A Design Space Exploration of Transmission-Line Links for On-Chip Interconnect

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

Department of Electrical and Computer Engineering

University of Rochester, Rochester, New York 14627

Email: {aaron.carpenter, jianyun.hu, michael.huang, hui.wu}@rochester.edu, liupeng@zju.edu.cn

Abstract—With increasing core count, chip multiprocessors (CMP) require a high-performance interconnect fabric that is energy-efficient. Well-engineered transmission line-based communication systems offer an attractive solution, especially for CMPs with a moderate number of cores. While transmission lines have been used in a wide variety of purposes, there lack comprehensive studies to guide architects to navigate the circuit and physical design space to make proper architecture-level analyses and tradeoffs. This paper makes a first-step effort in exploring part of the design space. Using detailed simulation-based analysis, we show that a shared-medium fabric based on transmission line can offer better performance and a much better energy profile than a conventional mesh interconnect.

Keywords: Transmission Line, On-chip Interconnect, Design Space Study

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 apparently accepted as the de facto solution [19], [28]. A packet switched network offers numerous advantages such as bandwidth scalability and modularity. However, it is not without drawbacks. Routers are complex structures that occupy significant chip real-estate and consume significant power [31]. 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 not need scalable bandwidth. 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 already been used in numerous ways including 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 tradeoffs 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 tradeoffs.

In this paper, we take a first-step effort exploring the design space of TL circuitry. As illustrated in Figure 1, this design space can be roughly broken down into three regions based on the transceiver circuitry. While we strive to evaluate optimal designs from each category, it is worth noting that future work will almost certainly push the envelope of all designs.

Fig. 1. Illustration of TL system design space.

The rest of the paper is organized as follows. Section II briefly describes background and related work. Section III discusses the design space study of the transmission line links. To understand the overall system-level impact, in Section IV we discuss a simple architecture using these links. Section V concludes.

II. BACKGROUND AND RELATED WORK

Transmission lines are common components in RF and microwave circuits. Their characteristics such as impedance, loss, propagation delay, dispersion and crosstalk depend on the structure, size, materials, and fabrication. For the application of transmission lines as global interconnect in CMOS circuits, electromagnetic (full-wave and quasi-TEM) analyses of on-chip transmission lines on silicon substrate (e.g., [18], [20], [26], [27], [29]) provided the groundwork. Circuit-level studies (e.g., [16], [17], [21], [34]) have been carried out to characterize the performance of transmission line based on-chip interconnect. Novel signaling and modulation schemes have been proposed [14], [25].

Much work has been done investigating particular implementations of transmission line based on-chip interconnect and the supporting transceiver circuitry [22], [30]. These studies each present a particular design choice, as well as the circuit level ramifications of these choices. System-level designs have explored transmission lines as a special-purpose interconnect for caches [5], [6] or as express lanes in a mesh system [11]–[13]. Similarly, a particular circuit design point is chosen in these studies.

For architects to use the right design to obtain *system-level* optimal tradeoffs, we need to go beyond isolated design point studies and better understand the tradeoffs of different circuit designs and their implications for overall system performance and energy efficiency. Our paper is a first-step effort towards this goal.

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 5) allowing packet-switching-free communication that is both energy-efficient and low-latency. In this paper, we focus on circuit- and system-level analyses in such a context. Clearly, transmission lines can be used in other ways in the on-chip interconnect.

From the system’s perspective, a channel’s latency, throughput, and energy efficiency are of primary 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.

A. Transmission Line Topology

While there are more transmission line structures, 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. Hence we focus

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.

on microstrips and CPS lines in this work. Figure 3 shows a cross section of each and the main parameters in their physical design. Microstrip lines are often chosen for their simplicity and are typically used with pure digital transmitters and receivers (inverters). In contrast, coplanar strips, paired with differential signaling provide extra robustness.

