

Designing Energy-Efficient Low-Diameter On-chip Networks with Equalized Interconnects

Ajay Joshi, Byungsub Kim and Vladimir Stojanović

Department of EECS, Massachusetts Institute of Technology, Cambridge, MA

Abstract

In a power and area constrained multicore system, the on-chip communication network needs to be carefully designed to maximize the system performance and programmer productivity while minimizing energy and area. In this paper, we explore the design of energy-efficient low-diameter networks (flattened butterfly and Clos) using equalized on-chip interconnects. These low-diameter networks are attractive as they can potentially provide uniformly high throughput and low latency across various traffic patterns, but require efficient global communication channels. In our case study, for a 64-tile system, the use of equalization for the wire channels in low-diameter networks provides $2\times$ reduction in power with no loss in system performance compared to repeater-inserted wire channels. The use of virtual channels in routers further reduces the power of the network by 25–50% and wire area by $2\times$.

1. Introduction

To keep scaling the performance, multicore processors rely on data and thread-level parallelism, improved energy-efficiency of smaller/simpler cores, and existence of easy-to-program, energy-efficient interconnect networks. While an order of magnitude in core energy-efficiency improvement is needed in the next decade to enable hundreds of cores on a power-limited die, to keep the system balanced, interconnect networks have to undergo a similar transformation. Low-diameter high-radix networks like crossbar [1], Clos [2], etc have uniform latency and throughput, which makes them attractive from the programming perspective. However, these topologies require global on-chip communication channels, which are expensive in terms of latency, power and area, when implemented with traditional repeater insertion. As a result, these schemes do not scale well to dozens or hundreds of cores under a global power constraint. On the other hand, high-diameter low-radix networks like mesh [3] use shorter channels and consume low power and area, but are difficult to program and their throughput and latency scale poorly with the size of the network.

In this paper we investigate the benefits of using equalized interconnects as global channels to build low-diameter on-chip networks with uniformly low la-

tency and high throughput. The energy-efficiency and bandwidth-density advantages of equalized interconnects over repeater insertion have been recently demonstrated in a few modeling and design studies [4–7]. We utilize the link efficiency models from these studies (especially [6]) to evaluate and quantify the benefits of these interconnects at the network level using a vertical design framework. In section 2, we describe our target system and a set of baseline on-chip networks. Section 3 gives a comparative overview of implementation costs of channels designed using pipelined repeater insertion versus equalization. A power–performance comparison of various network topologies designed using the two interconnect design techniques is presented in section 4.

2. Network topologies

We consider a 64-tile system designed using predictive 32 nm technology, operating at 5 GHz, with area of 400 mm². The proposed analysis is independent of the exact nature of the tile. Each tile could be heterogeneous and consist of one or more logic cores and a part of the on-chip cache, or the tiles could be homogeneous and have either only logic cores or only memory blocks. We adopt a message passing protocol for communication between the various tiles. We consider systems sized for ideal throughput per tile (TT) of 64 b/cyc and 128 b/cyc under uniform random traffic pattern. To cover the spectrum from high-diameter, low-radix to low-diameter, high-radix we will analyze *mesh*, concentrated mesh (*cmesh*), flattened butterfly (*flatFly*) [8] and *Clos* networks, shown in Figure 1. Table 1 shows some of the key design parameters for these network topologies sized for TT = 128 b/cyc.

Figure 1(a) shows a 2D 8x8 *mesh* network. This network uses distributed flow control and is commonly used in today’s multicore designs [3], [9] since it is relatively easy to implement. However, the *mesh* network has high average hop count resulting in high latency. For every hop, energy is consumed in the routers and wires. As core count increases, this energy could become significant, especially for global traffic patterns. In addition, programmer productivity is affected due to difficulty in uniform mapping of both local and global traffic patterns to a *mesh* network.

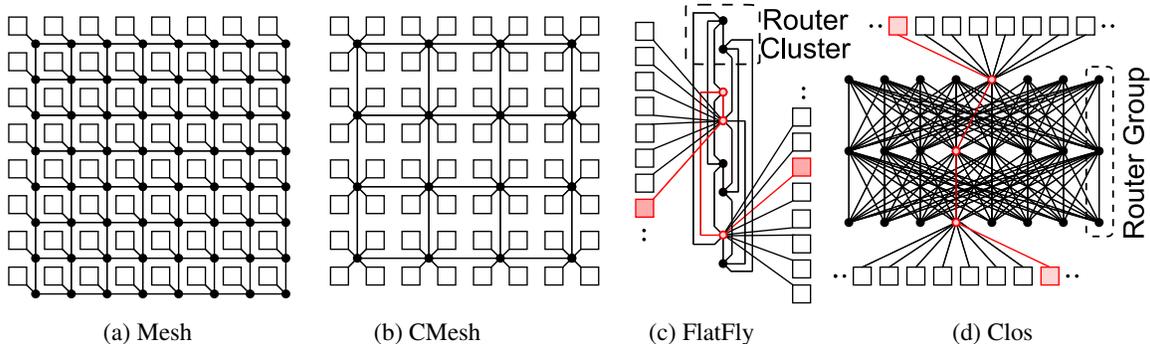


Figure 1: Logical View of 64-Tile Network Topologies – (a) 2D 8x8 *mesh*, (b) concentrated mesh (*cmesh*) with 4x concentration, (c) 8-ary 2-fly flattened butterfly (*flatFly*) (d) 8-ary 3-stage *Clos* network with eight middle routers. In all four figures: squares = tiles, dots = routers. In (a), (b) and (c) inter-dot lines = two opposite direction channels. In (d) inter-dot lines = uni-directional channels.

Topology	Channels				Routers		Latency					
	N_C	b_C	N_{BC}	$N_{BC} \cdot b_C$	N_R	radix	H	T_R	T_C	T_{TC}	T_S	T_0
<i>mesh</i>	224	256	16	4,096	64	5x5	2-15	2	1	0	2	7-46
<i>cmesh</i>	48	512	8	4,096	16	8x8	1-7	2	2	0	1	3-25
<i>flatFly</i>	32	512	8	4,096	8	12x12	1-2	2	2-6	0-1	1	3-16
<i>Clos</i>	128	128	64	8,192	24	8x8	3	2	3-5	0-1	4	16-20

Table 1: Example Network Configurations – Networks sized to support 128 b/cyc per tile under uniform random traffic. N_C = number of channels, b_C = bits/channel, N_{BC} = number of bisection channels, N_R = number of routers, H = number of routers along data paths, T_R = router latency, T_C = channel latency, T_{TC} = latency from tile to first router, T_S = serialization latency, T_0 = zero load latency.

