

Low Latency and Energy Efficient Multicasting Schemes for 3D NoC-based SoCs

Xiaohang Wang^{1,5}, Maurizio Palesi², Mei Yang³

¹ College of Computer Science and Technology

⁵ Department of Information Science and Electronic Engineering, Zhejiang University, China
{baikeina, liupeng}@zju.edu.cn

²Kore University, Enna, Italy
maurizio.palesi@unikore.it

Yingtao Jiang³, Michael C. Huang⁴, Peng Liu^{5*}

³Department of Electrical and Computer Engineering, University of Nevada Las Vegas, USA
{mei.yang, yingtao}@egr.unlv.edu

⁴Electrical and Computer Engineering department, Rochester University, USA
michael.huang@rochester.edu

Abstract—In this paper, two topology oriented multicast routing algorithms, MXYZ and AL+XYZ, are proposed to support multicasting in 3D Networks on Chips (NoCs). In specific, MXYZ is a dimension order multicast routing algorithm that targets 3D NoC systems built upon regular topologies, while AL+XYZ is applicable to NoCs with irregular topologies. If the output channel found by MXYZ is not available (i.e. in the same region), an alternative output channel is used to forward/replicate the packets in AL+XYZ. MXYZ is evaluated against a path based regular topology oriented multicast routing and AL+XYZ against an irregular region oriented multiple unicast routing algorithm. Our experimental results have demonstrated that the proposed MXYZ and AL+XYZ schemes have lower latency and energy consumption than the conventional path based multicast routing and the multiple unicast routing algorithms, meriting them to be more suitable for supporting multicasting in 3D NoC systems.

Keywords: Networks on Chip (NoC), multicast, routing algorithms

I. INTRODUCTION

With emerging three dimensional integration circuit technique (3D IC) [1], multi-core and many-core systems can now be connected through 3D Networks on Chips (NoCs) [2], which unquestionably offer high computation parallelism to vastly boost the performance. It has been observed that following the programming models developed for these 3D NoC-based SoCs, many applications SoCs exhibit significant one-to-many and one-to-all communications, which mandates efficient multicast support, especially at the hardware level, to help reduce both transmission latency and power consumption [3, 4].

The need of efficient hardware collective communication support is complicated by topological irregularity which might be caused by virtualization [5] or faulty components. To utilize the large number of computing resources in 3D NoC based multi-core and many-core systems, virtualization is used to facilitate multiple applications running in the system simultaneously. Unfortunately, the tile regions allocated for these applications may not necessarily be in regular shapes (e.g. 2D/3D mesh/torus). On the other hand, when there exist faulty links/nodes, the NoC topology may also end up being irregular. For example, Fig. 1 shows the topological irregularity caused

by virtualization in a 3D NoC based many-core system. In Fig. 1, three applications run simultaneously in the system. The regions allocated for applications 1, 2 and 3 are not regular cuboids.

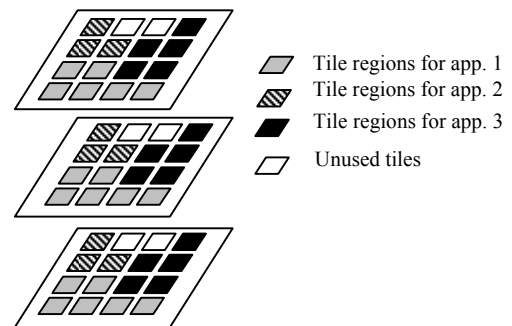


Figure 1. Illustration of irregular regions caused by virtualization in 3D NoC based many-core systems

In the open literature, there has been no work addressing irregular region oriented multicasting for 3D NoC systems. Even multicast routing algorithms for regular 3D NoCs are very rare. In general, the multicast routing algorithms can be broadly classified into two categories: path based multicast and tree based multicast [6]. In path based multicast routing, the multicast packets traverse along the Hamiltonian paths to reach the destinations. In tree-based multicast routing, packets have to be replicated to reach the destination when necessary. Tree based multicast routing might have lower transmission latency than the path-based approaches, as in path based multicast, packets might need to traverse non-minimal paths to form Hamiltonian paths [7]. Later in Fig. 3 (Section II.B), an example will be given to show the difference of the two kinds of multicast routing algorithms. In [8], a path based multicast routing algorithm is proposed, in which the network is partitioned, and a packet is generated following a Hamiltonian path in each partition. This idea can be borrowed for 3D regular NoC systems, but it is not applicable to systems with irregular topologies. Regular topology oriented multicast/unicast routing algorithms are not suitable for irregular topologies, as the tiles in path found by the routing algorithms may not be in the region. For example, neither could the above path based multicast routing or traditional

This work is supported by NSFC under grant 60873112 and also in part by grant 61028004.