Fig. 3. Cross-section of (a) microstrip lines (dotted line for interdigitated-MSL) and (b) coplanar strips

Attenuation & crosstalk: To understand the characteristics of the channel formed by these different transmission lines in isolation, we 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, a first-principle EM simulator [1] is used to obtain s-parameter profiles given the transmission line dimensions; and Advanced Design System (ADS), from Agilent Technologies 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. 3), varying the gap and spacing yields different attenuation and crosstalk. We sweep through the space to identify the optimal metal strip width and necessary spacing in each configuration (MSL or CPS). The results are plotted in Figure 4-a and Figure 4-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, we also plot the maximum throughput of MSL without crosstalk

Fig. 4. (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

(where the neighboring lines are not injected with any signals as noise sources). As we can see, the difference is significant: with crosstalk, the maximum capacity drops from about 60Gb/s to about 20Gb/s.

One simple approach to reduce crosstalk is to use an interdigitated organization of the strips, alternating signal lines and ground lines that provide some shielding.¹ Figure 4-a and 4-b suggest that I-MSL offers less protection against crosstalk and a somewhat lower throughput than CPS. To narrow the search, we will focus on CPS.

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 4-c, we limit the total pitch of all transmission lines and vary the number of lines to obtain the aggregate throughput of the system. Assuming a 2cm $\times$ 2cm CMP divided into sixteen 5mm $\times$ 5mm tiles, we limit the total width 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 bandwidth peaks at about 60 lines for both configurations and CPS offers a maximum of 1.9Tbps aggregate throughput. This is a substantial amount of raw bandwidth. 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 transmission 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 one paper. We focus on designs that are relatively simple and can be easily integrated with digital

¹Compared to the more generic notion of 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.

CMOS circuits. Note that transmission circuit design is not orthogonal to the design of the physical line. For instance, differential signaling naturally pairs with coplanar strips.

Figure 5 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.

Fig. 5. General schematic for the transmission line link interconnect.

Digital: Probably the simplest design is to use a chain of (large) inverters (Figure 6-a) to drive the TL (microstrip) “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). With a limited searching of inverter sizing, we found that even with this simple link design we can achieve a transmission rate of 10Gb/s over a 40mm 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 (Table I), 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

Propagation	Single Segment: 28.9 ps; Round-trip: 461.9 ps
Line Dimensions	56 lines, 45 μ m pitch; Length: 5mm per segment

Component	Bit-Rate (Gb/s)	Transmitter Side			Receiver Side			Total Energy/bit (pJ)
		Power (mW)	Latency (ps)	Area (μ m ²)	Power (mW)	Latency (ps)	Area (μ m 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 I

TRANSMISSION LINE LINK 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).

Fig. 6. (a) The digital transmitter design of digital inverters is also used for the transmitter and receiver in the fully-digital transceiver design, with larger size requirements.

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). A standard CMOS differential amplifier is used in our design. No special RF devices, like inductors, are used for better integration. The receiver is a chain of differential amplifiers scaled using inverse scaling [33], 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 results in more area and power. Nevertheless the overall per-bit energy is low (Table I).

One alternative to the chain of amplifiers is current-mode logic (CML) latched sampler, similar to the one presented in [10]. As shown in Figure 6-b, the latched sampler uses a cross-coupled latch immediately after a differential amplifier, which resulted 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 [36].

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 used as SerDes. These are designed using high-speed digital circuits but still introduce non-trivial delays as our simulations show (Table I).

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 clock. The magnitude of phase delta depends on the sender and can be quickly determined by sending and receiving a short test sequence in an initial calibration step. Data recovery circuit use the clock with the modified phase to ensure correct latching.

T/R switch: Because of the large metal area required to route TLLs, we need 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.

IV. SYSTEM-LEVEL IMPACT OF TRANSMISSION LINE LINKS

To understand the ultimate impact at the system level, we use the multi-drop transmission line links discussed above to build a bus-like global interconnect for a CMP. Compared to a conventional packet-switched interconnect, such a transmission-line link bus (TLLB) does not have packet relay or routing. But unlike a conventional bus (with implied broadcast capabilities), different nodes on a TLLB merely share the same transmission line medium for point-to-point communication. Such a bus needs a few architectural

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 [23].

elements to function. Note that the architecture issue is outside the scope of this paper and the design is kept simple and not optimized. More discussion of the architectural design of TLLB can be found in [9].

