

Using Transmission Lines for Global On-Chip Communication

Aaron Carpenter, Jianyun Hu, Jie Xu, Michael Huang, Hui Wu, and Peng Liu

Abstract—The growing number of cores in chip multiprocessors increases the importance of interconnection for overall system performance and energy efficiency. Compared to traditional distributed shared-memory architectures, chip-multiprocessors offer a different set of design constraints and opportunities. As a result, a conventional packet-relay multiprocessor interconnect architecture is a valid, but not necessarily optimal, design point. Worsening wire delays, energy-inefficient routers, and the decreased importance of *in-field* scalability, make the conventional packet-switched network-on-chip a less attractive option.

An alternative solution uses well-engineered transmission lines as communication links. These transmission lines, along with simple, practical circuits using modern CMOS technology, can provide low latency, low energy, high throughput channels which can be used as a shared-medium point-to-point link. The design of the transmission lines and transceiver circuits has important architectural impact. This paper includes a first-step design effort for these components, particularly when used for a globally shared-medium bus. For medium-scale CMPs, this interconnect backbone can eliminate the need for packet switching and provide energy, as well as performance benefits when compared to a conventional mesh interconnect. We will provide a design of such a system from the ground up, including design of the transmission lines, transceiver circuits, and a simple, yet effective, architectural design for a shared-medium interconnect, and show that such a design can be a compelling alternative to packet-switched networks for CMPs.

I. INTRODUCTION

As the number of cores integrated into a single chip steadily increases, an important component in chip multiprocessors (CMPs) is the on-chip interconnect. For a number of reasons, packet-switched interconnect is often accepted as the de facto solution [26], [42]. A packet switched network offers numerous advantages such as throughput scalability and modularity. However, it is not without drawbacks. Routers are complex structures that occupy significant chip real-estate and consume

This work is supported in part by NSF under the grants 0901701, 0829915, and 0747324, and by the NSFC under grant 61028004. Copyright (c) 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending an email to pubs-permissions@ieee.org.

A. Carpenter is with the Electrical & Computer Engineering Dept. at Binghamton University, carpente@binghamton.edu

J. Hu, J. Xu, M. Huang, and H. Wu are with the Electrical & Computer Engineering Dept. at the University of Rochester, [jiansun.hu, jie.xu, michael.huang, hui.wu}@rochester.edu](mailto:{jianyun.hu, jie.xu, michael.huang, hui.wu}@rochester.edu)

P. Liu is with the Information Science & Electronics Engineering Dept. at Zhejiang University, liupeng@zju.edu.cn

significant power [45]. Repeated packet relaying adds latency to communication and can be an important performance issue, especially for simpler topologies with large network diameters such as ring or mesh. These disadvantages are upfront costs paid even when the applications do no need scalable throughput. As such, alternative architectures should be explored. Transmission line based interconnects are a promising candidate.

A transmission line (TL) allows high signaling rate, speed-of-light propagation velocity, and can potentially provide sufficient throughput for a range of CMPs, such that packet relaying can be avoided altogether. TL-based designs have been used in the context of microprocessors, but the specific design used is often studied and described in an ad-hoc fashion. A TL link has a large degree of freedom in designing the channel medium, the coding scheme, and the circuitry in the signaling chain and offers a vast range of trade-offs between costs and benefits. There is a lack of comprehensive design space studies to help architects navigate the design space and make optimal system-wide trade-offs. However, the design choices made at the circuit level have a significant impact on the characteristics of the architectural implementation, and vice versa. Figure 1 qualitatively illustrates a TL and circuit design spaces.

This paper presents an exploration of the design space of TL circuitry, and provides a simple, yet effective architectural design, using the TL links as a shared interconnect backbone. The rest of the paper is organized follows: Section II gives some background and related work. Section III discusses the transmission line and transceiver circuit design spaces. Section IV describes the architectural design in depth, and Section V evaluates the design. Section VI concludes.




Fig. 1: Illustration of transmission line link system design space.

II. BACKGROUND & RELATED WORK

Transmission lines are common components in RF and microwave circuits. The characteristics of the transmission lines, such as impedance, loss, propagation delay, dispersion and

crosstalk depend on the structure, size, materials, and fabrication. For the application of transmission lines as a global interconnect in CMOS circuits, electromagnetic (full-wave and quasi-TEM) analyses of on-chip transmission lines on silicon substrate (e.g., [25], [28], [40], [41], [43]) provided the groundwork. Circuit-level studies (e.g., [11], [23], [24], [31], [32], [44], [54]) have been carried out to characterize the performance of transmission line based on-chip interconnect. Novel signaling and modulation schemes have been proposed [18], [37].

System-level analyses often chose a point design for use as a special-purpose interconnect for caches [6], [7] or as express lanes in a mesh system [15]–[17]. Similarly, a particular circuit design point is chosen in these studies.

For architects to use the right design to obtain *system-level* optimal trade-offs, we need to go beyond isolated design point studies and better understand the trade-offs of different circuit designs and their implications for overall system performance and energy efficiency. Our paper is an attempt to bridge the RF design exploration with an simple architectural implementation, in the form of a shared medium bus.

With the integration of multiple cores on a single die, proposals of advanced interconnection have emerged. These proposals range from networks-on-chip (NoC) [5], [21], [30], [48], [52], [58] to optical interconnects [20], [29], [38], [39], [53], [56], [60] or RF interconnects [6], [15]–[17], [50]. Even with the use of circuit- or device-level support for optics or RF circuitry, many designs still rely on packet-switching at the architecture level [6], [15]–[17], [20], [50].

Recently, bus designs have started to gain more attention as a supplement or alternative to pure packet-switched networks. Conventional digital buses are being explored as part of the interconnect design [22], [55]. These designs still rely on packet-switching to connect multiple buses either explicitly through routers [22] or implicitly via hubs connecting multiple bus segments [55]. With only buses in the system, it is argued that the coherence substrate can switch to a snoopy protocol that helps reduce transaction hops and thus overall latency. Transmission lines are used with wide-band communication circuits to provide a bus design with low latencies and high overall throughput, which in turn allows the bus to be the only fabric in a CMP and a purely circuit-switched fabric [13]. Most related work does not present a design exploration of the TL and circuitry, but chooses from a small subset of point designs.

Finally, using transmission lines for communication is a well-established technique in mixed-signal and analog systems. There is no need to rely on future development for devices and technologies, as in on-chip optical interconnect. In addition to leveraging transmission lines, packet latency can be reduced via various optimizations in a packet-switched interconnect. New topologies, such as flattened butterfly, use higher radix routers to reduce network diameter and thus the average number of hops [36]. Wiring between routers can also be optimized with customized sizing to trade off among latency, throughput density, and energy [46], [47].

III. PHYSICAL AND CIRCUIT DESIGN

With ever improving transistor performance, a communication system can achieve a data rate of tens of Gb/s per line and an aggregate data rate of Tb/s over on-chip global transmission lines. In medium-sized CMPs, the global network connecting different cores can be entirely based on a multi-drop transmission line system (illustrated in Figure 2 and Figure 3) allowing packet-switching-free communication that is both energy-efficient and low-latency.

