High Performance Collective Communication-Aware 3D Network-on-Chip Architectures

Biresh Kumar Joardar
School of EECS
Washington State University
Pullman, USA
bjoardar@eecs.wsu.edu

Karthi Duraisamy
School of EECS
Washington State University
Pullman, USA
kduraisa@eecs.wsu.edu

Partha Pratim Pande
School of EECS
Washington State University
Pullman, USA
pande@eecs.wsu.edu

***Abstract*— 3D Network-on-Chip (NoC) architectures are capable of achieving better performance and lower energy consumption compared to their planar counterparts. However, conventional 3D NoCs are not efficient in handling collective communication. Existing works mainly explore Path and Tree multicast distribution schemes for 3D NoCs. However, both these mechanisms involve high network latency and lack scalability. In this work, we propose a SMART (Single-cycle Multi-hop Asynchronous Repeated Traversal) 3D NoC architecture that is capable of achieving high-performance collective communication. The proposed High-Performance SMART (HP-SMART) 3D NoC achieves 65% and 31% latency improvements compared to the existing Path and Tree multicast-based 3D NoCs respectively. HP-SMART 3D NoC also achieves significant improvement in message latency compared to its 2D counterpart.**

I. INTRODUCTION

The emergence of three-dimensional (3D) Network-on-Chip (NoC) has revolutionized the design of high-performance and energy efficient manycore chips. Manycore platforms designed with 3D NoCs outperform their 2D counterparts in terms of latency, throughput and energy consumption [1]. The achievable performance of a 3D NoC can be improved further if the overall communication architecture is optimized according to the traffic patterns generated by the applications under consideration [2].

Collective communication refers to (a) distribution of a multicast/broadcast message from a single source node to multiple/all other nodes in the system and (b) aggregation of acknowledgments (ACKs) from multiple/all nodes in the system to a single destination node. This type of communication pattern is observed in various emerging applications. Neural-network based computing, real-time object recognition processing, neuromorphic computing, managing and configuring the on-chip networks and coherency protocols like AMD's Hammer generate significant multicast/broadcast data [3][4].

Collective communication traffic is more challenging to handle on conventional NoC platforms than the unicast traffic due to the following reasons: (1) collective communication requires long-range data exchanges and (2) dense collective communication can cause heavy congestion [4]. Without proper support from the underlying NoC architecture and routing protocol, collective communication can lead to high network and queueing latencies [4]. Hence, a 3D NoC, specifically designed to address collective communication is of paramount importance. Despite this fact, the design and analysis of a collective communication-aware 3D NoC is yet to be explored thoroughly. Most of the existing works on this topic have mainly focused on extending the conventional path and tree multicast distribution techniques employed in 2D mesh NoCs to the 3D architecture. However, it is well known that both the path and tree multicast have inherent limitations in scalability [4].

This work was supported, in part by the US National Science Foundation (NSF) grants CNS-1564014 and CCF 1514269 and USA Army Research Office grant W911NF-17-1-0485.

Recently, an efficient collective communication mechanism using the single-cycle multi hop asynchronous repeated traversal (SMART) paradigm has been proposed [5][6]. The SMART NoC employs an efficient router bypass control network along with the conventional 2D mesh interconnect utilizing multi-drop wires integrated with clockless repeaters. With the help of this control network, SMART NoC enables single cycle end-to-end transmissions leading to low-latency collective communication [6]. Despite its advantages, the conventional SMART routing is mainly suited for NoC architectures employing a relatively low clock frequency [5][7]. The physical link length between the communicating nodes is an important factor in determining the achievable clock frequency of SMART NoCs [7]. Compared to 2D NoCs, 3D NoCs reduce the physical link lengths between communicating cores by utilizing the vertical interconnects. The vertical links effectively act as long-range shortcuts. Therefore, compared to the conventional planar architectures, a SMART 3D NoC has the potential to benefit from the shorter path lengths and operate at higher clock frequencies.

