# HORNET: A Cycle-Level Multicore Simulator

Pengju Ren, Mieszko Lis, Myong Hyon Cho, Keun Sup Shim, Christopher W. Fletcher, Omer Khan, *Member, IEEE,* Nanning Zheng, *Fellow, IEEE,* and Srinivas Devadas, *Fellow, IEEE*

***Abstract*—We present HORNET, a parallel, highly configurable, cycle-level multicore simulator based on an ingress-queued wormhole router network-on-chip (NoC) architecture. The parallel simulation engine offers cycle-accurate as well as periodic synchronization; while preserving functional accuracy, this permits tradeoffs between perfect timing accuracy and high speed with very good accuracy. When run on six separate physical cores on a single die, speedups can exceed a factor of over 5, and when run on a two-die 12-core system with 2-way hyperthreading, speedups exceed 12×. Most hardware parameters are configurable, including memory hierarchy, interconnect geometry, bandwidth, crossbar dimensions, parameters driving power, and thermal effects. A highly parametrized table-based NoC design allows a variety of routing and virtual channel allocation algorithms out of the box, ranging from simple dimension-ordered routing to complex Valiant, ROMM, O1Turn or PROM schemes, BSOR, and adaptive routing. HORNET can run in network-only mode using synthetic traffic or traces, or directly emulate a MIPS-based multicore. HORNET is freely available under the open-source MIT license at http://csg.csail.mit.edu/hornet/.**

***Index Terms*—Multicore simulation, network-on-chip, parallel simulation.**

## I. Introduction

In recent years, architectures with several distinct CPU cores on a single die have become the standard: general-purpose processors now include as many as eight cores [1] and multicore designs with 64 or more cores are commercially available [2]. Experts predict that by the end of the decade we could have as many as 1000 cores on a single die [3].

For a multicore on this massive scale, connectivity is a major concern, and inefficient interconnects can severely limit performance. Current interconnects like buses, all-to-all point-to-point connections, and even rings clearly do not scale beyond a few cores. The relatively small scale of existing network-on-chip (NoC) interconnects has allowed plentiful on-chip bandwidth to make up for simple routing [4], but this will not last as scales grow from the $8 \times 8$ mesh of a 64-core chip to the $32 \times 32$ dimensions of a 1000-core: assuming all-to-all traffic and one flow per source/destination pair, a link in a $8 \times 8$ mesh with XY routing carries at most 128 flows, but in a $32 \times 32$ mesh, the worst link could be on the critical path of as many as 8192 flows.

These kinds of scales will expose phenomena not perceptible in current smaller-scale multicores, and future multicores will therefore require relatively high-performance on-chip networks and sophisticated routing. In such complex systems, complex interactions make real-world performance difficult to intuit, and designers have long relied on cycle-level simulations to guide algorithmic and architectural decisions; NoCs are no different. On a multicore scale, however, a cycle-level system simulator has high computation requirements, and taking advantage of the parallel execution capabilities of today's systems is critical.

With this in mind, we present HORNET, a highly configurable, cycle-level multicore simulator with support for a variety of memory hierarchies, interconnect routing and VC allocation algorithms, as well as accurate power and thermal modeling. Its multithreaded simulation engine divides the work equally among available host processor cores, and permits either cycle-accurate precision or increased performance at some accuracy cost via periodic synchronization. HORNET can be driven in network-only mode by synthetic patterns or application traces, or in full multicore mode using a built-in MIPS core simulator. Specifically, using HORNET, we:

1) show that results from small-scale NoC simulations cannot be used to guide architectural decisions on a 1000-core scale;
2) identify key factors for parallelizing NoC simulators and show how to take advantage of them for linear performance scaling as the number of host cores grows;
3) show that without cycle-level simulation, and, in particular, accurate modeling of congestion, various properties of the NoC being simulated (e.g., packet latencies) can suffer significant (e.g., 2×) errors in measurement, and that the detailed information provided by cycle-level simulation can drive architectural decisions;
4) demonstrate that end-to-end integration with a processor model is necessary for accurate modeling of application performance;
5) describe how simulated power and thermal profiles over the application's runtime and available on a per-tile granularity can drive such decisions as thermal constraint selection and sensor placement, as well as offer

Manuscript received May 20, 2011; revised September 28, 2011; accepted December 30, 2011. Date of current version May 18, 2012. This work was completed during the visit of P. Ren at the Massachusetts Institute of Technology. This paper was recommended by Associate Editor H.-H. S. Lee.

P. Ren and N. Zheng are with Xi'an Jiaotong University, Xi'an 710049, China (e-mail: pengjuren@gmail.com; nnzheng@mail.xjtu.edu.cn).

M. Lis, M. H. Cho, K. S. Shim, C. W. Fletcher, and S. Devadas are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: mieszko@csail.mit.edu; mhcho@csail.mit.edu; ksshim@csail.mit.edu; cwfletch@csail.mit.edu; devadas@csail.mit.edu).

O. Khan is with the University of Connecticut, Storrs, CT 06269 USA (e-mail: khan@uconn.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCAD.2012.2184760

0278-0070/$31.00 © 2012 IEEE

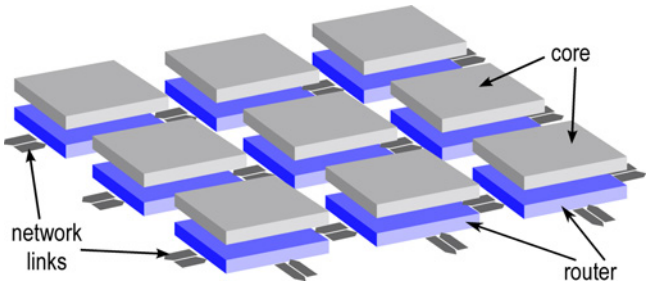


Fig. 1. Multicore system simulated by HORNET. The gray tiles (top) can be trace-driven packet injectors or cycle-level MIPS core models; the blue tiles (bottom) are cycle-level models of a flit-based virtual-channel wormhole router. While the illustration shows a 2-D mesh, HORNET can construct a system with any interconnect geometry.

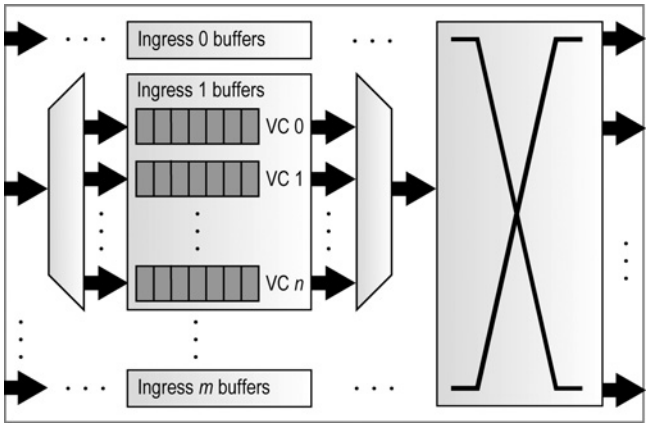


Fig. 2. Basic datapath of a NoC router modeled by HORNET. Packets arrive flit-by-flit on ingress ports and are buffered in ingress virtual channel (VC) buffers until they have been assigned a next-hop node and VC; they then compete for the crossbar and, after crossing, depart from the egress ports.

opportunities for power/thermal-aware routing algorithm design.

In the remainder of this paper, we first outline the design and features of HORNET and describe its parallelization and correctness in Section II. Next, in Section IV, we review the capabilities of HORNET and discuss speed versus accuracy tradeoffs using complete runs of selected SPLASH-2 applications [5] as well as simulations using synthetic traffic patterns. Finally, we review related research in Section V and offer concluding remarks in Section VI.

## II. Design and Features

In this section, we outline the range of systems that can be simulated by HORNET, and discuss the techniques used to parallelize simulations.

### A. Network Model

Fig. 2 illustrates the basic datapath of a NoC router modeled by HORNET. There is one ingress port and one egress port for each neighboring node, as well as for each injector (or CPU core) connected to the switch; each ingress port contains any number of virtual channel buffers (VCs), which buffer flits until they can traverse the crossbar into the next-hop node.

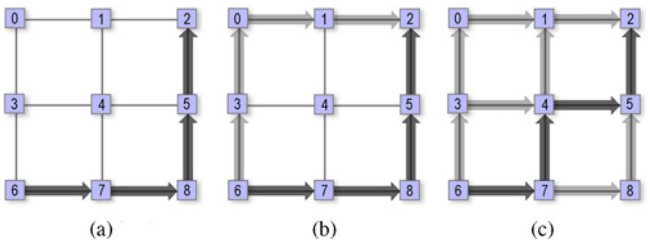


Fig. 3. Example routes for a flow between a source (node 6) and a destination (node 2) for three oblivious routing algorithms. A single path is highlighted in dark gray while other possible paths are shown in light gray. (a) XY. (b) O1TURN. (c) Two-phase ROMM.

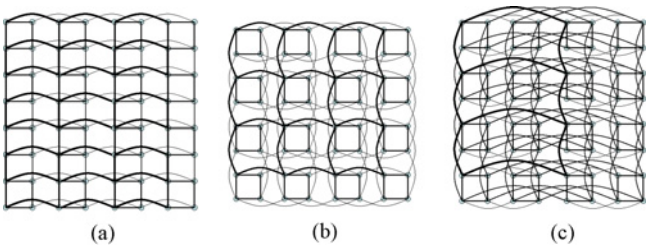


Fig. 4. Planar view of three example multilayer mesh interconnect geometries which can be directly configured in HORNET. (a) 3-D mesh–x1. (b) 3-D mesh–x1y1. (c) 3-D mesh–xcube.

As in any ingress-buffered wormhole router, packets arrive at the ingress ports flit-by-flit, and are stored in the appropriate virtual channel buffers. When the first flit of a packet arrives at the head of a VC buffer, the packet enters the route computation (RC) stage and the next-hop egress port is determined according to the routing algorithm. Next, the packet waits in the VC allocation (VA) stage until granted a next-hop virtual channel according to the chosen VC allocation scheme. Finally, in the switch arbitration (SA) stage, each flit of the packet competes for access to the crossbar and transits to the next node in the switch traversal (ST) stage. The RC and VA steps are active once per packet (to the head flit), while the SA and ST stages are applied per-flit.