*The correspondence author is Peng Liu

dimension ordered routing (DOR) work for the regions in Fig. 1. In the literature, routing tables are typically used to route packets in irregular topologies [9]. However, if the traffic could not be predicted (e.g. the traffic caused by the cache coherence protocol [4]), setting up routing tables may incur additional latency.

In our previous work [4, 10], an irregular region oriented multicasting strategy was proposed for 2D NoC systems. In this work, we will extend the work to deal with both regular and irregular region oriented multicasting problems for both regular and irregular 3D NoC systems. Correspondingly, two respective schemes, MXYZ for supporting multicasting in regular 3D NoC systems, and AL+XYZ for supporting multicasting in irregular regions in 3D NoC systems, are proposed.

The organization of the paper is as follows. Section II introduces the NoC architectural model and a multicast routing algorithm named MXYZ for regular 3D NoCs. Section III proposes a strategy to support multicast routing in irregular regions and an irregular region oriented multicast routing algorithm AL+XYZ based on MXYZ. Section IV evaluates the performance of MXYZ and AL+XYZ. Finally, Section V concludes the paper.

II. NOC ARCHITECTURAL MODEL AND MULTICAST SUPPORT FOR REGULAR 3D NOCS

A. Architectural and Energy Models

The target 3D NoC architecture in consideration is a stacked homogeneous system, where each tile is composed of a processor core and a router [11], as shown in Fig. 2. Each router has seven ports, i.e., East, West, North, South, Local, Up and Down. The 3D NoC system is composed by $N \times N \times N_z$ tiles. That is, the 3D NoC system has N_z layers and each layer has $N \times N$ tiles.




Figure 2. 3D NoC architecture

To support multicast routing, the replication unit is used at each input port to replicate the flits of a multicast packet according to the decision of the routing unit. Asynchronous replication scheme is assumed which allows multiple replicated flits to be forwarded independently [6]. If one replicated flit is blocked, other replicated flits can be forwarded asynchronously.

The average energy consumption for a unicast communication which sends one bit from source tile s to destination tile t can be represented as,

$$E_{Unicast}^{s,t} = \eta E_{Rbit} + \eta_H E_{LHbit} + \eta_V E_{LVbit} \quad (1)$$

where η is the number of routers traversed from tile s to tile t , E_{Rbit} is the bit energy consumed by the router, E_{LHbit} is the bit energy consumed on each horizontal link, and E_{LVbit} is the bit energy consumed on each vertical link. η_H and η_V are the numbers of horizontal and vertical links on the communication path. Following the wire model in [12], $E_{LHbit} = d_H V_{dd}^2 C_{wireH} / 2$ and $E_{LVbit} = d_V V_{dd}^2 C_{wireV} / 2$, where d_H and d_V are the lengths of the horizontal and vertical links, V_{dd} is the supply voltage, and C_{wireH} and C_{wireV} are the wire capacitance of horizontal and vertical links, respectively.

The average power consumption for a multicast communication which sends 1 bit from source tile s to the set of destination tiles \bar{D} can be represented as,

$$E_{Multicast}^{s,\bar{D}} = \eta E_{Rbit} + \eta_{MH} E_{LHbit} + \eta_{MV} E_{LVbit} \quad (2)$$

where η is the total number of routers, and η_{MV} and η_{MH} the total numbers of links that are on the multicast path from tile s to all tiles in \bar{D} , respectively.

B. Multicast Support for Regular 3D NoCs

In this section, a tree-based, dimension-ordered multicast routing algorithm MXYZ is proposed to support multicasting in regular 3D NoC systems. The routing algorithm of MXYZ is shown in Fig. 3.