From the system’s perspective, a channel’s latency, throughput, and energy efficiency are metrics of interest. In a transmission-line channel, the signal propagation latency is largely determined by the length of the line, as the propagation velocity is simply the speed of light in the medium ($c/\sqrt{\mu_r * \epsilon_r}$), which is roughly 6ps/mm for CMOS technologies where $\epsilon_r = 3.0$ is assumed, and likely decreases over time as low-K dielectric materials improve. Modern CMP dies are relatively stable in dimensions (about 2cm on each side). A multi-drop transmission line loop meandering through a 16-tile CMP therefore measures about 75mm in length, as in Figure 2, and a corresponding worst-case propagation delay of about 440ps. If a closed loop is used, the worst-case distance and delay becomes 40mm and 235ps, respectively. Transceiver circuitry will also add some delay. Nevertheless, the overall transmission latency is only a few cycles even for multi-GHz cores. As such, channel throughput is the key speed metric and can impact the serialization and queuing delay of the packet latency. Channel throughput and energy per bit in turn depend on the transmission line physical properties, as well as the transceiver circuitry.




Fig. 2: Top-level view of 16-core interconnect layout. The solid line is a physical bi-directional ring, and the dotted line is a bi-directional terminated bus.

A. Transmission Line Topology

While there are many transmission line structures, a few of the most common ones for on-chip interconnect are microstrip lines (MSL), coplanar waveguides (CPW) and coplanar strips (CPS). The latter two have similar characteristics, and CPS lines lead to higher interconnect density than CPW. Hence microstrips and CPS lines are the focus in this work. Figure 4 shows a cross section of each and the main parameters in their physical design. Microstrip lines are often chosen for their simplicity and can be used with pure digital transmitters and receivers (inverters).




Fig. 3: General schematic for the transmission line link interconnect.

In contrast, coplanar strips, driven by slightly more complex differential signaling, provide extra robustness.




Fig. 4: Cross-section of (a) microstrip lines and (b) coplanar strips. The dotted lines in (a) represent inter-digitated MSL.

Attenuation & crosstalk: To understand the characteristics of the channel formed by these different transmission lines in isolation, we can idealize the active circuitry and estimate the maximum channel throughput (bit-rate) purely based on the characteristics of the lines. This is performed using a pair of industrial grade simulators. Sonnet [1] is used to obtain S-parameter profiles, given the transmission line material and dimensions; and Advanced Design System (ADS) is used to take the resulting attenuation and crosstalk characteristics into account and perform transient analyses to estimate achievable data rate. All simulations were done using noisy environments, including aggressor lines to simulate crosstalk between neighboring lines.

Given the same pitch size ($W+G$ in Fig. 4-(a)), varying the gap and spacing yields different attenuation and crosstalk. Sweeping through the space to identify the optimal metal strip width and necessary spacing in each configuration (MSL or CPS) helps put these sizing decisions into broader context. The results are plotted in Figure 5-(a) and Figure 5-(b).

Clearly, as the pitch size increases, crosstalk lowers for both configurations. However, crosstalk remains high for MSL in absolute terms. In contrast, CPS is subject to much less crosstalk, thanks to the differential signaling. Without the cost of running a pair of differential strips, MSL potentially provides good throughput at the low end of the pitch scale ($< 25\mu m$), but the throughput saturates very fast. This saturation is mainly due to crosstalk. For illustration, the maximum throughput of MSL without crosstalk (where the neighboring lines are not injected

with any signals as noise sources, in this case, labeled as I-MSL, or inter-digitated-MSL) is also plotted. As we can see, the difference is significant: without crosstalk, the maximum capacity increases from about 20Gb/s to about 60Gb/s.

One simple approach to reduce crosstalk is to use an inter-digitated organization of the strips, alternating signal lines and ground lines that provide some shielding.¹ Figure 5-(a) and 5-(b) suggest that I-MSL offers less protection against crosstalk and a somewhat lower throughput than CPS, due to the single-ended signaling. CPS is chosen for this work, in order to narrow the search.

Aggregate throughput: Intuitively, wider metal strips (which lower attenuation) and larger spacing (which lowers crosstalk) both help improve single-channel throughput, but not necessarily throughput density. Since practical transmission lines are already much wider than typical digital (RC) wires, optimal use of metal space is important.

In Figure 5-(c), the total pitch of all transmission lines is limited and the number of lines is varied to obtain the aggregate throughput of the system. Assuming a $2\text{cm} \times 2\text{cm}$ CMP divided into sixteen $5\text{mm} \times 5\text{mm}$ tiles, the total width can be limited to 2.5mm, or half of the tile’s width. Note that this is a rather arbitrary limit and not a fundamental constraint.

As we can see, the throughput peaks at about 60 lines (each with a pitch of $45\mu m$) for both configurations and CPS offers a maximum of 1.9 Tbps aggregate throughput. This is a substantial amount of raw throughput. It is entirely conceivable that a medium-scale CMPs relies only on transmission lines to provide a shared-medium global interconnect. It is worth noting that when the transceiver circuitry is taken into account, the actual throughput can change in either direction: slower transistors can limit throughput, and equalization circuitry can compensate for the channel bandwidth limitation. The optimal number of lines, as a result, can also fluctuate.

B. Transmission Circuits

Transmitter and receiver: The transmission circuitry design space is equally vast and unlikely to be explored exhaustively in a single paper. This work focuses on designs that are relatively simple and can be easily integrated with CMOS circuits. Note that transceiver circuit design is not orthogonal to the design of the physical line. For instance, differential signaling naturally pairs with coplanar strips.

Figure 3 shows the general schematic of a single transmission link (surrounded by neighboring links) with transmission circuits. In general, the transmission circuit can be as simple as inverter-chain based fully digital circuits and as it becomes more sophisticated, it allows faster data rates at generally reduced per bit energy costs.

¹Compared to the more generic Co-Planar Waveguide (CPW) in which the width of the shielding line and its distance to a signal line are free variables, the inter-digitated organization places a shielding line equal in width to the signal line equal-distance to the two neighboring lines.




Fig. 5: (a) Per-line bit rate and (b) crosstalk as a function of wire pitch. (c) Aggregate bit-rate as a function of the number of lines in a 2.5mm space

Propagation	Single Segment: 28.9 ps; Worst-case: 442.5 ps
Dimensions	45 lines, $45\mu\text{m}$ pitch; Length: 5mm per segment

TABLE I: Transmission line characteristics.

Digital: Probably the simplest design is a chain of (large) inverters (Figure 6-(a)) to drive the TL (microstrip/CPW) “strongly” so that the attenuated signal still arrives at the receiver discernible by the same style of inverter chain (albeit with smaller sizes to reduce the load on the TL). Even with this simple link design, transmission lines can achieve a transmission rate of 10Gb/s over a 75mm TL. Unfortunately, when the line is used as a multi-drop medium and when other circuit elements are included in the simulation, the signal degradation is so severe that the system no longer works regardless of transistor sizing. A simple remedy is to repeat the transmitter at each node. Such repeated TL becomes uni-directional and adds significant gate delays on top of propagation delay. Indeed, the gate delay, at 30ps (Tables I and II), is comparable to propagation delay for each segment of the TL, and thus doubles the total latency. Note that at about 5mm apart, the repeaters are inserted far more sparsely than in typical digital wires.