1) *Interconnect Geometry:* The nodes in a system modeled by HORNET can be configured with pairwise connections to form any geometry, including rings, multilayer meshes (see Fig. 4), and tori. Each node may have as many ports as desired: for example, most nodes in the 2-D mesh shown in Fig. 1 have five ports (four facing the neighboring nodes and one facing the CPU); the number and size of virtual channels can be controlled independently for each port, allowing the CPU↔switch ports to have different VC configuration from the switch↔switch ports.

2) *Routing:* HORNET supports oblivious, static, and adaptive routing. A wide range of oblivious and static routing schemes is possible by configuring per-node routing tables. These are addressed by the flow ID and the incoming direction $\langle prev\_node\_id, flow\_id \rangle$, and each entry is a set of weighted next-hop results $\{\langle next\_node\_id, next\_flow\_id, weight \rangle, \cdots\}$. If the set contains more than one next-hop option, one is selected at random with propensity proportionate to the relevant $weight$ field, and the packet is forwarded to $next\_node\_id$ with its flow ID renamed to $next\_flow\_id$.

For example, in the case of simple XY routing, shown in Fig. 3(a), the routing tables for nodes 2, 5, 6, 7, and 8 would contain one entry for the relevant flow, addressed by the previous node ID (or 6 for the starting node 6) and the flow ID; the lookup result would direct the packet to the next node along the red path (or 2 for the terminal node 2) with the same flow ID and weight of 1.0. Static routing [6] is handled similarly. For O1TURN routing [7], illustrated in Fig. 3(b), the table at the start node (6) would contain two next-hop entries (one with next-hop node 3 and the other with next-hop node 7) weighted equally at 0.5, and the destination node (2) would have two entries (one arriving from node 1, and the other from node 5); the remaining tables do not differ from XY.

HORNET's table-driven routing directly supports probabilistic oblivious routing algorithms such as PROM [8], as well as probabilistic routing algorithms which first route the packet to a random intermediate node (say via XY routing) and only then to the final destination (e.g., Valiant [9] and its minimum-rectangle variant ROMM [10]). For example, the red path in Fig. 3(c) shows one possible route from node 6 to node 2 in a two-phase ROMM scheme: the packet is first routed to node 4 and then to its final destination. To fill the routing tables, we must solve two problems: 1) remember whether the intermediate hop has been passed; and 2) express several routes with different intermediate destinations but the same next hop as one table entry. The first problem is solved by *changing the flow ID* at the intermediate node, and renaming the flow back to its original ID once at the destination node; the second problem corresponds to sending the flow to one of two possible next-hop nodes weighted by the ratio of possible flows going each way regardless of their intermediate nodes. Consider, as an example, the routing entries at node 4 for a flow from node 6 to node 2. A packet arriving from node 3 must have passed its intermediate hop at node 4 (because otherwise XY routing to the intermediate node would have restricted it to arriving from node 7) and can only continue on to node 5 without renaming the flow. A packet arriving at node 4 from node 7 must not have passed its intermediate node (because otherwise it would be out of turns in its second XY phase and it could not get to its destination at node 2); the intermediate node can be either node 1 (with one path) or node 4 itself (also with one path), and so the table entry would direct the packet to node 1 (without flow renaming) or to node 5 (with flow renaming) with equal probability.

3) *Virtual Channel Allocation:* Like routing, virtual channel allocation (VCA) is table-driven. The VCA table lookup uses the next-hop node and flow ID computed in the route computation step, and is addressed by the four-tuple $\langle prev\_node\_id, flow\_id, next\_node\_id, next\_flow\_id \rangle$. As with table-driven routing, each lookup may result in a set of possible next-hop VCs $\{\langle next\_vc\_id, weight \rangle, \cdots\}$, and the VCA step randomly selects one VC among the possibilities according to the weights.

This directly supports dynamic VCA (all VCs are listed in the result with equal probabilities) as well as static set VCA [11] (the VC is a function of the flow ID). Most other VCA schemes used to avoid deadlock, such as that of O1TURN (where the XY and YX subroutes must be on different VCs), Valiant/ROMM (where each phase has a separate VC set), as well as various adaptive VCA schemes like the turn model [12], are easily implemented as a function of the current and next-hop flow IDs.

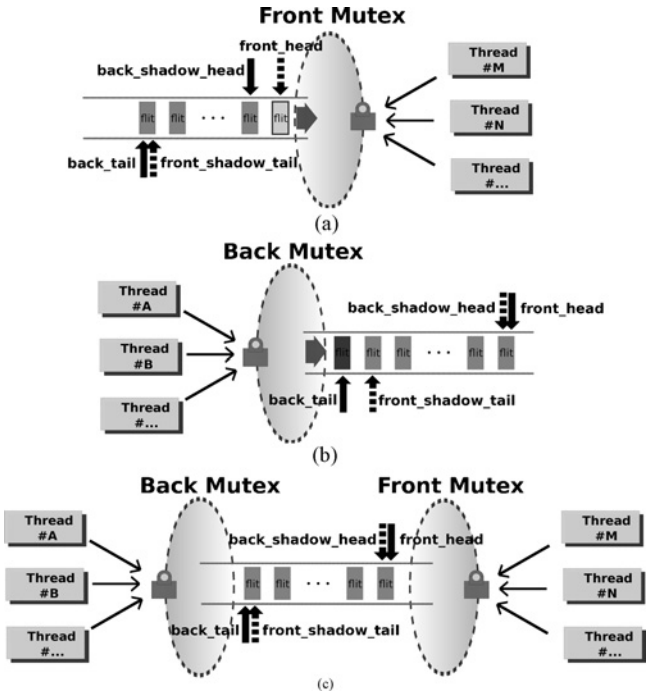


Fig. 5. Cross-thread synchronization of node-to-node data transfers in HORNET. (a) Dequeue operation: only front end locked. (b) Enqueue operation: only back end locked. (c) Negative-edge synchronization: both ends locked.

Finally, HORNET supports VCA schemes where the next-hop VC choice depends on the *contents* of the possible next-hop VCs, such as EDVCA [13] or FAA [14].

4) *Bidirectional Links:* HORNET allows inter-node connections to be bidirectional: links can optionally change direction as often as every cycle based on local traffic conditions, effectively trading off bandwidth in one direction for bandwidth in the opposite direction [15]. To achieve this, each link is associated with a modeled hardware arbiter which collects information from the two ports facing each other across the link (for example, number of packets ready to traverse the link in each direction and the available destination buffer space) and suitably sets the allowed bandwidth in each direction.

The performance of routing and VC allocation algorithms can be heavily affected by the regular nature of the synthetic traffic patterns often used for evaluation: for example, a simple round-robin VCA scheme can exhibit throughput unfairness and cause otherwise equivalent flows to experience widely different delays if the traffic pattern injects flits in sync with the round-robin period. Worse yet, a similarly biased crossbar arbitration scheme can potentially block traffic arriving from one neighbor by always selecting another ingress port for crossbar traversal.

While relatively sophisticated arbitration algorithms have been developed (e.g., *i*SLIP [16]), the limited area and power in an NoC, together with the requirement for fast line-rate decisions, restricts the complexity of arbitration schemes and, consequently, their robustness to adversarial traffic patterns.

TABLE I
*HORNET* CAN CONSTRUCT MANY ROUTERS FROM BASIC BUILDING BLOCKS LIST IN THE TABLE

| Components | Architecture Parameters | Statistic Parameter |
|---|---|---|
| Virtual buffers | $I_p$: ingresses from neighbor<br>$I_c$: ingresses from cores<br>$V_p$: neighbor ingress VCs<br>$V_c$: core ingress VCs<br>$B$: buffer size in flits<br>$F$: flit size in bits<br>$P_r$: virtual buffer read ports<br>$P_w$: virtual buffer write ports | $n_{br}$: flits read operation for each virtual buffer<br>$n_{bw}$: flits write operation for each virtual buffer |
| Crossbar | $X_C$: input ports from core<br>$X_P$: input ports from neighbor<br>$B_P$: output ports<br>$W$: port width in bits | $n_x$: number of flits go to the crossbar |
| Allocator | $A_1$: stage1 arbiter size<br>$N_1$: num of stage1 arbiters<br>$A_2$: stage2 arbiter size<br>$N_2$: num of stage2 arbiters<br>$L$ : package length in flits | $n_{r1}$: stage1 require flits<br>$n_{a1}$: stage1 active arbiters<br>$n_{r2}$: stage2 require flits<br>$n_{a2}$: stage2 active arbiters |
| LinkPower | $Link$: links for each direction<br>$W$ : bandwidth for each link<br>$Ll$: link length for each direction | $n_l$: flits transmit on the link<br>$l_a$: number of active links |
| Pipeline-reg | $N$: stages of router pipeline | $n_r$: switching register data |

Accurate statistic parameters are generated during simulation.

Instead of selecting one such algorithm, therefore, HORNET employs randomness to break arbitration ties: for example, the order in which next-in-line packets are considered for VC allocation, and the order in which waiting next-in-line flits are considered for crossbar traversal, are both randomized. While the pseudorandom number generators are by default initialized from an OS randomness source, the random seeds can be set by the user when exact reproducibility is required.

### B. Power and Thermal Modeling

Power dissipation has become a prominent constraint for architects, and is an important factor in routing algorithm selection for large-scale high-throughput on-chip networks. To enable power and thermal analysis, HORNET combines a dynamic power model based on ORION 2.0 [17] with a leakage power model; an accurate thermal model uses HOTSPOT 5.0 [18]. At runtime, various system configuration parameters (buffer sizes, port counts, and so on) and statistics (buffer reads/writes, crossbar transits) are passed to the ORION library for on-the-fly power estimation as shown in Table I and to HOTSPOT for thermal modeling: this enables not only the usual average and peak power and thermal analysis for the entire chip but also per-tile and per-time-period reporting.

### C. Concurrency, Synchronization, and Correctness

One aspect of correctness is the faithful modeling of the parallelism inherent in synchronous hardware, and applies even for single-threaded simulation. HORNET handles this by having a "positive-edge" stage (when computations and writes occur throughout the entire simulated device, but are not yet visible anywhere when read) and a separate "negative-edge" stage (when the written data are made visible) for every clock cycle; this correctly models the behavior of registered hardware where all data are captured in one fell swoop at the rising clock edge.

HORNET takes advantage of modern multicore processors by automatically distributing simulation work among the available cores; as we show in Section IV, this results in significant speedup of the simulations. The simulated system is divided into tiles comprising a single virtual channel router and any traffic generators connected to it (see Fig. 1), as well as a private pseudorandom number generator and any data structures required for collecting statistics. One execution thread is spawned for each available processor core (and restricted to run only on that core), and each tile is mapped to a thread; thus, some threads may be responsible for multiple tiles but a tile is never split across threads.