// The destination set \bar{D} is partitioned at the source tile s , separate copies of the // packets are sent to the output channels.

for each intermediate tile t ,

if ($t \in \bar{D}$)

A packet is replicated and ejected to the local processor.

else if t has the same Z coordinates as one of the destinations

A packet is replicated and sent in the corresponding Z direction.

else if t has the same Y coordinates as one of the destinations

A packet is replicated and sent in the corresponding Y direction.

else

t forwards the packet in the X direction.

Figure 3. MXYZ routing algorithm

Fig. 4(a) shows how MXYZ works. The source tile 6 sends multicast packets to these destinations: tiles 0, 4, 12, 28, 44.

- At the source tile 6, two packets are generated, i.e. one packet is sent to the West output channel (packet A) and the other to the East output channel (packet B).
- Packet A follows the XY routing path to reach the destination tile 0.
- Packet B reaches the destination tile 4 and then another packet is replicated (packet C) to the Up output channel to reach the destination tile 28. Packet B at tile 4 is forwarded to the North output channel to reach destination tile 12.
- Finally, at the destination tile 12, a packet is replicated to the Up output channel to reach the destination tile 44 (packet D).

To show the efficiency of tree based MXYZ over path based multicast routings, MXYZ is compared against the path based hybrid partition (HP) multicast routing algorithm [8] in Fig. 4. Fig. 4 shows an example where MXYZ could generate paths of fewer hop counts than HP. In Fig. 4 (b), there are totally 18 hop counts in the HP communication paths. In contrast, in Fig. 4 (a), there are only 9 hop counts in the MXYZ communication paths. The Hamiltonian paths in HP might lead to non-minimal communication path. Therefore the packets need to traverse in paths with more hop counts which might cause higher energy consumption and higher latency. For example, the packets could not be replicated to the west output at tile 6 to reach destination tile 0 in HP (Fig. 4 (b)). Instead, in Fig. 4 (a), a packet could be replicated to the west output at tile 6 directly to reach destination tile 0. Thus, the hop count of the paths could be reduced. Similarly, to reach the destination tile 44, the packet needs to traverse more hop counts in HP than that in MXYZ.

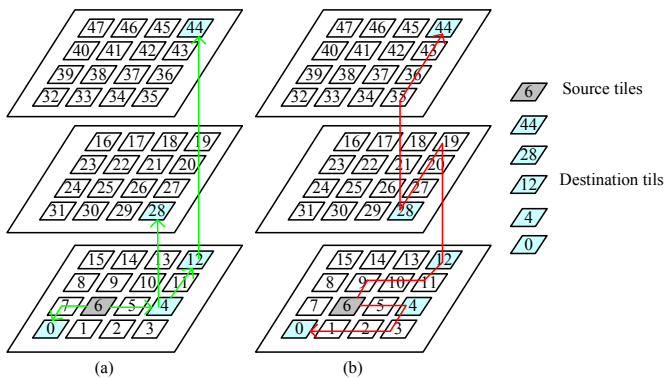


Figure 4. Comparison of HP and MXYZ in terms of hop counts by an example. (a) The communication paths following HP routing. (b) The communication paths following MXYZ routing.

III. IRREGULAR REGION ORIENTED MULTICASTING IN 3D NoC

A. Motivating Example

Before the irregular region oriented multicast routing algorithm is described in detail, a simple example is used to illustrate the basic idea of the irregular region oriented multicast routing strategy.

In Fig. 5, suppose the multicast routing algorithm is based on MXYZ. The red dashed arrow indicates the routing path from the source tile s to the destination tiles d_1 and d_2 assuming the region is regular. However, the shape of the region in consideration is not regular, i.e. the West output port of tile 11 does not connect to a tile inside the same region. Thus, the packets in the red arrow cannot reach the destinations. However, if the path turns to North at tile 11, the packet could reach both destinations (shown as the green solid arrow).

From this example, a simple yet effective irregular region oriented multicasting routing strategy can be derived. That is, a multicast routing algorithm supporting regular NoC systems is used as the basic routing algorithm. In case that the output

channel found by the basic multicast routing is not available (e.g., the tile connected to the output channel is not in the same region), an alternative output channel shall be selected. For example, at a tile, if the output channel in its X direction is not available, the packet can be forwarded to the output channel in its Y direction.