The hop latency could be reduced by moving to higher dimensional networks. This would improve the network performance, but these high-dimensional networks need high-radix routers and long channels when mapped to a 2D layout. These high-radix routers and long channels consume higher energy and area. Another way to lower hop count without too much overhead is by using concentration [10]. Figure 1(b) shows an example of a *mesh* with a concentration of 4. Here, the lower router and channel count balances with the higher router radices and channel lengths, in turn providing lower latency than *mesh* at comparable area and energy. However, like the *mesh* network, it is difficult to uniformly map the various applications to a *cmesh* network.

To provide low latency and high throughput, low-diameter networks like *flatFly* (Figure 1(c)) and *Clos* (Figure 1(d)) could be used. However, both networks require high-radix routers and long channels, which can increase the energy and area cost. The *flatFly* in Figure 1(c) is a variation of a 8-ary 2-fly butterfly and has 8 tiles connected to each 12×12 high-radix router, and when using dimension-ordered routing requires a maximum of two hops (intra-cluster and inter-cluster) for communication. Figure 2(a) shows a possible layout of a *flatFly* network. The *flatFly* network, however, lacks path diversity which

results in network congestion, making the uniform mapping of various applications harder for the programmer.

Compared to *flatFly*, the *Clos* network minimizes the congestion by providing larger path diversity while maintaining a hop count of 2. Figure 1(d) shows a 8-ary 3-stage *Clos* network with 8 possible paths between each pair of tiles. Figures 2(b) and (c) show two possible mappings of the *Clos* network to a 2D layout. In the layout in Figure 2(b) [11], the three 8×8 routers in a group are clustered together. On the other hand, for the layout in Figure 2(c) the intermediate set of routers are placed at the center of the chip. The first *Clos* layout requires channels of 5 different lengths, while the second *Clos* layout requires channels of 2 different lengths. In this paper we use the layout in Figure 2(c) since it consumes 15% lower power. In the *Clos* network, the data always travels through the intermediate set of routers, which results in uniform latency and throughput across various traffic patterns.

3. Interconnect technology

The on-chip communication networks are constrained by metal area and power dissipation. While designing the on-chip network channels, we want to minimize the channel power for a given area (width) of the channel, while

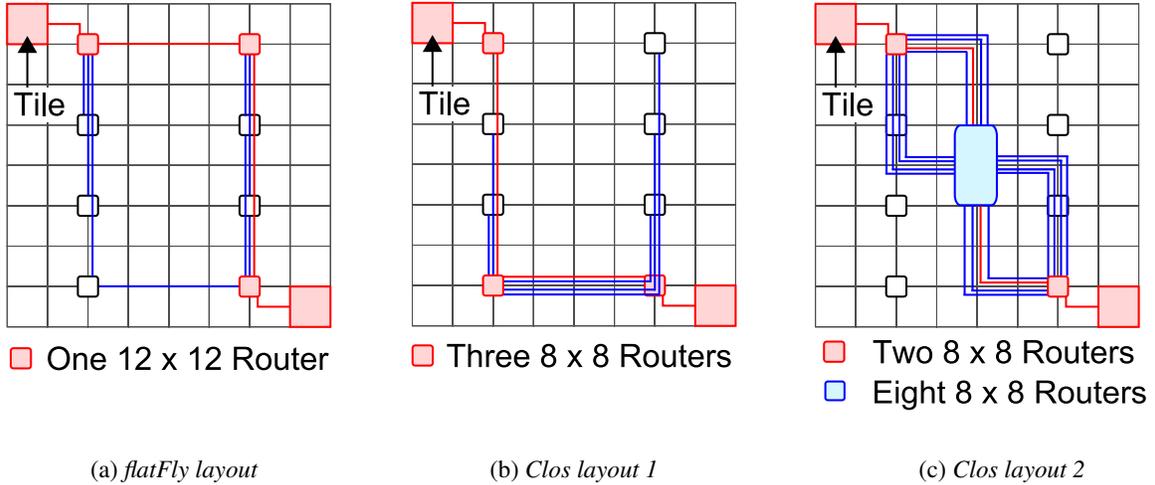


Figure 2: Physical layouts of *flatFly* and *Clos* networks. For *flatFly* layout and *Clos* layout 2, all bi-directional links (blue) associated with the highlighted routers are shown. For *Clos* layout 1, only the links associated with the intermediate router are shown. The data path is shown by the red links in all three layouts.

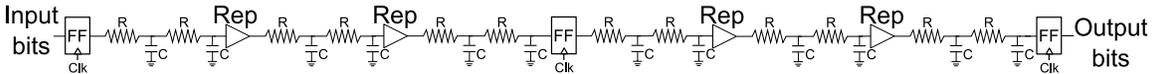


Figure 3: Repeater-inserted pipelined RC interconnect model. FF = Flip-flop.

satisfying the target latency and throughput constraints. In this section we provide a comparative overview of implementation costs of channels that are designed using pipelined repeater insertion (Figure 3) and using equalization (Figure 4), projected to a 32 nm technology generation. We assume that channels are designed using either semi-global (M6) or global routing layers (M9). The design parameters for these global/semi-global wires are obtained primarily from energy and bandwidth-density driven sizing optimization for both repeated and equalized interconnects [6] for target throughput of 5 Gbps. Table 2 shows the energy per bit (data-dependent and fixed) and the latency for 5 Gbps M6 and M9 links of various lengths. Here the fixed energy component includes bias/leakage energy and energy consumed in the clocked circuits in transmitter/receiver and interconnect flip-flops.

Repeater insertion is one of the most commonly used design techniques for on-chip interconnects. Here, re-

peated interconnects are pipelined to enable a more modular design and minimize latency uncertainty due to process variability and noise. For the various interconnect lengths listed in Table 2 that are designed as repeater-inserted pipelined interconnects, each pipeline segment is 2.5 mm long. The energy cost for the basic 2.5 mm pipeline segment is calculated using the data from [6] such that the latency is 200 ps and bandwidth is 5 Gbps. The energy cost for longer pipelined wire lengths is then calculated using the energy cost for the 2.5 mm wire. We include the cost associated with the flip-flops at the sender side, receiver side and intermediate pipeline stages.

