UC-PHOTON: A Novel Hybrid Photonic Network-on-Chip for Multiple Use-Case Applications

Shirish Bahirat and Sudeep Pasricha

ECE Department, Colorado State University, Fort Collins, CO

{shirish.bahirat, sudeep}@colostate.edu

Abstract

Multiple use-case chip multiprocessor (CMP) applications require adaptive on-chip communication fabrics to cope with changing use-case performance needs. Networks-on-chip (NoC) have recently gained popularity as scalable and adaptive on-chip communication fabrics, but suffer from prohibitive power dissipation. In this paper we propose *UC-PHOTON*, a novel hybrid photonic NoC communication architecture optimized to cope with the variable bandwidth and latency constraints of multiple use-case applications implemented on CMPs. Our detailed experimental results indicate that *UC-PHOTON* can effectively adapt to meet diverse use-case traffic requirements and optimize energy-delay product and power dissipation, with scaling CMP core count and multiple use-case complexity. For the five multiple use-case applications explored in this work, *UC-PHOTON* shows up to 46× reduction in power dissipation and up to 170× reduction in energy-delay product compared to traditional electrical NoC fabrics, highlighting the benefits of using the novel communication fabric.

1. Introduction

In recent years, rapid advances in technology scaling and increases in application complexity have given impetus to the design of chip multiprocessors (CMP) with multiple components (processors, memories, peripherals) integrated on a single chip. CMPs can support greater levels of parallelism and have been shown to provide significant improvements in performance-per-watt compared to uni-processor systems-on-chip (SoC) clocked at higher frequencies. Already, numerous CMP designs are commercially available today from several vendors for a wide range of computing systems from Blu-Ray recorders, car navigation systems, digital TVs, gaming consoles, to exa-flop supercomputers. Some examples include the Sony/IBM/Toshiba Cell [1], Fujitsu FR-1000V [2], NEC/ARM MPCore and MP211 [3], and Renesas SH-X3 [4], which have been developed for the consumer electronics domain.

One of the most challenging problems in CMP design today is the design of the on-chip communication fabric that inter-connects the multiple cores on a chip [5]. The on-chip communication fabric is responsible for satisfying strict latency and bandwidth constraints that are relentlessly increasing in stringency as application performance approaches the peta- and exa- flop levels. Unfortunately, on-chip interconnects have not scaled well with process technology. In ultra-deep submicron (UDSM) technology nodes below 65 nm, not only have interconnects become longer, but the signal delay on these long (global) interconnects has been steadily increasing with each successive technology generation, and now far exceeds gate delay. The International Technology Roadmap for Semiconductors (ITRS) acknowledges that delay on global interconnects has now become a major performance bottleneck, and the topmost challenge for the semiconductor industry [6]. In addition to becoming a potential source for performance bottlenecks, interconnects on a chip suffer from reduced reliability due to UDSM effects such as capacitive and inductive crosstalk, and higher dynamic and leakage power dissipation.

To cope with communication demands in emerging CMP designs, there has been a gradual shift away from circuit-switched bus-based communication architectures to packet-switched networks-on-chip (NoCs) [5][7]. Hierarchical and crossbar-based shared bus architectures lack the scalability to support high bandwidths, and are also more susceptible to intra-die process variations and interference due to crosstalk and external electromagnetic sources. NoCs generally offer a more regular structure and a layered design methodology, which leads to better predictability and improved reliability in UDSM technologies. Most importantly, NoCs can support higher bandwidths and are thus being hailed as a promising on-chip communication fabric for emerging CMP designs. However, in practice, the few existing implementations of NoCs have been found to have several drawbacks stemming from the large number of network interface (NI), router, link, and buffer components in NoCs that lead to a high area overhead and prohibitive power dissipation. For instance, recent NoC prototypes have shown NoCs taking a substantial portion of system power, e.g., ~40% in the MIT RAW chip [8] and ~30% in the Intel 80-core Teraflop chip [9]. Recent studies have also suggested that NoC power dissipation is much higher (by a factor of 10×) than what is needed to meet the performance needs of future CMPs [10][11]. Thus, radical new solutions are required to overcome the power brick wall facing NoCs in the next few years.

The relatively recent phenomenon of digital convergence [12] puts a further strain on communication architecture design in emerging CMPs. Applications in the digital convergence era have multiple operating modes, called *use-cases*, that have distinct workload and communication traffic characteristics. For instance, smart cellular phones today can support a combination of several functionalities including making and receiving calls, playback for MP3s, FM radio and video, GPS navigation, PDA support, wireless web browsing, games with 3D graphics, Bluetooth syncing, and so on. Some of the functionalities, such as 3D gaming have a much higher computational and communication workload to render sophisticated graphics and maintain high frame display rates. In contrast, MP3 playback requires simpler audio codec decoding and has a much lower computation and

11th Int'l Symposium on Quality Electronic Design

communication performance requirement. With the number of use-cases increasing rapidly in applications today (~tens to hundreds), designers are finding it extremely challenging to create CMP implementations that can support multiple heterogeneous use-cases. Figure 1 illustrates an example of a network routing application with multiple use-cases. The application is implemented as a CMP with two general purpose microprocessors, a digital signal processor (DSP), two on-chip memories and a network interface (I/O). Figure 1(a)-(c) show three use cases of the application that have vastly different bandwidth requirements between components. It is also possible for the application to execute multiple use cases simultaneously, as shown in Figure 1(d), where use cases 1 and 2 execute at the same time (called a *compound* use case). Figure 1(e) shows how the average workload traffic bandwidth varies for the use cases during application execution. Typically, switching between use cases takes place when users interact with the application or if there is a change in the environment (e.g., change in wireless signal strength, or battery level). Such a switch in use cases results in the temporal switch of application task and communication graphs. It is highly likely that a communication architecture customized for a single use case may not meet performance requirements for another use case. Thus there is a need to enhance traditional on-chip communication fabrics for multiple use-case applications.
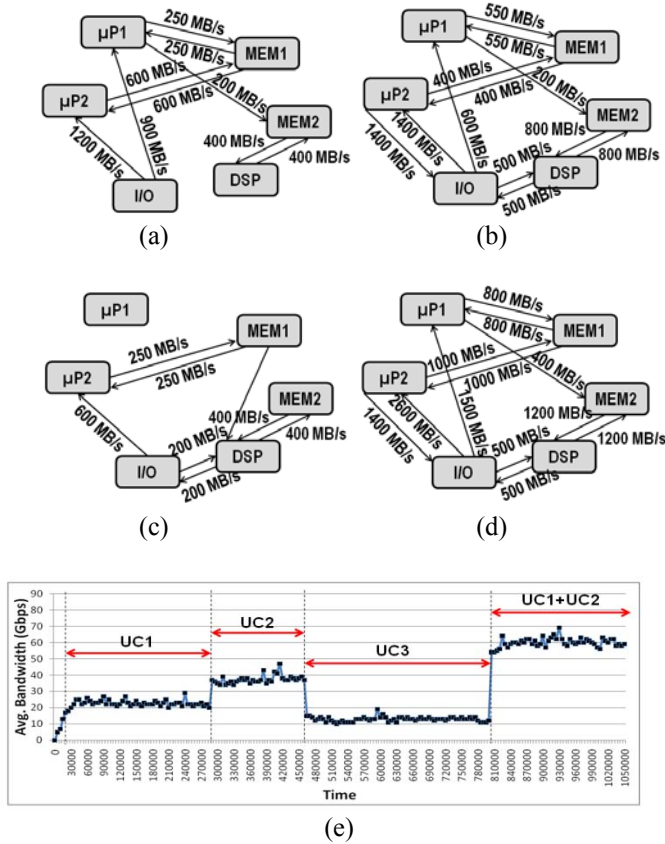
In this paper, we propose using a novel hybrid photonic NoC communication architecture called *UC-PHOTON* to cope with emerging multiple use-case applications and maximize performance-per-watt. *UC-PHOTON* is comprised of one or more photonic ring paths coupled to a traditional 2D electrical mesh NoC architecture. The photonic paths offload global communication from the electrical network, improving packet latency and reducing communication power dissipation. In addition, *UC-PHOTON* supports dynamic reconfiguration of the electrical and photonic networks. This enables runtime adaptation to changing traffic patterns, which allows network resources to be optimized for even lower power dissipation. Experimental results on several multiple use-case CMP benchmark implementations indicate that *UC-PHOTON* can scale with increasing use-case count and core count to save orders of magnitude power, reduce energy-delay product, and improve performance compared to traditional electrical NoC architectures.