A. Architecture Design

Partitioning the bus: We support two lengths of packets: data and meta (or control) packets. Data packets are larger (288 bits vs 72 bits). For better use the system bandwidth, separate meta- and data-packet buses are used (9 links for meta bus and 36 links for data bus).

Arbitration and receiver wake-up: A special TLL is used to connect nodes to the center of the chip where the arbiter lies. A requesting node sends over the destination ID over this special link to the arbiter. Upon granting the bus access, the arbiter sends back a grant token to the requester and a wake-up signal to the intended destination to power on receiver circuit and to prepare the PDR circuit. Both these feedbacks are also sent over TLLs.

To better hide certain latencies, the grant token can be sent a few cycles ahead of the actual cycle of availability. We do this by sending a token that encodes the number of cycles to wait before actual transfer. Upon receiving the token, the sender can count down and start preparative actions such as serialization to overlap with the waiting.

Signal draining: When the link forms a ring, we need to ensure that that a signal does not traverse the loop and overlap with a subsequent packet causing interference. This is achieved by having nodes that are outside the shorter path between the transmitter and the receiver turn their receiver to *drain* mode during transmission. In drain mode the T/R switch is turned on to siphon energy from the transmission line. The amplifiers, however, are turned off, since the information is useless. At each node, impedance tuning is done to minimize reflection. As a result, when the T/R switch is on, the node absorbs 50% of the energy propagated thus far. In other words, after 6 draining nodes, the signal energy in the transmission line is reduced to about 1% of the original signal, no longer a significant source noise. Furthermore, an extra cycle is padded to every transmission to allow the previous packet to drain out before starting the next packet.

B. Experimental Analysis

For brevity, the system-level experimental setup is summarized in Table II. For this experiment, we use the differential transmission circuit with latched sampler.

Performance analysis: Figure 7 compares the performance of a system using a TLLB with a system using a mesh interconnect. All results are normalized to that of an ideal interconnect, in which we do not model any routing delay, contention, or queuing delays. We model only the wire delay over the manhattan distance between the sender and receiver node (30ps/mm [32]). The proposed TLLB outperforms the mesh network for most applications. For 7 out of 18 benchmarks, the TLLB performs within 5% of ideal mesh

Simulator Environment	
Circuit Simulators	32-nm Predictive Tech. Model [3] used for ADS circuit modeling Sonnet [1] used for TL modeling
Architectural Simulator	SimpleScalar [8] 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], [15], [24] to reduce traffic
Benchmarks Used	
Splash-2 [35]	barnes (ba), cholesky (ch), fft (ff), fmm (fm) lu (lu), ocean (oc), radiosity (rs), radix (rx) raytrace (ry), water-spatial (ws)
Parsec [7]	blackscholes (bl), fluidanimate (fl)
Other	em3d (em), ilink (il), jacobi (ja) mp3d (mp), shallow (sh), tsp (ts)

TABLE II
SIMULATOR ENVIRONMENT & BENCHMARKS USED.

interconnect. Overall, TLLB achieves a relative performance of 80% (geometric mean) of ideal mesh, as compared to 76% for the conventional mesh network.

Fig. 7. System performance comparison.

The reason for TLLB's performance is its latency. In a medium-scale CMP like the one simulated here, the overall throughput demand seldom overwhelms the shared bus. As such, the superior latency characteristics from the links directly translate into lower end-to-end latency.

Energy savings: Other than providing high-speed communication, TLLB achieve dramatic energy reduction compared to a packet-switched interconnect where routers, buffers, and repeaters all contribute to significant energy overhead of the interconnect. Table I includes the power/energy of various components of the proposed TLLB. Including the overheads of arbitration and so forth, the per-bit energy is less than 1pJ, and more than 26x lower than that in a mesh. As such, the global interconnect energy is essentially negligible with TLLB, whereas in a mesh-based system, it accounts for about 20% of the entire chip energy consumption. Coupled with an overall speedup of 1.05, a TLLB-based CMP improves on system-wide energy-delay product by 1.38x over a mesh-based CMP.

