A QoS-Enabled On-Die Interconnect Fabric for Kilo-Node Chips

To meet rapidly growing performance demands and energy constraints, future chips will likely feature thousands of on-die resources. Existing network-on-chip solutions weren't designed for scalability and will be unable to meet future interconnect demands. A hybrid network-on-chip architecture called Kilo-NoC co-optimizes topology, flow control, and quality of service to achieve significant gains in efficiency.

Boris Grot
École Polytechnique Fédérale de Lausanne

Joel Hestness
University of Texas at Austin

Stephen W. Keckler
Nvidia

Onur Mutlu
Carnegie Mellon University

The semiconductor industry is rapidly moving toward rich, chip-level integration; in many application domains, chips with hundreds to thousands of processing elements are likely to appear in the near future. To address the communication needs of richly integrated chips, the industry has embraced structured, on-die communication fabrics called networks-on-chip (NoCs). However, existing NoC architectures have been designed for substrates with dozens of nodes, not hundreds or thousands; once scaled to tomorrow's kilonode configurations, significant performance, energy, and area overheads emerge in today's state-of-the-art NoCs.

We focus on NoC scalability from energy, area, performance, and quality of service (QoS) perspectives. Prior research indicates that richly connected topologies improve latency and energy efficiency in on-chip networks.[1,2] While our analysis confirms those results, it also identifies buffer overheads as a critical scalability obstacle that emerges once richly connected NoCs are scaled to configurations with hundreds of nodes. Large buffer pools adversely affect NoC area and energy efficiency. The addition of QoS support further increases storage overhead, virtual channel (VC) requirements, and arbitration complexity.

Our solution holistically addresses key sources of inefficiency in NoCs of highly integrated chip multiprocessors (CMPs) through a hybrid NoC architecture (called Kilo-NoC) that offers low latency, a small footprint, good energy efficiency, and strong service guarantees.

Kilo-NoC overview

Existing QoS approaches necessitate hardware support at every router, incurring network-wide costs and complexity overheads. We propose a QoS architecture that overcomes this limitation of previous designs by requiring QoS hardware at just a subset of the routers. Our approach consolidates shared resources, such as memory controllers,

0272-1732/12/$31.00 © 2012 IEEE Published by the IEEE Computer Society

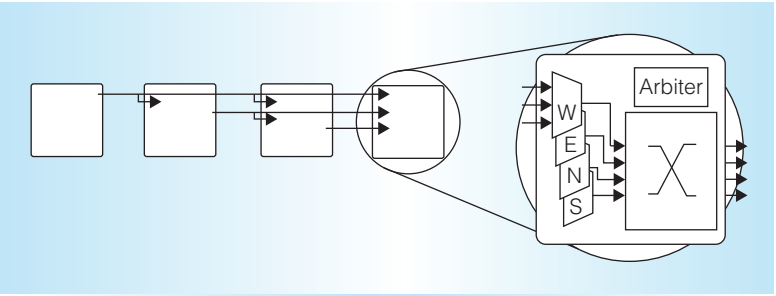

Figure 1. A conceptual view of the Multidrop Express Channels (MECS) architecture. MECS uses a point-to-multipoint interconnect model to enable rich connectivity with relatively few channels.

within a portion of the network and enforces service guarantees only within subnetworks that contain the shared resources. The enabling technology underlying the scheme is a richly connected topology that enables single-hop access to any QoS-protected subnetwork, effectively eliminating intermediate nodes as sources of interference. To our knowledge, our work is the first to consider the interaction between topology and QoS.

While topology-awareness offers a considerable reduction in QoS-related costs, it doesn't address the high buffer overheads of richly connected topologies. We eliminate much of this buffer expense by introducing a lightweight elastic buffer (EB) architecture that integrates storage directly into links. Again, our design leverages a feature of the topology to offer a single-network, deadlock-free EB architecture at a fraction of prior schemes' cost.

Together, these techniques synergistically work to enable a highly scalable kilo-node interconnect fabric. Our evaluation in the context of a thousand-terminal system reveals that the Kilo-NoC architecture is highly effective in relieving the scalability bottlenecks of today's NoCs. Compared to a state-of-the-art, QoS-enabled NoC, our proposed design reduces network area requirements by 45 percent and energy expenditure by 29 percent. The Kilo-NoC attains these benefits without sacrificing either performance or strength of service guarantees.