Mixed: The limitation of an all-digital link is that the signal at the receiver needs to maintain full swing. An analog receiver using current source amplifiers obviates the need of a full-swing signal and allows two benefits: First, the transmitter area and power can be decreased substantially. Second, the more forgiving receiver allows a faster bit rate.

Differential: Finally, the transmitter can adopt (analog) differential signaling over coplanar strips (Figure 6-(b)). It is worth noting that fully analog single-ended designs are also possible, but not fully explored in this work. A standard CMOS differential amplifier is used in this design. No special, hard-to-integrate RF devices, like inductors, are used. The receiver is a chain of differential amplifiers scaled using inverse scaling [51], allowing for high bandwidth and low power. The differential amplifiers are gated, and can be turned off when inactive, saving power/energy.

Differential signaling offers much better rejection of noise and permits faster data rate and lower power on the transmitter side. On the other hand, the receiver needs more amplification stages that result in more area and power. Nevertheless the overall per-bit energy is low (Table II).

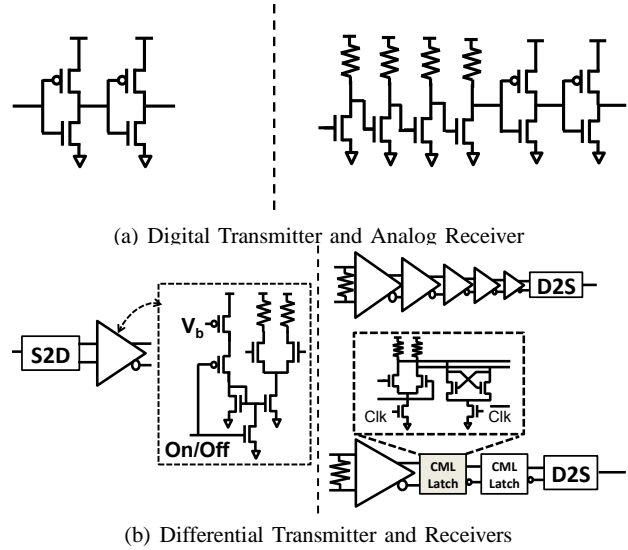


Fig. 6: (a) The digital transmitter design of digital inverters is also used for the transmitter and receiver in the fully-digital transceiver design. (b) The differential amplifier is used as a transmitter and as a component of the receiver design, forming a amplifier chain or a pre-amplification stage to drive a current-mode logic latch.

One alternative to the chain of amplifiers is current-mode logic (CML) latched sampler, similar to the one presented in [14]. As shown in Figure 6-(b), the latched sampler uses a cross-coupled latch immediately after a differential amplifier, which results in economy of circuit and still permits high data rate. Depending on the number of latches used, this circuit can subsume some of the deserialization functionality. In the extreme case, enough latches can be used to obviate any deserialization, greatly shortening the latency at some power cost. A latched sampler does require low-skew clocks, provided by circuit technologies such as injection locked clocking [61].

SerDes & PDR: Faster transistor speeds in modern and future generation CMOS technologies are an important contributor to the performance of a transmission line link bus (TLLB). On-chip TLL-based interconnect will operate at many times the core frequency, making serialization and deserialization (SerDes) necessary. Typically, multiple stages of 2:1 MUX/DEMUX are

Component	Bit-Rate (Gb/s)	Transmitter Side			Receiver Side			Total Energy/bit (pJ)
		Power (mW)	Latency (ps)	Area (μm^2)	Power (mW)	Latency (ps)	Area (μm^2)	
Digital	10	5	30	150	1.5	30	50	0.65-10.4
Mixed	17	20	30	250	8	35	60	1.65
Differential*	26.5	3.1	22	200	6.4	45	550	0.36
Latched Sampler	26.5	-	-	-	13	103	400	0.61
SERDES	-	1.6	750	220	1.15	650	165	0.1
PDR	-	-	-	-	0.4	150	60	0.02

TABLE II: Transceiver characteristics. Note that in the digital configuration, the transmitter latency is incurred every hop. The SERDES results are based on the fastest data rate (from analog transmission circuit). 32nm technology is used, as simulated using the predictive model in [3]. *This is the final design used for architectural analysis.

used as SerDes. These are designed using high-speed digital circuits but still introduce non-trivial delays as the simulations show (Table II). Often seen as a source of high power consumption for high-speed systems, we found that in our system SerDes does consume significant energy. Its small latency can also be hidden by pipelining in the steady state.

Phase and data recovery (PDR) is another necessary component to ensure the transmitters and receivers can properly communicate, and is independent of transceiver design. After a distance-dependent propagation delay, the transmitted pulses do not align with the receiver’s local clock. The magnitude of phase delta depends on the sender and can be quickly determined by sending and receiving a short test sequence in a initial, one-time calibration step. Data recovery circuits use the clock with the modified phase to ensure correct latching. Typically, clock recovery would also be necessary, but by using the injection-locked clocking scheme proposed in [61], we can exploit the globally synchronous clock, and rely only on phase recovery.

Isolation switch: Because of the large metal area required to route TLLs, it is necessary to share the lines among nodes. To prevent excessive loss and limit noise of inactive nodes, a switch is needed between the transceiver circuit and the transmission line tap.² When the switch is on, it must allow the signal to pass through with low loss and low distortion. When off, the switch must allow very little energy to be passed through in either direction. In 32nm technology, both of these goals can be accomplished reasonably well using a standard CMOS pass-gate structure. Additionally, the receivers and transmitters are power gated when not in use.

Final TLL design: Before exploring an architectural design, we summarize the final TL and circuit design, choosing the best design for the CMP environment. Coplanar strips are used as a final topology, as they utilize the space of the top metal layer more efficiently than the microstrips or coplanar waveguides; basic differential transmitters and receivers, scaled inversely, are also used without any equalization [51].

Our simulations show that a data rate of 26.4Gb/s can be achieved for a pair of transmission lines with a total pitch (including spacing) of $45\mu\text{m}$. Within 2.5mm of space, this pitch allows up to 55 pairs to be laid out (we use 45), totaling 1.45Tb/s of total throughput. All analysis for the full system

²Such a switch is also used in wireless systems to allow transmitter and receiver to time-share the antenna and is referred to as the T/R switch [33].

was done assuming a noisy environment.

A straightforward backbone interconnect based on transmission line links can be a good design option for general-purpose chip multiprocessors.

As we can see, transmission lines and associated circuits can be designed to provide low latency and high throughput, without the use of hard-to-integrate components (*i.e.*, inductors) or brute force throughput enhancement (*i.e.*, frequency division multiplexing, complex encoding). The architectural design can exploit these characteristics to improve the energy-efficiency of the interconnect backbone, removing the need for heavy-duty architectural solutions like packet-switched networks, particularly for medium-scaled CMPS.

IV. SHARED-MEDIUM ON-CHIP INTERCONNECT

Given these TLLs, the implementation of the links and how to allocate and use them for global communication falls to the architectural design space. This section will investigate the use of TLLs as a simple shared point-to-point link, and provide evidence that the traffic in a CMP can be relatively low, and thus the TLL bus design, which focuses on latency and energy efficiency, rather than scalability, can be a serious option for general-purpose chips.