Figure 5. An example showing the basic idea of the multicasting algorithm

B. Assumptions and Definitions

The shapes of the regions in consideration should be near convex. That is, a region is composed by N_z layers of planar sub-regions and each planar sub-region has the same shape. Each sub-region should be convex, i.e. the minimal path of any two tiles should be inside the same sub-region (Fig. 5). Regions with such shapes could be generated by the online incremental mapping algorithm described in our previous work [13]. We will consider regions with convex and non-convex polyhedron shapes in the future.

To support multicasting inside the irregular regions, a set of connectivity bits are added.

Definition 1 Connectivity bits. Each router has 4 connectivity bits, C_N , C_E , C_S , and C_W , and each of these bits defines the connectivity at the specific output direction. Suppose a tile has coordinate (x, y, z) , and C_N is 1 if tile (x, y, z) and its north neighbor tile $(x, y-1, z)$ are in the same region. Similarly, C_x is set to 1 if the tile and its neighbor tile in the x direction are in the same region.

Definition 2 Extended connectivity bits. To support up to M overlapped regions, each router located at the tile with coordinate (x, y, z) has $4 \times M$ connectivity bits [4], $\{C_{N[1]}, \dots, C_{N[M]}\}$, $\{C_{W[1]}, \dots, C_{W[M]}\}$, $\{C_{E[1]}, \dots, C_{E[M]}\}$, and $\{C_{S[1]}, \dots, C_{S[M]}\}$. Suppose a tile has coordinate (x, y, z) , $C_{x[q]} = 1$ ($q=1, \dots, M$) if tile (x, y, z) and its neighbor tile in the x direction are in the same region with ID q . Extended connectivity (EC) bits EC_N , EC_E , EC_S , and EC_W are defined as follows. Given the region ID q , $EC_N = C_{N[q]}$, $EC_W = C_{W[q]}$, $EC_E = C_{E[q]}$, $EC_S = C_{S[q]}$.

Definition 3 Virtual regions. At each tile with coordinate (x, y, z) , the network is partitioned into 10 virtual regions, R_0, R_1, \dots, R_9 , such that,

$$\text{Tiles with coordinate } (x_0, y_0, z) \in R_0, \text{ if } x_0 > x \text{ and } y_0 < y. \quad (3)$$

$$\text{Tiles with coordinate } (x, y_1, z) \in R_1, \text{ if } y_1 < y. \quad (4)$$

$$\text{Tiles with coordinate } (x_2, y_2, z) \in R_2, \text{ if } x_2 < x \text{ and } y_2 < y. \quad (5)$$

$$\text{Tiles with coordinate } (x_3, y, z) \in R_3, \text{ if } x_3 < x. \quad (6)$$

$$\text{Tiles with coordinate } (x_4, y_4, z) \in R_4, \text{ if } x_4 < x \text{ and } y_4 > y. \quad (7)$$

$$\text{Tiles with coordinate } (x, y_5, z) \in R_5, \text{ if } y_5 > y. \quad (8)$$

$$\text{Tiles with coordinate } (x_6, y_6, z) \in R_6, \text{ if } x_6 > x \text{ and } y_6 > y. \quad (9)$$

Tiles with coordinate $(x_7, y, z) \in R_7$, if $y_7 > y$. (10)

Tiles with coordinate $(x, y, z_8) \in R_8$, if $z_8 > z$. (11)

Tiles with coordinate $(x, y, z_9) \in R_9$, if $z_9 < z$. (12)

Some nodes at the boundary may have regions without nodes.

Fig. 6 shows the partition of the network into 10 virtual regions at tile 25. The virtual regions include eight planar quadrants and two vertical partitions. The partition of the network helps in identifying the destination locations, which actually determines how to replicate packets.