Background

Networks are commonly characterized along three dimensions: topology, routing, and flow control. Of these, topology is the single most important determinant of performance, energy efficiency, and cost (area).

To date, most NoCs that have been realized in silicon feature a ring or mesh topology. While such topologies are acceptable when interconnecting a modest number of nodes, their large average hop count in chips with hundreds of networked components represents a serious efficiency bottleneck. Each hop involves a router crossing, which often dominates the per-hop latency and energy cost due to the need to write and read the packet buffers, arbitrate for resources, and traverse a crossbar switch.

To bridge the scalability gap, researchers have proposed low-diameter NoC topologies that improve efficiency through rich internode connectivity. One such topology is the flattened butterfly, which fully interconnects the nodes in each dimension via dedicated point-to-point channels.[1] However, the flattened butterfly's channel requirements and crossbar complexity grow quadratically with the network radix and represent a scalability obstacle. An alternative organization, shown in Figure 1, uses multidrop express channels (MECS) to achieve the same degree of connectivity as the flattened butterfly but with fewer channels.[2] Each node in a MECS network has four output channels, one per cardinal direction. Lightweight drop interfaces integrated into the channel let packets exit the channel into one of the routers spanned by the link.

Unfortunately, long-range links require large buffer pools to cover the flight time of the data and the credit return. Such large buffer configurations carry a significant area and energy cost in richly connected NoCs with hundreds of nodes. For instance, in a 16 × 16 network with a low-diameter topology (for example, MECS or flattened butterfly), each router needs 30 network ports with up to 35 buffer slots per packet class per port. The chip-wide buffer requirements of such a NoC, assuming two packet classes and 16-byte links, exceed 8 Mbytes—an extraordinary amount from an area and energy standpoint, even in future process technologies.

Adding QoS support further increases buffer demands. Although recent work has

demonstrated NoC QoS architectures with smaller footprints than traditional networking schemes,[3,4] that work only considered mesh topologies. Our analysis shows that the benefits of these QoS architectures are greatly diminished in large-scale, low-diameter NoCs owing to the high VC requirements imposed by long link spans. For instance, in the hypothetical 16 × 16 network we mentioned earlier, a high-radix router outfitted with a state-of-the-art Preemptive Virtual Clock (PVC) NoC QoS architecture requires more than 700 VCs, compared to just 24 VCs in a PVC-enabled mesh router.

Improving the flow-control mechanism, which manages traffic flow by allocating resources to packets, can help mitigate high buffer overheads. One potential approach is bufferless flow control, recently examined by NoC researchers looking for ways to boost efficiency.[5] Unfortunately, existing bufferless architectures are unable to provide service guarantees. Integrating storage elements directly into links, a technique termed elastic buffering, is another promising direction recent research has taken.[6,7] While existing proposals offer some gains in efficiency, the serializing nature of EB links complicates deadlock avoidance and impedes the isolation of flows, which is necessary for QoS guarantees.

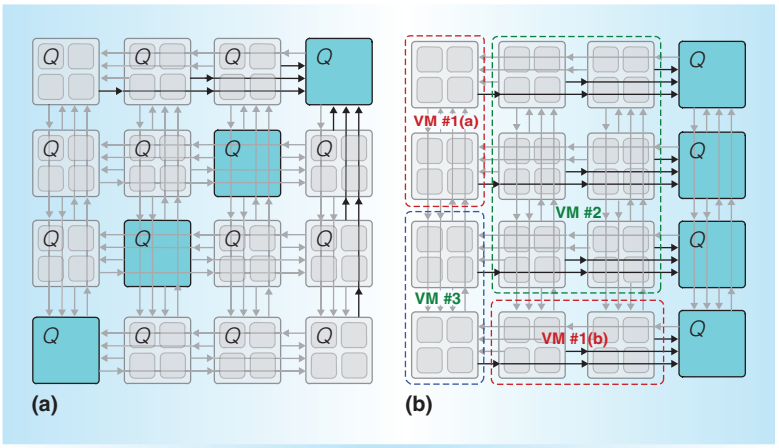

Figure 2. A 64-tile chip multiprocessor (CMP) with four-way concentration and MECS topology. Light nodes represent core and cache tiles; shaded nodes represent memory controllers; and Q indicates QoS hardware support. Dotted lines indicate virtual machine (VM) assignments in a topology-aware QoS architecture. Baseline QoS-enabled CMP (a); topology-aware QoS approach (b). Topology-awareness enables a reduction in the number of routers that require QoS hardware.