2. On-chip Photonic Interconnects: Overview

Recently, photonic interconnects have been proposed as a solution to overcome the on-chip communication power bottleneck [13]. Photonic interconnect technology is of interest because it has been shown to be much more energy efficient compared to copper (Cu) interconnects especially at high speeds and long distances [14]-[16]. The ability of photonic waveguides to carry many information channels simultaneously increases interconnect bandwidth density significantly, eliminating the need for a large number of wires to achieve adequate bandwidth. Photonic interconnects are becoming standard in data centers, and chip-to-chip photonic links have been demonstrated [17]. This trend will naturally bring photonic interconnects into the on-chip stack, particularly as a means to enable high bandwidth and low power data communication between hundreds of cores in future CMPs. Recent advances in the field of nanoscale silicon photonics have enabled highly integrated photonic interconnect-based components in CMOS-based ICs [18]-[20]. In fact, photonic elements have now become available as library cells in standard CMOS processes.
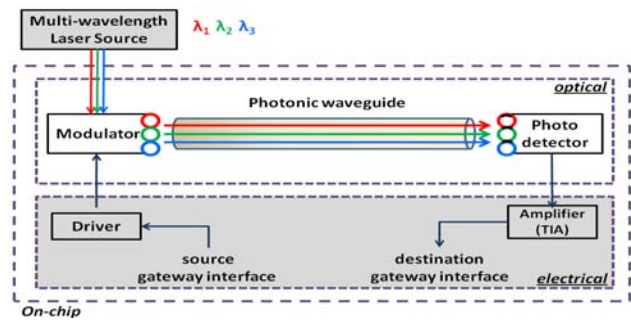

Figure 1: Multiple use case application with use cases (a) UC1, (b) UC2, (c) UC3, (d) UC1+UC2, (e) workloads

Figure 2: Building blocks of on-chip photonic interconnects

Figure 2 shows a high level overview of the key on-chip photonic transmission components as considered in this work: a multi-wavelength laser light source, modulator, photonic waveguide, and photodetector. Multiple wavelengths of light are coupled on to the chip using photonic fibers from off-chip (or on-stack) multi-wavelength

mode-locked lasers [21]. Several simultaneous data transfers over multiple wavelengths are possible in photonic interconnects, and such wave division multiplexing (WDM) is critical for ensuring high bandwidth transfers [22]. For a WDM scheme with λ wavelengths, allocation of wavelengths to traffic streams is done using 'multiplexing by core', with each of the n interfacing cores having exclusive access to λ/n wavelengths. This limits the number of transmitters, but provides substantial power savings. The photonic waveguide is made of CMOS-compatible silicon oxide, which has been shown to carry light with low losses (on the order of 2–3 dB/cm) and can be curved with bend radii on the order of 10 μm [23].
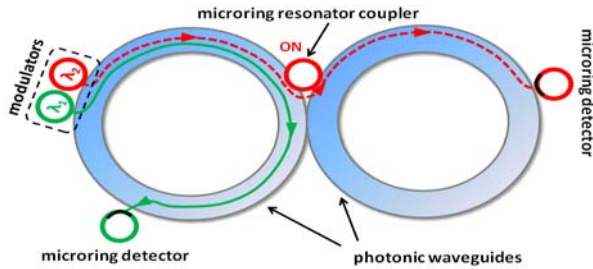

Figure 3: Microring resonator coupling

Wavelength-selective nanophotonic silicon modulators made out of microring resonator structures [24] are used to convert electrical signals into light at the source, for transmission. Microring resonators are also used at the destination as filter structures to "drop" the corresponding wavelength from the waveguide into a local photodetector device [25] that converts the light signal back to an electrical signal. Trans-impedance amplifier (TIA) circuits are used to amplify analog electrical signals at the receiver to digital voltage levels. Future CMPs with hundreds of cores will require multiple photonic waveguides to meet performance requirements. In such systems, microring resonators can be used as couplers to couple light between multiple waveguides. Figure 3 shows an example of this coupling, with light of wavelength λ_2 being coupled from the first to the second photonic waveguide. The coupling is enabled by injecting charge into the microring resonator coupler to change its index of refraction so that only a resonant wavelength λ_2 is coupled through it, while other non-resonant wavelengths (such as λ_1) remain unaffected. Ring filters and modulators must also be thermally tuned to maintain their resonance under on-die temperature variations [35]. We assume a single heater element per microring resonator structure in this work for this purpose.

An important consideration for WDM enabled photonic interconnects is the optical loss in its components. Optical loss impacts system design as it sets the required optical laser power and correspondingly the electrical laser power (at a roughly 30% conversion efficiency). In addition to the waveguide, optical losses exist for couplers, modulator/filter resonators, waveguide crossings, photodetectors, and also due to non-linearity. Section 6.1 summarizes the losses we take into account for our photonic system, to accurately characterize its power dissipation.