A. Traffic Demand

Typical microprocessors rely on packet-switched network for the on-chip communication because of the inherent scalability of the system. However, a small- or medium-sized CMP has an upper limit on the traffic demand, and thus an understanding of on-chip traffic is necessary before implementing any interconnect backbone, especially for a shared medium bus, which provides significant, but not scalable, throughput.

Node structure: With chip-multiprocessors, there is flexibility to determine what on-chip communication uses the packetized interconnect. A baseline assumption often made in literature is that a chip consists of tiles, each with a core, an L1 cache, and a slice of a globally shared L2 (last-level) cache.

Sometimes a small number of cores and L2 slices are clustered into a node (concentrating interconnect demand). In such a system, the backbone network only makes a stop at every node. This organization of cores requires an intra-node fabric (*e.g.*, crossbar) that connects multiple L1 caches and the L2 cache banks in the node [10], [11], [13].

Minimizing horizontal traffic: To sustain high-speed processing, each core demands sufficient “vertical” throughput to fetch data from lower levels in the memory hierarchy all the way up to the core. Ideally, this vertical throughput is being provided by dedicated links between different levels of caches in the core’s node. However, depending on the address mapping, the data may be physically located on a cache in a remote node, incurring demand for “horizontal” throughput. Much research has been done to optimize the location of data to avoid unnecessary horizontal traffic. For instance, data can be mapped either statically or dynamically to the node where it is most often accessed or migrated there at run-time [4], [19], [34]. Such optimizations are important in their own right and will, as a side effect, significantly reduce the demand on the backbone, further strengthening the appeal of shared-medium, relay-free solutions.

In summary, communication in a chip-multiprocessor is carried out on a collection of fabrics; many architectural factors impact how much traffic depends on the backbone. Hence, sacrificing scalability of the backbone to achieve better energy efficiency and latency can be a viable alternative.

B. Bus Architecture

Figure 7 shows an overview of such an interconnect subsystem. Each node uses the proposed transceiver circuits to deliver packets over the shared transmission lines connecting all nodes. Note that unlike the conventional notion of a bus that often implies broadcast capability, our bus is merely a shared medium that allows point-to-point communication. Prior to the transfer of payload data on the bus, two setup operations are performed: arbitration and receiver wake-up.




Fig. 7: Overview of the bus-based communication subsystem.

Arbitration: The use of a shared-medium bus structure requires an arbitration mechanism. While any implementation of a permission granting system works, this design included a centralized system which can be thought of a centralized token ring. Because the ring is centralized, the “token” can quickly pass to the next requester. This arbiter is essentially a priority encoder for, say, 16 bits in a 16-node system. Larger, far more complex priority encoders are used in the timing critical store-forwarding circuit inside the core. We have measured a straightforward, unoptimized synthesis of a 16-node arbiter and compared it to the synthesized router used in a packet-switched interconnect [49]. The router’s overall delay is 4.3x that of the arbiter (1.65ns vs 0.38ns). The router is also much larger (10x), consumes far more power (20x), and is used more frequently (per flit-hop).

Receiver wake-up: For energy efficiency, the receivers operate in two modes. When the message is intended for a node,

its receiver transfers energy from the transmission line to the detector. On the other hand, when the message is intended for another node, the node is set to cause minimum loss for the through signal. For this reason, a setup step is performed immediately before payload data transmission to “wake up” the intended receiver, while other receivers remain in the off (and high isolation) mode. This setup is done in a pipelined fashion.

The request and grant signals are transferred over transmission lines similar to those used to build the bus. Such transfers take additional latency (modeled faithfully in this study) that will only be exposed when the bus is lightly loaded.

Turn-around time and bundling: After the transmission of the payload, the bus will be idle for a period of time to allow the signal to “drain” from the links. Even in the short distance of on-chip transmission lines, the wave’s propagation delay is not negligible. The amount of time needed to wait before another node can start to use the bus to transmit depends on the distance between the current transmitting node and the next scheduled to transmit. In most cases, a full cycle of turn-around time is enough. In the extreme case, a two-cycle turn-around delay is needed.

Note that in the special case of the same node transmitting another packet there is no need for such a turn-around period. Thus for better utilization of the bus throughput, this design uses a policy that allows *bundling*: sending multiple packets for each bus arbitration. When consecutive packets are sent from the same node, only the last packet will incur any turn-around time penalty. The impact of bundling is quantified in Section V-C.

Partitioning the bus: A simple way to get high throughput out of the bus structure is to use a wide bus that minimizes serialization latency. For example, a 32-byte cache line payload can be sent in one processor cycle over a bus with 32 data links operating at a data rate 8 times the computing clock speed. Clearly, a wide bus is wasteful for smaller payloads such as requests. In a shared-memory architecture, meta packets are common (about 60% in our suite of applications). Having another, smaller bus for meta packets is a clear option. In fact, with relatively small costs, it is possible to have multiple buses for meta packets. They can be used to increase throughput, or to support different types of requests such as in Alpha GS320 [27] (which prevents fetch deadlocks and eliminates the need to use NACK in their protocol). For simplicity, for this work, the bus consists of a single bus for meta packets and another one for data packets.

V. ANALYSIS OF SMALL TO MEDIUM CMPS

A. Experimental Setup

Architectural simulations of the proposed design were performed using an extensively modified version of SimpleScalar [9]. PopNet [2] is used to model the packet-switched network, while extra support was added to model the transmission line-based bus. The processor setup is described in Table III.

Table III also lists the benchmarks used to test the design space, including Splash-2 [59] and Parsec [8], using the respectively cited input sizes. Abbreviations are used in the data figures,

and the corresponding abbreviation is in parentheses in the table. Each benchmark is fast-forwarded according to the requirements of the binary of that benchmark. An offline profile is used to determine data page mapping, which is a common technique to reduce traffic by localizing data. The profile assigns a data page to the core which will access its contents most frequently [4], [19], [34].


Simulator Environment	
Circuit Simulators	32-nm Predictive Tech. Model [3] used for ADS circuit modeling Sonnet [1] used for TL modeling
Architectural Simulator	SimpleScalar [9] extensively modified for CMP Popnet [2] to model conventional mesh network
System Specifications	3.3GHz, 16-core, 8-fetch, 64-entry LSQ 128-entry ROB, 16KB private L1 cache per core 2MB shared L2 cache w/ 15 cycle latency 72-bit flit, 1-flit meta-packet, 4-flit data-packet Page-coloring [4], [19], [34] to reduce traffic
Benchmarks Used	
Splash-2 [59]	barnes (ba), cholesky (ch), fft (ff), fmm (fm) lu (lu), ocean (oc), radiosity (rs), radix (rx) raytrace (ry), water-spatial (ws)
Parsec [8]	blackscholes (bl), fluidanimate (fl)
Other	em3d (em), ilink (il), jacobi (ja) mp3d (mp), shallow (sh), tsp (ts)

TABLE III: Simulator environment & benchmarks used.