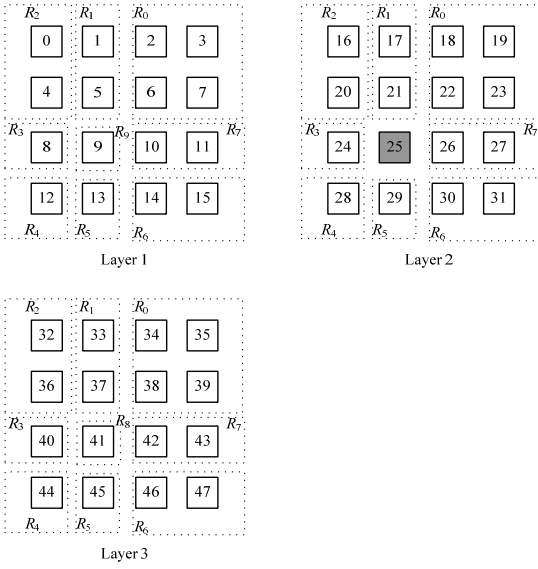


Figure 6. The virtual regions of tile 25.

C. Irregular Region Oriented Multicasting Routing Strategy and Algorithm

Based on the observation in Fig. 5, an irregular region oriented multicasting routing framework can be derived as follows. 1) A multicast routing supporting regular 3D NoC systems is used as the basic multicast routing algorithm. 2) In case that the output channel found by the basic multicast routing algorithm is not available (e.g. the tile connected to the output channel is not in the same region), an alternative output channel is selected to reach the destinations.

Based on this strategy, an irregular region oriented multicasting routing algorithm named as Alternative XYZ (AL+XYZ) is proposed for 3D NoC systems, which is based on MXYZ algorithm in section II. However, it shall be noted that, apart from MXYZ, any multicast routing algorithm that works for regular 3D NoC systems can be used as the basic multicast routing algorithm to develop irregular region oriented multicasting routing algorithms using the above strategy. Fig. 7 shows the multicast routing unit. The pseudo code of AL+XYZ is shown in Fig. 8 which corresponds to the MR module in Fig. 7. In Fig. 8, the following two sets of vectors are used.

- 1) Destination bits vectors for each region. Ten bit vectors $IN_{R_0}, \dots, IN_{R_9}$, each with $N \times N \times N_z$ bits, are used to represent the destinations within one region.

- 2) Output destination bit vectors. Seven bit vectors $N_DestSet$, $S_DestSet$, $E_DestSet$, $W_DestSet$, $Up_DestSet$, and $Down_DestSet$, each of $N \times N \times N_z$ bits, are used to represent the destinations in each output channel.

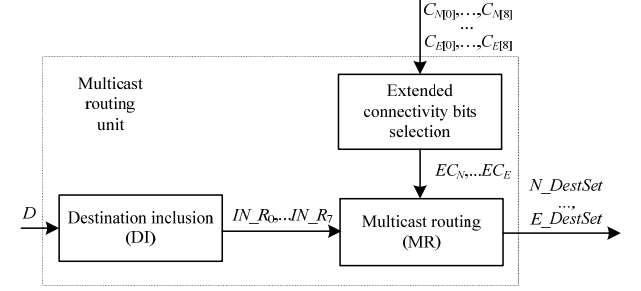


Figure 7. Structure of the multicast routing unit.

// Alternative multicast XYZ for irregular regions (AL+XYZ)
// Bit vectors $N_DestSet$, $E_DestSet$, $W_DestSet$, $S_DestSet$, $Up_DestSet$,
// and $Down_DestSet$ are initialized to 0
// Bit vectors IN_{R_i} ($0 \leq i \leq 9$) is 1 if R_i includes any multicast destination.

```
if (IN_R0) begin// vector IN_R0 is non-zero
if (EC_E == 0) N_DestSet= IN_R0 OR N_DestSet
else E_DestSet= IN_R0 OR E_DestSet
end
```

```
if (IN_R1) N_DestSet= IN_R1 OR N_DestSet
```

```
if (IN_R2) begin// vector IN_R2 is non-zero
if (EC_W == 1) W_DestSet= IN_R2 OR W_DestSet
else N_DestSet= IN_R2 OR N_DestSet
end
if (IN_R3) W_DestSet= IN_R3 OR W_DestSet
```

```
if (IN_R4) begin// vector IN_R4 is non-zero
if (EC_W == 1) W_DestSet= IN_R4 OR W_DestSet
else S_DestSet= IN_R4 OR S_DestSet
end
```

```
if (IN_R5) S_DestSet= IN_R5 OR S_DestSet
```

```
if (IN_R6) begin// vector IN_R6 is non-zero
if (EC_E == 1) E_DestSet= IN_R6 OR E_DestSet
else S_DestSet= IN_R6 OR S_DestSet
end
```

```
if (IN_R7) E_DestSet= IN_R7 OR E_DestSet
if (IN_R8) Up_DestSet= IN_R8 OR Up_DestSet
if (IN_R9) Down_DestSet= IN_R9 OR Down_DestSet
```