Kilo-NoC architecture

We describe the proposed Kilo-NoC architecture in the context of a 1,024-tile CMP implemented in 15-nm technology.

Baseline design

Figure 2a shows the baseline NoC organization, scaled down to 64 tiles for clarity. We employ concentration[8] to reduce the number of network nodes from 1,024 to 256 by integrating four terminals at each router via a crossbar switch. A node refers to a network node, while a terminal is a discrete system resource (such as a core, cache tile, or memory controller) with a dedicated port at a network node. The nodes are interconnected via a richly connected MECS topology. We chose MECS due to its low diameter, scalable channel count, modest switch complexity, and the unique capabilities multidrop offers. PVC enforces QoS guarantees for the virtual machines (VMs) sharing a die.[4]

We arrange the 256 concentrated nodes in a 16 × 16 grid. Each MECS router integrates 30 network input ports (15 per dimension). With one cycle of wire latency between adjacent nodes, maximum channel delay from one edge of the chip to another is 15 cycles. The round-trip credit time is 35 cycles, once router pipeline delays are included. This round-trip latency establishes a lower bound for per-port buffer requirements in the absence of any location-dependent optimizations. To guarantee freedom from protocol deadlock, each port needs a dedicated VC per packet class. With two priority levels (*request* at low priority and *reply* at high priority), a pair of 35-deep VCs affords deadlock-free operation while covering the maximum round-trip credit delay. The total buffer requirements are 70 flits at each input port and 2,100 flits for the entire 30-port router. With 16-byte flits, total required storage is 32 Kbytes per router and 8.4 Mbytes network wide.

To guarantee QoS, packets from different nodes require separate VCs to prevent priority inversion within a VC buffer. To accommodate a worst-case pattern consisting of single-flit packets from different nodes, an

unoptimized router would require 35 VCs per port. Several optimizations, such as location-dependent buffer sizing, can be used to reduce the VC and buffer requirements at additional design expense and arbitration complexity. Here, we assume a 25 percent reduction in per-port VC requirements. Assuming a maximum packet size of four flits, a baseline QoS-enabled architecture requires 25 four-deep VCs per port, 750 VCs and 3,000 flit slots per router, and 12 Mbytes of storage network wide.

Topology-aware QoS architecture

Our first optimization target is the QoS mechanism, which imposes significant VC and buffer overheads. In contrast to existing network QoS architectures that demand dedicated QoS logic and storage at every router, we seek to limit the number of nodes requiring hardware QoS support. Our proposed scheme, called Topology-Aware QoS (TAQ), accomplishes this goal by isolating shared resources into dedicated regions of the network, called *shared regions* (SRs), with hardware QoS enforcement within each SR. The rest of the network is freed from the burden of hardware QoS support and enjoys reduced cost and complexity.

The TAQ architecture leverages the rich intradimension connectivity provided by MECS (or other low-diameter topologies) to ensure single-hop access to any shared region, which we achieve by organizing the SRs into columns spanning the entire width of the die. Single-hop connectivity guarantees interference-free transit into an SR. Once inside the SR, a packet is regulated by PVC (or another QoS mechanism) as it proceeds to its destination. To prevent unregulated contention for network bandwidth at concentrated nodes outside of the SR, we require the OS or hypervisor to coschedule only threads from the same VM onto a node. Figure 2b shows our proposed organization, including a sample assignment of nodes to VMs. Note that though the figure's SR column is on the edge of the die, TAQ doesn't require such placement.