Equalization can be used to improve the bandwidth of a point-to-point interconnect. In an equalized interconnect, a feed-forward equalizer filter (FFE) is used to shape the transmitted pulse to minimize the intersymbol-interference (ISI) at the receiver side, allowing higher data rate. A decision-feedback equalizer (DFE) is also some-

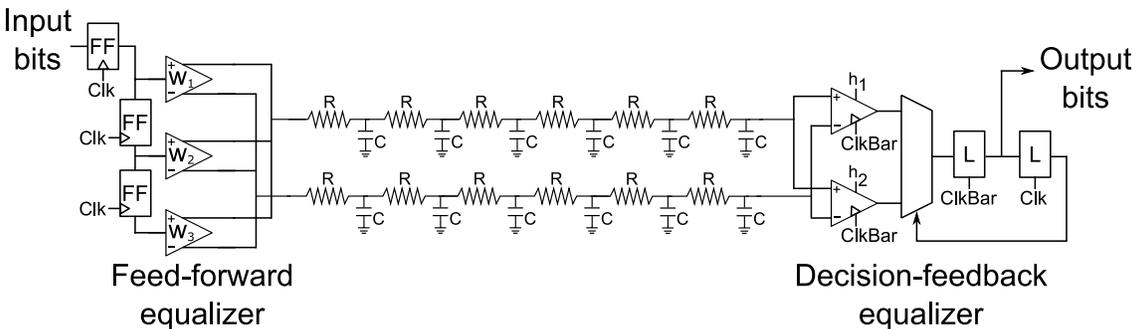


Figure 4: Equalized RC interconnect model. FF = Flip-flop, L = Latch [6].

Interconnect length (mm)	Semi-global wires (M6)						Global wires (M9)					
	DDE (fJ/bt)		FE (fJ/bt)		LAT (cyc)		DDE (fJ/bt)		FE (fJ/bt)		LAT (cyc)	
	Rep	Eqz	Rep	Eqz	Rep	Eqz	Rep	Eqz	Rep	Eqz	Rep	Eqz
2.5	110	–	20	–	1	–	90	–	20	–	1	–
5	215	34	35	39	2	2	175	24	35	37	2	2
7.5	320	101	50	51	3	2	260	34	50	43	3	2
10	425	63	65	73	4	4*	345	58	65	50	4	2
12.5	530	130	80	85	5	4*	430	108	80	60	5	2
15	635	92	95	107	6	6 [#]	515	63	95	81	6	4*

Table 2: Energy and latency for repeater-inserted pipelined and equalized interconnects of various lengths designed using M6 (wire thickness = 403 nm, interlevel dielectric thickness = 400 nm) and M9 (wire thickness = 576 nm, interlevel dielectric thickness = 800 nm) parameters for 32 nm technology. DDE = Data-dependent energy, FE = Fixed energy, fJ/bt = femtojoules per bit-time, LAT = Latency, Rep = Repeater-inserted pipelined interconnects – all repeater-inserted interconnects are pipelined with each pipeline stage having a length of 2.5 mm. Eqz = Equalized interconnects, *2 pipeline segments with 2 cycles per pipeline segment, [#]3 pipeline segments with 2 cycles per pipeline segment.

times used at the receiver to cancel the first trailing ISI tap, thus relaxing the FFE. Equalization improves the latency close to speed of light [12] and also the energy efficiency of the transmission. The low latency results from the fact that equalization uses the fast propagation velocity at high frequencies. The FFE transmitter shapes the phase of the received signal, so the receiver observes constant delay over the data bandwidth. The low energy cost of the equalized interconnect results from output voltage swing attenuation along the wire, which effectively reduces the total amount of charge that the FFE driver must inject into the wire. The low latency and energy cost of equalized interconnect provides better performance–energy tradeoff than repeater-inserted pipelined interconnect. For equalized interconnects of various lengths, the energy and latency listed in Table 2 are obtained by combining the results from [6] with projections on the receiver costs and latency from [7]. The energy-efficiency of equalized interconnects is a strong function of distance. In order to achieve best energy-efficiency with reasonable latency, it is sometimes beneficial to pipeline a long equalized interconnect using shorter equalized segments. For example, for a channel length of 15 mm in lossy M6, the most efficient design uses three pipelined 5 mm segments that are individually equalized.

Table 2 also shows that equalized interconnects consume 4–8 \times lower data-dependent energy (DDE) than repeater-inserted pipelined interconnects, and have comparable fixed energy (FE). As a result, at low channel utilization (interconnect activity), the energy consumed by both repeater-inserted pipelined and equalized channels is comparable and overall scales well with utilization. On the other hand, RF and photonic interconnects have high energy efficiency only at high utilization due to significant FE component (due to carrier or laser/heating power) [11, 13].

4. Evaluation

In this section we compare the performance and power of *mesh*, *cmesh*, *flatFly* and *Clos* networks with repeater-inserted pipelined and equalized channels. We use the Booksim [14] simulator to simulate these network topologies for our target 64-tile system. The simulator models router pipeline latencies, router contention, credit-based flow control and serialization overheads. Warm-up, measure and drain phases of several thousand cycles are used to accurately determine the average latency at a given injection rate. For our analysis, we consider synthetic traffic patterns based on partitioned application model [11]. The entire set of tiles is divided into partitions, and tiles communicate only with other tiles in the same partition. Depending on the mapping of these logical partitions to the actual physical layout, the tiles in a partition could be co-located or distributed across the die. A uniform random (UR) traffic pattern corresponds to a single partition case. A P8C traffic pattern divides the tiles into 8 groups (8 tiles per partition) with the tiles within a partition being co-located. The P8D partition traffic pattern also consists of 8 partitions with the tiles in a partition striped across the chip. In case of the P2D (2 tiles per partition) traffic pattern, the tiles are located in the diagonally opposite quadrants of the chip.