3. Related Work

Several researchers have compared the physical properties of on-chip electrical and photonic interconnects [13][16][26][27] at the circuit level and highlighted the signal speed and power benefits of photonic interconnects. Other work from industry and academia has been focusing on photonic device fabrication, for components such as gigascale

modulators (e.g., [24], [28]), photodetectors (e.g., [25], [29]), couplers, switches (e.g., [30], [31]), buffers, on-chip waveguides and on-chip WDM devices (e.g., [32]). A few recent works have explored the system-level impact of using hybrid photonic-electric interconnect architectures and proposed hybrid Cu-photonic crossbars [33][34], Clos networks [35], fat-trees [36] and meshes [37]. These works use WDM to increase bandwidth in photonic waveguides and non-blocking micro-resonator based photonic switches [38] for routing photonic messages. Our previous work [39] explored a simpler WDM-enabled hybrid Cu-photonic architecture with a parallel photonic ring waveguide interfaced to an electrical mesh NoC. While this architecture is more area and cost effective than other hybrid photonic-electric topologies (e.g., ~60× lower photonic layer area footprint than hybrid photonic mesh [37]), it is not scalable for large CMP designs.

Most of the work on designing NoCs has focused on optimization for an application with only a single operating mode (use case), e.g. [40]-[42]. These techniques lead to sub-optimal designs for today's applications with multiple use cases [43]. Several works [44]-[49] (including our own previous work [50]) have explored dynamic reconfiguration to adapt to changing traffic patterns for a single use case application, using dynamic voltage scaling/dynamic frequency scaling (DVS/DFS), adaptive routing schemes, and adaptive arbitration to improve performance and power dissipation. Only recently have a few approaches started to focus on designing and optimizing NoCs to meet performance constraints of multiple use-cases. Murali et al. [51][52] proposed using DVS/DFS, adaptive routing, and adaptive time division multiple access (TDMA) slot allocation to reduce NoC power dissipation. Hansson et al. [43][53] described an optimization technique to reduce switching time between use cases on a mesh NoC fabric. However, none of these works have explored runtime reconfiguration for hybrid photonic NoC architectures to optimize power dissipation.
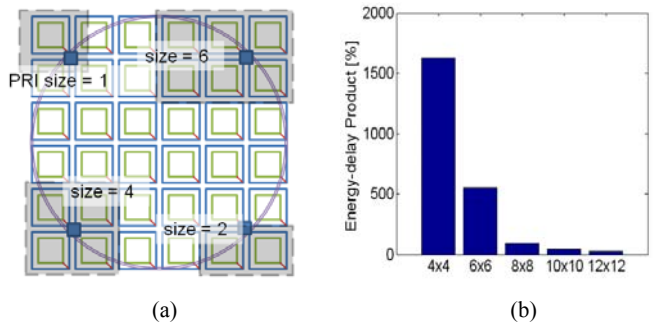

Figure 4: (a) 6x6 hybrid NoC with photonic ring and various sizes of photonic regions of influence (PRI), (b) % improvement in energy-delay product for hybrid NoC with photonic ring compared to conventional 2D mesh NoC, with scaling CMP complexity.

This paper makes two novel contributions. First, we extend our previously proposed low-cost photonic ring topology [39] and explore novel multi-ring topologies to improve performance scalability for emerging CMPs with hundreds of cores. Second, we explore runtime reconfiguration of both the electrical and photonic networks

to significantly improve communication performance and reduce power dissipation for multiple use-case applications. Our proposed *UC-PHOTON* architecture enables several new techniques for runtime optimization, in addition to existing techniques such as DVS/DFS, clock gating, adaptive TDMA slot allocation, and adaptive arbitration.

4. UC-PHOTON Overview

4.1. Background

Our proposed *UC-PHOTON* communication architecture utilizes 3D integration with separate planes for logic and silicon photonics. The basic architecture comprises of an electrical NoC interfaced to a silicon photonics layer that has photonic waveguide-based interconnect paths. In general, photonic waveguides and components for complex topologies such as mesh, torii, and fat trees can be prohibitively expensive in terms of fabrication cost and area overhead. Consequently, in our previous work [39], we proposed a low overhead hybrid photonic NoC architecture with a parallel ring-based photonic waveguide interfaced to a traditional 2D electrical mesh NoC. *Gateway interface* routers provided the connectivity between the electrical layer and the modulators and photodetectors in the photonic layer. The photonic ring was shown to provide a faster and more energy-efficient path for on-chip global communication compared to traditional electrical NoCs.

To improve scalability, a photonic region of influence (PRI) was also defined in [39], which refers to the number of cores around the gateway interface that can utilize the photonic path for communication. Figure 4(a) shows a 6×6 CMP with varying PRI sizes at the four gateway interfaces. If a router falls under a PRI, it is modified to additionally consider the photonic path for global communication for incoming flits. Note that while the sizes for the PRIs are shown as different at each gateway interface in the figure, this is for illustration purposes only, and in practice we assume a fixed PRI size for all gateway interfaces. For smaller sized systems (e.g., 4×4 CMPs), limiting the number of cores interfacing with each gateway interface to one (i.e., PRI size = 1) may be sufficient to offload a majority of the global communication from the electrical network. For more complex systems (e.g., 8×8 CMPs) a larger PRI size may be more appropriate. However, we have found that as the system size increases (e.g., 10×10, 12×12 cores etc), increasing the PRI size provides rapidly diminishing returns.

Figure 4(b) shows the % improvement in the energy-delay product of the hybrid photonic-ring NoC compared to a conventional electrical mesh NoC, with scaling CMP complexity. For each CMP configuration, results were averaged over several SPLASH-2 benchmark [54] implementations, and optimal PRI sizes (to get the lowest energy-delay product) were used. It is clear from the figure that the benefits of using the photonic ring become insignificant for large CMP sizes. This is primarily because the photonic ring is under-utilized due to its limited coverage area on chip, even though the global communication requirements are higher. Thus, more scalable hybrid communication fabrics are needed to cope with the more stringent latency and power constraints of future CMPs.

4.2. Topology

To overcome the scalability limitations discussed above, we propose a more general hybrid photonic NoC architecture called *UC-PHOTON* that employs multiple photonic ring paths. The topology can be characterized by a 6-tuple $<k,b,n,r,w,c>$ where k is number of photonic rings, b is the bitwidth of the photonic waveguides, n is number of uplinks/downlinks (i.e., gateway interfaces), r is the PRI size, w is the number of WDM channels, and c is the number of cores in the CMP. For the purpose of this work, we analyzed various 6-tuple spaces to better understand the

scalability of various configurations. Figure 5 shows four configurations of the proposed architecture with varying numbers of photonic ring paths (k = 2,3,4,5; with n = 8,12,16,16 gateway interfaces respectively) that can better span an IC chip. The configurations shown are for a fixed sized CMP (c = 36), have a bitwidth b = 256, PRI size r = 2, and w = 16 WDM channels. The building blocks of the photonic communication, such as the modulators, waveguides, photo-detector receivers, and microring resonator switches are the same as described in Section 2. The electrical mesh NoC requires modifications in the router architecture to incorporate additional traffic to/from the photonic layer. This is discussed in detail next.
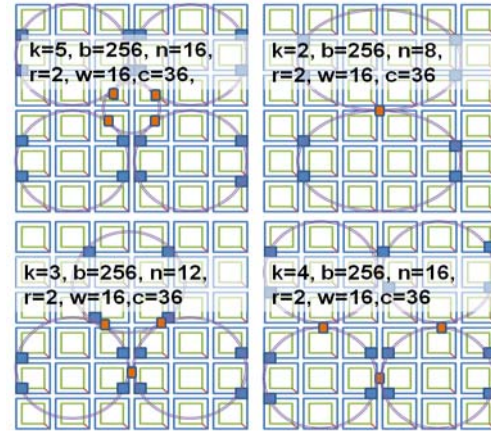

