

## Per-Packet Global Congestion Estimation for Fast Packet Delivery in Networks-on-Chip

Pejman Lotfi-Kamran

**Abstract** Networks-on-chip (NOCs) are becoming the de facto communication fabric to connect cores and cache banks in chip multiprocessors (CMPs). Routing algorithms, as one of the key components that influence NOC latency, are the subject of extensive research. Static routing algorithms have low cost but unlike adaptive routing algorithms, do not perform well under non-uniform or bursty traffic. Adaptive routing algorithms estimate congestion levels of output ports to avoid routing traffic over congested ports. As global adaptive routing algorithms are not restricted to local information for congestion estimation, they are the prime candidates for balancing traffic in NOCs. Unfortunately, destinations of packets are not considered for congestion estimation in existing global adaptive routing algorithms. We will show that having identical congestion estimates for packets with different destinations prevents global adaptive routing algorithms from reaching their peak potential. In this work, we introduce *Fast*, a low-cost global adaptive routing algorithm that estimates congestion levels of output ports on a per-packet basis. The simulation results reveal that *Fast* achieves lower latency and higher throughput as compared to those of other adaptive routing algorithms across all workloads examined. *Fast* increases the throughput of an  $8 \times 8$  network by 54%, 30%, and 16% as compared to DOR, Local, and RCA on a synthetic traffic profile. On realistic benchmarks, *Fast* achieves 5% average, and 12% maximum latency reduction on SPLASH-2 benchmarks running on a 49-core CMP as compared to the state of the art.

**Keywords** Adaptive routing · chip multiprocessor · congestion estimation · network-on-chip

---

P. Lotfi-Kamran  
School of Computer Science  
Institute for Research in Fundamental Sciences (IPM)  
Tel.: +98-21-24509404  
Fax: +98-21-24509444  
E-mail: plotfi@ipm.ir

## 1 Introduction

Today's information centric world is powered by processors. The unprecedented growth of processors performance in the past 30 years has enabled many useful applications and services to emerge. Many of these applications and services have inherent parallelism [46,4,6,12]. Taking advantage of the parallelism in applications, and driven by the need to increase performance of processors, the industry and academia have introduced multi-core processors (e.g., [18,43,3,39,30,16]).

In light of known scalability limitations for crossbar- and bus-based designs, existing many-core chip multiprocessors (CMPs), such as Tiler's Tile series, have featured a network-on-chip (i.e., a mesh-based interconnect fabric) and a tiled organization. Typically, each tile integrates a core, a slice of the shared last-level cache (LLC) with directory, and a router. The resulting organization enables cost-effective scalability to high core counts.

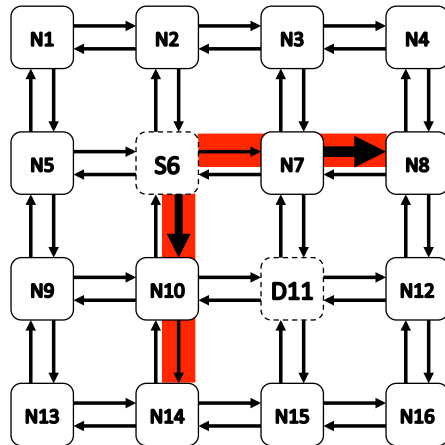
In tiled processors, the network-on-chip is responsible to carry cores' requests for data and instructions to the LLC, and to bring back the responses to the requesting cores. As the delay of delivering instructions and data is critical to the performance of parallel applications, the design of a fast network-on-chip is the subject of many research activities [38,33,45,29,1,22,25,34,2,44].

One of the key components that has a major impact on the speed of a network-on-chip is routing algorithm [17,26,23,28,31,42,15]. At each router, the routing algorithm decides how to route incoming packets to the destinations. The simplest routing algorithm is dimension order routing (DOR) which statically routes packets along one of the dimensions before moving to the next one. Static routing algorithms are simple and have low complexity, but cannot adapt routing decisions to the traffic in the network, and hence, result in high network delay and low performance under non-uniform or bursty traffic.

To address the shortcomings of static routing algorithms, researchers have proposed adaptive routing algorithms. An adaptive routing algorithm associates a congestion estimate to each output port that a packet can be forwarded to. When a packet can be routed over more than one output port, an adaptive routing algorithm picks the port that has the lowest congestion.

A local adaptive routing algorithm uses metrics that are accessible in a router or in the neighboring routers (e.g., number of free virtual channels (VCs), number of free slots in the downstream buffer, etc.) to estimate congestion levels of output ports. Local adaptive routing algorithms reduce latency of the network over static routing algorithms. However, as they only observe the status of a small fraction of the network, local adaptive routing algorithms greedily rely on local information to route packets, and hence, frequently take suboptimal decisions [15].

Global adaptive routing algorithms benefit from global (or regional) metrics to estimate congestion levels of output ports, and hence, can take better decisions. Regional Congestion Awareness (RCA) [15], which is one of the state-of-the-art global adaptive routing algorithms, aggregates the congestion estimates of links in a region (e.g., a row or a column) of a NOC to estimate

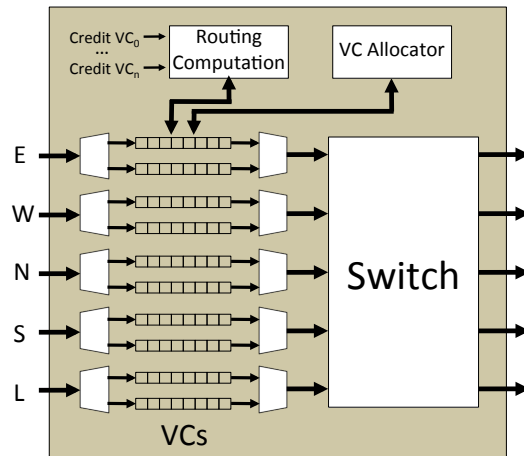


**Fig. 1** Illustration of the inefficiency of existing global adaptive routing algorithms in estimating the congestion levels of output ports. The thickness of a link in the shaded area is proportional to the traffic passing through it.

congestion levels of routers' output ports. While using an aggregated congestion estimate enables RCA to improve NOC performance over local adaptive routing algorithms, existing aggregation methods introduce noise into the congestion estimate, and prevent global adaptive routing algorithms to reach their peak potential.

Figure 1 illustrates the inefficiency of existing global adaptive routing algorithms in estimating congestion levels of routers' output ports. Assume that *Node 6* has a packet that needs to be sent to *Node 11*. *Node 6* can either forward the packet on EAST port to *Node 7*, or on SOUTH port to *Node 10*. A global adaptive routing algorithm combines the congestion status of the EAST port of *Node 6* and *7*, and the congestion status of the SOUTH port of *Node 6* and *10* to estimate congestion levels of the EAST and the SOUTH port, respectively. While the destination of the packet is *Node 11*, and definitely, this packet will not pass through the EAST port of *Node 7* or the SOUTH port of *Node 10*, nevertheless, the congestion status of these ports are part of the congestion estimates that *Node 6* will use to pick the favorite port. If one of these ports is congested (e.g., EAST port of *Node 7*), and consequently, contributes significantly to the congestion estimate, the global adaptive routing algorithm's decision will be suboptimal.