In a multithreaded simulation, therefore, *functional* correctness requires that inter-tile communication be lossless, and, in the context of a pair of tiles, in-order: this way, all flits along the same path are eventually delivered in the correct order. Because in HORNET a tile is never divided across threads, inter-thread communication is limited to flits crossing from one node to another, and some fundamentally sequential but rarely used features (such as writing VCD dumps). *Timing* accuracy additionally requires that the effects of events in a given tile on its neighbors occur promptly and in the correct order: in other words, that clock cycles be *globally* synchronized.

In the HORNET implementation, virtual channel queues (VCs, see Fig. 2) constitute the sole interface where two tiles (and therefore two threads) interact. These structures consist of a fixed-length buffer together with a head pointer and a tail pointer. Since modeling the synchronous clock-driven updates requires that flits freshly enqueued at the tail (as well as free entries caused by dequeues at the head) are not visible at the opposite end until the next clock cycle, both pointers have "shadow" copies. As shown in Fig. 5, the head reads from a shadow (and possibly stale) copy of the tail pointer to determine the queue length, and the tail reads from a shadow copy of the head pointer to determine the available free space; at the "negative" clock edge the shadow pointers are reconciled with their originals, which makes any changes visible at the next "positive" edge. The decoupled implementation also makes it straightforward to employ separate head and tail locks to preserve correctness while maximizing parallelism: dequeues only lock the head lock [Fig. 5(a)] and enqueues only lock the tail lock [Fig. 5(b)], making it possible to enqueue flits (in one thread) and dequeue flits (in another thread) concurrently; the shadow copy synchronization at the negative edge updates both shadow pointers, and must hold both locks [Fig. 5(c)].

In concurrent simulation, *timing* inaccuracies can arise from when a simulated tile instantaneously (in one clock cycle) observes a set of changes effected by another tile over several clock cycles. This has two consequences: one is that the "local" clocks in each tile may be out of sync and any values computed using both may be inaccurate; the other is that a tile may observe the effects of too many (or too few) flit arrivals and different relative flit arrivals, leading to altered local congestion conditions and therefore different flit delivery times. HORNET avoids the first problem by keeping most collected statistics with the flits being transferred and

updating them on the fly: for example, a flit's latency is updated incrementally at each node as the flit makes progress through the system, and is therefore immune to variation in the relative clock rates of different tiles.

The remaining inaccuracy is controlled by periodically synchronizing all threads on a barrier. 100% timing accuracy demands that threads be synchronized twice per clock cycle (once on the positive edge and once on the negative edge), and, indeed, simulation results in that mode precisely match those obtained from sequential simulation. Less frequent synchronizations are also possible: the synchronizing barriers then occur every few cycle instead of twice every cycle, and, as discussed in Section IV, result in significant speed benefits. Even with loose synchronization the *functional* correctness is conserved, since flits along the same route will arrive in the correct order; the only cost is some degradation in *timing* accuracy, which, as discussed in Section IV, turns out to be fairly minor.

### D. Processor Core Integration

Fig. 1 shows the system simulated by HORNET. Each tile contains a flit-based NoC router, connected to other routers via point-to-point links with any desired interconnect geometry, and, optionally, one of several possible traffic generators. These can be either trace-driven injectors or cycle-level MIPS simulators. A common bridge abstraction presents a simple packet-based interface to the injectors and cores, hiding the details of DMA transfers and dividing the packets into flits and facilitating the development of new core types.

1) *Trace-Driven Injector:* The simple trace-driven injector reads a text-format trace of the injection events: each event contains a timestamp, the flow ID, packet size, and possibly a repeat frequency (for periodic flows). The injector offers packets to the network at the appropriate times, buffering packets in an injector queue if the network cannot accept them and attempting retransmission until the packets are injected. When packets reach their destinations they are immediately discarded.

2) *MIPS Simulator:* Each tile can be configured to simulate a built-in single-cycle in-order MIPS core; the core can be loaded with statically linked binaries compiled with a MIPS cross-compiler such as GCC.

The MIPS core is connected to a configurable memory hierarchy which supports an arbitrary number of private or shared cache levels backed by a shared main memory. Memory coherence among the caches is ensured either by an implementation of the MSI cache coherence protocol or via a NUCA-style distributed shared memory with remote-access reads and stores; either option uses the configured on-chip network to communicate with main memories, directories, and other caches.

To directly support MPI-style applications, the network can also be directly exposed to the processor core via a system call interface: the program can send packets on specific flows, poll for packets waiting at the processor ingress, and receive packets from specific queues. The sending and receiving process models a DMA, freeing the processor while the packets are being sent and received.

TABLE II
SYSTEM CONFIGURATIONS USED IN SIMULATION

| Characteristic | Configuration |
|---|---|
| Topology | 32×32 2-D mesh, 16×16 2-D mesh, 8×8 2-D mesh; |
| Routing | XY, O1TURN, ROMM, Valiant |
| VC allocation | Dynamic, EDVCA |
| Link bandwidth | 1 flit/cycle |
| VCs per port | 4, 8 |
| VC buffer size | 4, 8 flits |
| Avg. packet size | 8 flits |
| Traffic workloads | Transpose, bit-complement, shuffle, H.264 decoder profile; Splash-2 traces: fft, RADIX, swaptions, water; natively executed Parsec applications: blackscholes |
| Warmup cycles | 200 000 for synthetic traffic, 0 for applications |
| Analyzed cycles | 2 000 000 for synthetic traffic, full running time for applications |
| Server CPUs used | 2× Intel Xeon X5680 6-core with HT, Intel Core i7 4-core with HT |
| # HT cores used for simulation | 1...24 |
| Sync period | clock-accurate, multiple cycles |

## III. METHODS

To support our claims with concrete examples, we ran HORNET simulations with various system configurations. The salient configuration features used in various combinations in our experiments are listed in Table II.

PARSEC benchmarks were scaled for 1024 cores and run directly on the integrated MIPS model. SPLASH-2 traces were obtained by running the benchmarks [5] in the distributed x86 multicore simulator Graphite [19] with 64 application threads; all network transmissions were logged and the traces were then replayed in HORNET. To obtain significant network congestion, the x86 core was assumed to run on a clock ten times faster than the network. This was necessary because the SPLASH benchmarks were written for a multiprocessor environment where the cost of inter-processor communication was much higher, and thus particular attention was paid to frugal communication; the plentiful bandwidth and relatively short latencies available in NoC-based multicores make this kind of optimization less critical today.

Although each simulation collects a wide variety of statistics, most reports below focus on average in-network latency of delivered traffic—that is, the number of cycles elapsed from the time a flit was injected into a network router ingress port to the time it departed the last network egress port for the destination CPU—as most relevant to current and future cache-coherent shared-memory NoC multicores. For speedups, we measured elapsed wall-clock times with HORNET as the only significant application running on the relevant server. Finally, to quantify the accuracy of the loosely synchronized simulations, we first ran HORNET with full clock-accurate synchronization to obtain a baseline; we then repeated the experiment with different synchronization periods (but the same random number seed, and so on) and compared the reported average latencies as an accuracy measurement.

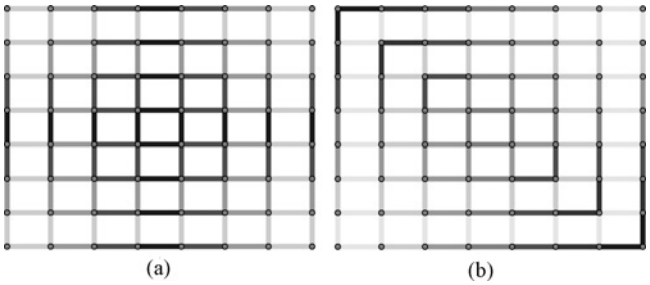


Fig. 6. Link usage for synthetic BITCOMP and TRANSPOSE traffic using the XY routing algorithm. For clarity, a 64-node $8 \times 8$ 2-D mesh is shown, and the most encumbered link in each case is shared by 8 and 7 flows, respectively; for a 1024-node $32 \times 32$ mesh, these figures grow to 32 and 31. Even in the $8 \times 8$ case there are some effects on the latency for many of these flows (see Fig. 2), but in the $32 \times 32$ case the effects are so severe that many flows cannot deliver any packets in reasonable time. (a) Bitcomplement. (b) Transpose.

To enable power and thermal analysis, we integrated HORNET with a power model based on ORION 2.0 [17] and a thermal model uses HOTSPOT 5.0 [18].

## IV. DISCUSSION

### A. Simulation Challenges for Large-Scale Multicores

Scaling multicores and their on-chip networks to thousand-core levels presents challenges that do not arise in existing systems with fewer than one hundred cores. On the one hand, there is the simple challenge of significantly more traffic concentrated on few nodes: the off-chip bandwidth grows much more slowly than on-chip transistor counts [20], and the resulting higher core-to-memory ratio will raise traffic centered around the memory controller to unprecedented levels.

On the other hand, various congestion effects present but not significantly detrimental in smaller networks are radically amplified in on-chip interconnects on a 1000-core scale. For example, while a single one-way link in an $8 \times 8$ mesh with dimension-ordered routing (DOR) might at worst be the bottleneck for 128 distinct flows (assuming all-to-all traffic and one flow per source/destination node pair), the most encumbered link in a 1024-core $32 \times 32$ mesh could be on the critical path of as many as 8192 flows.[1] Worse yet, local congestion can cause long-distance flows to experience exponentially long latencies (see Fig. 8): indeed, in long-running high-traffic simulations of a 1024-core, $32 \times 32$ mesh network, we observed that some flows delivered very few or even no packets precisely because of this effect, whereas in a 64-core, $8 \times 8$ network this was never a problem.

Fig. 6 shows the number of flows on each link for the BITCOMP and TRANSPOSE benchmarks using XY routing. For clarity, the figure shows a $8 \times 8$ mesh, where the most encumbered link (darkest) is shared by eight flows; in a 1024-node $32 \times 32$ mesh, the most encumbered link would be shared by 32 flows. This means that, in each clock cycle, packets contending for those links have a 1 : 32 chance on the average to traverse the next link, which for flows with many links to traverse grows exponentially and quickly amounts to a significant delay. Fig. 7 shows how this affects packet delivery rate and the latency the packets actually delivered: because average latencies grow exponentially with the number of hops a flow must travel, the number of packets delivered during the 2 000 000 simulated cycles drops off rapidly, and at some point some flows receive no packets at all before simulation ends. The same figures show that a different virtual

[1]The number of flows on the most encumbered link for DOR on a $n \times n$ mesh is $\frac{n^3}{4}$.