Figure 5: Four configurations of the proposed hybrid photonic NoC architecture for a 6x6 CMP

4.3. Routing, Flow Control and Deadlock Avoidance

The *UC-PHOTON* communication architecture supports wormhole switching, XY dimension order routing for routing flits in the electrical NoC, and a modified PRI-aware XY routing scheme for selective data transmission through the photonic links. Communicating cores lying within the same photonic region of influence communicate using the electrical NoC (intra-PRI transfers). Cores that need to communicate and reside in different photonic regions of influence communicate using the photonic paths, if the size of the data is above a certain user-defined threshold (inter-PRI transfers). In this way large data messages can be offloaded from the electrical NoC and sent over a faster and more energy efficient photonic path. Transfers between cores lying outside photonic regions of influence occur normally via the electrical network. Network interfaces (NIs) ensure that header flits contain coordinates of the source and destination of the packet being injected into the NoC, as well as a flag indicating that the message size is large enough to potentially traverse a photonic path (for inter-PRI transfers). All routers have 'region validation' units that use XY routing for intra-PRI transfers or for transfers to cores not residing in any PRIs via the electrical NoC. Otherwise if an inter-PRI transfer is detected by the 'region validation' unit at the source router, the packets are re-routed to the gateway interface of the local PRI using XY routing, traverse the photonic link to the destination gateway interface, and then are routed to the destination router using XY routing. The data traversing the photonic path is not buffered in the photonic layer, and the photonic transfer can therefore be considered to be a form of *photonic circuit switching*. Flow control for these transfers is implemented using the switch-to-switch ACK/NACK scheme in both the electrical and photonic links, to ensure that the destination is able to accept the transmitted data. Flits traversing a photonic link receive ACK/NACK information from the destination via the electrical NoC, just like with electrical links.
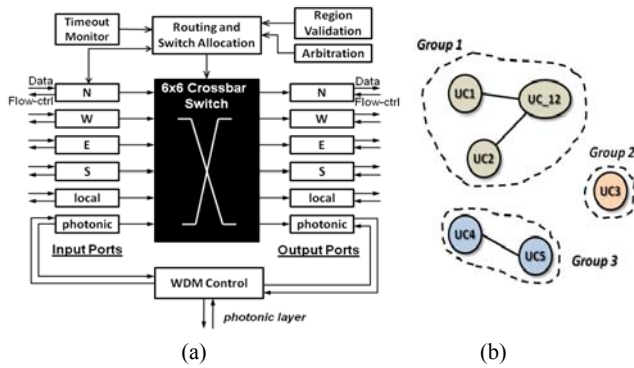
Figure 6: (a) gateway interface electrical router architecture, (b) use-case critical transition graph (CTG)

There are two types of routers used in UC-PHOTON: *(i)* regular electrical mesh routers that have 5 I/O ports (N, S, E, W, local core) with the exception of the boundary routers that have fewer ports, and *(ii)* gateway interface routers that have six I/O ports (N, S, E, W, local core, photonic link) and are responsible for sending/receiving flits to/from photonic interconnects in the photonic layer. Both types of routers have an input and output queued crossbar with a 4-flit buffer on each input/output port. Figure 6(a) shows a high level block diagram of the gateway interface router. Note the additional ports for interfacing with the photonic rings. These ports are connected to a 'WDM control' module that controls wavelength assignment to different traffic flows, to enable WDM for high bandwidth photonic communication. If multiple requests contend for access to the photonic waveguide at a gateway interface, then the request with the furthest distance to the destination is given preference (other schemes can be used here as well).

While XY routing has been proven to be deadlock-free for mesh-like regular NoCs (as no channel dependency cycles can be formed between dimensions), the modifications made to this routing scheme in our architecture to accommodate photonic transfers may end up creating deadlock conditions. We extensively studied deadlocks in the proposed architecture when packets traverse the photonic ring paths. To overcome a potential deadlock, we arrived at using low overhead timeout flits sporadically interleaved with the flits for the long data messages traversing the photonic paths. This is a form of regressive deadlock recovery [55]. If a timeout flit reaches a router where flits are blocked, a 'timeout monitor' module in the router can detect a timeout event and recognize potential cases where flits are blocked due to deadlock, and drop the blocked flits, while sending a NACK signal in the reverse direction to indicate the flits being dropped. This allows the system to unblock and recover from potential deadlock. While the method has the overhead of the additional flits in long messages intended for photonic links and monitoring module in the routers, this is still simpler than other potential deadlock resolution alternatives such as keeping a reserved deadlock free channel and draining deadlocked packets through this channel until the deadlock condition clears.

4.4. Dynamic Reconfiguration

A key feature of *UC-PHOTON* is the support for dynamic reconfiguration for low power operation. Since different use-cases have different traffic flows, and bandwidth/latency requirements, runtime reconfiguration strategies can be employed at the transition point between use-case executions to save power. There is always a timing overhead incurred when transitioning between use-cases, primarily to load the new use-case data and code, distribute control signals across the chip, and gracefully shut down the current use-case. Some use cases are critical and need to be loaded and run quickly. For other use cases, the transitioning time can be much longer, from hundreds of micro-seconds to several milli-seconds. In

this time we can implement strategies to reduce power dissipation. Obviously, the runtime reconfiguration strategies possible during critical use-case switching are much more restricted than for regular use-cases switching. We define a *critical transition graph* as an undirected graph $CTG(V,E)$ where each vertex $v_i \in V$ is a use case, and an undirected edge (v_i, v_j) represents the need for fast switching between use cases v_i and v_j. Figure 6(b) illustrates a *CTG* for a multiple use-case application, where for instance use cases *UC4* and *UC5* need to implement fast switching.