In this paper, we introduce *Fast*, a global adaptive routing algorithm that eliminates noise in the process of global congestion estimation. Unlike existing global routing algorithms that use an identical congestion estimate for all packets that may go through an output port at a given time, *Fast* estimates, on a per-packet basis, a congestion level for each port, taking into consideration the destinations of the packets. *Fast* uses a novel mechanism that requires only 1-bit of global information for each router in the row and column. We will show that *Fast* improves network latency by 5%, on average, and 12%, at best, over



**Fig. 2** Organization of a mesh router.

the state of the art on SPLASH-2 benchmarks, while requiring less hardware overhead.

The rest of this paper is organized as follows. Section 2 presents background information on networks-on-chip. Section 3 reviews existing routing algorithms, and Section 4 discusses why existing global adaptive routing algorithms cannot reach their peak potential. Section 5 introduces *Fast* global adaptive routing algorithm. In Section 6, we evaluate our proposal for global adaptive routing. Finally, we conclude in Section 8.

## 2 Background

Except for the boundary routers, every router in a 2D mesh, has five input/output ports (four connected to the neighboring routers and one for the local core). Organization of a wormhole router includes the input FIFO for each port, route computation unit, VC allocation unit, crossbar control logic, and the crossbar. The packet transmission is performed by dividing it to smaller pieces called flits. A flit enters into the router through one of the ports and is stored in its FIFO. If the flit is a header, indicating the start of a new packet, it proceeds to the routing unit, which determines the output port that the packet should use. Then, the header flit attempts to acquire a virtual channel for the next hop. Upon a successful VC allocation, the header flit enters the switch arbitration stage, where it competes for the output port with other flits from the other router input ports. Once the crossbar passage is granted, the flit traverses the switch and enters the channel. Subsequent flits belonging to the same packet can proceed directly to the crossbar and go to the output port. The main architectural element of a router is shown in Figure 2.

To eliminate the delay of routing from the critical path of a router, researchers proposed look-ahead routing, in which routing decisions are performed one hop in advance (i.e., For each packet, a router performs routing decision for the next hop, as the previous router has already performed the routing for this hop) [13]. The delay of the routing being off the critical path, researchers proposed dynamic routing methods, which are more complex, for better distribution of traffic in the network.

The main difference between a static and dynamic router is that, depending on the network condition, in a dynamic router, the route computation unit may select different paths at different times for the same source-destination pair. Various factors may be used for choosing an output port like (1) the number of free VCs, (2) the crossbar pressure, and (3) the number of free buffers at the corresponding input port in the downstream router. We use (2) to estimate network condition for adaptive routers as we experimentally verified that this metric works best across the range of workloads that we examined (Our results also corroborate prior work [15]).

Packets can be routed in both X and Y directions without restriction, adaptive routing algorithms need a mechanism to guarantee deadlock avoidance. In networks having virtual channels (general case), usually the following method, which is called virtual sub-network, is used to guarantee deadlock avoidance. Virtual channels in Y dimensions are divided into two parts. The network is partitioned into two sub-networks called +X sub-network and -X sub-network each having half of the channels in the Y dimension. If the destination node is to the right of the source, the packet will be routed through the +X sub-network. If the destination node is to the left of the source, the packet will be routed through the -X sub-network. Otherwise that packet can be routed using either sub-network [36].

In this work, we only consider minimal routing algorithms, as they are by far the most dominant method in the literature and real-machine implementations. Investigating the usefulness of the techniques that we present in this work under non-minimal routing is part of our future work.

### 3 Existing Routing Algorithms

In this section, we review three main classes of routing algorithms: dimension order routing (DOR), local adaptive routing, and global adaptive routing.

#### 3.1 Dimension Order Routing (DOR)

The simplest type of routing is static routing, in which the path that a packet travels is determined only based on the source and the destination of the packet. One of the widely used static routing mechanism is dimension order routing (DOR). Having a static order of the dimensions (e.g., X, Y, etc.), with DOR, packets are routed to the correct position in the highest-order dimension before moving to the next dimension.

DOR is easy to implement and has low complexity. However, as network conditions do not have any roles in the routing decisions, this routing algorithm is incapable of balancing non-uniform traffic in the NOC. Consequently, DOR leads to large network latency when traffic is bursty (e.g., traffic of a CMP), and hence, is not well suited for chip multiprocessors.

### 3.2 Local Adaptive Routing

To address the shortcomings of static routing algorithms, researchers proposed using adaptive routing algorithms. In adaptive routing, the path that a packet travels not only depends on the source and the destination of the packet, but also on the network conditions. For this purpose, adaptive routers estimate congestion level of each output port. When a packet can be routed over two ports, the port with a smaller congestion level is taken.

Local adaptive routing algorithms take advantage of metrics that are accessible in a router to estimate congestion level of an output port. Various metrics have been used in the literature for estimating congestion level of an output port. The number of free VCs [7], the number of free buffer slots in the downstream router [23], and the active demand that each output port experiences (i.e., crossbar demand) [15] are among such metrics. Corroborating prior work [15], we experimentally verify that the crossbar demand works best across all of the workloads that we examined.


A local adaptive routing algorithm distributes traffic just by observing the status of a single router in the downstream path. If a link or router is congested further away, a local adaptive routing algorithm has no knowledge of it, and hence, cannot avoid the congested link or router.

### 3.3 Global Adaptive Routing

Global adaptive routing algorithms address the shortcoming of local adaptive routing algorithms by estimating congestion level of a router's output port, not just by the status of the router itself, but also by taking into consideration the status of other routers. For this purpose, a global adaptive routing algorithm requires a monitoring network to send congestion status among routers. Every router aggregates (takes weighted average of) its local congestion estimate and that of downstream routers. The resulting global congestion estimate is being used for the output port selection, and also gets propagated to upstream routers. Figure 3 shows how in a global adaptive routing algorithm, congestion status is propagated in the EAST direction [15].

## 4 Why not existing global adaptive routing algorithms?

Global adaptive routing algorithms have the potential to reduce the network latency beyond what is possible by local adaptive routing algorithms. However,




**Fig. 3** Congestion status flow in a global adaptive routing algorithm [15].

as mentioned in the Introduction Section, existing global adaptive routing algorithms cannot reach their peak potential because they do not estimate congestion level of an output port on a per-packet basis: congestion estimate of an output port is independent of packets' destinations. Therefore, it is possible that a further-away congested link that is not on the path to the destination contributes significantly to the congestion estimate, preventing global adaptive routing algorithms from balancing traffic in NOCs.