Figure 8. The pseudo code of AL+XYZ

If MXYZ is used as the basic routing algorithm, the alternative output direction is in the Y direction. Thus, if there are destinations in R_0 and R_2 , North (Y+) is selected as alternative output direction. If there are destinations in virtual regions 4 and 6, South (Y-) is selected as alternative output direction. Note that, if destinations are in $R_1, R_3, R_5, R_7, R_8, R_9$ i.e., the X+, X-, Y+, Y-, Z-, Z+ directions, there is no alternative output direction. The reason is that, according to the assumption in Section II, the region must be near convex which ensures that between any pair of tile there exists at least a minimal path inside the region. It is clear that for each destination in $R_1, R_3, R_5, R_7, R_8, R_9$ there is only one minimal

path to these destinations. Hence, there is no alternative output direction in those virtual regions. Thus, only the alternative output directions for R_0, R_2, R_4, R_6 are found.

To avoid deadlock, two virtual networks are used, VN_0 and VN_1 . VN_0 does not allow packets to turn to North while VN_1 does not allow packets to turn to South. The virtual network to be used is decided for each packet at the source router and cannot be changed at the intermediate routers.

D. Hardware cost

In order to investigate the overhead of hardware cost incurred by the multicast routing unit, a 7×7 unicast router configured with two virtual channels and without multicast routing unit is compared with the 7×7 AL+XYZ router, in terms of area and power consumption. Synopses Design Compiler 2009.06 is used as the synthesis tool, with TSMC 65nm CMOS library. The power consumption and area of AL+XYZ router are 50 mW and 153776 um^2 . As a reference, the power consumption and area of the 7×7 unicast router are 43 mW and 141775 um^2 . Thus, the area overhead of AL+XYZ over the unicast router is less than 7%.

IV. PERFORMANCE EVALUATION

A. Experimental setup

TABLE I. SIMULATION CONFIGURATION

Network size	4x4x3
Destination Size	8
Packet Length	8
Flit Size(bits)	75
VC depth	8
VC number	2
Distribution of the Destinations	Uniform
C_{Hwire} (fF/mm)	212.12
C_{Vwire} (fF/mm)	600
Vias length(um)	50




Figure 9. Two regions in the 4x4x3 NoC system

Table I lists the simulation settings. multicast to unicast ratio (MUR) is used to measure the ratio of the multicast over the entire communication. MUR is defined to be the packet number of multicasting to that of unicasting. Noxim [14] is used as the NoC system simulator.

Two sets of experiments are performed. The first set of experiments is used to compare MXYZ against a path based hybrid partition (HP) multicast routing algorithm [8] in regular 3D NoC systems, in terms of power consumption and latency.

The second set of experiments is used to compare AL+XYZ and multiple unicast (MUC) in irregular regions in 3D NoC systems. In case of unavailable output channel, MUC selects an alternative channel for output. Fig. 9 shows the two regions visualized in the 4x4x3 NoC system.

B. The Efficiency of MXYZ in Regular 3D NoC Systems

Fig. 10 (a) shows the normalized energy consumption of HP over MXYZ, i.e. the energy consumption of HP divided by that of MXYZ. Fig. 10 (a) shows that, when the injection rate is high (e.g. above 0.09), the energy consumption of HP is about $1.7x \sim 1.9x$ over that of MXYZ. Fig. 10 (b) shows the latency of HP and MXYZ. As the inject rate increases, the latency of HP increases more rapidly than that of MXYZ. The reason is that, the Hamiltonian paths in HP result in non-minimal path for the packets. Thus, both the energy consumption and latency are increased.

