol. 5, 2003, pp. V-217–V-220.
- [12]. P. Kapur and K. C. Saraswat (2002). Comparisons between Electrical and Optical Interconnects for On-Chip Signaling. In Proc. IEEE Int. Interconnect Technol. Conf, pp. 89–91
- [13]. J. Dally and B. Towles (2001). Route packets, not wires: On-chip interconnection networks, in Proc. Design Autom. Conf. pp. 684–689.
- [14]. R. Yoshimura, M. Hikita, M. Usui, S. Tomaru, and S. Imamura (1997). Polymeric Optical Waveguide Films with 45° Mirrors Formed with a 90° V-shaped Diamond Blade, *Electron. Lett.*, vol. 33, no. 15, pp. 1311–1312.