The channel widths (see Table 3) are calculated using the bisection bandwidth criteria such that the ideal network throughput for all topologies under uniform random traffic is the same. For *cmesh* and *flatFly*, instead of using a single network we used two parallel networks with half the calculated channel widths (*cmeshX2* and *flatFlyX2*). This is required as in some cases the calculated channel widths evaluate to be larger than the message size. An added advantage of two parallel networks is good path diversity. The *mesh*, *cmeshX2* and *flatFlyX2* use dimension-

ordered routing, while *Clos* uses randomized oblivious routing where the intermediate router is randomly chosen. All four topologies have credit-based backpressure [14]. For configurations with no virtual channels (VCs), each flit experiences a 2-cycle latency within the router – first cycle for arbitration and second cycle for router crossbar traversal. The route computation is done only for the head flit. For the configurations with VCs, the head flit has a latency of four cycles – one each for route computation, virtual-channel allocation, switch allocation, and switch traversal, while the remaining body/tail flits have a latency of two cycles – one for switch allocation, and second for switch traversal [14]. The latency of the channels is dependent on the topology and physical layout. A message passing interface is adopted and each message size was assumed to be 512 bits. We assume that the input buffer for each port of the router can store 4 messages. The following events - channel utilization and switch traversal are counted while simulating each traffic pattern on each topology. These event counts are then used for calculating power dissipation. While calculating power, we considered both M6 and M9 channels. Both M6 and M9 channels showed similar trends in power savings with equalization and VCs. Here, we present the results for M9 channels. We consider two types of configurations – networks with no VCs and networks with 4 VCs.

4.1. Networks with no virtual channels

Figure 5 shows the latency versus offered bandwidth plot for the four topologies, running different traffic patterns. We consider the two channel widths listed in Table 3. For the *mesh* network (Figure 5(a)), the UR and P8D traffic patterns have comparable saturation through-

Throughput per tile (b/cyc)	Channel width (b)			
	<i>mesh</i>	<i>cmesh</i>	<i>flatFly</i>	<i>Clos</i>
64	128	128×2	128×2	64
128	256	256×2	256×2	128

Table 3: Channel widths for various system types and topologies. For *cmesh* and *flatFly* we use two parallel networks.

put as the source and destination pairs are distributed across the chip. In case of the P2D traffic pattern, all communication paths are between diagonally opposite quadrants of the chip, which increases contention among the network resources, in turn reducing the saturation throughput. For the P8C traffic patterns all the paths are local, resulting in lower contention and higher saturation throughput. As we double the channel width (TT = 64 b/cyc → TT = 128 b/cyc), there is a 50% increase in saturation throughput. On an average, the *cmeshX2* (Figure 5(b)) network exhibits higher saturation throughput than the *mesh* network due to path diversity and it has lower latency due to lower hop count. The *flatFlyX2* (Figure 5(c)) exhibits similar behavior to the *mesh* and *cmeshX2* networks across various traffic patterns. The saturation throughputs for P2D, UR and P8D are comparable to the *cmeshX2*, but the saturation throughput for the P8C traffic pattern is much higher as none of the router-to-router channels are required for communication. The *flatFlyX2* network, however, provides lower and more uniform average latency than *mesh* and *cmeshX2*. The *Clos* (Figure 5(d)) network has a much more uniform latency and throughput across various traffic patterns, due to the availability of path diversity that reduces congestion.

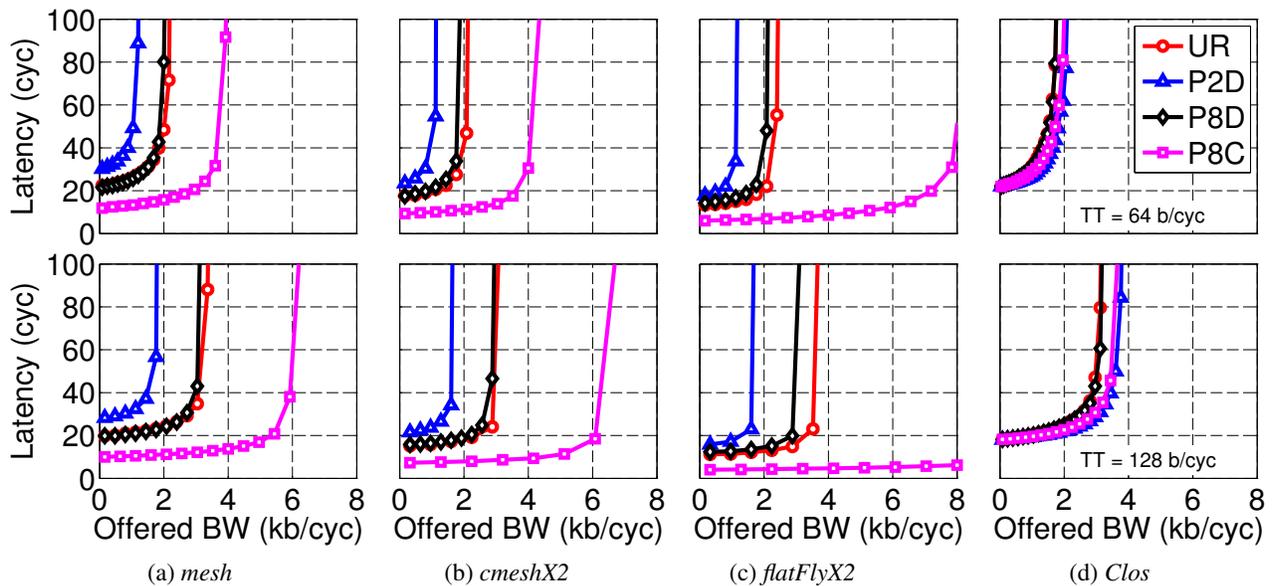


Figure 5: Latency vs. Offered Bandwidth – No VCs, TT = Throughput per tile (for that row).