Congestion estimate being independent of a packet's destination, congested links that are outside the path to the destination may influence the routing decision while they ideally should not. To quantify the importance of taking into consideration the path to the destination while estimating congestion level, we count the number of routing decisions where all congested links are outside the path to the destination but still influence the routing decision (these congested links should not affect the routing decision as they are outside the path to the destination but existing global routing algorithms have no way of identifying such links) and normalize it to the number of routing decisions where at least one congested link is on the path to the destination. Figure 4 shows the results of this experiment across all SPLASH-2 benchmarks. We consider a link congested if two or more packets compete for that link in a given cycle. When congested links are NOT on the path to the destination, and links that are on the path to the destination are not congested, the congestion estimate gets heavily influenced by the congested links that are outside the path, misleading global adaptive routing algorithms. Figure 4 shows that the number of such misleading cases is large across SPLASH-2 benchmarks: on average for every three useful congestion notifications, there is a misleading one.

Congestion estimation being independent of packets' destinations, existing global adaptive routing algorithms are incapable of sifting useful congestion notifications from misleading ones. As the number of misleading congestion notifications is high, existing global adaptive routing algorithms cannot fully balance traffic in the network; missing the opportunity to minimize latency.



**Fig. 4** Number of misleading global congestion notifications of existing global adaptive routing algorithms, normalized to the number of useful global congestion notifications across SPLASH-2 benchmarks.

## 5 Fast Global Adaptive Routing Algorithm

To exclusively take into consideration the congestion status of links that are on the path to the destination, a global adaptive routing algorithm should include destinations of packets in the process of estimating congestion levels of output ports. Destinations of packets enable global routing algorithms to dynamically include congestion status of downstream routers that are on the path to the destination in the routing decision.

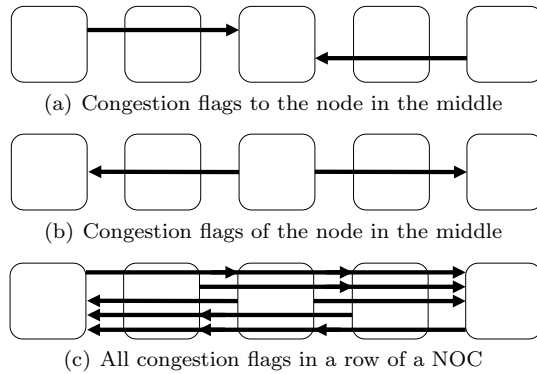
Unfortunately, existing global routing algorithms are inherently incapable of selectively excluding congestion status of not-on-the-path-to-the-destination downstream routers from global congestion status. This is due to the fact that as congestion status of downstream routers gets mixed with each other (i.e., the weighted average) to form global congestion status, it is no longer possible to extract congestion status of a particular downstream router from the global congestion status.

To address this shortcoming, we rely on concatenation, instead of weighted average, to form global congestion status. As with concatenation, the correspondence between bits in global congestion status and status of downstream routers is preserved, a global routing algorithm can take advantage of packets' destinations to dynamically include congestion status of routers that are on the path to the destination in the routing decision.

In this work, we associate a 1-bit congestion flag per routers' output port. This 1-bit flag is ONE when the port is congested and ZERO otherwise. Each router propagates its congestion flags to downstream routers in the same row and column. As a result, every router receives a congestion flag for each router in the same row and column. Figure 5(a) shows the congestion flags that a node in the middle of a row receives from other nodes in the same row.

For every output port, a router propagates the congestion flag of the port to the neighbor that is connected to the port. The neighbor has access to the congestion flag (i.e., can use it in the routing process) and also propagates





**Fig. 5** Congestion flags of *Fast* global adaptive routing algorithm in a row of a network-on-chip.

the flag further in the row or column for the use of other upstream routers. Figure 5(b) shows the wires that carry two of the four congestion flags of a node to the rest of the nodes in the same row (there are other wires propagating the other two congestion flags to the nodes in the column, but are not shown).

Every router requires  $n$  1-bit wire segments (i.e., each wire segment connects two neighboring nodes) to propagate the two congestion flags of EAST and WEST ports to all the nodes in the row, where  $n$  is the number of nodes in the row. Likewise, there will be another  $n$  1-bit wire segments to propagate the congestion flags of NORTH and SOUTH ports to all the nodes in the column. In total, we need  $n \times n$  1-bit wire segments in a row (or a column) to propagate all congestion flags of all the nodes in the row (or the column) as shown in Figure 5(c).

In this work, we introduce *Fast*, a simple and low-cost global routing algorithm to demonstrate the potentials of per-packet congestion estimation. *Fast* leverages the following heuristic to minimize NOC delays: as links become further away from the current node, their congestion status becomes less important. The heuristic states that a congested link near the current router is more important (i.e., has more negative impact) than a congested link that is further away.

Having the congestion flags of all the routers in the same row and column, the routing logic can be as simple as the following. If the destination is on the same row or column as the current node (i.e., there is only one option for routing the packet), pick the only possible output port. Otherwise, if local congestion metrics favor a port, pick that port. Otherwise, for every possible output port, calculate the length of the continuously uncongested path to the destination in the row or column using global congestion flags. Pick the port with longer uncongested path to the destination.

As a near congested link has more negative impact on performance than a further away congested link, *Fast* relies on local congestion metrics to pick the favorite port. If local congestion metrics do not favor an output port over

the other one, *Fast* uses global congestion flags to break the tie. Using congestion flags and for each possible output port, *Fast* calculates the length of the straight path to the destination that starts from the output port and ends with the first congested link. *Fast* picks the output port that enables a packet to travel more hops before facing a congested link.

---

**Algorithm 1** Pseudocode of Fast Global Adaptive Routing
 

---

```

1: procedure FAST
2:   if  $((diff\_x <> 0) \ \& \ (diff\_y <> 0))$  then
3:     if  $local\_x <> local\_y$  then
4:       Local adaptive routing based on  $local\_x$  and  $local\_y$ 
5:     else
6:       if  $diff\_x > 0$  then
7:          $dist\_x \leftarrow + diff\_x$ 
8:          $dir\_x \leftarrow + 1$ 
9:       else
10:         $dist\_x \leftarrow - diff\_x$ 
11:         $dir\_x \leftarrow - 1$ 
12:       if  $diff\_y > 0$  then
13:          $dist\_y \leftarrow + diff\_y$ 
14:          $dir\_y \leftarrow + 1$ 
15:       else
16:         $dist\_y \leftarrow - diff\_y$ 
17:         $dir\_y \leftarrow - 1$ 
18:        $cnt\_x \leftarrow 0$ 
19:        $cnt\_y \leftarrow 0$ 
20:       while  $((dist\_x <> 0) \ \& \ (g\_cong\_x[pos\_x + cnt\_x] == 0))$  do
21:          $cnt\_x \leftarrow cnt\_x + dir\_x$ 
22:          $dist\_x \leftarrow dist\_x - 1$ 
23:       while  $((dist\_y <> 0) \ \& \ (g\_cong\_y[pos\_y + cnt\_y] == 0))$  do
24:          $cnt\_y \leftarrow cnt\_y + dir\_y$ 
25:          $dist\_y \leftarrow dist\_y - 1$ 
26:       if  $cnt\_x >= cnt\_y$  then
27:         if  $dir\_x == +1$  then
28:           Pick EAST
29:         else
30:           Pick WEST
31:       else
32:         if  $dir\_y == +1$  then
33:           Pick SOUTH
34:         else
35:           Pick NORTH
36:     else
37:       Just one option: pick the only possible direction