Traffic impact of page placement: A significant body of research exists to reduce unnecessary remote accesses by trying to map data close to the threads that frequently access the data. The solutions range from simple heuristics to map pages (*e.g.*, first-touch) to sophisticated algorithms that migrate data on the fly. Such optimizations not only improve performance on their own by reducing average latencies, but also serve to reduce horizontal traffic. This research uses a simple model as a proxy of a “middle-of-the-road” solution to localize data. Specifically, the last-level cache is shared and page interleaved. Off-line profiling assigns pages the color that matches the color of the node where the pages are accessed most frequently.

Figure 8-(a) shows that simple techniques can already cut down on unnecessary horizontal traffic. Without data mapping optimizations, using round-robin data distribution in an n -node system, each L1 miss has a 1 in n chance of being served locally. Hence, one would expect remote traffic to be roughly 94%, 88%, and 75% respectively for 16, 8, and 4 node systems. With even a simple profiling technique, the percentage of remote accesses drops to 53%, 46%, and 35%, respectively.

The performance impact of such data mapping on a canonical mesh interconnect is shown in Figure 8-(b). Note that the 16-node organization has 1 core linked to its own L2 slice. The 8-node organization clusters 2 cores into a single node. The result is a longer latency for using the intra-node fabric to access the cache slices local to the node, but a decrease in the number of remote accesses that use the backbone interconnect. The decrease in horizontal traffic and increased locality results in a speedup of more than 2x over a baseline with round-robin page allocation. Clearly, better data placement is an important optimization in its own right, and the sophistication and effect of the technique will only increase over time. The important side effect of traffic reduction alleviates a problem for the simpler shared-medium



(a) Remote Accesses



(b) Cluster Performance

Fig. 8: (a) Percentage of L2 accesses that are remote. The 3 configurations are 1, 2, and 4 cores per node. (b) Speedup due to profiling and clustering. The bar on the left is for 1 core per node, the right bar is for 2 cores per node. The baseline in this case is a 16-core mesh with round-robin data distribution.

relay-free interconnect, such as our design.

Performance comparison: While the TLL bus has a more limited aggregate throughput, it offers a better latency in general and in particular for packets between far apart nodes. Figure 9 compares the execution speed of this interconnect (with a bundling factor of 3) with a mesh. In this experiment, the chip-multiprocessor has 16 cores and is organized into 16 or 8 nodes. At this scale, the limit in throughput is seldom a problem for any application and, in general, more than compensated for by the superior latency. Even the more throughput demanding applications, such as *em3d*, *mp3d*, and *ocean*, perform comparably to mesh, especially in an 8-node configurations. On average, applications run faster on the TLL bus than on the mesh by 1.15x in the 16-node and 1.17x in the 8-node configurations, respectively.

An idealized interconnect system was also designed, and it was verified that the TLL bus performs close to this upperbound (more later). For instance, the 8-node system can achieve 91% performance of the ideal system.

As can be seen in Figure 8-(b), even though the intra-node fabric becomes slower as the node size increases, the benefit of having a smaller network in general outweighs the cost of slower intra-node accesses. In a mesh-based system, clustering helps improve performance by 4%. Just as with the case of better data placement, these optimizations reduce the demand on the backbone interconnect and has a slightly more significant benefit (6%) in the TLL bus system.

To summarize, even though bus architectures face throughput scalability challenges, in modest-scale chip-multiprocessors




Fig. 9: Speedup of TLL bus system over the respective (16- or 8-node) mesh-based system. The left bar in each group represent 16 node configuration and the right bar, 8 node.

and when natural steps are taken to improve performance, the disadvantages of TLL bus are much mitigated and the benefit becomes more pronounced.

B. Power Savings

One of the main disadvantages of canonical mesh networks is the high power and energy consumption [22], [35], [46], [55]. On average, the network power accounts for around 20% of the total system’s power. In contrast, the TLL bus uses no relay or energy-intensive routing. The power consumption of TLL bus is low in both absolute and relative terms. An entire link consumes 12.7mW while active (Table II shows power of individual components). Even when all lines are working all the time, the total power is around 600mW. When idling, the power consumption is even lower. Leakage in the communication circuit is estimate to be around $10\mu W$ per node [3], essentially negligible.


Comparing the energy consumed by the TLL bus to the power statistics from the network power model, Orion [57], there is a reduction in network energy of about 26x. With this reduction, the energy spent in the interconnect is less than 1% of the total energy consumption.

C. The Impact of Bundling

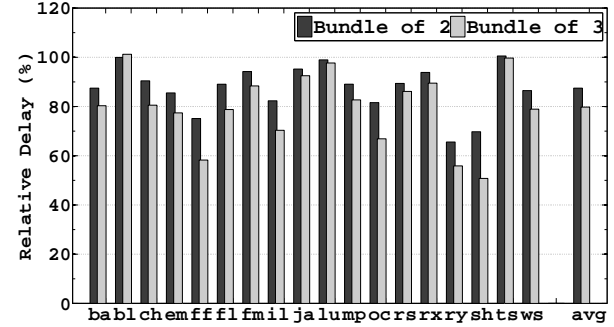
As discussed in Section IV-B, the turn-around time also wastes bus throughput and can be mitigated with bundling. So far, the design has used a bundling factor of 3, *i.e.*, each node can send up to 3 packets before yielding the bus. Figure 10 shows the impact of varying the bundling factor from 1 (no bundling) to 3. As we can see, the performance generally increases when the bundling factor increases. Without bundling, much throughput is wasted due to turn-around, so there is a noticeable performance increase with a bundling of 2. However, too much bundling can be detrimental to performance as well (*e.g.*, in the case of *tsp*). Figure 10-(b) shows the average overall packet latency for a bundling of 2 and 3 compared to no bundling. On average, bundling of 2 and 3 saves 13% and 20% respectively of the latency and improves performance by 2.0% and 3.4% respectively.

D. Scaling Up

While many-core chips will fill a certain market niche, a significant fraction of general-purpose chip-multiprocessors may



(a) Speedup



(b) Overall latency

Fig. 10: (a) Speedup of the 16-node system with bundling of 2 and 3, over the system without bundling. (b) Overall packet latency relative to a non-bundled system. The left and right bar correspond to a bundling of 2 and 3 respectively.

have only a relatively modest number of cores. The proposed design works well in such an environment. As the number of cores increases beyond a threshold, the viability of our current design will decrease. A limited scalability test is conducted with a 64-core system organized into 2- or 4-core nodes (32 nodes, 2 cores each; and 16 nodes, 4 cores each), using the exact same bus design as before. Figure 11 summarizes the performance result compared to the (scaled-up) mesh-based design with the same clustering.

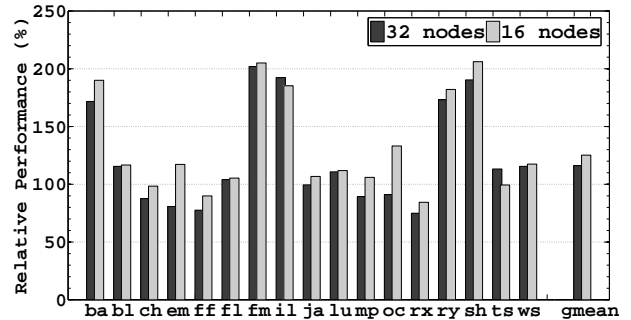


Fig. 11: Relative performance of a 64-core system. For the TLL bus configurations, a bundle of 3 is used.