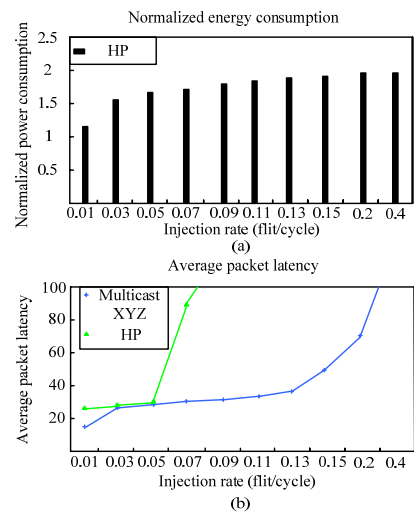


Figure 10. Comparison of HP and MXYZ in terms of (a) energy consumption and (b) latency when MUR = 0.3.

C. The Efficiency of AL+XYZ in Irregular Regions in 3D NoC Systems

Fig. 11 (a) shows the normalized energy consumption of MUC over AL+XYZ when the MUR is low (MUR=0.05) which means that the multicast communication has a small ratio over the entire traffic. From Fig. 11 (a), we can see that, when the injection rate is low (e.g. 0.01), the energy consumption of MUC is the same as that of AL+XYZ. When the injection rate increases, the energy consumption of MUC is about $1.3x \sim 1.4x$ of that of AL+XYZ.

Fig. 11 (b) shows the energy consumption when the multicasting communication has a large ratio over the entire traffic. When the injection rate increases (e.g. above 0.03), the energy consumption of MUC is higher than that of AL+XYZ by a factor of $1.7x \sim 2.2x$. This can be attributed to the fact that MUC generates more packets than AL+XYZ. Multiple unicast packets are generated to each destination. Thus, the extra packets contribute to higher energy consumption. The higher the ratio of multicast traffic, the higher is MUC energy consumption overhead.

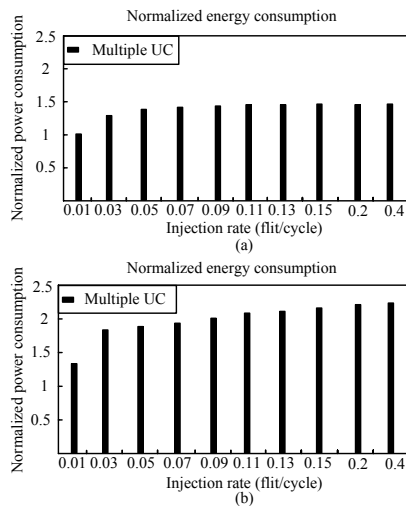


Figure 11. Comparison of MUC and AL+XYZ in terms of energy consumption with (a) MUR = 0.05 and (b) MUR = 0.3.

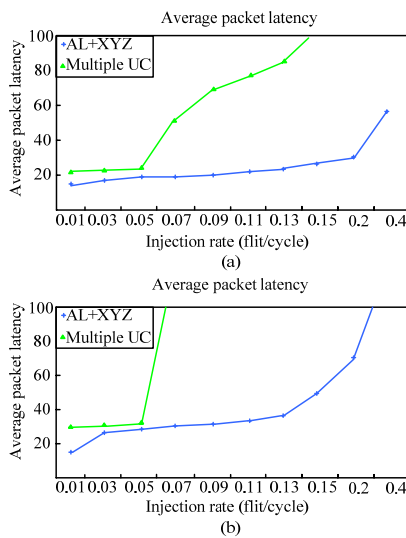


Figure 12. Comparison of MUC and AL+XYZ in terms of latency with (a) MUR = 0.05 and (b) MUR = 0.3.

Fig. 12 (a) shows the latency when MUR is low (MUR=0.05). We can see that, when the injection rate is low (below 0.07), the latency of MUC is relatively low. When the injection rate is above 0.07, the latency of MUC increases rapidly. The increase in the latency of AL+XYZ is much slower until the injection rate is 0.4.