```

---

The pseudocode of the *Fast* adaptive routing algorithm is shown as Algorithm 1. The inputs of this algorithm are  $pos\_x$ ,  $pos\_y$ ,  $diff\_x$ ,  $diff\_y$ ,  $local\_x$ ,  $local\_y$ ,  $g\_cong\_x$ , and  $g\_cong\_y$ .  $pos\_x$  and  $pos\_y$  determine the X and Y position of current router.  $diff\_x$  and  $diff\_y$  show both the distance (i.e., the absolute value) to the destination and the direction (i.e., the sign) of the destination in each of the X and Y dimensions, respectively.  $local\_x$  and  $local\_y$  are the

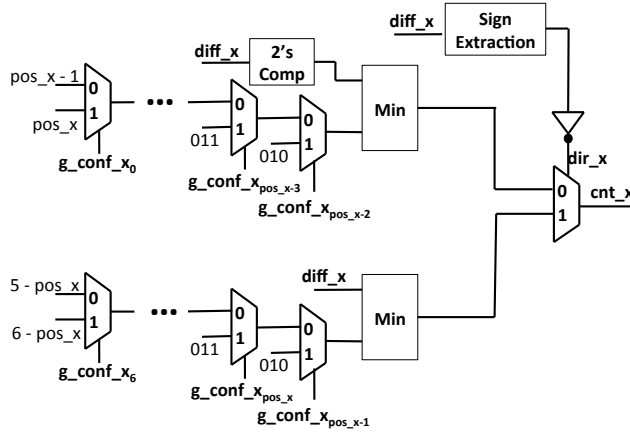


Fig. 6 Hardware realization of  $cnt_x$  and  $cnt_y$ .

congestion estimates of neighboring routers, and  $g\_cong\_x$ , and  $g\_cong\_y$  are the global congestion flags of routers in the X and Y dimensions, respectively.

Line 2 of the pseudocode checks if there are two options for routing the packet. If the condition does not hold, the control goes to Line 37 and picks the only available option using standard routing techniques. Otherwise if the condition holds and consequently two options are available for routing the packet, Line 3 checks if local congestion metrics (congestion level of neighbors) are different. If they are different, *Fast* relies on these local metrics to route the packet (i.e., no global congestion metric will be used). Otherwise, Lines 6–19 determine the direction and the distance to the destination in each of the X and Y dimensions. Lines 20–25 count how many hops the packet passes through on the path to the destination before encountering a congested link in each of the X and Y dimensions. Finally, Lines 26–35 pick the option (out of the two available options) in which the packet faces a congested link further away.

### 5.1 Hardware Realization of Fast Adaptive Routing Algorithm

In this part, we propose a hardware implementation for *Fast* adaptive routing algorithm, and evaluate its critical path delay. To this goal, we suggest a hardware implementation for computing  $cnt_x$  and  $cnt_y$  (See Algorithm 1 for details of  $cnt_x$  and  $cnt_y$ ), and then provide a complete hardware implementation for *Fast* algorithm using hardware implementation of  $cnt_x$  and  $cnt_y$ .

$cnt_x$  and  $cnt_y$  store the distance to the first congested link in the X and Y directions, respectively. We use a chain of multiplexers to implement this computation. Figure 6 shows the hardware realization of  $cnt_x$  (Hardware

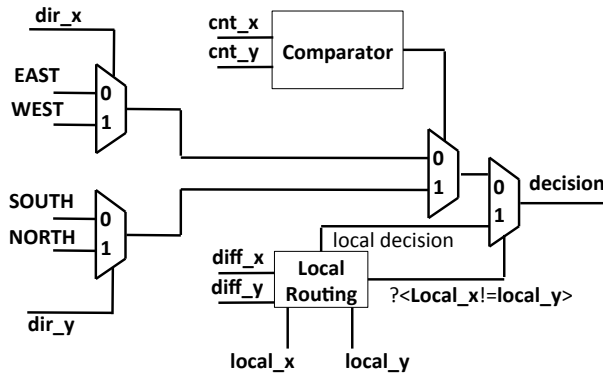


Fig. 7 Hardware realization of *Fast* adaptive routing algorithm.

realization of  $cnt_y$  is identical). As the chain of multiplexers do not take into consideration the destination of the packet, we take the minimum of the output of the mux-chain and the distance to the destination as the value of  $cnt_x$  (The same is true for  $cnt_y$ ). The suggested hardware realization of  $cnt_x$  is simple, because it only requires standard digital components (2-input multiplexer, 2's complement, comparator, and sign extractor). Moreover, it is fast, because it operates on values with few bits (4 to 8 bits).

Having a hardware realization for  $cnt_x$  and  $cnt_y$ , the hardware implementation of the whole *Fast* routing algorithm is trivial. Figure 7 shows the hardware realization of *Fast* routing algorithm. It essentially checks the condition to determine if local or global routing should be used. In case local routing is appropriate, the output comes from the standard local router. Otherwise, the output is determined using  $cnt_x$  and  $cnt_y$  (the one that is larger wins).

We use detailed technology models to derive component latencies. The worst case delay of the routing algorithm happens in a router in the corner of an  $8 \times 8$  network-on-chip that has the highest number of multiplexers in the mux-chain. At 32 nm technology, the routing easily fits in one clock cycle for frequencies up to 5 GHz given the critical-path delay of the *Fast* routing algorithm.

While *Fast* is simple and has low cost, it estimates congestion levels of output ports on a per-packet basis. Consequently, it can showcase the importance of per-packet congestion estimation by global routing algorithms.

## 6 Evaluation

To assess the effectiveness of the proposed routing algorithm (i.e., *Fast*), we compare it against DOR, Local, and RCA (the latter is the state-of-the-art

**Table 1** Baseline network configuration and variation.

Characteristic	Baseline	Variation
Topology	$7 \times 7$ 2D Mesh	$4 \times 4$ , $15 \times 15$
Routing	Minimal, fully-adaptive, virtual sub-network deadlock avoidance	–
Per-hop latency	3	–
VC/port	2	4
Flit buffers/VC	6	–
Packet size (flits)	5	10
Traffic workload	Bit Rotate, Hotspot, Transpose, Uniform	SPLASH-2
Packets/node	10000	–

global adaptive routing algorithm). We use synthetic and real workloads for the evaluation.