Depending on how virtual machines on a die are placed, certain intra-VM and inter-VM transfers might have to flow through a router at a node mapped to an unrelated VM. Because such scenarios can result in inter-VM contention at routers lacking QoS support, we use simple routing rules that exploit the combination of a richly connected topology and QoS-enabled regions to avoid inter-VM interference. The rules can be summarized as follows.

- Communication within a dimension is unrestricted, since a low-diameter topology provides interference-free, single-hop communication in a given row or column.
- Dimension changes are unrestricted if the turn node belongs to the same VM as the packet's source or destination. For example, all cache-to-cache traffic associated with VM #2 in Figure 2b stays within a single convex region and never needs to transit through a router in another VM.
- Packets requiring a dimension change at a router associated with a node of an unrelated VM must flow through one of the shared regions. Depending on the locations of the communicating nodes and the SRs, the resulting routes may be nonminimal. For instance, in Figure 2b, traffic from partition (a) of VM #1 that is transiting to partition (b) must take the longer path through the shared column to avoid turning at a router associated with VM #2.

Our proposal preserves service guarantees for all VMs, regardless of the locations of communicating nodes. However, placing all of a VM's resources in a contiguous region can improve both performance and energy efficiency by reducing communication distance and minimizing accesses to the SRs.

Elastic-buffer flow control

Freed from the burden of enforcing QoS, routers outside of shared regions enjoy a significant reduction in VCs to just one VC per packet class. Yet, despite this reduction, a MECS kiloterminal network with two packet priority classes still requires a prohibitive 8 Mbytes of buffer capacity. In an effort to further reduce buffer overheads, we turn to elastic buffering.

Although conventional elastic-buffered networks are incompatible with QoS due to the serializing nature of EB flow control (which can cause priority inversion within a channel), our proposed TAQ architecture enables elastic buffering outside of the shared regions by eliminating interference among nodes from different VMs. A point-to-multipoint MECS topology also greatly reduces overall storage requirements, because all downstream destination nodes effectively share each buffer slot in a channel. In contrast, the point-to-point topologies considered in earlier EB studies offer limited storage savings, because buffer capacity is simply shifted from routers to links.

One drawback of existing EB architectures is the additional complexity associated with guaranteeing freedom from protocol deadlock. An earlier proposal ensured deadlock freedom through a conventional VC architecture, thus negating the benefits of elastic buffering.[6] A more recent study advocates pure, elastic-buffered NoCs with no VCs.[7] While a pure EB architecture is leaner than alternative designs, it requires a dedicated physical network for each packet class for deadlock avoidance, increasing NoC area and wire pressure.