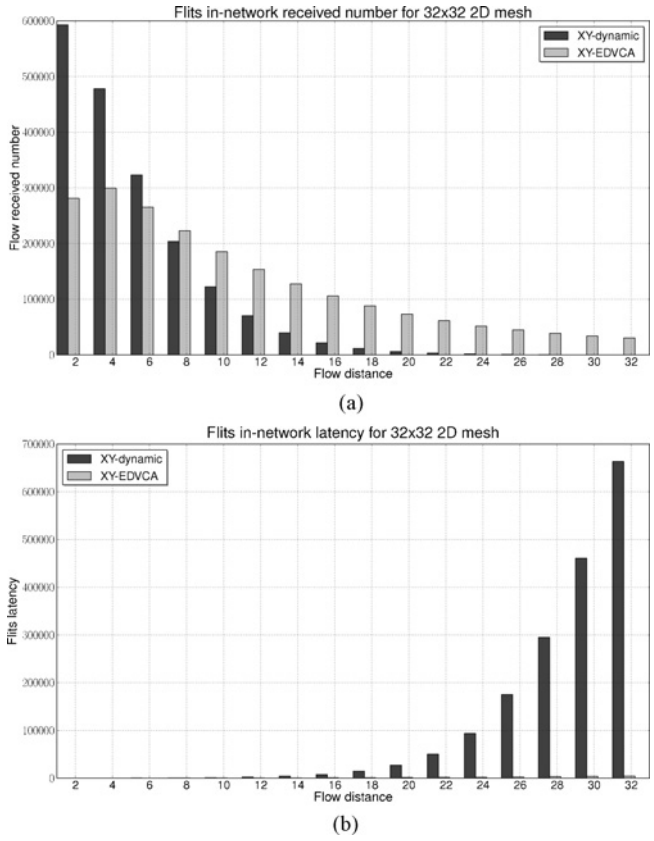


Fig. 7. In large meshes, the XY routing algorithm (shown in black) fails to scale because per-hop fairness leads to (a) decreasing packet delivery rates and (b) exponentially increasing delays for longer flows (see Fig. 8); indeed, after the number of hops that must be traversed exceeds about 32, there only tens of flits received, when distance up to 48, some delays exceed the simulation time and some flows never receive any packets. Routing algorithms like EDVCA (shown in gray) mitigate this effect and are more appropriate for large-scale NoCs. To observe this effect, accurate simulation of NoC router buffers and contention is necessary. (2 000 000 cycles on a $32 \times 32$ 2-D mesh using the synthetic bit-complement benchmark). (a) Packet delivery rate in flows as a function of the distance packets must travel. (b) Average packet latency in flows as a function of the distance packets must travel.

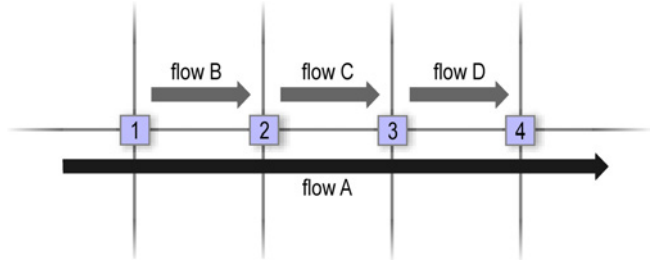


Fig. 8. In a heavily loaded network, traffic on long-path routes can suffer significantly more latency than those on short routes. In this case, flow *A* must compete for the link with a short flow (*B*, *C*, and *D*) at *every step* of its route; assuming locally fair arbitration between any two flows, this effect can result in delays for *A* that are exponential in its path length.

channel allocation scheme (EDVCA [13]), which combines the benefits of dynamic and static virtual channel allocation schemes, ameliorates head-of-line blocking, significantly improving throughput compared with equivalent baseline, and show its great advantage for large-scale networks.

Clearly, extrapolating architectural decisions for large-scale on-chip networks from small-scale simulations runs severe risks of missing significant performance bottlenecks, and accurate simulation of large networks is a necessary step in the design process.

### B. Parallelization and Performance

Since large-scale designs must be simulated directly, scalability is a key consideration in simulator design. One possibility is to abstract away detail and give up cycle-level simulation, but as discussed is Section IV-C below, this is undesirable for on-chip network design. Another is to implement the system on an FPGA directly or via a time-multiplexing system like HaSIM [21]; while those approaches offer excellent performance, they require a very low-level, time-consuming design and verification process. In HORNET, we instead take advantage of today's commodity processors featuring several cores on a single die: a single simulation can be split into multiple threads running in parallel on as many cores as are available.

Efficient parallel simulation requires designing the simulator for concurrency from ground up. The key factor limiting performance is inter-thread communication, which can be divided into: 1) communication within the simulated network itself; 2) synchronization barriers for clock-accurate results; and 3) any shared data structures. As discussed in Section II-C, HORNET threads share no structures other than the synchronized queues that carry traffic among the simulated routers, so communication only involves 1) and 2).

Clock synchronization (twice per cycle in cycle-accurate mode) incurs the most performance cost because all threads must wait on the same barrier. While this is inexpensive and allows linear scaling when all cores are on the same die, barrier communication across separate processor dies becomes time-consuming and limits performance (see Fig. 9), especially when systems with few cores are being simulated and thread workload per simulated cycle is relatively light. To allow further speedup, HORNET allows barrier synchronization to be performed periodically instead. From a functional correctness standpoint, this makes no difference, since all traffic will still arrive subject to the original ordering constraints and any deterministic algorithm running on the CPU cores will have the same results. The loose synchronization does imply some loss of fidelity in reported timing, but there HORNET ensures high accuracy by accumulating statistics separately in each thread, carrying measurements within each transmitted packet, and never basing measurements on relative values from two cores. As a result, timing measurements retain near very high fidelity even when scaling across separate dies and hyperthreaded cores (Fig. 9).

While communication due to simulated network traffic is unavoidable, HORNET employs fine-grained locking to ensure maximum parallelism. The virtual channel buffers—the only communication points between any two tiles—have front and back locks which can be separately held by different threads: this allows HORNET to ensure that results from cycle-accurate parallel simulations are *identical* to those from an equivalent single-thread simulation (given the same randomness seeds), and that intertile communication does not limit performance (see Fig. 9).

To further improve performance in the network-only and application trace configurations, HORNET can fast-forward the clocks in each tile when there are no flits buffered in the network and no flits about to be injected for some period of time. Because in that situation no useful work can possibly result, HORNET advances the clocks to the next injection event and continues cycle-by-cycle simulation from that point without altering simulation results. Clearly, heavy traffic loads will not benefit from fast-forwarding because the network buffers are never drained and HORNET never advances the clock by more than one cycle. Fig. 10 shows that the benefit on low-volume traffic depends intimately on the traffic pattern: an application which, like BLACKSCHOLES and FFT, has long pauses between traffic bursts, will benefit significantly; an application which spreads the small amount of traffic it generates evenly over time like the H.264 profile will rarely allow the network to fully drain and therefore will benefit little from fast-forwarding.

### C. Congestion and Cycle-Level Simulation

High-level architectural simulators tend to assume an idealized interconnect network and generally either do not consider congestion or approximate it with an analytical model. For interconnect network design itself, however, congestion effects are of prime importance, as they dictate, for example, what routing algorithms should be employed. To estimate the effect of congestion, we performed simulations of the SPLASH benchmark suite in the congestion-accurate configuration and in a congestion-oblivious configuration where injection bandwidth was limited as in the accurate model but the transit latencies were simple hop-counts. As Fig. 12 shows, ignoring congestion effects can cause the simulation to significantly underestimate simulation-time measurements: depending on the amount of network traffic generated by the benchmark, the effect ranged from $2\times$ to negligible.

When congestion must be modeled accurately, cycle-accurate simulation is indispensable. For example, network congestion can have significant effects on how the network configuration—say the number and size of the virtual channels—affects in-network latency (i.e., the latency incurred after the relevant flit is seen by the processor as successfully sent). Intuitively, adding more virtual channels should generally allow more packets from different flows to compete for transmission across the crossbar, increase crossbar efficiency, and therefore reduce observed packet latency. While this holds when traffic is light, it may have an opposite effect in a relatively congested network: as Fig. 13 illustrates on five SPLASH-2 applications, doubling the number of VCs while holding the size of each VC constant causes the observed in-network latency in a relatively congested network to actually *increase*. This is because the total amount of buffer space in the network also doubled, and, when traffic is heavy and delivery