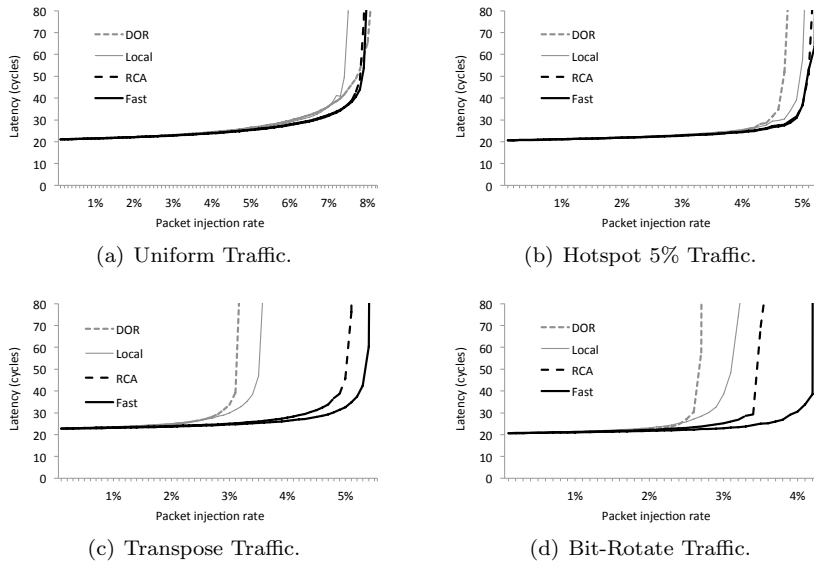
## 6.1 Methodology

To determine latency-throughput characteristics, we use a detailed cycle-accurate NOC simulator for the virtual-channel routers. Each input virtual channel has a buffer (FIFO) depth of six flits. The congestion threshold value for *Fast* routing is set to two, meaning that an output port of a router is considered to be congested when at least two packets compete for the same output port in a given cycle. The per-hop latency for packets is 3 cycles and congestion flags can pass one hop in a single cycle (as they do not need routing or arbitration). In all of the simulations, the latency is measured by averaging the latency of the packets when each local core generates 10000 packets. The routers use the minimally fully adaptive virtual sub-network deadlock avoidance technique discussed in [36]. Table 1 shows the baseline network configuration, and the variations used in the sensitivity studies.

### 6.1.1 Technology Parameters

We use publicly available tools and data to estimate the area and energy of global routing algorithms. Our study targets a 32 nm technology node with an on-die voltage of 0.9 V and a 2 GHz operating frequency.

We use custom wire models, derived from a combination of sources [2, 20], to model links and router switch fabrics. For links, we model semi-global wires with a pitch of 200 nm and power-delay-optimized repeaters that yield a link latency of 125 ps/mm. On random data, links dissipate 50 fJ/bit/mm, with repeaters responsible for 19% of link energy. For area estimates, we assume that link wires are routed over logic or SRAM and do not contribute to network area; however, repeater area is accounted for in the evaluation. Our buffer models are taken from ORION 2.0 [21]. We model flip-flop based buffers.



**Fig. 8** Latency versus packet injection rate for DOR, Local, RCA, and *Fast* on a  $7 \times 7$  2D mesh with two virtual channels per port and 5-flit packets. (a) uniform random traffic, (b) Hotspot 5% traffic, (c) Transpose traffic, and (d) Bit-rotate traffic.

### 6.1.2 Workloads

We evaluate *Fast* using four standard synthetic traffic patterns: *bit-rotate*, *hotspot 5%*, *transpose*, and *uniform random*. These synthetic traffic patterns provide insight into the relative strengths and weaknesses of the competing routing algorithms.

Moreover, we evaluate *Fast* using traces of SPLASH-2 benchmarks [46]. The traces are obtained from a 49-core CMP, arranged in a  $7 \times 7$  2D mesh topology [25].

## 6.2 Synthetic Traffic Patterns

We evaluate various routing algorithms on a  $7 \times 7$  2D mesh network with two virtual channels using standard synthetic traffic patterns. The packet length is set to five flits. The packet latency for different traffic profiles are shown in Figure 8.

### 6.2.1 Uniform Random Traffic Profile

In this traffic profile, every node sends packets to other nodes while using a uniform distribution to construct the destination set of packets [27]. Figure 8(a) shows the average communication delay as a function of the average packet

injection rate. For this traffic profile, DOR, RCA, and *Fast* routing algorithms lead to lower latencies (uniform traffic is balanced under DOR routing [15]).

### 6.2.2 Hotspot Traffic Profile

The same as uniform random, each node sends packets to other nodes in the network. But nodes at positions (3, 3), (3, 5), (5, 3), and (5, 5) receive 5% more packets. The average communication delays as a function of average packet injection rate for this traffic pattern is shown in Figure 8(b). For this traffic pattern, adaptive routing algorithms perform much better than DOR, while *Fast* and RCA leading to the lowest latencies. This is due to the non-uniform nature of the traffic pattern.

### 6.2.3 Transpose Traffic Profile

In transpose traffic profile within an  $n \times n$  mesh network, a node at position  $(i, j)$  ( $i, j \in [0, n - 1)$ ) only sends packets to the node at position  $(n - i - 1, n - j - 1)$ . This traffic pattern is similar to the concept of transposing a matrix [14]. This traffic profile leads to a non-uniform traffic distribution with heavy traffic in the center of the mesh. As the results shown in Figure 8(c) reveal, if the packet injection rate is very low, the routing techniques behave similarly. As the injection rate increases and links get congested, the *Fast* and RCA routing algorithms lead to smaller average delays, with *Fast* being the absolute best.

### 6.2.4 Bit-Rotate Traffic Profile

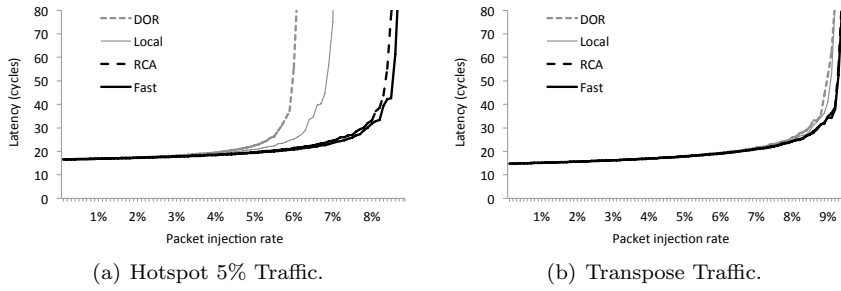
In an  $n \times n$  mesh network, a node at position  $(i, j)$  ( $i, j \in [0, n - 1)$ ) is assigned an address  $i * n + j$ . In bit-rotate traffic profile, a node with address  $s$  sends packets only to the node with address  $d$ , where  $d_i = s_{i+1 \bmod n}$ . This traffic profile leads to a non-uniform traffic distribution. With this traffic profile, *Fast* routing algorithm leads to the lowest latencies, beating DOR, Local, and RCA.

## 6.3 Sensitivity Analysis


In this section, we change some of the NOC parameters to evaluate the sensitivity of the routing algorithms to the change in the network parameters.

### 6.3.1 NOC Size

Figure 9 and 10 show the latency versus packet injection rate for a small (i.e.,  $4 \times 4$ ) and a large (i.e.,  $15 \times 15$ ) mesh NOC under the hotspot 5% (9(a) and 10(a)), and transpose (9(b) and 10(b)) traffic profiles with 5-flit packets. For hotspot 5%, nodes at positions (1, 1), (1, 2), (2, 1), and (2, 2) for the  $4 \times 4$ , and nodes at positions (3, 3), (3, 12), (12, 3), and (12, 12) for the



**Fig. 9** Latency versus packet injection rate for DOR, Local, RCA, and *Fast* on a  $4 \times 4$  2D mesh with two virtual channels per port and 5-flit packets. (a) Hotspot 5% traffic, and (b) Transpose traffic.