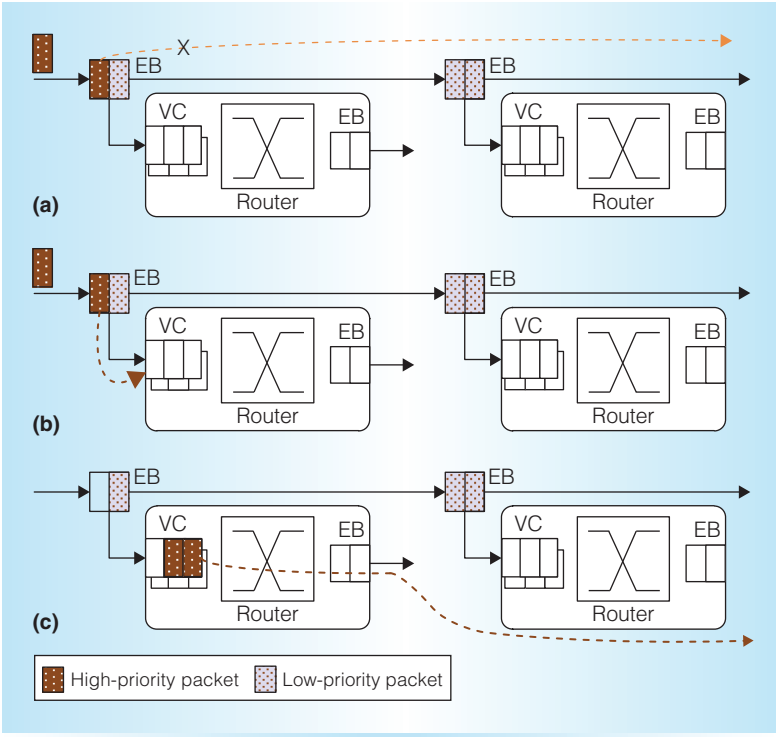

Figure 3. An example of deadlock avoidance in a network with elastic buffers (EB) via the Just-in-Time (JIT-VC) mechanism. A high-priority packet in a MECS channel is obstructed by a low-priority packet (a). The high-priority packet acquires a buffer at a router associated with the EB (b). The high-priority packet switches to a new MECS channel and proceeds toward its destination (c). Our architecture avoids deadlock by providing escape paths for blocked packets.

Low-cost elastic buffering

We propose an EB organization that affords considerable area savings over earlier schemes. Our approach combines elastic-buffered links with minimal VC support, enabling a single-network architecture with hybrid EB/VC flow control. The key to minimizing VC costs comes through a novel flow-control mechanism called Just-in-Time VC binding (JIT-VC), which enables a packet in the channel to allocate a VC from an EB adjacent to the router. In doing so, our design essentially eliminates the buffer credit loop, whose length determines VC and buffer requirements. The resulting organization represents a scalable alternative to traditional VC architectures in which buffer requirements are proportional to the link delay, necessitating large buffer pools to cover long link spans.

Because packets regulated by JIT-VC flow control don't reserve downstream buffer space before entering the channel, they leave the network susceptible to protocol deadlock. We assure deadlock freedom by providing an escape path for blocked packets into intermediate routers along their direction of travel by exploiting the multidrop feature of MECS channels in concert with the JIT-VC mechanism. Under normal operation, a packet will allocate a VC once it reaches the EB at the target (turn or destination) node. However, should a high-priority (for example, reply) packet be blocked in the channel, it can escape into a JIT-allocated VC at another node. Once buffered at an escape router, a packet will switch to a new MECS channel by traversing the router pipeline like any other packet. To prevent circular deadlock, we don't let packets switch dimensions at an escape node. Figure 3 shows a high-level depiction of our approach.

Forward progress in the proposed EB-enabled network is guaranteed through simple

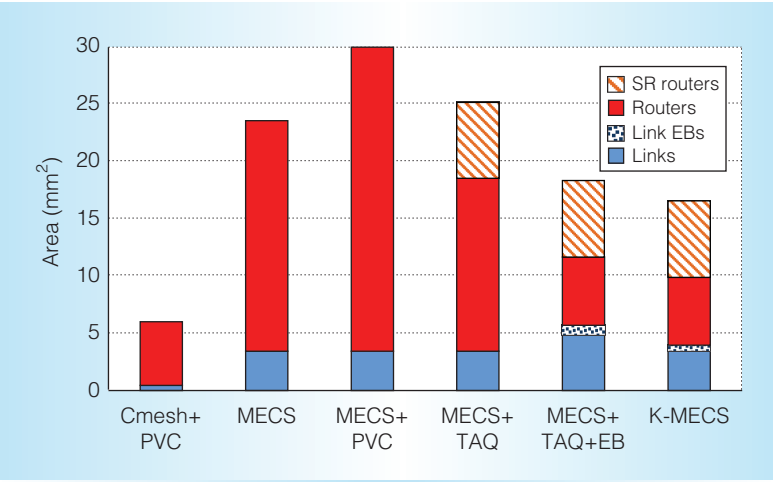

Figure 4. Area breakdown of various NoC configurations. K-MECS has the lowest area among organizations with similar or greater bisection bandwidth.

microarchitectural mechanisms that segregate packets into dedicated VCs according to priority class at each router, allocate resources in strict priority order (as determined by packet class), and ensure that high-priority packets eventually escape into a router along their paths by bypassing or draining lower-priority packets using JIT-VC. For details of the scheme, along with a proof showing its freedom from deadlock, see our paper for the 38th International Symposium on Computer Architecture (ISCA 2011).[9]

A single-network EB scheme, as we describe here, enables significant reduction in storage requirements for nodes outside of shared regions. Given a maximum packet size of four flits and two priority classes, a pair of four-deep VCs suffices at each router input port. Compared to a baseline PVC-enabled MECS router with 25 VCs per port, our approach reduces both VC and buffer requirements by a factor of 12.