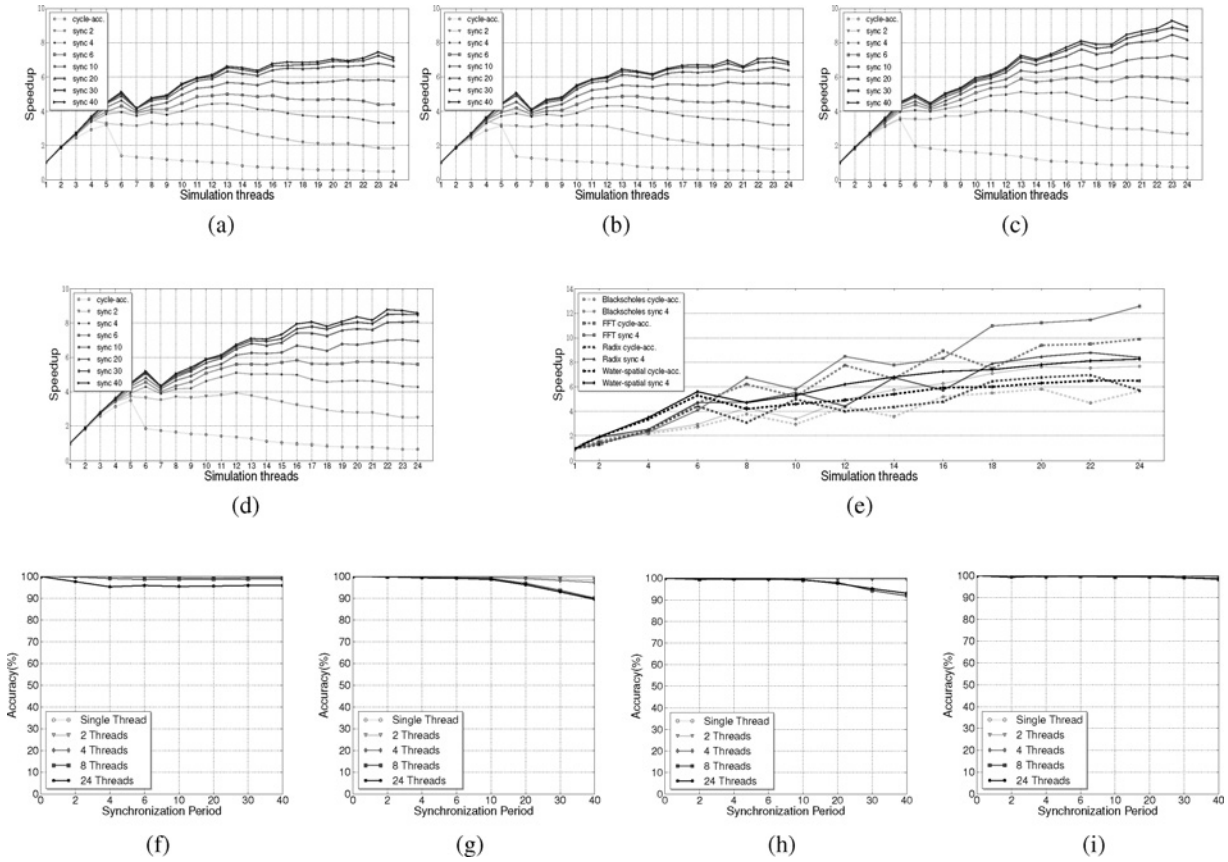


Fig. 9. Parallelization speedup for cycle-accurate and loosely-synchronized simulations of 64-core and 1024-core trace simulations of the Parsec BLACKSCHOLES benchmark and the SPLASH-2 banchmarks FFT, RADIX, and WATER-SPATIAL on a system with 24 hyperthreaded cores. The figures show that simulation speeds scale linearly up to the six physical cores on the same die (4–5 × speedup depending on benchmark and synchronization level); the overheads of pairing up threads using hyperthreading (7–12 threads) and communication between two separate dies (13+ threads) result in slightly slower speedup gains. Because each thread has more work to do between synchronizations for the 1024-core simulations, speedups are high even for cycle-accurate simulations (sync 0); for 64-core simulations, increased synchronization periods help speed up simulations significantly, especially when the simulation uses many threads. At the same time, simulation fidelity (as measured by average packet latency deviation from cycle-accurate simulation) is high even for loosely synchronized simulations. (For all the simulations, we used a 12× Intel Xeon X5680 at 3.33 GHz, on two 6-core dies, each hyperthreaded 2-way for a total of 24 cores (cores 1–12 on core 0, with 7–12 virtual; cores 13–24 on core 1 with 19–24 virtual). (a) BLACKSCHOLES 64-core speedup. (b) FFT 64-core speedup. (c) RADIX 64-core speedup. (d) WATER-SPATIAL 64 core speedup. (e) SPLASH-2 1024-core speedup. (f) BLACKSCHOLES 1024-core accuracy. (g) FFT 1024-core accuracy. (h) RADIX 1024-core accuracy. (i) WATER-SPATIAL 1024-core accuracy.

rates are limited by the network bandwidth, the flits at the tails of the VC queues must compete with flits from more VCs and thus experience longer delays. Indeed, when the VC queue sizes are halved to keep the total amount of buffer space the same, the 4-VC setup exhibits shorter latencies than the 2-VC equivalent as originally expected.

While in a lightly loaded network almost any routing and VC allocation algorithm will perform well, heavier loads lead to different congestion under different routing and VC algorithms and performance is significantly affected; again, accurately evaluating such effects calls for a cycle-level simulator and real applications. Fig. 11 shows the effect of routing and VC allocation scheme on performance of the SPLASH-2 benchmarks in a relatively congested network. While the algorithms with more path diversity (O1TURN and ROMM) do lower observed in-network latency, the performance increase is not as much as might be expected by considering the increased bandwidth available to each flow.

Modern multicore designs can reduce on-chip network congestion by placing several independent memory controllers in different parts of the network. Since in a cache-coherent system a memory controller generally communicates with all processor cores, modeling congestion is critical in evaluating the tradeoff between adding memory controllers and controlling chip area and pin usage. For example, Fig. 11 shows in-network latency for two cache-coherent systems: one with one memory controller (MC) and the other with five. While the 5-MC configuration improves performance across the board, the improvement varies across applications, from an order of magnitude (e.g., FFT, SWAPTIONS) to negligible (e.g., in the case of BLACKSCHOLES, which has relatively little memory traffic). Such detailed simulations enable the chip architect to consider the exact tradeoff between the increased area and performance improvement. More significantly, the two choices impose different constraints on selecting the routing and VC allocation logic: while the congestion around a central memory controller makes controlling congestion via routing and VC allocation, the average latency in a system with five memory controllers does not vary significantly under different routing algorithms and EDVCA, and the designer might choose to save area and reduce implementation complexity in the network switch.

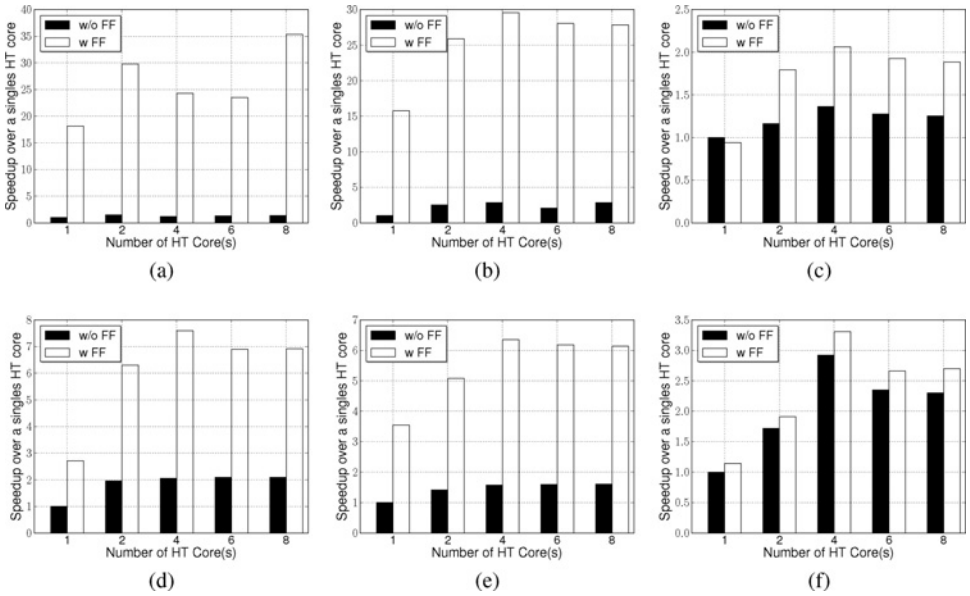


Fig. 10. Performance benefits from fast-forwarding. Unlike BLACKSCHOLES, FFT, RADIX, SWAPTIONS, and H.264 see significant speedups when the network is idle, the low-traffic H.264 profile gains little because packets are sent with relatively constant frequency and the network is rarely fully drained. (a) BLACKSCHOLES. (b) FFT. (c) RADIX. (d) SWAPTIONS. (e) WATER-SPATIAL. (f) H.264.

### D. Processor Model Integration

Much of the research in network-on-chip microarchitecture relies on synthetic traffic patterns or application traces collected under the assumption of an ideal interconnect network. This approach generally ignores the interdependencies among the various flows and the delays caused by instructions that must wait until network packets are delivered (for example, memory accesses that miss the per-core cache hierarchy). For these reasons, a precise evaluation of the performance of NoC-based multicores in running real applications requires that the CPU core and the network be simulated together.

To quantify the differences between trace-driven and real application traffic, we implemented Cannon's algorithm for matrix multiplication [22] in C using message-passing and targeting the MIPS core simulator that ships with HORNET. We ran the simulation on 64 cores and applied it to a $128 \times 128$ matrix; to stress the network, cores were mapped randomly, per-cell data sizes were assumed to be large, and computations were taken to be relatively fast. For the trace version, we assumed an ideal single-cycle network, logged each network transmission event, and later replayed the traces in HORNET; for the combined core+network version, we ran the benchmark with the MIPS cores simulated by HORNET directly interacting with the on-chip network.

The results, shown in Fig. 14, illustrate that the processor cores may have to spend significant amounts of time waiting for the network. On the one hand, a destination node waiting for a packet may block until the packet arrives. On the other hand, the sending node may have to wait for the destination core to make progress: when the destination is nearby (e.g., adjacent), even a relatively short packet can exceed the total buffer space available in the network, and the sending core may have to stall before starting the following packet until the current packet has been at least somewhat processed by the destination core and network buffers have freed up.

### E. Power Analysis

As on-chip network complexity and size increases, concerns about cooling, thermal budgets, and battery lifetime are at the forefront of NoC design [23]. Achieving good power performance requires not only optimizing specific system components using techniques such as clock and voltage gating, but also choosing the right routing algorithms, virtual channel configuration, and other network parameters as appropriate for the expected traffic patterns. In this section, we briefly discuss how HORNET helps evaluate power and thermal as part of the design process.

1) *Power Efficiency and Routing/VC Allocation Algorithm Choice:* While common wisdom approximates power dissipation in high-bandwidth networks as directly related to the amount of traffic, in reality routing algorithms and architectural considerations are key. Fig. 15(a) shows that the choice of routing algorithm can have significantly more effect on power efficiency than the amount of traffic sent: the overall energy-delay product per flit (EDPPF) for less power-efficient routing schemes can be many times greater than for the most efficient one, BSOR static routing.