For critical use-case transitions, low overhead runtime reconfiguration schemes are desirable to save power. *UC-PHOTON* employs low-overhead runtime optimization techniques *(i)-(iii)* described below for critical use-case transitions, and all the techniques *(i)-(v)* for non-critical use-cases transitions:

(i) DVS/DFS: Dynamic supply voltage and clock frequency scaling (DVS/DFS) is one of the most widely used runtime optimization techniques to reduce power dissipation. In our approach, NoC link and router frequencies are dynamically adapted to meet performance requirements while consuming the minimum power. Since use-case performance requirements are known in advance, the available slack can be precisely utilized to achieve maximum power dissipation. An almost quadratic reduction in dynamic power dissipation can be achieved using this approach. We use a conservative model for voltage scaling, where we assume that the square of the voltage scales linearly with the frequency [56].

(ii) Clock gating: For some use cases, not every link or router of the NoC needs to be active to implement all communication flows. In such a scenario, to reduce dynamic power dissipation, clock gating can be employed. Clock gating is the most effective solution for optimizing the dynamic power, and is supported by most commercial synthesis and optimization tools. The main idea behind clock gating is to shut-down a circuit's blocks that are not performing useful computations during some particular clock cycles. In our approach, clock gating is used to shut down links and routers that do not need to be accessed for a use-case.

(iii) Dynamic WDM: Wavelength division multiplexing allows several photonic signals to be transmitted simultaneously in a single photonic waveguide using different wavelengths which do not interfere with each other. WDM can thus significantly improve photonic interconnect bandwidth density over electrical interconnects. We assume that each of the photonic waveguides has λ available wavelengths for WDM, thus creating a λ-way WDM photonic path. The value of λ has significant implications for performance, cost and power since using a larger number of wavelengths improves bandwidth but requires additional modulators and receivers, which increases area, cost and power overhead in the photonic layer. In our work, we utilize a practically achievable conservative λ value of 16. Since the dissipated power in the modulators and receivers is typically a linear function of the number of WDM channels employed, reducing the number of WDM channels can save power. Our hybrid photonic NoC supports rapidly varying the number of WDM channels during use-case transitions, by shutting off channels (modulators/receiver pairs) when data bandwidth requirements are low to save power, and enabling the channels when bandwidth requirements become high, to maintain performance goals.

(iv) Runtime PRI reconfiguration: The size of the region of photonic influence impacts the photonic path utilization. Small region sizes promote more transfers via the electrical NoC, while large region sizes increase the traffic flows eligible for transfer via the photonic rings. However, increasing the PRI size beyond a certain point can be counter-productive, increasing latency and power dissipation, as shown in our previous work [39]. Since application traffic characteristics can change across use-cases, a single PRI size may not be adequate in optimally distributing traffic flows between the electrical and photonic paths to minimize power

dissipation. Thus *UC-PHOTON* supports dynamically varying the PRI size at runtime during non-critical use-cases transitions to track changing application traffic requirements, and achieve low power operation. The reconfiguration step involves updating region boundary coordinates in tables in the 'region validation' units of the NoC routers, which can take several hundreds of cycles, and hence is only applied during non-critical use-case transitions.

(v) Adaptive TDMA slot allocation: The TDMA slot allocation in a router for different traffic flows controls the bandwidth and also the average latency of packets for the flows in a NoC. Since different use-cases have different bandwidth and latency requirements, changing the TDMA slot allocation during use-case transition is a way to adapt to the new use-case. The goal is to give more slots to critical traffic flows in a use-case with more stringent constraints. Since the TDMA slot allocation information is assumed to be stored in a separate memory, it takes several hundreds of cycles to update the TDMA slot allocations in all the NoC routers. Thus, this technique is only applied during non-critical use-case transitions.

6. Experiments

6.1. Experimental Setup

Photonic 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 photonic interconnects, electrical signals must be converted into light and then back into an electrical signal. This process requires a performance and power overhead that must be taken into account for an accurate analysis. To explore the impact of using *UC-PHOTON* in CMPs, we modeled it by extensively modifying our in-house cycle accurate SystemC-based NoC simulator [57][60]. Five multiple use-case benchmark applications were selected and implemented on multiple cores in the simulator model. The goal was to explore applications with a wide spectrum of core count and use-case complexity. Table 1 shows the characteristics of these applications, with core complexity varying from 36 to 100 cores and use-case complexity varying from 5 to 30. SPL2xA and SPL2xB are multiple use-case applications derived from combining SPLASH-2 benchmarks (*Cholesky, FFT, Fmm, Lu, Ocean, Radix*). PNET1, PNET2, and PNET3 are multiple use-case networking applications used for packet processing, and forwarding [50]. We targeted a 32 nm process technology, and assume a 400 mm^2 CMP die area. A high level floorplanner [58] is used to determine core placement and link lengths.

The operating frequency of the photonic rings was estimated by calculating the time needed for the light to travel from any node to the farthest node, so that data can be transmitted to all nodes in one cycle. Through geometric calculations for the ring, using delay values from Table 2, and incorporating latching delays (using ITRS data [6]) we obtained a maximum operating frequency of greater than 3 GHz for the different sizes of CMPs we considered. Thus the photonic rings (and the communication network) were safely clocked at 2.3 GHz. Delay estimates for the various photonic interconnect-centric components were obtained from [14] and device fabrication results (e.g., [59]), and are shown in Table

2, for the 32 nm node. The delay of an optimally repeated and sized electrical (Cu) wire at 32 nm was assumed to be 42 ps/mm [13]. Delays for other electrical NoC components (routers, NI, buffers) were obtained from post-synthesis gate-level models after layout.

Table 1: Multi use-case application characteristics

Application	Cores	Use cases	Critical Trans.
SPL2x A	6x6	5	16
SPL2x B	6x6	10	32
PNET1	8x8	15	44
PNET2	8x8	20	54
PNET3	10x 10	30	98

Table 2: Delay of PHOTON components at 32nm

Component	32 nm
Modulator driver (ps)	9.5
Modulator (ps)	3.1
Waveguide (ps/mm)	15.4
Photo Detector (ps)	0.22
Receiver (ps)	4.0

The power dissipated in *UC-PHOTON* can be categorized into two components: electrical network power and photonic ring network power. The static and dynamic power dissipation of electrical routers and links in this work is based on results obtained from Orion 2.0 [61] incorporated into our simulator. For calculating power dissipation of the modulator driver and TIA power we used ITRS device projections [6] and standard circuit procedures. The energy consumption of each transmitter and receiver is 20 fJ/bit (dynamic) and 5 fJ/bit (static), as derived from [35]. A thermal tuning energy of 16 fJ/bit is considered for each heater element. In addition, an electrical laser power of 3.3 W (with 30% laser efficiency) is also considered in the power calculations. The laser power value accounts for the per component optical losses for the coupler/splitter (1.2 dB), non-linearity (1 dB at 30mW), waveguide (3 dB/cm), waveguide crossings (0.05 dB), ring modulator (1 dB), receiver filter (1.5 dB) and photodetector (0.1 dB).
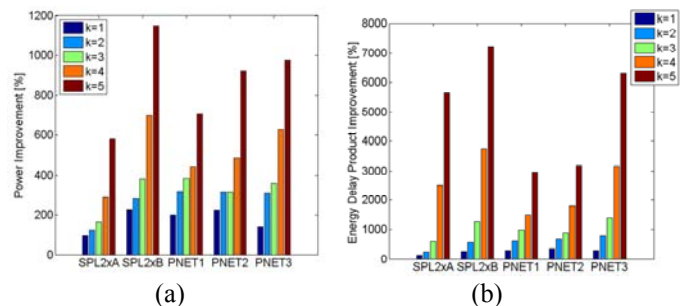