**Fig. 10** Latency versus packet injection rate for DOR, Local, RCA, and *Fast* on a  $15 \times 15$  2D mesh with two virtual channels per port and 5-flit packets. (a) Hotspot 5% traffic, and (b) Transpose traffic.


$15 \times 15$  mesh receive 5% more traffic. For the small network, the results in Figure 9 show that *Fast* performs very well on Hotspot 5%, achieving 42% and 26% better throughput than DOR and Local and slightly exceeding that of RCA. On Transpose traffic profile, *Fast* is 6% better than Local and DOR and slightly better than of RCA. On smaller networks, *Fast* provides excellent visibility into the congestion state of the network.

For the large network (i.e., Figure 10), adaptive approaches do not perform as well versus DOR. This is due to the fact that noise in congestion estimate is increased. However, under both traffic profiles, *Fast*, which eliminates noise, is the absolute best routing algorithm, with RCA and DOR being the second and third, respectively.


### 6.3.2 Packet Length

Figure 11 shows load-latency graphs for long packets (i.e., packets that are 10 flits long) in a  $7 \times 7$  2D mesh under the hotspot 5% (11(a)) and transpose (11(b)) traffic profiles. The average packet latencies for all routing algorithms are higher than those of short packets. The increased latency is a known characteristic of wormhole routing with long packets. The reason is that for long





**Fig. 11** Latency versus packet injection rate for DOR, Local, RCA, and *Fast* on a  $7 \times 7$  2D mesh with two virtual channels per port and 10-flit packets. (a) Hotspot 5% traffic, and (b) Transpose traffic.



**Fig. 12** Latency versus packet injection rate for DOR, Local, RCA, and *Fast* on a  $7 \times 7$  2D mesh with four virtual channels per port and 5-flit packets. (a) Hotspot 5% traffic, and (b) Transpose traffic.

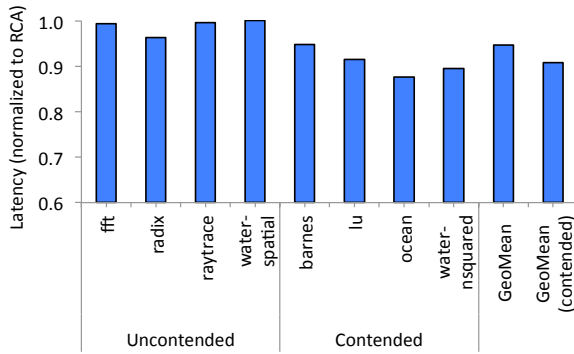
packets, an imbalance in the resource utilization occurs because packets are likely to hold resources over multiple routers. Similar to the case of 5-flit packets, *Fast* routing algorithm performs the best under both traffic profiles, with RCA and Local being the second and third, respectively.

### 6.3.3 Network with Four Virtual Channels

Figure 12 shows latency versus packet injection rate for DOR, Local, RCA, and *Fast* routers with four virtual channels. In this case, we consider a  $7 \times 7$  2D mesh with 5-flit packets. As the results show, *Fast* continues to perform better than other routing algorithms in a network with four virtual channels.

## 6.4 SPLASH-2 Benchmarks

Figure 13 shows the average packet latency of *Fast* routing algorithm across SPLASH-2 benchmarks, normalized to RCA. Contention is the cause of significant packet latency in *barnes*, *lu*, *ocean*, and *water-nsquared*; thus global adaptive routing algorithms have an opportunity to improve performance. Although *Fast* provides equal or lower latency than RCA for all benchmarks, the



**Fig. 13** Average latency of *Fast* across SPLASH-2 benchmarks normalized to the latency of RCA.

greatest benefit is on *ocean* with over 12% reduction in latency. On average, *Fast* provides a latency reduction of 5% across all benchmarks, and 10% across contended benchmarks versus the state of the art.

### 6.5 Hardware Overhead

RCA requires nine bits of information per output port to propagate global congestion status to upstream routers, and *Fast* only requires an average of three bits per port. Consequently, *Fast* requires 3x fewer wires for global congestion propagation, which represents a 8% reduction in the number of wires for a 64-bit flit network.

The 8% reduction in the number of wires is important for wire-intensive designs that are wire limited. But, as most of the area of a mesh network is due to the buffers and crossbars, in absolute terms, the difference between the area of RCA and *Fast* is small (1.29 mm<sup>2</sup> versus 1.27 mm<sup>2</sup>).

### 6.6 Power Dissipation

Our analysis confirms findings of prior work that CMP NOC is not a significant consumer of power at the chip level [24, 29]: NOC power is below 1 W under SPLASH-2 workloads. Good L1 cache performance is the main reason for the low power consumption at the NOC level.

## 7 Related Work

In [19], a static routing algorithm for two-dimensional meshes which is called XY is introduced. In this routing algorithm, each packet first travels along the X and then the Y dimension to reach the destination. For this method,

deadlock never occurs but no adaptivity exists in this algorithm. An adaptive routing algorithm named turn-model is introduced in [14] based on which another adaptive routing algorithm called OddEven turn is proposed in [5]. To avoid deadlock, OddEven method restricts the position that turns are allowed in the mesh topology. Another algorithm called DyAD is introduced in [17]. This algorithm is a combination of a static routing algorithm called oe-fix, and an adaptive routing algorithm based on the OddEven turn algorithm. Depending on the congestion condition of the network, one of the routing algorithms is invoked.

Another adaptive routing is hot potato or deflection routing [37,11,35] which is based on the idea of delivering a packet to an output channel at each cycle. If all the channels belonging to minimal paths are occupied, then the packet is misrouted. When contention occurs and the desired channel is not available, the packet, instead of waiting, will pick any alternative available channels (minimal or non-minimal) to continue moving to the next router; therefore the router does not need buffers. In hot potato routing, if the number of input channels is equal to the number of output channels at every router node, packets can always find an exit channel and the routing is deadlock free. However, livelock is a potential problem in this routing. Also, hot potato increases message latency even in the absence of congestion and bandwidth consumption. Accordingly, performance of hot potato routing is not as good as other wormhole routing methods [8].

There are adaptive routing algorithms for increasing fault tolerance of the on-chip network. Stochastic communication method has been proposed to deal with permanent and transient faults of network links and nodes [9]. This method has the advantage of simplicity and low overhead. The selection of links and of the number of redundant copies to be sent on the links is stochastically done at runtime by the network routers. As a result, the transmission latency is unpredictable and, hence, it cannot be guaranteed. Also, stochastic communication is not efficient in terms of power dissipation and latency.