V. CONCLUSIONS

Packet-switched on-chip interconnects often provides scalable bandwidth but at the expense of latency and can consume significant energy. A properly designed communication link

based on transmission line can achieve ultra-low latency and power consumption and can thus be a serious candidate for on chip interconnect, especially for moderately-sized CMPs.

In this paper, we have navigated part of the design space for transmission line links. Our simulation-based study shows that (1) advances in technology allows very high data rates and low energy even with only simple transmission circuitry; (2) a much higher data rate and better energy efficiency can be achieved with some analog circuitry and differential signaling; and (3) the superior latency and energy characteristics of links translate to significant system-level improvements: In a 16-core CMP, a rudimentary TLLB-based version outperforms a mesh-based system by 5% while saving more than 26x in interconnect energy. This translates into an improvement of 1.38x in system-wide energy-delay product.

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. Balasubramanian, 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] B. Beckmann and D. Wood. TLC: Transmission Line Caches. In *Proc. Int'l Symp. on Microarch.*, pages 43–54, December 2003.
- [6] 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.
- [7] 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.
- [8] D. Burger and T. Austin. The SimpleScalar Tool Set, Version 2.0. Technical report 1342, Computer Sciences Department, University of Wisconsin-Madison, June 1997.
- [9] A. Carpenter, J. Hu, J. Xu, M. Huang, and H. Wu. A Case for Globally Shared-Medium On-Chip Interconnect. In *Proceedings of the International Symposium on Computer Architecture*, June 2011.
- [10] T. Chalvatzis, K. Yau, R. Aroca, P. Schvan, M. Yang, and S. Voinigescu. Low-Voltage Topologies for 40-Gb/s Circuits in Nanoscale CMOS. *IEEE Journal of Solid-State Circuits*, 42(7):1564–1573, July 2007.
- [11] 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.
- [12] 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.
- [13] 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, 2008.
- [14] 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.
- [15] 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.
- [16] A. Deutsch. Electrical characteristics of interconnections for high-performance systems. *Proceedings of the IEEE*, 86(2):315–357, February 1998.
- [17] A. Deutsch, P.W. Coteus, G.V. Kopcsay, H.H. Smith, C.W. Surovic, B.L. Krauter, D.C. Edelstein, and P.L. Restle. On-chip wiring design challenges for gigahertz operation. *Proceedings of the IEEE*, 89(4):529–555, April 2001.
- [18] 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.
- [19] S. Furber and J. Bainbridge. Future trends in soc interconnect. In *IEEE International Symposium on System-on-Chip*, pages 183–186, November 2005.
- [20] H. Hasegawa, M. Furukawa, and H. Yanai. Properties of Microstrip Line on Si-SiO₂ System. 19(11):869–881, Nov. 1971.
- [21] 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.
- [22] 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.
- [23] Y. Jin and C. Nguyen. Ultra-Compact High-Linearity High-Power Fully Integrated DC-20-GHz 0.18- μ m CMOS T/R Switch. *IEEE Transactions on Microwave Theory and Techniques*, 55(1):30–36, Jan. 2007.
- [24] R. Kessler and M. Hill. Page Placement Algorithms for Large Real-Indexed Caches. *ACM Transactions on Computer Systems*, 10(4):338–359, 1992.
- [25] 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.
- [26] 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.
- [27] Y. Kwon, V. Hietala, and K. Champlin. Quasi-TEM Analysis of “Slow-Wave” Mode Propagation on Coplanar Microstructure MIS Transmission Lines. 35(6):545–551, Jun. 1987.
- [28] 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.
- [29] V. Milanovic, M. Ozgur, D.C. DeGroot, J.A. Jargon, M. Gaitan, and M.E. Zaghloul. 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.
- [30] 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.
- [31] 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.
- [32] N. Muralimanohar and R. Balasubramanian. Interconnect Design Considerations for Large NUCA Caches. In *Proc. Int'l Symp. on Comp. Arch.*, pages 369–380, June 2007.
- [33] 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.
- [34] C. Svensson. Electrical interconnects revitalized. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 10(6):777–788, December 2002.
- [35] 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.
- [36] 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.