Evaluation highlights

We rigorously evaluated the set of proposed optimizations in terms of their effect on network efficiency, performance, and QoS. We used detailed technology models for area and energy and simulation-based studies for performance and QoS analysis. Here, we present a sampling of the results and refer interested readers to our ISCA 2011 paper for additional findings, insights, and details regarding our methodology.[9]

We model a 1,024-tile CMP in 15-nm technology with an on-chip voltage of 0.7 V and a die area of 256 mm^2, excluding peripheral circuitry. At the network level, four-way concentration reduces the number of routers to 256, of which 64 correspond to various shared resources—potentially including memory controllers, fixed-function accelerators, and I/O interfaces.

We evaluated the following NoC organizations:

- Cmesh+PVC: a concentrated mesh topology with PVC-based QoS support.
- MECS: baseline MECS network with no QoS support.
- MECS+PVC: QoS-enabled MECS network with PVC-based QoS support.
- MECS+TAQ: MECS network with the proposed topology-aware QoS architecture. PVC enforces QoS inside four shared regions and no QoS support exists elsewhere.
- MECS+TAQ+EB: TAQ-enabled network augmented with a pure, elastic-buffer flow-control architecture. Deadlock freedom is ensured through separate *request* and *reply* networks. Elastic buffering is deployed only outside shared regions, with conventional buffering and PVC inside SRs.
- K-MECS: the proposed Kilo-NoC configuration, featuring TAQ and hybrid EB/VC flow control with JIT-VC allocation (outside SRs).

Area analysis

Figure 4 breaks down the total network area into four resource types: links, link-integrated EBs, regular routers, and SR routers (TAQ-enabled topologies only). For links, we account for the area of drivers and receivers and anticipate that wires are routed over logic in a dedicated layer.

TAQ proves to be an effective optimization for reducing network area. Compared to a baseline, QoS-enabled MECS network (MECS+PVC), TAQ enables a 16-percent area reduction (MECS+TAQ bar) due to diminished buffer requirements. The pure,

elastic-buffered NoC further reduces the area footprint by 27 percent (MECS+TAQ+EB), but at the cost of a 56-percent increase in wire requirements precipitated by the need for a second network. K-MECS offers an additional 10-percent area reduction and curtails wire pressure compared to a pure EB architecture by not requiring a second network to guarantee deadlock freedom. To put the optimizations in perspective, the conventionally buffered, QoS-enabled SR routers in K-MECS account for more than one half of the total router area but make up just a quarter of the network nodes.

The smallest network area is found in the Cmesh topology, due to its modest bisection bandwidth. The Cmesh NoC occupies 2.8 times less area than the K-MECS network, but offers 8 times less network bandwidth. Link area represents just 7 percent of the Cmesh network area, while accounting for 21 percent of the richly connected K-MECS network. Among design with comparable bandwidth, K-MECS represents the most area-efficient configuration.

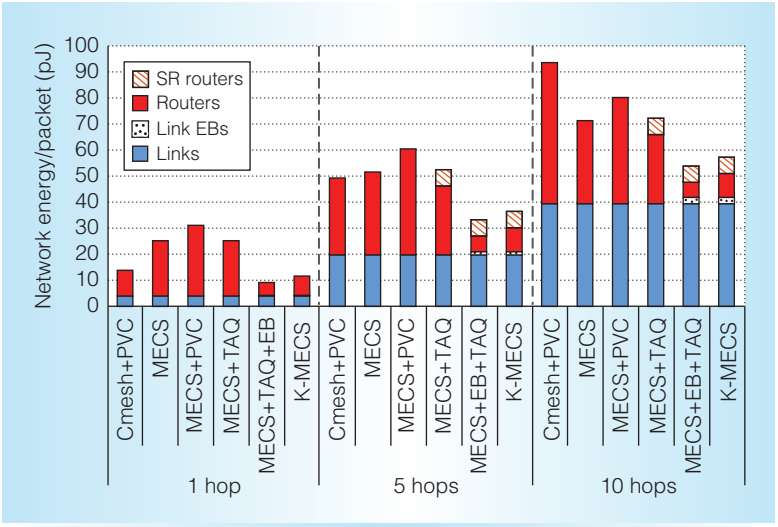

Figure 5. NoC energy breakdown as a function of the communication distance. Topology awareness improves NoC energy efficiency by reducing QoS overheads and by enabling QoS-friendly elastic buffering.