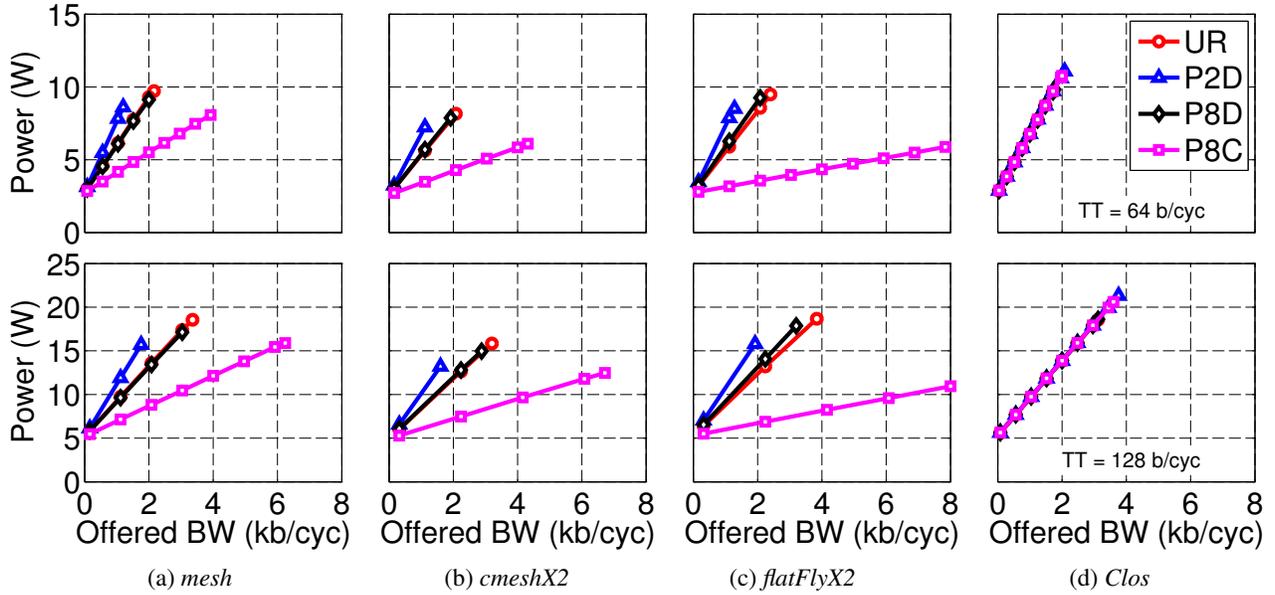


Figure 6: Power vs. Offered Bandwidth – Interconnects designed using pipelining through repeater insertion, no VCs, TT = Throughput per tile (for that row).

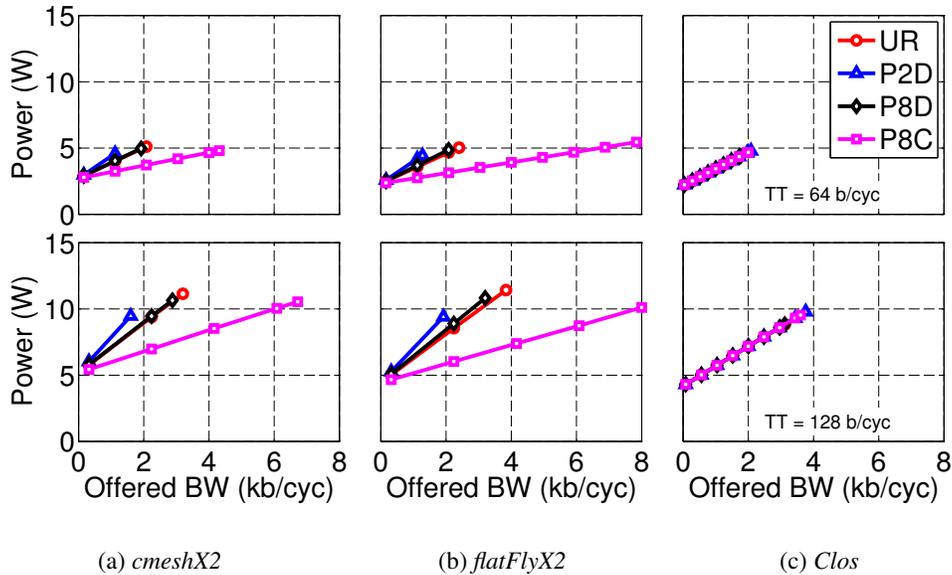


Figure 7: Power vs. Offered Bandwidth – Interconnects designed using equalization, No VCs, TT = Throughput per tile (for that row).

Figure 6 shows a plot of the power dissipation versus the offered bandwidth for the various topologies and traffic patterns. The power dissipation is calculated using channel energy/bit-time values in Table 2, and router models adapted from [15]. The *cmeshX2* has $2\times$ less routers than *mesh*. Hence, even with a higher router radix ($1.6\times$), *cmeshX2* consumes lower power in routers than *mesh* at comparable throughput. The number of channels in *cmeshX2* is $2\times$ lower, but each channel is $2\times$ longer compared to *mesh*. Hence, the channel power is compara-

ble. At saturation, the *flatFlyX2* network consumes power comparable to the *mesh* at their respective saturations for the UR, P2D and P8D traffic pattern, with roughly half the latency. For the P8C traffic pattern, the power dissipation is lower in *flatFlyX2*, as the global channels are not used and power is dissipated only in the routers and the static components of the wires. For the *Clos* network, the power consumed across various traffic patterns is uniform and is comparable to UR/P8D in *mesh* network. For a given topology, sizing-up the network from TT = 64

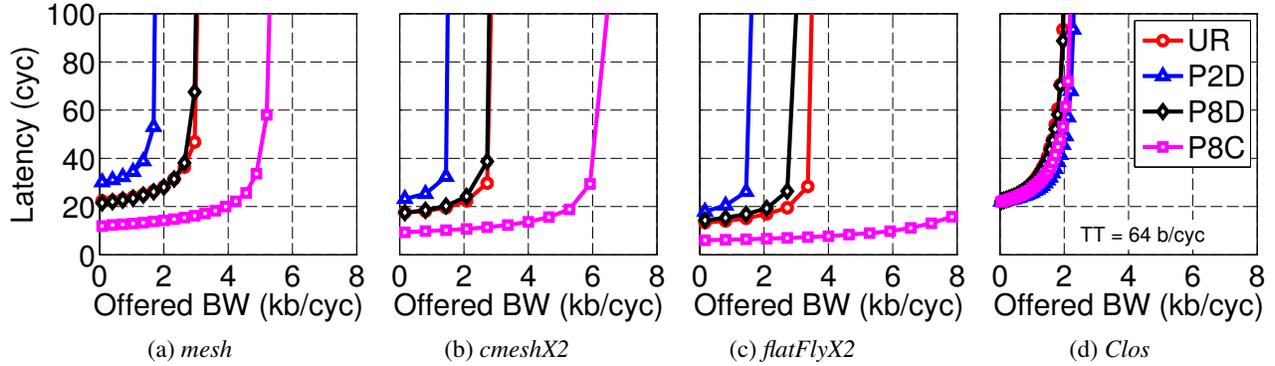


Figure 8: Latency vs. Offered Bandwidth – 4 VCs per router, TT = Throughput per tile (for that row).