2) *Power Efficiency and Router Microarchitecture Choice:* Power-efficiency effects of microarchitectural decisions also depend on the traffic pattern and routing algorithm choice, and must be carefully evaluated via simulation at design time. For example, intuition suggests that a router with more virtual channels will burn more power; on the other hand, however, adding more VCs might make routing more efficient and cut the total time required to deliver all of the packets, resulting in overall power savings. As Fig. 15(b) shows, the most efficient configuration depends on the traffic pattern and the routing algorithm used: for most combinations we tested, a 4-VC router was most power-efficient, but for a significant fraction a 2-VC setup was the best choice. Similarly, given the same number of VCs, a design can use a full crossbar connecting all

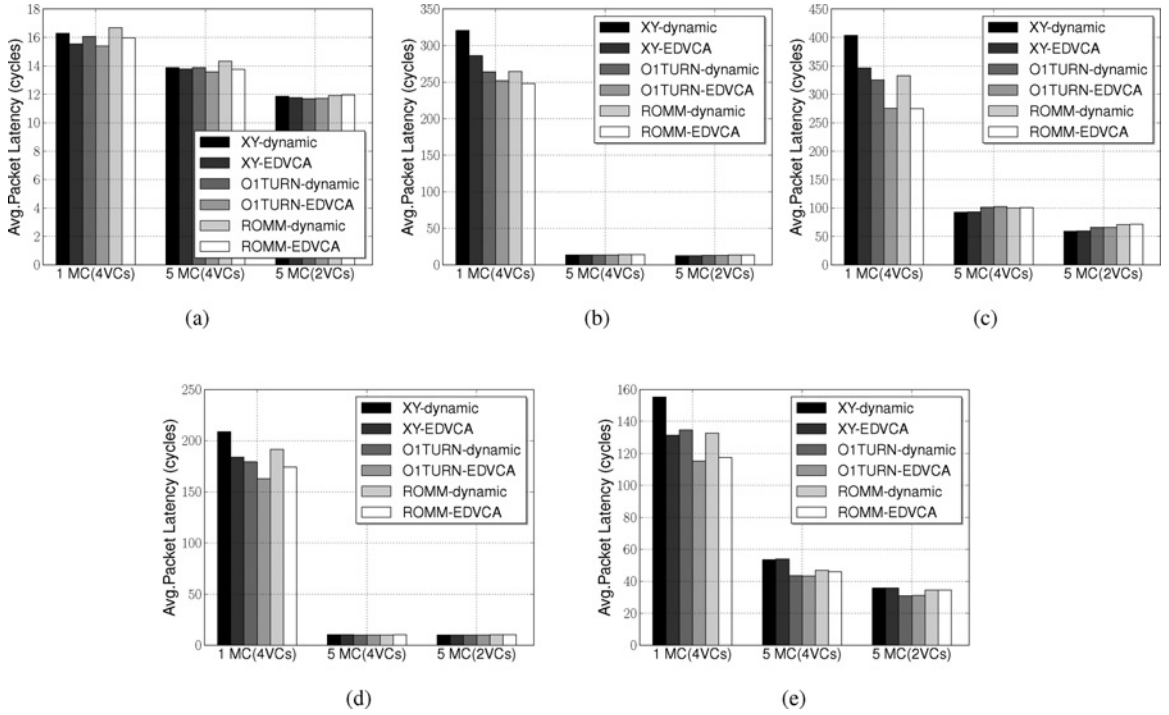


Fig. 11. Effect of routing, VC configuration, and varying the number of memory controllers on in-network latency (and therefore memory system performance) in a 64-core system running traces from BLACKSCHOLES, FFT, RADIX, SWAPTIONS, and WATER-SPATIAL benchmarks. Multiple memory controllers significantly reduce congestion, and the resulting improvements vary depending on the memory traffic patterns of each application. (a) BLACKSCHOLES. (b) FFT. (c) RADIX. (d) SWAPTIONS. (e) WATER-SPATIAL.

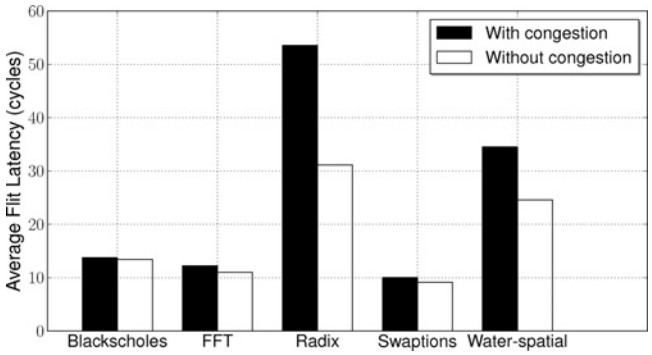


Fig. 12. Effect of congestion on flit latency for five SPLASH applications: for RADIX and WATER-SPATIAL, which both generate a lot of network traffic, not modeling congestion results in nearly 2× and 1.5× network latency underestimate; for BLACKSCHOLES, FFT and SWAPTIONS, which both have much less traffic, the difference, although present, is not significant. The results for the remaining SPLASH applications were similar: the congestion effect for high-traffic applications was similar to RADIX and for low-traffic applications resembled SWAPTIONS. (64-core system with 4 VCs, buffer depth is 4 flits).

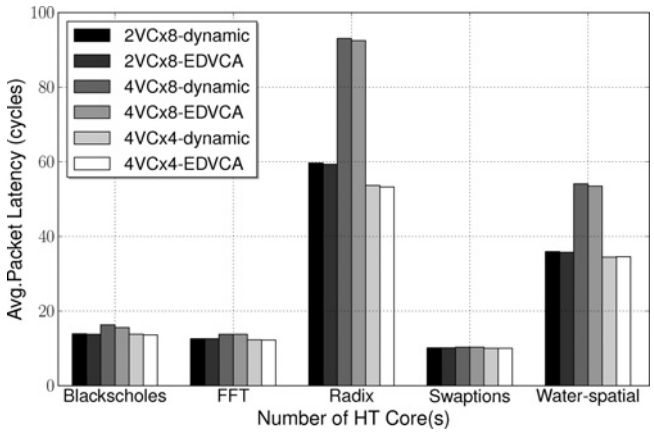


Fig. 13. In-network latency for different VC buffer configurations. Counter-intuitively, for heavy traffic loads like RADIX and WATER-SPATIAL, increasing the number of VCs from 2 to 4 while keeping the VC sizes constant at 8 flits actually *increases* in-network latency because packets can be buffered inside the network. When total VC memory size is held constant, doubling the number of VCs to 4 (and correspondingly halving their capacities to 4 flits) decreases latency as expected. for light traffic loads like the others, there are not so much congestion requirement for more virtual channels.

ingress VCs to all next-hop VCs, or can use per-port muxes to select one candidate among the VCs before entering a much smaller crossbar (see Fig. 2); as Fig. 15(c) illustrates, smaller crossbars tend to be more power-efficient for most routing schemes, but for some algorithms (e.g., O1turn) enlarging the crossbar results in shorter run times and net power savings.

### *F. Thermal Effects*

While processor core and cache thermal effects have been extensively studied, available interconnect network models report only steady-state averages for the entire chip. Fig. 16 shows that choosing thermal constraints based merely on average or peak temperature data can be misleading: for applications in which network load varies significantly over time, basing interconnect design decisions on the mean values runs the risk of thermal runaways when the application enters a heavy-traffic phase, while using worst-case peak values may result in over-provisioned, expensive thermal packaging.

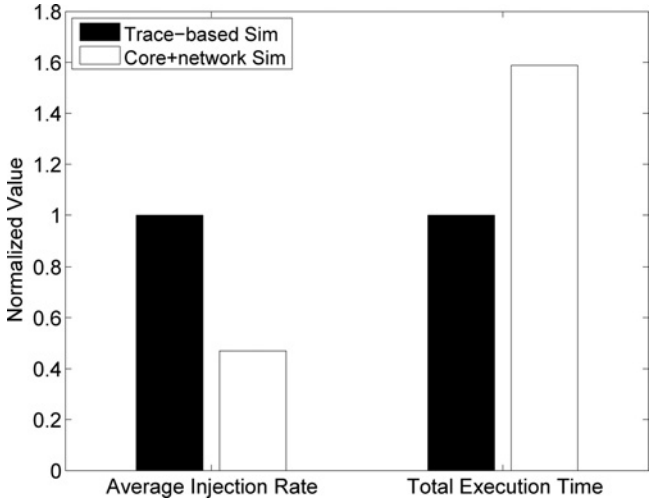


Fig. 14. Trace-based simulation lacks the feedback loop from the network to the sending (or receiving) core; this allows cores to inject packets unrealistically fast and permits the application to finish much earlier than realistically possible.

Instead, the designer might choose a design point based on the temperature profiles of the target applications, and ensure that application execution—and hence network traffic—are throttled when temperature rises above some maximum.

Such throttling requires attaching thermal sensors to the die itself; although placing more sensors on the die would provide a more accurate thermal picture, the sensors themselves are relatively expensive and power-hungry, and generally very few are present on a chip. We reasoned that, since our SPLASH runs were done with one memory controller in one corner of the mesh, the switches bordering might become a thermal hotspot and the memory controller would be a good place for a sensor. As illustrated in Fig. 17, however, the thermal hotspot in our simulations varied in magnitude but remained in the center of the chip regardless of the benchmark and routing algorithm: this is because the XY routing algorithm we used (and, indeed, nearly all available algorithms) route a greater proportion of the traffic via the central region of the mesh. This result suggests that placing a sensor in the central area of the die should suffice.

The availability of time-resolved and space-resolved thermal measurements within HORNET allows us to investigate routing algorithms which can reroute traffic on possibly longer paths (e.g., near the edges of the mesh) instead of throttling down performance when temperature rises; the development of such power-adaptive routing schemes remains an interesting topic of future research.

## V. RELATED WORK

One NoC simulator that stands out among the many simple, limited-purpose software NoC simulators is Garnet [24]. Like HORNET, Garnet models an NoC interconnection network at the cycle-accurate level: the model allows either a standard ingress-queued virtual channel router with a rigid five-stage pipeline or a flexible egress-queued router. Integration with GEMS provides a full-system simulation framework and a