(a)   (b)

Figure 7: % improvement for non-reconfigurable *UC-PHOTON* vs. 2D electrical mesh NoC, (a) power, (b) energy-delay product

6.2. Results

Our first set of experiments compares the *UC-PHOTON* architecture with a traditional electrical 2D mesh NoC. Figure 7(a)-(b) show the improvement in average power and energy-delay product for our proposed *UC-PHOTON* architecture without enabling any runtime optimizations. The results are shown for the five different applications (from Table 1) implemented on the five different configurations of our proposed architecture ($k(n)$=1(4), 2(8), 3(12), 4(16), 5(16); b=256; r=optimal; w=16) and the 2D electrical mesh NoC. All implementations represented by bars in Figure 7 satisfy the bandwidth and latency constraints of all use-cases in the respective applications. It can be seen that even without enabling any runtime reconfiguration, *UC-PHOTON* achieves improvements from 2.1× to 12.6× in power saving and a reduction from 3.1× to 73.2× in energy-delay product over a traditional 2D electrical mesh NoC. These significant improvements are due to the offloading of multi-hop communication from the electrical mesh path to the more energy efficient and lower latency photonic path. Increasing the number of photonic rings (value of k) proportionally improves both the power and energy-delay characteristics of the on-chip communication infrastructure (at the cost of photonic layer complexity) due to more easily accessible photonic paths across the chip, and more efficient photonic path utilization. Thus hybrid photonic NoC fabrics like *UC-PHOTON* have a clear advantage over traditional 2D electrical NoC fabrics when it comes to reducing power and energy-delay product in future CMP designs.
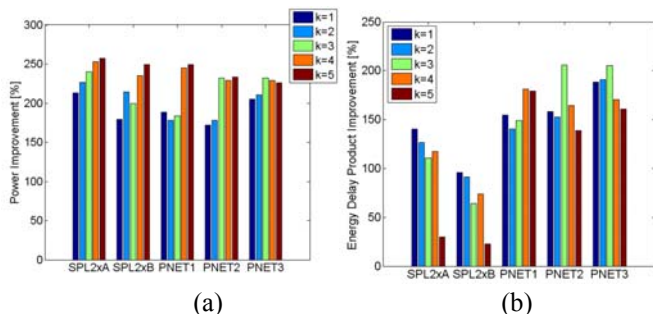

Figure 8: % improvement for runtime reconfigurable *UC-PHOTON* vs. non-reconfigurable *UC-PHOTON*, (a) power, (b) energy-delay product

In the next set of experiments, we were interested in exploring the impact of enabling runtime optimizations described in Section 4.4 in the *UC-PHOTON* architecture. Figure 8(a)-(b) show the improvements in average power and energy-delay product for the case when runtime reconfiguration optimizations are enabled compared to the case when runtime reconfiguration options are not utilized in *UC-PHOTON*. Results are shown for the same application implementation scenarios as in the previous set of experiments, with implementations represented by bars in Figure 8 satisfying the bandwidth and latency constraints of all use-cases in the respective applications. It can be seen that enabling runtime reconfiguration has a notable impact on reducing power dissipation (from 2.7× to 3.6×) as well as energy-delay product (from 1.2× to 3.1×). For smaller

benchmarks like SPL2xA and SPL2xB, an interesting phenomenon can be noticed – increasing the number of photonic rings can significantly reduce the improvements in energy-delay product. Even for larger and more complex benchmarks such as PNET1, PNET2, and PNET3, the improvements in energy-delay product saturate after a point and subsequently increasing the number photonic rings reduces the energy-product improvements. This happens because all runtime reconfiguration optimizations have a latency overhead (and also a small energy overhead) associated with their implementations. For more complex configurations (e.g., k = 4, 5) the latency overhead of runtime optimizations such as dynamic WDM and PRI reconfiguration can easily increase average packet latency, and consequently the energy-delay product. Overall however, it can be seen from the results that there is always a benefit in performing runtime reconfiguration, especially for power dissipation and energy-delay product.

Our final set of experiments compares the runtime reconfiguration strategies enabled in the *UC-PHOTON* architecture with the approaches proposed in the only two NoC-based works for coping with multiple use-case applications that have been previously proposed in literature. The first approach [51] proposes creating a synthetic worst case use-case and performing runtime reconfigurations that include DVS/DFS, adaptive TDMA slot allocation, and adaptive routing during use-cases transitions. The second approach [52], also proposes the same runtime reconfigurations, but does not create a worst-case use-case; instead optimizing on a per-use-case basis. In the interest of fairness for the comparison study, we adapt these runtime reconfiguration approaches and implement them on hybrid photonic NoC topologies similar to *UC-PHOTON*. This ensures that we are comparing the effectiveness of the runtime reconfiguration strategies.