A local adaptive deadlock free routing algorithm called Dynamic XY (DyXY) has been proposed in [26]. In this algorithm, which is based on the static XY algorithm, a packet is sent either to the X or Y direction depending on the congestion condition. It uses local information which is the current queue length of the corresponding input port in the neighboring routers to decide on the next hop. It is assumed that the collection of these local decisions should lead to a near-optimal path from the source to the destination. The main weakness of DyXY is that the use of the local information in making routing decision could forward the packet in a path which has congestion in the routers farther than the current neighbors. The technique described in [15] overcomes this problem. It uses global information in making a routing decision. The technique requires a mechanism to mix local and global congestion information. The main problem of this work, as is discussed in this paper, is the lack of correspondence between a router in the downstream path and the global congestion value. *Fast* resolves this problem by having a dedicated single-bit congestion wire for each router in the downstream path. Several other techniques also attempted

to address this problem [10,40,41,32,31], but they either use low-cost mechanism which cannot reach the full potential of global adaptive routing [31] or significantly increase the complexity and hence the hardware overhead of the routers [10,40,41,32]. In this work, we showed that by using a few wires in a smart way, a low-cost fast adaptive routing algorithm for networks-on-chip emerges.

## 8 Conclusion

Chip-multiprocessors require fast networks-on-chip to maximize performance. A routing algorithm is one of the key components that influences network's latency. A static/local adaptive routing algorithm has no/limited visibility into the network, and hence, cannot minimize network's latency. A global routing algorithm, which has more visibility into the network, is an alternative that has the potential to minimize latency and maximize performance. One of the key components of a global routing algorithm is congestion estimation unit, which measures the likelihood of a packet experiencing congestion if it gets forwarded to a particular router's output port. Unfortunately, existing congestion estimation units ignore destinations of packets, and hence, introduce noise into the congestion estimate, preventing existing global routing algorithms from minimizing network's latency.

This work introduced *Fast*, a global adaptive routing algorithm that does congestion estimation on a per-packet basis. *Fast* benefits from a low-cost congestion status network (1-bit per port) that informs upstream routers of congestion in downstream routers. Taking advantage of the congestion status network, *Fast* dynamically forwards packets to output ports, where the first congested link is further away. We showed that while *Fast* requires fewer wires, and hence silicon area, it reduces network's latency by up to 12% over the state-of-the-art global adaptive routing algorithm.

## References

1. Bakhoda, A., Kim, J., Aamodt, T.M.: Throughput-Effective On-Chip Networks for Manycore Accelerators. In: Proceedings of the 43rd Annual IEEE/ACM International Symposium on Microarchitecture, pp. 421–432. NY, NY, USA (2010)
2. Balfour, J.D., Dally, W.J.: Design Tradeoffs for Tiled CMP On-Chip Networks. In: Proceedings of the 20th Annual ACM International Conference on Supercomputing, pp. 187–198. Cairns, Queensland, Australia (2006)
3. Barroso, L.A., Gharachorloo, K., McNamara, R., Nowatzyk, A., Qadeer, S., Sano, B., Smith, S., Stets, R., Verghese, B.: Piranha: A Scalable Architecture Based on Single-Chip Multiprocessing. In: Proceedings of the 27th Annual International Symposium on Computer Architecture, pp. 282–293. Vancouver, British Columbia, Canada (2000)
4. Bienia, C., Kumar, S., Singh, J.P., Li, K.: The PARSEC Benchmark Suite: Characterization and Architectural Implications. In: Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques, pp. 72–81. New York, New York, USA (2008). DOI 10.1145/1454115.1454128
5. Chiu, G.M.: The odd-even turn model for adaptive routing. IEEE Transactions on Parallel and Distributed Systems **11**(7), 729–738 (2000)

6. Council, T.P.P.: <http://www.tpc.org/default.asp>
7. Dally, W.J., Aoki, H.: Deadlock-Free Adaptive Routing in Multicomputer Networks Using Virtual Channels. *IEEE Transactions on Parallel and Distributed Systems* **4**(4), 466–475 (1993)
8. Duato, J., Yalamanchili, S., Lionel, N.: *Interconnection Networks: An Engineering Approach*, 1st edn. Morgan Kaufmann Publishers Inc. (2002)
9. Dumitras, T., Marculescu, R.: On-Chip Stochastic Communication. In: *Proceedings of the Conference on Design, Automation and Test in Europe - Volume 1*, p. 10790 (2003)
10. Ebrahimi, M., Daneshmand, M., Farahnakian, F., Plosila, J., Liljeberg, P., Palesi, M., Tenhunen, H.: HARAQ: Congestion-Aware Learning Model for Highly Adaptive Routing Algorithm in On-Chip Networks. In: *Proceedings of the sixth IEEE/ACM International Symposium on Networks-on-Chip*, pp. 19–26 (2012)
11. Feige, U., Raghavan, P.: Exact Analysis of Hot-potato Routing. In: *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*, pp. 553–562 (1992)
12. Ferdman, M., Adileh, A., Kocherber, O., Volos, S., Alisafae, M., Jevdjic, D., Kaynak, C., Popescu, A.D., Ailamaki, A., Falsafi, B.: Clearing the Clouds: A Study of Emerging Scale-Out Workloads on Modern Hardware. In: *Proceedings of the 17th International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 37–48. London, England, UK (2012)
13. Galles, M.: Spider: A High-Speed Network Interconnect. *IEEE Micro* **17**(1), 34–39 (1997)
14. Glass, C.J., Ni, L.M.: The Turn Model for Adaptive Routing. In: *Proceedings of the 19th Annual International Symposium on Computer Architecture*, pp. 278–287. Queensland, Australia (1992)
15. Gratz, P., Grot, B., Keckler, S.W.: Regional Congestion Awareness for Load Balance in Networks-on-Chip. In: *Proceedings of the 14th International Symposium on High-Performance Computer Architecture*, pp. 203–214. Salt Lake City, UT, USA (2008)
16. Grot, B., Hardy, D., Lotfi-Kamran, P., Falsafi, B., Nicopoulos, C., Sazeides, Y.: Optimizing Data-Center TCO with Scale-Out Processors. *IEEE Micro* **32**(5), 52–63 (2012)
17. Hu, J., Marculescu, R.: DyAD: Smart Routing for Networks-on-Chip. In: *Proceedings of the 41st Annual Design Automation Conference*, pp. 260–263. San Diego, CA, USA (2004)
18. Intel: Intel Xeon Processor X5670. <http://ark.intel.com/products/47920/>
19. Intel: A Touchstone DELTA System Description. In: *Technical report, Supercomputer Systems Division, Intel Corporation* (1991)
20. International Technology Roadmap for Semiconductors (ITRS), 2011 Edition. URL <http://www.itrs.net/Links/2011ITRS/Home2011.htm>
21. Kahng, A.B., Li, B., Peh, L.S., Samadi, K.: ORION 2.0: A Fast and Accurate NoC Power and Area Model for Early-Stage Design Space Exploration. In: *Proceedings of the Conference on Design, Automation, and Test in Europe*, pp. 423–428. Nice, France (2009)
22. Kim, J., Dally, W.J., Abts, D.: Flattened Butterfly: A Cost-Efficient Topology for High-Radix Networks. In: *Proceedings of the 34th Annual International Symposium on Computer Architecture*, pp. 126–137. San Diego, California, USA (2007)
23. Kim, J., Park, D., Theodorides, T., Vijaykrishnan, N., Das, C.R.: A Low Latency Router Supporting Adaptivity for On-chip Interconnects. In: *Proceedings of the 42nd Annual Design Automation Conference*, pp. 559–564. Anaheim, California, USA (2005)
24. Kumar, A., Kundu, P., Singh, A.P., Peh, L.S., Jha, N.K.: A 4.6Tbits/s 3.6GHz Single-cycle NoC Router with a Novel Switch Allocator in 65nm CMOS. In: *Proceedings of the 25th International Conference on Computer Design*, pp. 63–70 (2007)
25. Kumar, A., Peh, L.S., Kundu, P., Jha, N.K.: Express Virtual Channels: Towards the Ideal Interconnection Fabric. In: *Proceedings of the International Symposium on Computer Architecture*, pp. 150–161. San Diego, California, USA (2007)
26. Li, M., Zeng, Q.A., Jone, W.B.: DyXY: A Proximity Congestion-aware Deadlock-free Dynamic Routing Method for Network on Chip. In: *Proceedings of the 43rd Annual Design Automation Conference*, pp. 849–852. San Francisco, CA, USA (2006)
27. Lin, X., Ni, L.: Multicast Communication in Multicomputer Networks. *IEEE Transactions on Parallel and Distributed Systems* **4**(10), 1105–1117 (1993)