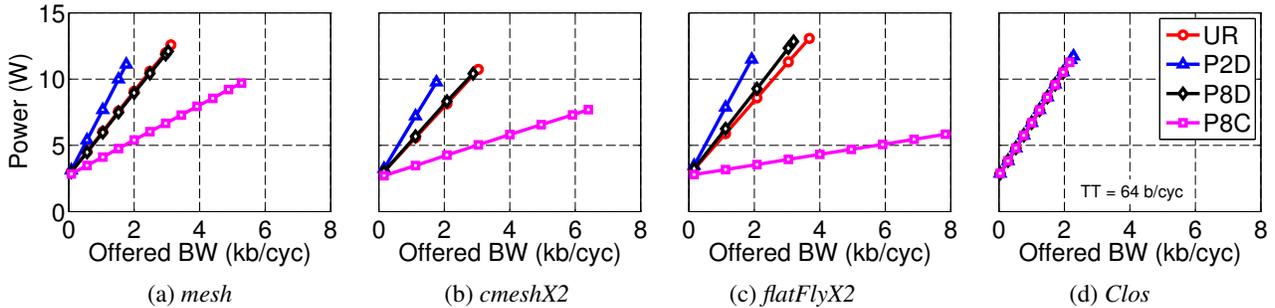


Figure 9: Power vs. Offered Bandwidth – Interconnects designed using pipelining through repeater insertion, 4 VCs per router, TT = Throughput per tile (for that row).

b/cyc to TT = 128 b/cyc increases the fixed power by $2\times$ while the data-dependent power is comparable for same throughput.

Figure 7 shows the power versus bandwidth plots for *cmeshX2*, *flatFlyX2* and *Clos* networks with equalized channels. The larger energy-cost advantage of equalized interconnects over repeated interconnects is partially amortized at the network level by the router energy costs. As described earlier, the power consumed in repeater-inserted pipelined and equalized interconnects is comparable at low offered bandwidths, and that translates to similar network power. As offered bandwidth increases, the data-dependent component starts to dominate and at saturation, we observe $1.5\text{--}2\times$ better energy-efficiency in networks with equalized interconnects.

4.2. Networks with virtual channels

Virtual channels [16] can be used to improve the performance of all four topologies. Figure 8 shows the latency versus bandwidth plots for the four topologies with TT = 64 b/cyc with four VCs. While implementing VCs, the total size (in bits) of the input buffer per port and crossbar size is maintained to be equal to the no-VC case. For the *mesh*, *cmeshX2* and *flatFlyX2* networks, $1.5\text{--}2\times$ improvement in the saturation throughput is observed by using 4 VCs. On the other hand, for the *Clos* network the improvement in the saturation throughput is minimal

because of the inherent path diversity in the topology. An interesting result of using VCs is that a system with *mesh/cmshX2/flatFlyX2* and TT = 64 b/cyc with 4 VCs roughly matches the saturation throughput of a system with TT = 128 b/cyc with no-VCs. There is a small increase in the latency when VCs are used due to increase in the router latency and credit-loop latency.

Figure 9 shows the power versus offered bandwidth plots for systems with various topologies and TT = 64 b/cyc and with 4 VCs. At comparable throughput, the system with TT = 64 b/cyc and 4 VCs consumes a 25–50% lower power than the corresponding system with TT = 128 b/cyc and no VCs. This is because the overhead due to VCs is lower than the fixed energy savings due to narrower channel widths and smaller crossbar in the router. A similar savings in power consumption can be observed when equalized interconnects are used (see Figure 7 and Figure 10).

Overall, compared to using a *cmeshX2 / flatFlyX2 / Clos* with TT = 128 b/cyc with no VCs and repeater-inserted pipelined interconnects, corresponding networks with TT = 64 b/cyc, 4 VCs and equalized interconnects consumes around $2\text{--}3\times$ lower power at comparable throughput. In addition, using $2\times$ narrower channels in system with TT = 64 b/cyc saves the network wire area. If we compare a *mesh* network with no VCs designed using pipelining through repeater insertion to

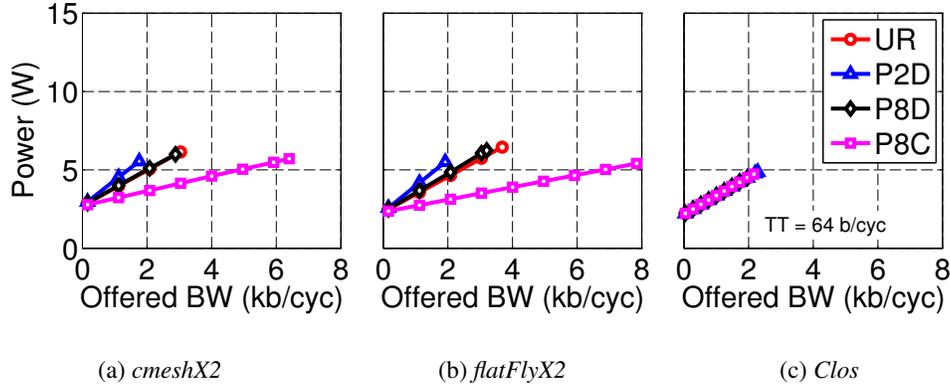


Figure 10: Power vs. Offered Bandwidth – Interconnects designed using equalization, 4 VCs per router, TT = Throughput per tile.

a *cmeshX2/flatFlyX2* designed using equalized interconnects and VCs, then *cmeshX2/flatFlyX2* consumes close to $3\times$ lower power and $2\times$ lower wire area at comparable throughput. The *Clos* network cannot match the saturation throughputs of *mesh* for local traffic patterns, but for global traffic patterns, the *Clos* network can match the throughput, while consuming $3\times$ lower power and $2\times$ lower area. In addition, *Clos* network also provides uniform latency and throughput across various traffic patterns. Results for M6 show identical trends with 15–25% larger power at saturation depending on the network.

Figure 11 shows a plot of the power split between the channel data-dependent power, channel fixed power, router data-dependent power and router fixed power for two saturation-throughput matched designs. The top figure corresponds to TT of 128 b/cyc with no VCs, while the bottom figure corresponds to TT of 64 b/cyc with 4 VCs. Here the router fixed power component consists of the power dissipated in the clocked circuits in the flip-flops in the router. For both designs, the power dissipated in the routers progressively increases from *Clos* to *flatFlyX2* to *cmeshX2* to *mesh* in that order. This can be attributed to the change in router count and router sizes across the various topologies.