Figure 9(a)-(f) compares our runtime reconfiguration enabled *UC-PHOTON* architecture with the runtime reconfiguration approaches from [51] and [52] implemented on hybrid photonic NoCs. The figures show % improvement for the reconfigurable *UC-PHOTON* architecture in average power dissipation (compared to (a) approach from [51], (b) approach from [52]), average throughput performance (compared to (c) approach from [51], (d) approach from [52]), and energy-delay product (compared to (e) approach from [51], (f) approach from [52]). Note that as in previous experiments, the implementations represented as bars in the figures satisfy bandwidth and latency constraints for all application use-cases. It can be seen that the reconfiguration strategies proposed with *UC-PHOTON* result in a reduction in power dissipation and energy-delay product, while also providing higher throughput performance, compared to approaches from both [51] and [52]. These results highlight the effectiveness of the runtime PRI reconfiguration, dynamic WDM, and clock gating optimizations that are enabled only in our proposed *UC-PHOTON* communication architecture, and not proposed in the approaches from [51] and [52].
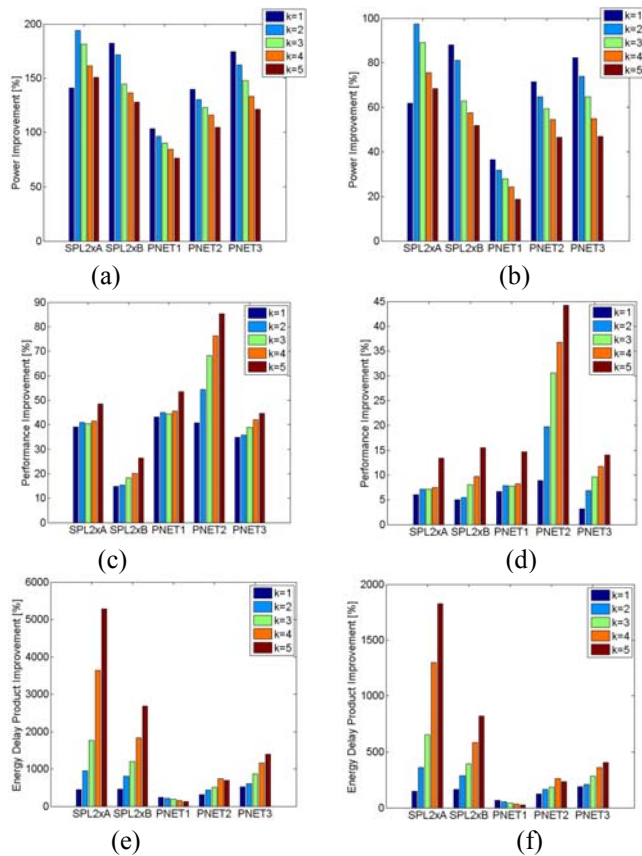
Figure 9: % improvement for runtime reconfigurable UC-PHOTON vs. approaches from [51] and [52], (a) avg. power w.r.t. [51], (b) avg. power w.r.t. [52], (c) performance w.r.t. [51], (d) performance w.r.t. [52], (e) energy-delay product w.r.t. [51], (f) energy-delay product w.r.t. [52]

Ultimately, the multi-ring *UC-PHOTON* configurations provide much lower average power and energy-delay products compared to single-ring configurations, but at the cost of increased complexity in the photonic layer and more overhead in the electrical layer due to the additional gateway interfaces. Our analysis of the area overhead of the different *UC-PHOTON* configurations in the electrical layer indicate that the absolute area overhead due to router enhancements in the electrical layer to implement dynamic reconfiguration schemes, deadlock recovery, and photonic interfaces increases with core count and as the number of photonic rings is increased, but is still minimal, at less than 2% chip area. In the future, as electrical NoC power dissipation is expected to be higher by a factor of around 10× than what is needed to enable tera- and petaflop performance levels of future CMPs [10][11], innovative on-chip communication paradigms are sorely needed. Reconfigurable hybrid photonic NoC architectures like *UC-PHOTON* proposed in this work can enable up to 46× reduction in power dissipation and up to 170× reduction in energy-delay product compared to traditional electrical NoC fabrics, while still satisfying bandwidth and latency constraints for all application use-cases. This is a very promising result that motivates the need for hybrid photonic NoCs to enable high performance-per-watt communication infrastructures in future CMP architectures that implement multiple use-case applications.

7. Conclusion

Emerging CMP applications today have tens to hundreds of cores and numerous use-cases with different communication bandwidth and latency requirements. Designing an on-chip communication fabric for these large systems that can satisfy the requirements of multiple use-cases at runtime is a challenging problem facing system designers today. While networks-on-chip (NoCs) are a promising solution that can provide scalable performance for large CMP designs, their performance-per-watt characteristics are not satisfactory for future CMP designs. In this paper we proposed a novel hybrid photonic NoC communication architecture called *UC-PHOTON* that can achieve scalable performance and performance-per-watt characteristics for small as well as large CMP designs. Our proposed hybrid photonic NoC architecture utilizes photonic ring paths interfaced with an electrical mesh NoC, and provides low-latency, high bandwidth, and power efficient data transfers. The novel hybrid architecture also supports various runtime reconfiguration optimizations to adapt to changing use-case performance requirements. Results of our experiments indicate that *UC-PHOTON* provides significantly better power, performance, and energy-delay characteristics compared to traditional 2D electrical NoCs designed to cope with multiple use-case applications.

8. References

[1]   D. Pham et al., "The design and implementation of a first-generation CELL processor", Proc. ISSCC, 2005.

[2]   A. Suga and K. Matsunami, "Introducing the FR 500 embedded microprocessor", IEEE MICRO, 20:21–27, 2000.

[3]   J Cornish, "Balanced energy optimization", Proc. ISLPED, 2004.

[4]   M. Ito et al., "An 8640 MIPS SoC with Independent Power-Off Control of 8 CPUs and 8 RAMs by An Automatic Parallelizing Compiler", Proc. ISSCC 2008.

[5]   S. Pasricha, N. Dutt. "On-Chip Communication Architectures", Morgan Kauffman, Apr 2008

[6]   International Technology Roadmap for Semiconductors (ITRS) 2007 Edition, http://public.itrs.net.

[7]   W.J. Dally and B. Towles, "Route Packets, Not Wires: On-Chip Interconnection Networks", Proc. DAC, pp. 684-689, June 2001.

[8]   M. Taylor et al., "The Raw Microprocessor", IEEE Micro, Apr 2002.

[9]   S. Vangal et al., "An 80-Tile 1.28 TFLOPS Network-on-Chip in 65 nm CMOS", Proc. IEEE ISSCC, Feb. 2007.

[10]   J.D. Owens, et al., "Research challenges for on-chip interconnection networks", Proc. IEEE Micro, 27(5), 2007.

[11]   R. G. Beausoleil, et al., "A Nanophotonic Interconnect for High-Performance Many-Core Computation", Proc. HOTI, 2008.

[12]   A. Covell, F, Whyte, "Digital Convergence: How the Merging of Computers, Communications and Multimedia is Transforming Our Lives", Aegis Publishing Group, 1999.

[13]   M. Haurylau, et al., "On-chip Optical Interconnect Roadmap: Challenges and Critical Directions", IEEE JQE 12(6), Nov 2006.

[14]   G. Chen, et al., "Predictions of CMOS Compatible On-Chip Optical Interconnect", Proc. SLIP, pp. 13-20, 2005.

[15] M.J. Kobrinsky, et al., "On-Chip Optical Interconnects", Proc. Intel Technology Journal, 8(2): 129-142, May 2004.

[16] S. Pasricha, N. Dutt, "Trends in Emerging On-Chip Interconnect Technologies", IPSJ Transactions on System LSI Design Methodology, Vol. 1, Sep 2008.

[17] L. Schares, et al., "Terabus: Terabit/Second-Class Card-Level Optical Interconnect Technologies", JQE, 12(5), Sep/Oct 2006.

[18] A. Biberman, et al., "Demonstration of All-Optical Multi-Wavelength Message Routing for Silicon Photonic Networks", Proc. OFCC, Mar 2008.

[19] F. Xia, et al., "Ultra-Compact High Order Ring Resonator Filters Using Submicron Silicon Photonic Wires for On-Chip Optical Interconnects", Proc. Optics Express, 15(19), Sep 2007.