Energy analysis

Figure 5 shows network-level energy efficiency for three different access patterns: nearest neighbor (1 mesh hop), semilocal (5 mesh hops), and random (10 mesh hops). The nearest-neighbor pattern incurs one link and two router traversals in all topologies. In contrast, 5-hop and 10-hop patterns require three router accesses (which represents the worst case) in low-diameter MECS networks, while requiring six and 11 router crossings, respectively, in Cmesh. We assume that a quarter of all accesses in the multihop patterns are to shared resources, necessitating transfers to and from the shared regions in TAQ-enabled networks.

In general, EB-enabled networks have better energy efficiency than other organizations. K-MECS is the second most efficient design among the evaluated alternatives, reducing NoC energy by 16 to 63 percent on local traffic and by 20 to 40 percent on nonlocal patterns. A pure EB architecture (MECS+TAQ+EB) is 22 percent more efficient than K-MECS on local traffic and 6 to 9 percent better on nonlocal routes, due to a reduction in buffer and switch input power; however, these reductions come at greater area expense and lower throughput as compared to K-MECS.

Links are responsible for a significant fraction of overall energy expense, limiting the benefits of router energy optimizations. For instance, links account for 69 percent of energy expended on random traffic in K-MECS. PVC-enabled routers in the shared regions also diminish the energy efficiency of K-MECS and other TAQ-enabled topologies.

Results summary

Table 1 summarizes the area, power requirements, zero-load latency, and throughput (maximum sustained injection rate before the network saturates) of different topologies in a kilo-terminal network in a 15-nm technology. Power numbers are derived for a 2 GHz clock frequency and random (10-hop) traffic at average network loads of 1 and 10 percent. Latency and throughput values are also for random traffic, with 50 percent of the nodes communicating.

Our proposed topology-aware QoS optimization effectively reduces network area and power consumption without compromising performance. Compared to a baseline MECS network with PVC support (MECS+PVC), TAQ reduces network area

Table 1. Area, power, and performance characteristics of various NoC architectures.

	Area (mm^2)	Power @ 1% (W)	Power @ 10% (W)	Zero-load latency (cycles)	Throughput (%)
Cmesh+PVC	6.0	3.8	38.3	36	9
MECS	23.5	2.9	29.2	20	29
MECS+PVC	29.9	3.3	32.9	20	29
MECS+TAQ	25.1	3.0	29.6	20	29
MECS+TAQ+EB	18.2	2.2	22.2	20	24
K-MECS	16.5	2.3	23.5	20	29

by 16 percent and power consumption by 10 percent (MECS+TAQ). Furthermore, TAQ enables elastic-buffered flow control outside of the shared regions, which further reduces area by 27 percent and power consumption by 25 percent, but degrades throughput by over 17 percent (MECS+TAQ+EB). The throughput reduction is caused by a severe shortage of network buffer capacity, aggravated by the shared nature of MECS links. Finally, K-MECS combines TAQ with the hybrid EB/VC flow-control architecture, which we also propose in this work. The resulting organization restores throughput and improves area efficiency at a small power penalty when compared to a pure elastic-buffered NoC.

An important contribution of our work lies in our topology-aware approach to QoS, which represents a new direction for scalable network QoS architectures. Whereas all prior schemes have focused on minimizing per-router cost and complexity, our research suggests that router optimizations might be secondary to architectural mechanisms that reduce the need for QoS support in the first place.