The channel fixed power for all four networks is comparable. Table 2 shows that for repeater-inserted pipelined interconnects the data-dependent energy increases as channel length increases. As a result of this, for all traffic patterns the channel data-dependent power consumed in the *Clos* network is the highest as it always uses global channels. For UR, P2D and P8D traffic pattern, the *flatFlyX2* network consumes higher channel data-dependent power than *mesh* and *cmeshX2* as it needs to use long global channels. On the other hand, for the P8C traffic pattern, the *flatFlyX2* consumes no channel data-dependent power as no channels are utilized. Between *mesh* and *cmeshX2*, for all traffic patterns, the *cmeshX2* consumes lower channel data-dependent power due to lower channel utilization.

The channel fixed energy and data-dependent energy for equalized interconnects is comparable and $4\text{--}8\times$ lower, respectively, to that of repeater-inserted pipelined interconnects (from Table 2). Hence, for networks designed using equalized interconnects, the total fixed and data-dependent channel power dissipation is comparable and lower, respectively, compared to the networks designed using repeater-inserted pipelined interconnects. The power dissipated in the routers is similar for both pipelined repeater-inserted and equalized design.

For design with TT = 64 b/cyc and 4 VCs the router power is lower than that for design with TT = 128 b/cyc and no VCs for corresponding topologies. The lower router power is due to a significant reduction in crossbar power compared to the minimal power increase due to the additional buffer interface circuits and VC allocation unit required for implementing virtual channels. The channel fixed power in design with TT = 64 b/cyc and 4 VCs is $2\times$ lower than that of design that has TT = 128 b/cyc and no VC as the channel widths are $2\times$ lower. The channel data-dependent power for both designs is comparable as though the channel width reduces by $2\times$, the channel utilization factor increases by $2\times$.

5. Related work

Previous research work has looked extensively at both issues of low-latency and improved energy-efficiency [4, 5, 7, 12] for point-to-point on-chip interconnect. However, very little work has been done to evaluate the impact of these new signaling techniques at the network/system level, compared to the extensive research and evaluation of repeated-interconnect based networks-on-chip. Beckman et al [17] proposed taking advantage of the transmission line behavior of the on-chip wires to design long point-to-point channels with low latency. The proposed technique was primarily applied for communication between the cores and cache in a multicore system. Compared to a dynamic non-uniform cache architecture, the

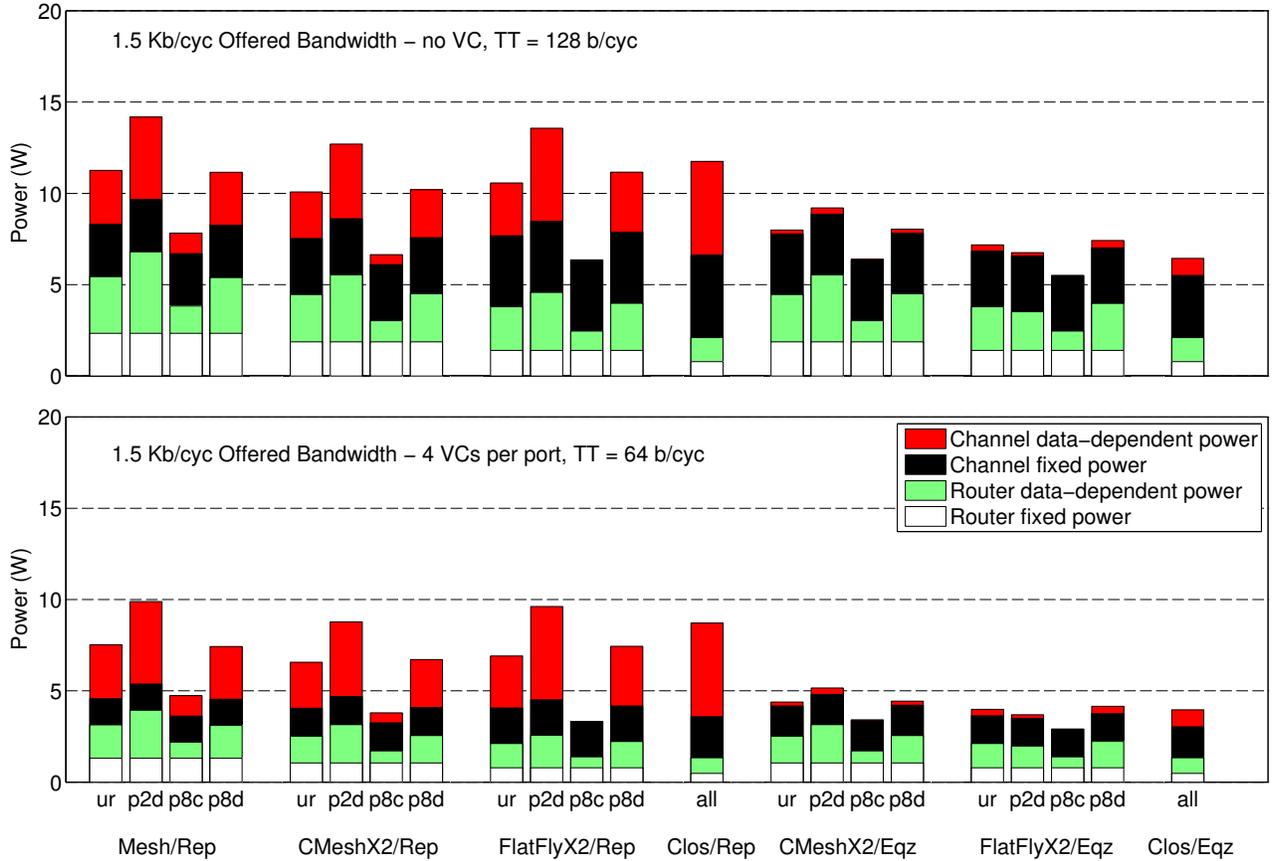


Figure 11: Power breakdown Power for the Clos network did not vary significantly across various traffic patterns.

use of transmission line behavior reduced the network’s dynamic power dissipation and cache area at comparable performance.