Fig. 12 (b) shows the latency when MUR is high (MUR = 0.3). MUC saturates at injection rate = 0.7. AL+XYZ has much lower latency compared to that of MUC. The reason is that, MUC generates extra packets to support multicasting. Thus, the increase in the number of packets contributes to severe congestion and earlier saturation point. AL+XYZ, on the other hand, avoids extra packets generation which can help alleviate the congestion situation.

V. CONCLUSION

In this paper, two efficient multicast routing algorithms MXYZ and AL+XYZ were proposed to support 3D NoC systems, respectively. MXYZ is a dimension order multicast routing algorithm (DOR) for regular 3D NoC systems. To support multicast routing in irregular regions, AL+XYZ was developed based on the following strategy. If an output channel found by a multicast routing algorithm for regular NoC systems (e.g. MXYZ) is not available, an alternative output channel is selected. In regular 3D NoC systems, MXYZ could outperform a path based multicast routing scheme HP in terms of latency and energy consumption. In irregular regions in 3D NoC systems, AL+XYZ was shown to outperform another multicast scheme, i.e. multiple unicast routing (MUC), in both latency and energy consumption. Due to lower latency and energy consumption than their counterparts, MXYZ and AL+XYZ are justified to be used as effective multicast routing algorithms in regular or virtualized irregular regions in 3D NoC systems.

REFERENCES

- [1] A. Sheibanyrad, 3D Integration for NoC-based SoC Architectures: Springer Verlag, 2010.
- [2] S. Borkar, "Thousand core chips: a technology perspective," in Proc. Design Automation Conf, pp. 746-749, 2007.
- [3] L. Wang, Y. Jin, H. Kim, and E. J. Kim, "Recursive partitioning multicast: a bandwidth-efficient routing for on-chip," in Proc. ACM/IEEE Int'l Symp. Networks-on-Chip, pp. 64-73, 2009.
- [4] X. Wang, M. Yang, Y. Jiang, and P. Liu, "On an efficient NoC multicasting scheme in support of multiple applications running on irregular sub-networks," *Microprocessors and Microsystems*, vol. 35, no. 2, pp. 119-129, 2011.
- [5] F. Triviño, J. L. Sánchez, F. J. Alfaro, and J. Flich, "Virtualizing network-on-chip resources in chip-multiprocessors," *Microprocessors and Microsystems*, vol. 35, no. 2, pp. 230-245, 2011.
- [6] J. Duato, S. Yalamanchili, and L. Ni, *Interconnection networks an engineering approach*: Morgan Kaufmann, 2002.
- [7] N. E. Jerger, L. S. Peh, and M. Lipasti, "Virtual circuit tree multicasting: a case for on-chip hardware multicast support," in Proc. Int'l Symp. Computer Architecture, pp. 229-240, 2008.
- [8] M. Ebrahimi, M. Daneshlab, P. Liljeberg, and H. Tenhunen, "Partitioning methods for unicast/multicast traffic in 3D NoC architecture," in Proc. IEEE Int'l Symp Design and Diagnostics of Electronic Circuits and Systems, pp. 127-132, 2010.
- [9] A. Mejia, M. Palesi, J. Flich, S. Kumar, P. López, R. Holmark, and J. Duato, "Region-based routing: a mechanism to support efficient routing algorithms in NoCs," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 17, no. 3, pp. 356-369, 2009.
- [10] X. Wang, M. Yang, Y. Jiang, and P. Liu, "Efficient multicasting scheme for irregular mesh-based NoCs," in Proc. IEEE Int'l SoC Conf., pp. 384-387, 2010.
- [11] V. F. Pavlidis and E. G. Friedman, "3-D topologies for networks-on-chip," *IEEE Trans. Very Large Scale Integration Systems*, vol. 15, no. 10, pp. 1081 - 1090, 2007.
- [12] H. Matsutani, M. Koibuchi, and H. Amano, "Tightly-coupled multi-layer topologies for 3-D NoCs," in Proc. Int'l Conf. Parallel Processing, pp. 75 - 85, 2007.
- [13] X. Wang, M. Palesi, M. Yang, Y. Jiang, M. Huang, and P. Liu, "Power-aware run-time incremental mapping for 3-D networks-on-chip," accepted by IFIP Int'l Conf. Network and Parallel Computing.
- [14] Noxim. [Online]. Available: www.sourceforge.net/Noxim.