Our work also points to a growing problem of NoC energy consumption in communication-intensive chips, as evidenced by the data in Table 1 (column "Power @ 10%"). As process scaling pushes the limits of on-die integration, future substrates will either restrict the extent of internode communication to save interconnect power or employ NoC architectures that push the envelope of energy efficiency. While the former approach is clearly undesirable as it limits the benefits of integration, the latter calls for significant innovation in the interconnect space. Specialization and tight integration of NoC components is one promising direction for improving the interconnect fabric's energy efficiency, and this work represents a step in that direction. MICRO

Acknowledgments

This research was supported by NSF CISE Infrastructure grant EIA-0303609 and NSF grant CCF-0811056.

References

1. J. Kim, J. Balfour, and W. Dally, "Flattened Butterfly Topology for On-Chip Networks," *Proc. 40th Int'l Symp. Microarchitecture* (Micro 40), ACM, 2007, pp. 172-182.
2. B. Grot, J. Hestness, S.W. Keckler, and O. Mutlu, "Express Cube Topologies for on-Chip Interconnects," *Proc. 15th Int'l Symp. High-Performance Computer Architecture* (HPCA 2009), IEEE CS, 2009, pp. 163-174.
3. J.W. Lee, M.C. Ng, and K. Asanović, "Globally-Synchronized Frames for Guaranteed Quality-of-Service in On-Chip Networks," *Proc. 35th Int'l Symp. Computer Architecture* (ISCA 2008), IEEE CS, 2008, pp. 89-100.
4. B. Grot, S.W. Keckler, and O. Mutlu, "Preemptive Virtual Clock: A Flexible, Efficient, and Cost-Effective QoS Scheme for Networks-on-Chip," *Proc. 42nd Int'l Symp. Microarchitecture* (Micro 42), ACM, 2009, pp. 268-279.
5. T. Moscibroda and O. Mutlu, "A Case for Bufferless Routing in On-Chip Networks," *Proc. 36th Int'l Symp. Computer Architecture* (ISCA 2009), IEEE CS, pp. 196-207.

6. A.K. Kodi, A. Sarathy, and A. Louri, ''iDEAL: Inter-Router Dual-Function Energy and Area-Efficient Links for Network-on-Chip (NoC) Architectures,'' *Proc. 35th Int'l Symp. Computer Architecture* (ISCA 2008), ACM, 2008, pp. 241-250.
7. G. Michelogiannakis, J. Balfour, and W. Dally, ''Elastic-Buffer Flow Control for On-Chip Networks,'' *Proc. 15th Int'l Symp. High-Performance Computer Architecture* (HPCA 2009), IEEE CS, 2009, pp. 151-162.
8. J.D. Balfour and W.J. Dally, ''Design Tradeoffs for Tiled CMP On-Chip Networks,'' *Proc. 23rd Int'l Conf. Supercomputing* (ICS 2006), ACM, 2006, pp. 187-198.
9. B. Grot, J. Hestness, S.W. Keckler, and O. Mutlu, ''Kilo-NOC: A Heterogeneous Network-on-Chip Architecture for Scalability and Service Guarantees,'' *Proc. 38th Int'l Symp. Computer Architecture* (ISCA 2011), IEEE CS, 2011, pp. 268-279.

Boris Grot is a postdoctoral researcher at École Polytechnique Fédérale de Lausanne. His research focuses on processor architectures, memory systems, and interconnection networks for high-throughput, energy-aware computing. Grot has a PhD in computer science from the University of Texas at Austin.

Joel Hestness is a PhD student in computer science at the University of Texas at Austin. His research interests include future highly integrated chips, on-chip networks, and communication. Hestness has a BS in computer science and in mathematics from the University of Wisconsin-—Madison.

Stephen W. Keckler is the senior director of architecture research at Nvidia and a professor in the Department of Computer Science at the University of Texas at Austin. His research interests include parallel computer architectures, memory systems, and interconnection networks. Keckler has a PhD in computer science from the Massachusetts Institute of Technology.

Onur Mutlu is an assistant professor in the Electrical and Computer Engineering Department at Carnegie Mellon University. His research interests include computer architecture, hardware/software cooperation, and memory and communication systems. Mutlu has a PhD in electrical and computer engineering from the University of Texas at Austin.

Direct questions and comments about this article to Boris Grot, EPFL IC ISIM PARSA, INJ 238 (Batiment INJ), Station 14, CH−1015, Lausanne, Switzerland; boris.grot@epfl.ch.