28. Lotfi-Kamran, P., Daneshtalab, M., Lucas, C., Navabi, Z.: BARP-A Dynamic Routing Protocol for Balanced Distribution of Traffic in NoCs. In: Proceedings of the Conference on Design, Automation and Test in Europe, pp. 1408–1413. Munich, Germany (2008)
29. Lotfi-Kamran, P., Grot, B., Falsafi, B.: NOC-Out: Microarchitecting a Scale-Out Processor. In: Proceedings of the 45th Annual IEEE/ACM International Symposium on Microarchitecture, pp. 177–187. Vancouver, BC, Canada (2012)
30. Lotfi-Kamran, P., Grot, B., Ferdman, M., Volos, S., Kocberber, O., Picorel, J., Adileh, A., Jevdjic, D., Idgunji, S., Ozer, E., Falsafi, B.: Scale-Out Processors. In: Proceedings of the 39th Annual International Symposium on Computer Architecture, pp. 500–511. Portland, Oregon, USA (2012)
31. Lotfi-Kamran, P., Rahmani, A.M., Daneshtalab, M., Afzali-Kusha, A., Navabi, Z.: EDXY - A Low Cost Congestion-Aware Routing Algorithm for Network-on-Chips. *Journal of Systems Architecture* **56**(7), 256–264 (2010)
32. Ma, S., Enright Jerger, N., Wang, Z.: DBAR: An Efficient Routing Algorithm to Support Multiple Concurrent Applications in Networks-on-chip. In: Proceedings of the 38th Annual International Symposium on Computer Architecture, pp. 413–424 (2011)
33. Marculescu, R., Ogras, U.Y., Peh, L.S., Jerger, N.E., Hoskote, Y.: Outstanding Research Problems in NoC Design: System, Microarchitecture, and Circuit Perspectives. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **28**(1), 3–21 (2009)
34. Michelogiannakis, G., Balfour, J., Dally, W.J.: Elastic-Buffer Flow Control for On-Chip Networks. In: Proceedings of the 15th IEEE International Symposium on High-Performance Computer Architecture, pp. 151–162. Raleigh, NC, USA (2009)
35. Moscibroda, T., Mutlu, O.: A Case for Bufferless Routing in On-Chip Networks. In: Proceedings of the 36th Annual International Symposium on Computer Architecture, pp. 196–207 (2009)
36. Ni, L.M., McKinley, P.K.: A Survey of Wormhole Routing Techniques in Direct Networks. *Computer* **26**(2), 62–76 (1993)
37. Nilsson, E., Millberg, M., Oberg, J., Jantsch, A.: Load Distribution with the Proximity Congestion Awareness in a Network on Chip. In: Proceedings of the Conference on Design, Automation and Test in Europe - Volume 1, p. 11126 (2003)
38. Ogras, U.Y., Hu, J., Marculescu, R.: Key Research Problems in NoC Design: A Holistic Perspective. In: Proceedings of the 3rd International Conference on Hardware/Software Codesign and System Synthesis, pp. 69–74. Jersey City, NJ, USA (2005)
39. Ozer, E., Flautner, K., Idgunji, S., Saidi, A., Sazeides, Y., Ahsan, B., Ladas, N., Nicopoulos, C., Sideris, I., Falsafi, B., Adileh, A., Ferdman, M., Lotfi-Kamran, P., Kuu-lusa, M., Marchal, P., Minas, N.: EuroCloud: Energy-Conscious 3D Server-on-Chip for Green Cloud Services. In: Proceedings of the Workshop on Architectural Concerns in Large Datacenters in conjunction with ISCA (2010)
40. Ramanujam, R.S., Lin, B.: Destination-based Adaptive Routing on 2D Mesh Networks. In: Proceedings of the 6th ACM/IEEE Symposium on Architectures for Networking and Communications Systems, pp. 19:1–19:12 (2010)
41. Ramanujam, R.S., Lin, B.: Destination-based Congestion Awareness for Adaptive Routing in 2D Mesh Networks. *ACM Transactions on Design Automation of Electronic Systems* **18**(4), 60:1–60:27 (2013)
42. Schonwald, T., Zimmermann, J., Bringmann, O., Rosenstiel, W.: Fully Adaptive Fault-Tolerant Routing Algorithm for Network-on-Chip Architectures. In: Proceedings of the 10th Euromicro Conference on Digital System Design Architectures, Methods and Tools, pp. 527–534. Lubeck, Germany (2007)
43. Shin, J.L., Tam, K., Huang, D., Petrick, B., Pham, H., Hwang, C., Li, H., Smith, A., Johnson, T., Schumacher, F., Greenhill, D., Leon, A.S., Strong, A.: A 40nm 16-Core 128-Thread CMT SPARC SoC Processor. In: Proceedings of the IEEE International Solid-State Circuits Conference, pp. 98–99. San Francisco, CA, USA (2010)
44. Singh, A., Dally, W.J., Gupta, A.K., Towles, B.: GOAL: A Load-Balanced Adaptive Routing Algorithm for Torus Networks. In: Proceedings of the 30th Annual International Symposium on Computer Architecture, pp. 194–205. Tel-Aviv, Israel (2003)
45. Vangal, S.R., Howard, J., Ruhl, G., Dighe, S., Wilson, H., Tschanz, J., Finan, D., Singh, A., Jacob, T., Jain, S., Erraguntla, V., Roberts, C., Hoskote, Y., Borkar, N., Borkar,

- 
- S.: An 80-Tile Sub-100-W TeraFLOPS Processor in 65-nm CMOS. *IEEE Journal of Solid-State Circuits* **43**(1), 29–41 (2008)
46. Woo, S.C., Ohara, M., Torrie, E., Singh, J.P., Gupta, A.: The SPLASH-2 Programs: Characterization and Methodological Considerations. In: *Proceedings of the 22nd International Symposium on Computer Architecture*, pp. 24–36. S. Margherita Ligure, Italy (1995). DOI 10.1145/223982.223990