As the system grows in size, the probability of the bus becoming a bottleneck increases. In a few cases (*e.g.*, *fft* and *radix*), the performance of the TLL bus is significantly worse than the conventional mesh interconnect (Figure 11). On the other hand, when the throughput is not a bottleneck resource, the

latency advantage over mesh becomes even more pronounced. As a result, the performance gap between the bus-based and mesh-based systems widens for many applications (*e.g.*, *fmm* and *shallow*). On average, the TLL bus performs 16% and 25% better than mesh for a 32- and 16-node system, respectively. Clearly, simply having better aggregate throughput scalability is not enough. A packet-switched interconnect (including hierarchical bus) segments wires to allow simultaneous traffic, improving overall throughput at the expense of latency. The result can also be a serious performance issue for chip-multiprocessors.

In other words, a bus architecture should not be written off as a possible solution for on-chip interconnect. After all, no design is truly scalable in all respects. The sacrifice in latency in some packet-switched interconnects can be an even more serious performance problem, not to mention the significantly higher energy cost. Additionally, there are potential optimization opportunities for transmission line link buses, including circuit-switched segments, coherence optimizations, and extracting better utilization out of the TLLB architecture, all which make the interconnect more scalable. There are more details on the scalability of the TLLB system in [12].

To better understand the limitation of bus-based system, the TLL bus is also compared it to an idealized interconnect system using conventional digital wires. In this system, no throughput limitation or contention is modeled for the interconnect. A packet's delay is calculated as 0.03mm/ps based on the latency-optimized wires in [46].




Fig. 12: Performance of TLL bus relative to idealized contention-free, low-latency interconnect.

Figure 12 shows the performance of the TLL bus in 32-node and 16-node configurations (both have 64 cores) normalized to that of the ideal interconnect. As we can see, while 7 out of 18 benchmarks perform within 10% of the idealized case, the limited throughput shows significant limitation in a number of applications where performance can be improved several folds. Nevertheless, the bus system achieves 67% and 72% of the idealized performance, for 32- and 16-nodes respectively, showing a somewhat graceful degradation beyond its intended usage range. Recall, in a 16-core, 8-node system, the bus can achieve 91% of the ideal's performance.

VI. CONCLUSIONS

Packet-switched interconnect, using simplistic digital wires, is often accepted by many as the default solution for on-chip

communication for future chips. While the superior scalability certainly carries significant advantages, there are, nonetheless, non-trivial issues such as the area cost of the router, the latency impact, and power overhead of repeated packet relays. While continued research will undoubtedly mitigate some of the issues, we should also investigate alternative solutions.

In this paper, we make a case for a different type of design. Our design space exploration lends insight to the co-design of the circuit-level and system-level design decisions. The simulation-based study shows that (1) advances in technology allows very high data rates and low energy even with only simple transceiver circuits; (2) a much higher data rate and better energy efficiency can be achieved with some analog circuits and differential signaling; (3) the superior latency and energy characteristics of the links translate to potential improvement at the system level; and (4) with this underlying capability, a truly packet-switching-free interconnect is both easy to build and quite competent to support the traffic demand for modestly sized chip-multiprocessors. Experimental analyses have shown that in a medium-scale 16-core system, this design achieves 91% of that in an *idealized* wire-based interconnect. The performance degrades rather gracefully, still achieving 72% performance of the ideal configuration in a 64-core system. Compared with a canonical mesh interconnect, the transmission line link bus provides advantages in latency, resulting in better average performance (1.17x in a 16-core system and 1.25x in a 64-core system).

Another important benefit of avoiding packet switching and relaying is the inherent energy efficiency of the communication system. The energy reduction in the backbone network is more than an order of magnitude compared to a mesh. This energy advantage of the TLL bus is important in itself and also provides capital for future optimizations that compensate for the throughput limitation.

REFERENCES


- [1] <http://www.sonnetsoftware.com/>.
- [2] PoPNet. <http://www.princeton.edu/~peh/orion.html>.
- [3] Predictive Technology Modeling. <http://ptm.asu.edu/>.
- [4] M. Awashiti, K. Sudan, R. Balasubramonian, and J. Carter. Dynamic Hardware-Assisted Software-Controlled Page Placement to Manage Capacity Allocation and Sharing within Large Caches. In *Proc. Int'l Symp. on High-Perf. Comp. Arch.*, pages 250–261, February 2009.
- [5] J. Balfour and W. J. Dally. Design Tradeoffs for Tiled CMP On-Chip Networks. In *Proc. Int'l Conf. on Supercomputing*, pages 187–198, June 2006.
- [6] B. Beckmann and D. Wood. TLC: Transmission Line Caches. In *Proc. Int'l Symp. on Microarch.*, pages 43–54, December 2003.
- [7] B. Beckmann and D. Wood. Managing Wire Delay in Large Chip-Multiprocessor Caches. In *Proc. Int'l Symp. on Microarch.*, pages 319–330, November 2004.
- [8] C. Bienia, S. Kumar, J. Singh, and K. Li. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In *Proc. Int'l Conf. on Parallel Arch. and Compilation Techniques*, September 2008.
- [9] D. Burger and T. Austin. The SimpleScalar Tool Set, Version 2.0. Technical report 1342, Computer Sciences Department, University of Wisconsin-Madison, June 1997.
- [10] A. Carpenter. *The Design and Use of High Speed Transmission Line Links for Global On-Chip Communication*. PhD thesis, Dept. of Electrical & Computer Engineering, University of Rochester, March 2012.