Similarly, the use of multiple RF carriers over on-chip transmission lines is proposed in [13]. These point-to-point channels are provided as low-latency short-cuts in addition to the underlying mesh network for on-chip communication. While RF interconnects can be used to provide low-latency paths across the chip, due to the transmission line nature of the channel, proposed polymorphic adaptive bandwidth allocation and reconfiguration as well as multicast are difficult to achieve due to stub effects on channel transfer function and varying termination impedance conditions. Given the required size and pitch of transmission lines, significantly more complex transceivers and large number of carriers per wire are needed to match the bandwidth-density of equalized interconnects. Equalized interconnects avoid these issues by operating wires in the RC regime.

The use of silicon photonics for on-chip communication has been proposed by many researchers [11, 18, 19]. The key idea here is that the electrical power cost associated with an optical channel is independent of the length of the channel. In addition, the photonic channels

drastically improve the bandwidth density through wavelength division multiplexing and reduce latency. Similar to RF interconnects, the photonic channels, however, have a fixed overhead in terms of laser power, which is consumed by the laser source and thermal tuning power, which is consumed to maintain the resonant frequency of various photonic devices. Unless the on-chip photonic network is carefully designed, at low network utilization, these costs can be prohibitively expensive and any potential advantage in latency, throughput and power gained through photonic channels can be lost due to fixed laser and thermal tuning power overhead.

6. Conclusion

In this paper, we presented a framework to quantify the benefits of using equalized on-chip interconnects as global channels for energy-and-area efficient low-diameter on-chip networks like *flattened butterfly* and *Clos*. Compared to the high-diameter networks like *mesh* that are designed using repeater-inserted pipelined interconnects, the use of interconnect equalization in *flattened butterfly* and *Clos* networks reduces power dissipation by 1.5-2 \times depending on traffic. Furthermore, the use of vir-

tual channels significantly improves the energy-efficiency (25–50%) and area (2×). An added advantage of using the *Clos* network is that it has a very narrow distribution of latency and throughput across various traffic patterns, which eases the multicore programming challenges. The *cmesh* network provides a power-performance trade-off comparable to the *flattened butterfly/Clos*, but the difficulty in programming the *cmesh* network makes it unattractive. The power of equalized networks scales well with network utilization, which is one of critical issues in proposed RF and photonic interconnects alternatives, which are potentially energy-efficient only at high network utilizations.

Acknowledgments

This work was supported by Intel Corp. The authors would like to thank Christopher Batten for discussions on interconnect network topologies, Ian Young and Alexandra Kern for discussions on equalized interconnects and process technology, and Imran Shamim for help with router modeling.

References

- [1] U. Nawathe *et al.*, “Implementation of an 8-core, 64-thread, power-efficient SPARC server on a chip,” *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 1, pp. 6–20, Jan. 2008.
- [2] C. Clos, “A study of non-blocking switching networks,” *Bell System Technical Journal*, vol. 32, pp. 406–424, 1953.
- [3] S. Bell *et al.*, “TILE64 processor: A 64-core SoC with mesh interconnect,” *Int’l Solid-State Circuits Conf.*, Feb. 2008.
- [4] R. Ho *et al.*, “High speed and low energy capacitively driven on-chip wires,” vol. 43, no. 1, pp. 52–60, Jan. 2008.
- [5] E. Mensink *et al.*, “A 0.28pJ/b 2Gb/s/ch transceiver in 90nm CMOS for 10 mm on-chip interconnects,” in *Proc. Digest of Technical Papers. IEEE International Solid-State Circuits Conference ISSCC 2007*, 11–15 Feb. 2007, pp. 414–612.
- [6] B. Kim and V. Stojanović, “Characterization of equalized and repeated interconnects for NoC applications,” *IEEE Design and Test of Computers*, vol. 25, no. 5, pp. 430–439, 2008.
- [7] B. Kim *et al.*, “A 4Gb/s/ch 356fJ/b 10mm equalized on-chip interconnect with nonlinear charge-injecting transmitter filter and transimpedance receiver in 90nm cmos technology,” in *Proc. Digest of Technical Papers. IEEE International Solid-State Circuits Conference ISSCC 2009*, 8–12 Feb. 2009, pp. 66–67.
- [8] J. Kim, J. Balfour, and W. J. Dally, “Flattened butterfly topology for on-chip networks,” in *Proc. 40th Annual IEEE/ACM International Symposium on Microarchitecture MICRO 2007*, 1–5 Dec. 2007, pp. 172–182.
- [9] S. Vangal *et al.*, “80-tile 1.28 TFlops network-on-chip in 65 nm CMOS,” *Int’l Solid-State Circuits Conf.*, Feb. 2007.
- [10] J. Balfour and W. J. Dally, “Design tradeoffs for tiled cmp on-chip networks,” in *ICS ’06: Proceedings of the 20th annual International Conference on Supercomputing*, 2006, pp. 187–198.
- [11] A. Joshi *et al.*, “Silicon-photonic clos networks for global on-chip communication,” in *Proc. International Symposium on Networks-on-Chip NOCS 2009*, 10–13 May 2009.
- [12] A. Jose, G. Patounakis, and K. Shepard, “Near speed-of-light on-chip interconnects using pulsed current-mode signalling,” June 2005, pp. 108–111.
- [13] M. C. F. Chang *et al.*, “Power reduction of CMP communication networks via RF-interconnects,” in *Proc. MICRO-41 Microarchitecture 2008 41st IEEE/ACM International Symposium on*, 8–12 Nov. 2008, pp. 376–387.
- [14] W. J. Dally and B. P. Towles, *Principles and Practices of Interconnection Networks*. Elsevier, 2004.
- [15] H. Wang, L.-S. Peh, and S. Malik, “A technology-aware and energy-oriented topology exploration for on-chip networks,” *Design, Automation and Test in Europe, 2005. Proceedings*, pp. 1238–1243 Vol. 2, March 2005.
- [16] W. J. Dally, “Virtual-channel flow control,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 3, no. 2, pp. 194–205, 1992.
- [17] B. M. Beckmann and D. A. Wood, “TLC: transmission line caches,” in *Proc. 36th Annual IEEE/ACM International Symposium on MICRO-36 Microarchitecture*, 2003, pp. 43–54.
- [18] A. Shacham, K. Bergman, and L. P. Carloni, “On the design of a photonic network-on-chip,” in *Proc. First International Symposium on Networks-on-Chip NOCS 2007*, 7–9 May 2007, pp. 53–64.
- [19] D. Vantrease *et al.*, “Corona: System implications of emerging nanophotonic technology,” in *Proc. 35th International Symposium on Computer Architecture ISCA ’08*, 21–25 June 2008, pp. 153–164.