[20] C.L. Schow, et al., "160-Gb/s, 16-Channel Full-Duplex, Single-Chip CMOS Optical Transceiver", Proc. OFCC, Mar 2007.

[21] B. R. Koch, et al., "Mode-locked silicon evanescent lasers", Proc. Opt. Express 15(18), Sep 2007.

[22] W. Bogaerts, et al., "Compact Wavelength-Selective Functions in Silicon-on-Insulator Photonic Wires", JQE, 12(6), Dec. 2006.

[23] Y. Vlasov et al., "Losses in single-mode silicon-on-insulator strip waveguides and bends", Proc. Opt. Express, 12(8), 2004.

[24] Q. Xu, et al., "12.5 Gbit/s Carrier-Injection-Based Silicon Microring Silicon Modulators", Proc. Opt. Express, 15(2), 2007.

[25] A. Gupta, et al., "High-Speed Optoelectronics Receivers in SiGe", Proc. VLSI Design, pp. 957-960, Jan 2004.

[26] G. Tosik, et al., "Power dissipation in optical and metallic clock distribution networks in new VLSI technologies", Proc. IEE EL, 40(3):198-200, Feb 2004.

[27] Ian O'Connor, "Optical solutions for system-level interconnect", Proc. SLIP, pp. 79-88, Feb 2004.

[28] W.M.J. Green, et al., "Ultra-Compact, Low RF Power, 10 Gb/s Silicon Mach-Zehnder Modulator", Opt. Express, 15(25), 2007

[29] S. Sahni et al. "Junction Field-effect-transistor based Germanium Photodetector on Silicon-on-Insulator", Proc. Optics Letters, May 2008.

[30] A. K. Okyay et al., "Silicon Germanium CMOS Optoelectronic Switching Device: Bringing Light to Latch", IEEE TED, Dec 2007.

[31] Y. Vlasov et al, "High-throughput Silicon Nanophotonic Wavelength-insensitive Switch for On-Chip Optical Networks" Proc. Nature Photonics, Apr 2008.

[32] A. Biberman et al., "First Demonstration of On-Chip Wavelength Multicasting", Proc. OFCC, March 2009

[33] M. Brie´re, et al., "System Level Assessment of an Optical NoC in an MPSoC Platform", Proc. DATE, Mar 2007.

[34] D. Vantrease et al., "Corona: System Implications of Emerging Nanophotonic Technology", Proc. ISCA, pp. 153-164, 2008.

[35] A. Joshi et al., "Silicon-Photonic Clos Networks for Global On-Chip Communication", Proc. NOCS 2009.

[36] H. Gu et al. "A Low-Power Fat Tree-based Optical Network-on-Chip for Multiprocessor System-on-Chip", Proc. NOCS 2009.

[37] A. Shacham, K. Bergman, and L.P. Carloni, "The Case for Low- Power Photonic Networks on Chip", Proc. DAC, 2007.

[38] S. Koohi, S. Hessabi, "Contention-Free on-Chip Routing of Optical Packets", Proc. NOCS 2009.

[39] S. Bahirat, S. Pasricha, "Exploring Hybrid Photonic Networks-on-Chip for Emerging Chip Multiprocessors", Proc. IEEE/ACM CODES+ISSS Oct 2009.

[40] A. Hansson et al., "A unified approach to constrained mapping and routing on network-on-chip architectures", Proc. ISSS 2005.

[41] J. Hu, R. Marculescu, "Exploiting the Routing Flexibility for Energy/Performance Aware Mapping of Regular NoC Architectures", Proc. DATE 2003.

[42] S. Murali, G. De Micheli, "Bandwidth-Constrained Mapping of Cores onto NoC Architectures", Proc. DATE 2004.

[43] A. Hansson, K. Goossens, "Trade-Offs in the Configuration of a Network on Chip for Multiple Use-Cases", Proc. NOCS 2007.

[44] I. Loi et al., "Synthesis of Low-Overhead Configurable Source Routing Tables for Network Interfaces", Proc. DATE 2009.

[45] M. A. A. Faruque, "Configurable Links for Runtime Adaptive On-chip Communication", Proc. DATE 2009.

[46] M. Li, et al., "DyXY - A Proximity Congestion-Aware Deadlock-Free Dynamic Routing Method for Network on Chip", Proc. DAC 2006.

[47] K. Lahiri, et al., "Design of high-performance system-on-chips using communication architecture tuners", IEEE TCAD, May 2004.

[48] T.D. Richardson, et al., "A hybrid SoC interconnect with dynamic TDMA-based transaction-less buses and on-chip networks", Proc. VLSI Design (VLSID), 2006. pp. 8-15.

[49] S. Pandey, T. Murgan, M. Glesner, "Energy Conscious Simultaneous Voltage Scaling and On-Chip Communication Bus Synthesis", Proc. VLSI-SoC 2006.

[50] S. Pasricha, N. Dutt, F. Kurdahi, "Dynamically Reconfigurable On-Chip Communication Architectures for Multi Use-Case Chip Multiprocessor Applications", Proc. ASPDAC, Jan 2009.

[51] S. Murali et al. "Mapping and configuration methods for multi-use-case networks on chips", Proc. ASPDAC, 2006.

[52] S.Murali et al. "A methodology for mapping multiple use-cases on to networks on chip", Proc. DATE, 2006.

[53] A. Hansson, M. Coenen, K. Goossens, "Undisrupted Quality-of-Service during Reconfiguration of Multiple Applications in Networks on Chip", Proc. DATE 2007

[54] S.C. Woo et al. "The SPLASH-2 programs: Characterization and methodological considerations", ISCAS, 1995.

[55] H. Al-Awwami, et al., "ZOMA: A Preemptive Deadlock Recovery Mechanism for Fully Adaptive Routing in Wormhole Networks", Proc. ICCNMC, pp. 519-525, 2001.

[56] J. Rabaey et al., "Digital Integrated Circuits", Prentice Hall, 2002.

[57] SystemC initiative. www.systemc.org.

[58] S. N. Adya, I. L. Markov, "Fixed-outline Floorplanning: Enabling Hierarchical Design", IEEE Trans TVLSI, Dec. 2003.

[59] I.-W. Hsieh, et al., "Ultrafast-Pulse Self-Phase Modulation and Third-Order Dispersion in Si Photonic Wire-Waveguides", Optics Exp., 14(25), Dec. 2006.

[60] S. Pasricha, N. Dutt, M. Ben-Romdhane, "Extending the Transaction Level Modeling Approach for Fast Communication Architecture Exploration", Proc. DAC 2004.

[61] A. Kahng, et al., "ORION 2.0: A Fast and Accurate NoC Power and Area Model for Early-Stage Design Space Exploration" Proc. DATE, 2009.