- [11] A. Carpenter, J. Hu, M. Huang, H. Wu, and P. Liu. A design space exploration of transmission-line links for on-chip interconnect. In *Proc. Int'l Symp. on Low-Power Electronics and Design*, June 2011.
- [12] A. Carpenter, J. Hu, O. Kocabas, M. Huang, and H. Wu. Enhancing effective throughput for transmission-line based bus. In *Proc. Int'l Symp. on Comp. Arch.*, June 2012.
- [13] A. Carpenter, J. Hu, J. Xu, M. Huang, and H. Wu. A case for globally shared-medium on-chip interconnect. In *Proc. Int'l Symp. on Comp. Arch.*, June 2011.
- [14] T. Chalvatzis, K. Yau, R. Aroca, P. Schvan, M. Yang, and S. Voignescu. Low-Voltage Topologies for 40-Gb/s Circuits in Nanoscale CMOS. *IEEE Journal of Solid-State Circuits*, 42(7):1564–1573, July 2007.
- [15] M. Chang, J. Cong, A. Kaplan, C. Liu, M. Naik, J. Premkumar, G. Reinman, E. Socher, and S. Tam. Power Reduction of CMP Communication Networks via RF-Interconnects. In *Proc. Int'l Symp. on Microarch.*, pages 376–387, November 2008.
- [16] M. Chang, J. Cong, A. Kaplan, M. Naik, G. Reinman, E. Socher, and R. Tam. CMP Network-on-Chip Overlaid With Multi-Band RF-Interconnect. In *Proc. Int'l Symp. on High-Perf. Comp. Arch.*, pages 191–202, February 2008.
- [17] M. Chang, E. Socher, S. Tam, J. Cong, and G. Reinman. RF Interconnects for Communications On-chip. In *Proc. Int'l Symp. on Physical Design*, pages 78–83, April 2008.
- [18] R. Chang, N. Talwalkar, C. Yue, and S. Wong. Near Speed-of-Light Signaling Over On-Chip Electrical Interconnects. *IEEE Journal of Solid-State Circuits*, 38(5):834–838, May 2003.
- [19] S. Cho and L. Jin. Managing Distributed, Shared L2 Caches through OS-Level Page Allocation. In *Proc. Int'l Symp. on Microarch.*, pages 455–468, December 2006.
- [20] M. Cianchetti, J. Kerekes, and D. Albonesi. Phastlane: a rapid transit optical routing network. In *Proc. Int'l Symp. on Comp. Arch.*, pages 441–450, 2009.
- [21] W. Dally and B. Towles. Route Packets, Not Wires: On-Chip Interconnection Networks. In *Proc. Design Automation Conf.*, pages 684–689, June 2001.
- [22] R. Das, S. Eachempati, A. Mishra, V. Narayanan, and C. Das. Design and Evaluation of a Hierarchical On-Chip Interconnect for Next-Generation CMPs. In *Proc. Int'l Symp. on High-Perf. Comp. Arch.*, February 2009.
- [23] A. Deutsch. Electrical characteristics of interconnections for high-performance systems. *Proceedings of the IEEE*, 86(2):315–357, February 1998.
- [24] A. Deutsch, P. Coteus, G. Kopcsay, H. Smith, C. Surovic, B. Krauter, D. Edelstein, and P. Restle. On-chip wiring design challenges for gigahertz operation. *Proceedings of the IEEE*, 89(4):529–555, April 2001.
- [25] A. Deutsch, G. Kopcsay, V. Ranieri, K. Cataldo, E. Galligan, W. Graham, R. McGouey, S. Nunes, J. Paraszczak, J. Ritsko, R. Serino, D. Shih, and J. Wilczynski. High-Speed Signal Propagation on Lossy Transmission Lines. *IBM Journal of Research and Development*, 34(4):601–615, July 1990.
- [26] S. Furber and J. Bainbridge. Future trends in soc interconnect. In *IEEE International Symposium on System-on-Chip*, pages 183–186, November 2005.
- [27] K. Gharachorloo, M. Sharma, S. Steely, and S. Van Doren. Architecture and design of AlphaServer GS320. In *Proc. Int'l Conf. on Arch. Support for Prog. Lang. and Operating Systems*, pages 13–24, November 2000.
- [28] H. Hasegawa, M. Furukawa, and H. Yanai. Properties of Microstrip Line on Si-SiO₂ System. *IEEE Transactions on Microwave Theory and Techniques*, 19(11):869–881, Nov. 1971.
- [29] G. Hendry, J. Chan, S. Kamil, L. Olifer, J. Shalf, L. Carloni, and K. Bergman. Silicon Nanophotonic Network-On-Chip Using TDM Arbitration. In *Hot Interconnect*, pages 88–95, August 2010.
- [30] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar. A 5-GHz Mesh Interconnect for a Teraflops Processor. *IEEE Micro*, 27(5):51–61, 2007.
- [31] H. Ito, J. Inoue, S. Gomi, H. Sugita, K. Okada, and K. Masu. On-chip Transmission Line for Long Global Interconnects. In *IEEE International Electron Devices Meeting. IEDM Technical Digest*, pages 677–680, December 2004.
- [32] H. Ito, M. Kimura, K. Miyashita, T. Ishii, K. Okada, and K. Masu. A Bidirectional- and Multi-Drop-Transmission-Line Interconnect for Multipoint-to-Multipoint On-Chip Communications. *IEEE Journal of Solid-State Circuits*, 43(4):1020–1029, April 2008.
- [33] Y. Jin and C. Nguyen. Ultra-Compact High-Linearity High-Power Fully Integrated DC-20-GHz 0.18-um CMOS T/R Switch. *IEEE Transactions on Microwave Theory and Techniques*, 55(1):30–36, Jan. 2007.
- [34] R. Kessler and M. Hill. Page Placement Algorithms for Large Real-Indexed Caches. *ACM Transactions on Computer Systems*, 10(4):338–359, 1992.
- [35] J. Kim. Low-Cost Router Microarchitecture for On-Chip Networks. In *Proc. Int'l Symp. on Microarch.*, pages 255–266, December 2009.
- [36] J. Kim, C. Nicopoulos, D. Park, R. Das, Y. Xie, V. Narayanan, M. S. Yousif, and C. R. Das. A Novel Dimensionally-decomposed Router for On-chip Communication in 3D Architectures. In *Proc. Int'l Symp. on Comp. Arch.*, pages 138–149, June 2007.
- [37] J. Kim, I. Verbauwhede, and M. Chang. Design of an Interconnect Architecture and Signaling Technology for Parallelism in Communication. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 15(8):881–894, August 2007.
- [38] N. Kirman, M. Kirman, R. Dokania, J. Martinez, A. Apsel, M. Watkins, and D. Albonesi. Leveraging Optical Technology in Future Bus-based Chip Multiprocessors. In *Proc. Int'l Symp. on Microarch.*, pages 492–503, December 2006.
- [39] N. Kirman and J. Martinez. A Power-Efficient All-Optical On-Chip Interconnect Using Wavelength-Based Oblivious Routing. In *Proc. Int'l Conf. on Arch. Support for Prog. Lang. and Operating Systems*, pages 15–28, March 2010.
- [40] T. Kitazawa and T. Itoh. Propagation characteristics of coplanar-type transmission lines with lossy media. *Microwave Theory and Techniques, IEEE Transactions on*, 39(10):1694–1700, October 1991.
- [41] Y. Kwon, V. Hietala, and K. Champlin. Quasi-TEM Analysis of "Slow-Wave" Mode Propagation on Coplanar Microstructure MIS Transmission Lines. *IEEE Transactions on Microwave Theory and Techniques*, 35(6):545–551, Jun. 1987.
- [42] R. Marculescu and P. Bogdan. The chip is the network: Toward a science of network-on-chip design. *Foundations and Trends in Electronic Design Automation*, 2(4):371–461, 2009.
- [43] V. Milanovic, M. Ozgur, D. DeGroot, J. Jargon, M. Gaitan, and M. Zaghoul. Characterization of broad-band transmission for coplanar waveguides on cmos silicon substrates. *Microwave Theory and Techniques, IEEE Transactions on*, 46(5):632–640, May 1998.
- [44] K. Miyashita, T. Ishii, H. Ito, N. Ishihara, and K. Masu. An Over-12-Gbps On-Chip Transmission Line Interconnect with a Pre-Emphasis Technique in 90nm CMOS. In *Electrical Performance of Electronic Packaging, 2008 IEEE-EPEP*, pages 303–306, October 2008.
- [45] S. Mukherjee, P. Bannon, S. Lang, A. Spink, and D Webb. The Alpha 21364 Network Architecture. *IEEE Micro*, 22(1):26–35, January/February 2002.
- [46] N. Muralimanohar and R. Balasubramonian. Interconnect Design Considerations for Large NUCA Caches. In *Proc. Int'l Symp. on Comp. Arch.*, pages 369–380, June 2007.
- [47] N. Jouppi N. Muralimanohar, R. Balasubramonian. Optimizing NUCA Organizations and Wiring Alternatives for Large Caches With CACTI 6.0. In *Proc. Int'l Symp. on Microarch.*, pages 3–14, December 2007.
- [48] B. Nayfeh, K. Olukotun, and J. Singh. The Impact of Shared-Cache Clustering in Small-Scale Shared-Memory Multiprocessors. In *Proc. Int'l Symp. on High-Perf. Comp. Arch.*, pages 74–84, February 1996.
- [49] L. Peh and W. Dally. A Delay Model and Speculative Architecture for Pipelined Routers. In *Proc. Int'l Symp. on*