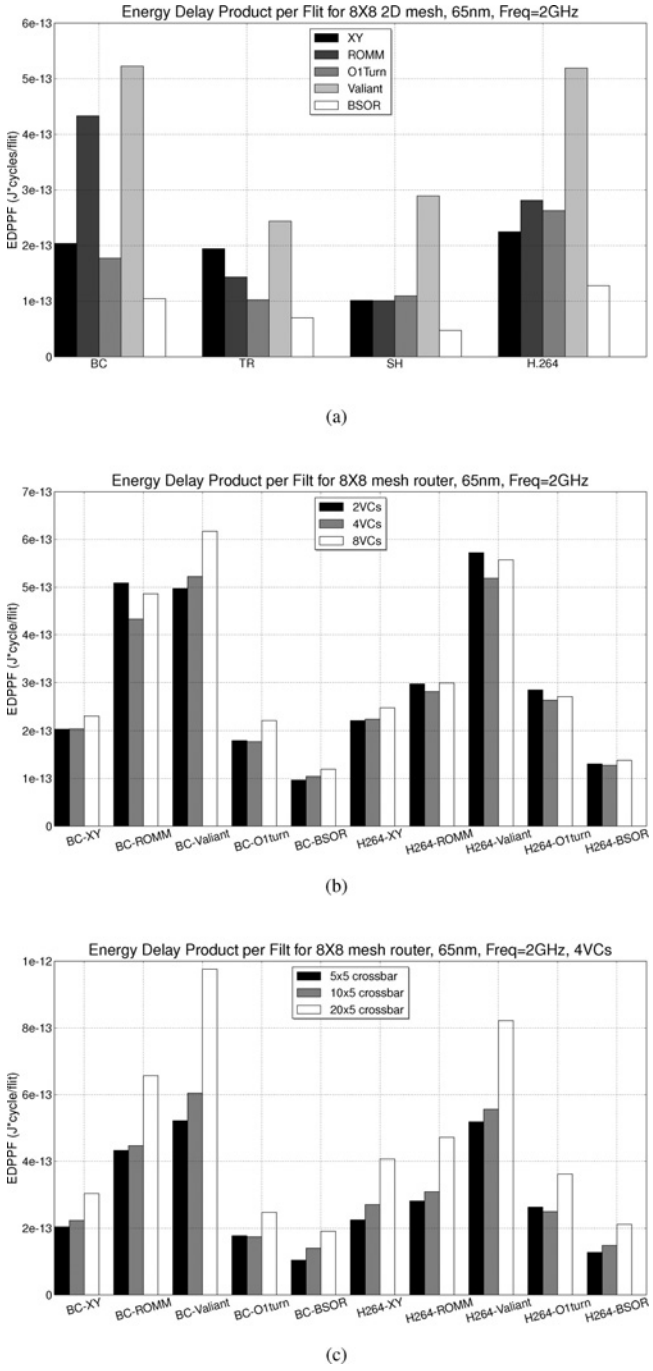


Fig. 15. (a) Routing algorithm choice has significantly more effect on power efficiency than the amount of traffic (100 000 flits were sent and delivered in each case) or architectural choices such as amount of (b) buffering, or (c) crossbar size. Nevertheless, the most power-efficient architectural choices depend on the combination of (b) traffic patterns and routing: for some loads a 2-VC configuration is the most power-efficient, while for others a 4-VC setup is best. The effects of crossbar size also depend on the (c) traffic pattern and routing scheme: while generally a smaller crossbar is more power-efficient, in some cases (e.g., H.264-O1turn) the extra VCs significantly cut the total time required to deliver the traffic and outweigh the additional per-cycle cost of a larger crossbar. ($8 \times 8$ 2-D mesh under synthetic traffic patterns: BITCOMP, TRANSPOSE, SHUFFLE, and parallel H.264 traffic profile, virtual channel buffer size is 8 flits). (a) EDPPF for different routing algorithm with same configuration. (b) EDPPF for different number of VCs. (c) EDPPF for different crossbar sizes, all using 4 VCs.

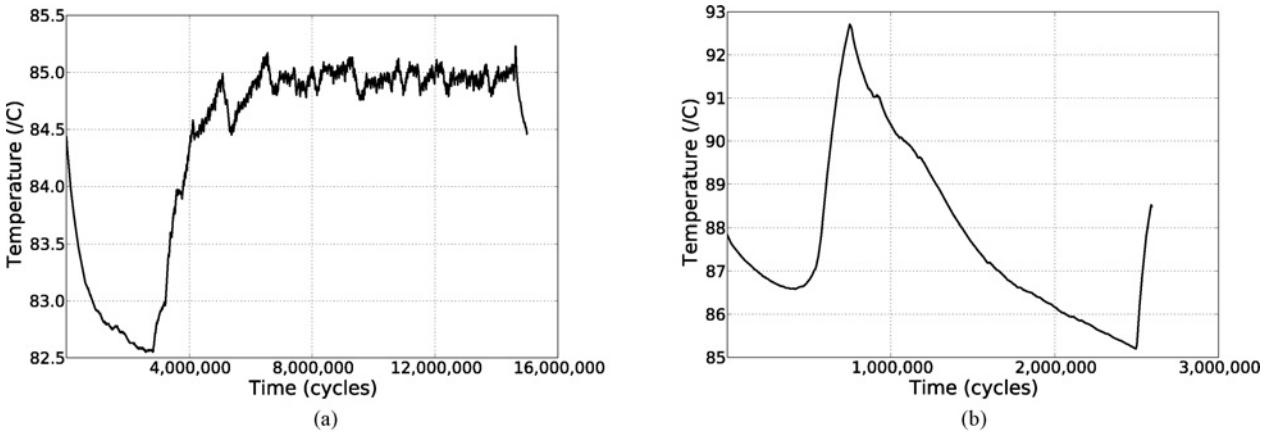


Fig. 16. Temperature traces over the runtime of different SPLASH applications. While for OCEAN a peak (or, indeed, mean) temperature estimate might be used to choose thermal constraints, the activity-dependent temperature variation in radix RADIX means that neither the mean nor the peak provides the best architectural tradeoff. (a) OCEAN. (b) RADIX.

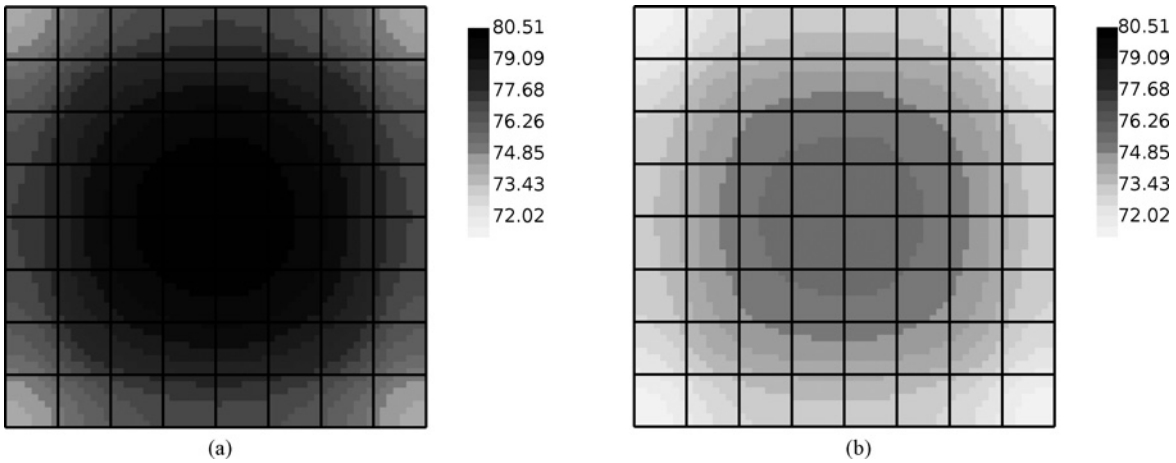


Fig. 17. Steady-state temperature distribution over a $8 \times 8$ mesh NoC for two SPLASH applications. While the overall magnitude varies significantly (in this case by over 5 °C), the overall distribution remains the same: even though the memory controller is located in the lower-left corner, the central nodes suffer the highest temperatures (the remaining SPLASH benchmarks and routing algorithms other than XY show similar temperature profiles). (a) RADIX. (b) WATER-SPATIAL.

memory model, while integration with ORION [25] provides power estimation. Table III compares Hornet with Garnet. RSIM [26] simulates shared-memory multiprocessors and uniprocessors designed for high instruction-level parallelism; it includes a multiprocessor coherence protocol and interconnect, and models contention at all resources. SICOSYS [27] is a general-purpose interconnection network simulator that captures essential details of low-level simulation, and has been integrated in RSIM. Noxim [28] models a mesh NoC and, like HORNET, allows the user to customize a variety of parameters like network size, VC sizes, packet size distribution, routing scheme, and so on; unlike HORNET, however, it is limited to 2-D mesh interconnects and is traffic-pattern-driven rather than integrated with a processor frontend. Booksim [29] allows for more network geometries but is also driven by synthetic traffic patterns. None of these simulators significantly exploits available multicore parallelism.

Highly configurable, parallelized architectural modeling is not a new idea. The Simplescalar toolset [30] can model a variety of processor architectures and memory hierarchies, and enjoys considerable popularity among computer architecture researchers. Graphite [19] is a Pin-based multicore simulator that stands out for its ability to model thousands of cores by dividing the work among not just multiple cores on the same die but multiple networked computers; it does not, however, interface with a cycle-level network model and its latency and congestion models are probabilistic. Finally, the growth in complexity and the need for ever-increasing amounts of verification has led to the development of FPGA-based

TABLE III
COMPARISON OF HORNET AND GARNET

| Features | GARNET | HORNET |
|---|---|---|
| Simulation Model | Integrated With GEMS Traffic Traces | With Built-in MIPS Traffic Traces Graphite Based Traces |
| Cycle-level accuracy | Yes | Yes |
| Scalability via paralellization | No | Yes |
| Configurable microarchitecture | Yes | Yes |
| Configurable routing | Yes | Yes |
| Power model | Yes | Yes |
| Thermal model | No | Yes |

simulators like HaSIM [21] and FPGA-level emulator platforms like RAMP [31], which, though far more difficult to configure, are much faster than software solutions.

## VI. Conclusion

We introduced HORNET, a highly configurable, cycle-accurate network-on-chip simulator that can be driven by network traces, or with a built-in MIPS simulator. HORNET's parallelized simulation engine can scale nearly linearly with the number of physical cores in the processor while preserving cycle-accurate behavior, and allows the user to obtain even more speed via loose synchronization, which preserves correctness but can introduce some inaccuracy in performance measurements.

## References


[1] S. Rusu, S. Tam, H. Muljono, D. Ayers, J. Chang, R. Varada, M. Ratta, and S. Vora, "A 45 nm 8-core enterprise Xeon® processor," in *Proc. IEEE Asian Solid-State Circuits Conf.*, Nov. 2009, pp. 9–12.

[2] S. Bell, B. Edwards, J. Amann, R. Conlin, K. Joyce, V. Leung, J. MacKay, M. Reif, L. Bao, J. Brown, M. Mattina, C. Miao, C. Ramey, D. Wentzlaff, W. Anderson, E. Berger, N. Fairbanks, D. Khan, F. Montenegro, J. Stickney, and J. Zook, "TILE64-processor: A 64-core SoC with mesh interconnect," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2008, pp. 88–598.

[3] S. Borkar, "Thousand core chips: A technology perspective," in *Proc. Des. Automat. Conf.*, 2007, pp. 746–749.