- High-Perf. Comp. Arch.*, pages 255–266, 2001.
- [50] A. Roy and M. Chowdhury. RS/Wireless Interconnects in Future On-Chip and Board-Level Clock Distribution Network. In *Proc. Int'l Conf. Electro/Information Technology*, pages 542–545, May 2007.
 - [51] E. Sackinger and W. Fischer. A 3-GHz 32-dB CMOS Limiting Amplifier for SONET OC-48 Receivers. *IEEE Journal of Solid-State Circuits*, 35(12):1884–1888, December 2000.
 - [52] D. Sanchez, G. Michelgeannakis, and C. Kozyrakis. An Analysis of On-Chip Interconnection Networks for Large-Scale Chip Multiprocessors. *ACM Transactions on Architecture and Code Optimization*, 7(1), 2010.
 - [53] A. Shacham, K. Bergman, and L. Carloni. On the Design of a Photonic Network-on-Chip. In *First Proc. Int'l Symp. on Networks-on-Chip*, pages 53–64, May 2007.
 - [54] C. Svensson. Electrical interconnects revitalized. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 10(6):777 – 788, December 2002.
 - [55] A. Udipi, N. Muralimanohar, and R. Balasubramonian. Towards Scalable, Energy-Efficient, Bus-Based On-chip Networks. In *Proc. Int'l Symp. on High-Perf. Comp. Arch.*, pages 1–12, January 2010.
 - [56] D. Vantrease et al. Corona: System Implications of Emerging Nanophotonic Technology. In *Proc. Int'l Symp. on Comp. Arch.*, June 2008.
 - [57] H. Wang, X. Zhu, L. S. Peh, and S. Malik. Orion: A Power-Performance Simulator for Interconnection Networks. In *Proc. Int'l Symp. on Microarch.*, pages 294–305, November 2002.
 - [58] D. Wentzlaff et al. On-Chip Interconnection Architecture of the Tile Processor. *IEEE Micro*, 27(5):15–31, 2007.
 - [59] S. Woo, M. Ohara, E. Torrie, J. Singh, and A. Gupta. The SPLASH-2 Programs: Characterization and Methodological Considerations. In *Proc. Int'l Symp. on Comp. Arch.*, pages 24–36, June 1995.
 - [60] J. Xue, A. Garg, B. Ciftcioglu, J. Hu, S. Wang, I. Savidis, M. Jain, R. Berman, P. Liu, M. Huang, H. Wu, E. Friedman, G. Wicks, and D. Moore. An Intra-Chip Free-Space Optical Interconnect. In *Proc. Int'l Symp. on Comp. Arch.*, pages 94–105, June 2010.
 - [61] L. Zhang, A. Carpenter, B. Ciftcioglu, A. Garg, M. Huang, and H. Wu. Injection-Locked Clocking: A Low-Power Clock Distribution Scheme for High-Performance Microprocessors. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 2008.




Aaron Carpenter received the B.S degree in 2005, M.S. in 2006, and Ph.D. in 2012, each from the Electrical and Computer Engineering department at the University of Rochester. He is currently an Assistant Professor in the Electrical and Computer Engineering department at Binghamton University, USA. His research interests include many aspects of computer architecture, including chip multiprocessors, on-chip interconnect design and energy-efficient systems.




Jianyun Hu (S'05) received B.Sc. degree in Electrical Engineering and M.Sc. degree in Microelectronics from Fudan University, China, in 2003 and 2006, respectively. He is currently working towards the Ph.D. degree in the Department of Electrical and Computer Engineering at the University of Rochester.

He was with Qualcomm Inc. in the summer and fall of 2011, where he worked on circuit design and simulation for cellular RF

IC. His research interests include wideband high-speed RF/analog IC design.




Jie Xu Jie Xu received his BS in Electrical Engineering from University of Science and Technology of China, Hefei, China, in 2006. He received his MS in Electrical Engineering from University of Rochester in 2011, and is currently pursuing PhD. In 2006-2009, he was with Chinese Academy of Sciences, researching on DSP and RF systems. His research interests include RF and analog circuits and systems.



Michael Huang received the BS degree in computer science and engineering from Tsinghua University, Beijing, in 1994, the MS and the PhD degree in computer science from University of Illinois at Urbana-Champaign in 1999 and 2002, respectively. From 1994 to 1997, he was a lead architect in building a 32-processor hierarchical shared-memory multiprocessor research prototype. He joined the faculty of the Electrical and Computer Engineering department in 2002. In 2010, he was on sabbatical at IBM T. J. Watson Research Center working on future POWER processor concept development.

His research interests include various aspects of high-performance computer architecture such as processor microarchitecture, communication and memory substrate, reliability, and energy-efficient and complexity-effective design. He is particularly interested in addressing emerging issues and exploring new capabilities in the underlying device, circuit, and manufacturing technology. He is a recipient of the NSF CAREER award and a member of the IEEE and the ACM.



Hui Wu received the B.Sc. degree in electrical engineering and M.Sc. degree in microelectronics from Tsinghua University in 1996 and 1998, and the Ph.D. degree in electrical engineering from California Institute of Technology in 2003, respectively. He was a co-op researcher at IBM T. J. Watson Research Center in 2001. In 2002-2003, he was with Axiom Microdevices.

In 2003, Dr. Wu joined the faculty of the University of Rochester, where he is an

Associate Professor of Electrical and Computer Engineering. His current research interests are in inter-and intra-chip optical/electrical interconnects, silicon photonics, electronic-phonic integrated circuits (EPIC), wideband RF and high-speed integrated circuits, high performance clocking, and nanoelectronics using emerging technologies.



Peng Liu received his BS in Optical Engineering and MS in Optical Engineering from Zhejiang University, Hangzhou, China, in 1992 and 1996, respectively, and PhD in Communication and Electrical Engineering from Zhejiang University in 1999. In 1999, he joined the Faculty of the Information Science and Electronic Engineering department at Zhejiang University, where he was promoted to Associate Professor in 2002. He spent the 2009-2010 academic year as

a Visiting Scholar at University of Rochester working on high-performance computer architectures. His research interests include embedded processor microarchitecture, multiprocessor system-on-chip architectures, on-chip interconnection networks, parallel computer architectures, and VLSI design.