[4] D. Wentzlaff, P. Griffin, H. Hoffmann, L. Bao, B. Edwards, C. Ramey, M. Mattina, C. C. Miao, J. Brown, and A. Agarwal, "On-chip interconnection architecture of the tile processor," *IEEE Micro*, vol. 27, no. 5, pp. 15–31, Sep.–Oct. 2007.

[5] S. Woo, M. Ohara, E. Torrie, J. Singh, and A. Gupta, "The SPLASH-2 programs: Characterization and methodological considerations," in *Proc. Int. Symp. Comput. Architect.*, 1995, pp. 24–36.

[6] M. A. Kinsy, M. H. Cho, T. Wen, E. Suh, M. van Dijk, and S. Devadas, "Application-aware deadlock-free oblivious routing," in *Proc. Int. Symp. Comput. Architect.*, 2009, pp. 208–219.

[7] D. Seo, A. Ali, W. Lim, and N. Rafique, "Near-optimal worst-case throughput routing for two-dimensional mesh networks," in *Proc. Int. Symp. Comput. Architect.*, 2005, pp. 432–443.

[8] M. H. Cho, M. Lis, K. S. Shim, M. A. Kinsy, and S. Devadas, "Path-based, randomized, oblivious, minimal routing," in *Proc. Int. Workshop Netw. Chip Architect.*, 2009, pp. 23–28.

[9] L. G. Valiant and G. J. Brebner, "Universal schemes for parallel communication," in *Proc. ACM Symp. Theory Comput.*, 1981, pp. 263–277.

[10] T. Nesson and S. L. Johnsson, "ROMM routing: A class of efficient minimal routing algorithms," in *Proc. Int. Workshop Parallel Comput. Routing Commun.*, 1994, pp. 185–199.

[11] K. S. Shim, M. H. Cho, M. A. Kinsy, T. Wen, M. Lis, E. Suh, and S. Devadas, "Static virtual channel allocation in oblivious routing," in *Proc. Int. Symp. Netw. Chip*, 2009, pp. 38–43.

[12] C. J. Glass and L. M. Ni, "The turn model for adaptive routing," in *Proc. Int. Symp. Comput. Architect.*, 1992, pp. 278–287.

[13] M. Lis, K. S. Shim, M. H. Cho, and S. Devadas, "Guaranteed in-order packet delivery using exclusive dynamic virtual channel allocation," Comput. Sci. Artif. Intell. Lab., Massachusetts Inst. Technol., Cambridge, MA, Tech. Rep. MIT-CSAIL-TR-2009-036, 2009.

[14] A. Banerjee and S. Moore, "Flow-aware allocation for on-chip networks," in *Proc. Int. Symp. Netw. Chip*, 2009, pp. 183–192.

[15] M. H. Cho, M. Lis, K. S. Shim, M. A. Kinsy, T. Wen, and S. Devadas, "Oblivious routing in on-chip bandwidth-adaptive networks," in *Proc. Int. Conf. Parallel Architect. Compilation Tech.*, 2009, pp. 181–190.

[16] N. McKeown, "The iSLIP scheduling algorithm for input-queued switches," *IEEE/ACM Trans. Netw.*, vol. 7, no. 2, pp. 188–201, Apr. 1999.

[17] A. Kahng, B. Li, L.-S. Peh, and K. Samadi, "Orion2.0: A fast and accurate NoC power and area model for early-stage design space exploration," in *Proc. Des. Automat. Test Eur. Conf. Exhibit.*, 2009, pp. 1530–1591.

[18] K. Skandron, M. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, "Temperature-aware microarchitecture," in *Proc. Int. Symp. Comput. Architect.*, 2003, pp. 2–13.

[19] J. E. Miller, H. Kasture, G. Kurian, C. Gruenwald, N. Beckmann, C. Celio, J. Eastep, and A. Agarwal, "Graphite: A distributed parallel simulator for multicores," in *Proc. Int. Symp. High Performance Comput. Architect.*, 2010, pp. 1–12.

[20] *Assembly and Packaging*, International Technology Roadmap for Semiconductors, 2007.

[21] M. Pellauer, M. Vijayaraghavan, M. Adler, Arvind, and J. Emer, "A-port networks: Preserving the timed behavior of synchronous systems for modeling on FPGAs," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 2, no. 3, pp. 1–26, 2009.

[22] L. E. Cannon, "A cellular computer to implement the Kalman filter algorithm," Ph.D. dissertation, Dept. Electric. Eng. Comput. Sci., Montana State Univ., Bozman, 1969.

[23] R. Marculescu, U. Y. Ogras, L. S. Peh, N. E. Jerger, and Y. Hoskote, "Outstanding research problems in NoC design: System, microarchitecture, and circuit perspectives," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 28, no. 1, pp. 3–21, Jan. 2009.

[24] N. Agarwal, T. Krishna, L. Peh, and N. Jha, "GARNET: A detailed onchip network model inside a full-system simulator," in *Proc. Int. Symp. Performance Anal. Syst. Softw.*, 2009, pp. 33–42.

[25] H. Wang, X. Zhu, L. Peh, and S. Malik, "Orion: A power-performance simulator for interconnection networks," in *Proc. Int. Symp. Microarchitect.*, 2002, pp. 294–305.

[26] V. S. Pai, P. Ranganathan, and S. V. Adve, "RSIM: Rice simulator for ILP multiprocessors," *SIGARCH Comput. Archit. News*, vol. 25, no. 5, p. 1, 1997.

[27] V. Puente, J. Gregorio, and R. Beivide, "Sicosys: An integrated framework for studying interconnection network performance in multiprocessor systems," in *Proc. Euromicro Conf. Parallel, Distrib. Netw.-Based Process.*, vol. 0. 2002, p. 0015.

[28] Noxim. (2010). *The NoC Simulator* [Online]. Available: http://noxim.sourceforge.net

[29] N. Jiang, G. Michelogiannakis, D. Becker, B. Towles, and W. J. Dally, *Booksim 2.0 User's Guide*. Palo Alto, CA: Stanford University, Mar. 2010.

[30] T. Austin, E. Larson, and D. Ernst, "Simplescalar: An infrastructure for computer system modeling," *Computer*, vol. 35, no. 2, pp. 59–67, 2002.

[31] Arvind, K. Asanovic, D. Chiou, J. C. Hoe, C. Kozyrakis, S.-L. Lu, M. Oskin, D. Patterson, J. Rabaey, and J. Wawrzynek, "RAMP: Research accelerator for multiple processors—a community vision for a shared experimental parallel HW/SW platform," Dept. Electric. Eng. Comput. Sci., Univ. California, Berkeley, CA, Tech. Rep. UCB/CSD-05-1412, Sep. 2005.


**Pengju Ren** received the B.S. degree in electrical engineering from Xi'an Jiaotong University, Xi'an, China, in 2004. He is currently pursuing the Ph.D. degree with the School of Electronic and Information Engineering, Xi'an Jiaotong University.

He was a Visiting Scholar with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, from October 2009 to January 2011. His current research interests include on-chip networks, scalable many-core designs, and very large scale integration architectures for digital video processing.

**Mieszko Lis** is currently pursuing the Ph.D. degree with the Massachusetts Institute of Technology (MIT), Cambridge. Before coming to MIT for doctoral studies in computer architecture and computational biology, he received the Bachelors and Masters degrees from MIT.

He accumulated extensive industry experience as a cofounder of a fabless semiconductor company and a high-level hardware synthesis startup. His current research interests include massive-scale multicores and the advanced coherent memory hierarchies required to support them.

**Myong Hyon Cho** received the Bachelors degree from Seoul National University, Seoul, Korea, and the Masters degree from the Massachusetts Institute of Technology, Cambridge. He is currently pursuing the Ph.D. degree with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology.

His current research interests include many-core computer architectures, memory subsystems, and on-chip networks.

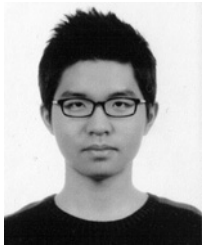

**Keun Sup Shim** received the B.S. degree in electrical engineering from the Korea Advanced Institute of Science and Technology, Daejeon, Korea, in 2006, and the M.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, in 2010. He is currently pursuing the Ph.D. degree in electrical engineering and computer science with MIT.

His current research interests include high-performance computer architectures, scalable many-core designs, and on-chip networks.

**Christopher W. Fletcher** received the B.S. degree in electrical engineering and computer science from the University of California, Berkeley. He is currently a second-year Graduate Student studying computer science with the Massachusetts Institute of Technology, Cambridge.

His current research interests include energy efficiency, performance, and scalability concerns in multicore systems.

**Omer Khan** (M'09) received the Bachelors degree from Michigan State University, East Lansing, in 2000, and the Ph.D. degree in electrical and computer engineering from the University of Massachusetts, Amherst, in 2009.

He is currently an Assistant Professor of electrical and computer engineering with the University of Connecticut (UConn), Storrs. Prior to joining UConn, he was a Post-Doctoral Fellow with the Massachusetts Institute of Technology, Cambridge. His current research interests include computer architecture. He has co-authored numerous papers in this area.

**Nanning Zheng** (SM'93–F'06) graduated from the Department of Electrical Engineering, Xi'an Jiaotong University, Xi'an, China, in 1975, and received the M.S. degree in information and control engineering from Xi'an Jiaotong University in 1981 and the Ph.D. degree in electrical engineering from Keio University, Yokohama, Japan, in 1985.

He joined Xi'an Jiaotong University in 1975, and is currently a Professor and the Director of the Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University. His current research interests include computer vision, pattern recognition, machine vision and image processing, neural networks, and hardware implementation of intelligent systems.

Dr. Zheng became a member of the Chinese Academy of Engineering in 1999, and has been the Chief Scientist and the Director of the Information Technology Committee of the China National High Technology Research and Development Program since 2001. He was the General Chair of the International Symposium on Information Theory and Its Applications and the General Co-Chair of the International Symposium on Nonlinear Theory and Its Applications, both in 2002. He is a member of the Board of Governors of the IEEE ITS Society and the Chinese Representative on the Governing Board of the International Association for Pattern Recognition. He also serves as an Executive Deputy Editor of the *Chinese Science Bulletin*.

**Srinivas Devadas** (F'98) is currently a Professor of electrical engineering and computer science with the Massachusetts Institute of Technology (MIT), Cambridge, and has been on the faculty of MIT since 1988. He served as the Associate Head with responsibility for computer science from 2005 to 2011. He has worked on the areas of computer-aided design, testing, formal verification, compilers for embedded processors, computer architectures, computer security, and computational biology, and has co-authored numerous papers and books in these areas.