Considering all the above facts, here we propose a SMART control network integrated 3D NoC architecture capable of high performance (supporting high clock frequency and low latency) collective communication. The salient features of this work are:

- We design a multi-drop TSV-based vertical link integrated with clock-less buffers. This vertical link enables seamless extension of the SMART control network to 3D NoC.
- We propose a novel routing scheme that can distribute/aggregate the collective communication traffic efficiently in a 3D NoC operating with high clock frequency. This routing mechanism is also capable of distributing (or aggregating) multiple broadcasts (or acknowledgements) simultaneously without causing any congestion.
- We demonstrate the effectiveness of the proposed SMART 3D NoC in handling collective communication through comparative performance evaluations against the path and tree multicast based 3D NoCs and the planar SMART NoC.

II. RELATED WORKS

There are a few lines of research that are relevant to this paper. We summarize them under two categories.

Collective communication-aware 2D NoCs: To meet the requirements of collective communication, a mesh NoC performing broadcast over uncongested trees and incorporated with ACK aggregation and single-cycle-multiport replication mechanisms is proposed [8]. Wireless [4] and surface-wave NoCs [9] enable low-latency collective communication. However, on-chip wireless and surface wave communication are yet to be adopted widely.

3D NoCs for multicast and broadcast: In [10], the authors have proposed a set theory-based model to study the multicast

communication in 3D NoCs. Design of a 3D NoC router to support efficient multicast transmission and congestion and hotspot elimination has been proposed in [11]. A 3D NoC-Bus Hybrid architecture supporting collective communication is discussed in [12]. Path multicast mechanism for 3D NoCs employing multiple network partitions is proposed in [14]. The partition-based methodology introduces parallelization in the Hamiltonian path based routing scheme. Hence, this methodology reduces the time required to complete a broadcast by distributing the message simultaneously across multiple partitions. For the 3D NoCs, it has been shown that a Hybrid Partitioning method comprising of Dual partitioning and Vertical Partitioning can perform better than the common 2D partitioning methods [14].

Tree multicast in 3D has been widely studied. In [13], the authors have proposed two multicast aware routing mechanisms, a regular XYZ tree based routing for the conventional mesh and a fault-tolerant multicast distribution scheme for NoCs with irregular topologies. Unlike path multicast, single or two stage routers can be used to implement tree multicast as demonstrated in [8]. This reduces the impact of intra-router delays in latency.

Both path and tree multicast involve high latency and do not scale with system size [4]. To overcome these limitations, we exploit the benefits introduced by both 3D integration and SMART control mechanism to design a high performance NoC.

III. Collective Communication-Aware SMART 3D NoC Design

In this section, first we elaborate the proposed SMART 3D NoC connectivity followed by the associated routing scheme.

A. Proposed NoC architecture

For the collective communication-aware SMART 3D NoC we employ the control network proposed in [15]. Specifically, we use 1D-SMART control network, which only allows multi-hop traversals in one dimension at a time. The 1D-SMART control network supports higher clock frequency than 2D-SMART [5] that allows bypass at turn routers. Following [15], the proposed NoC consists of two parallel networks, a normal 3D mesh data-transfer network and a control network distributing SMART hop Setup Requests (SSR). It has already been shown that the SSR control does not introduce significant area overheads [15].

Fig. 1 illustrates the proposed NoC connectivity. Each processor tile in the NoC consists of two routers, the primary data-transfer router and an SSR router. In addition to the SSR mesh connectivity, the control network also incorporates single-bit multi-drop pre-SSR wires. The pre-SSR wires span from each SSR router-port to all the other SSR routers that are in same dimension and within HPC_{max} hops, where HPC_{max} is the number of intermediate hops bypassed in one clock cycle. It should be noted that for SMART 3D NoCs, the HPC_{max} values are different for vertical and the planar links (denoted as $VHPC_{max}$ *and* $PHPC_{max}$, respectively) depending on the physical parameters of the TSVs and the planar metal wires employed.

Each *SMART hop* consists of four router pipeline stages, namely: local arbitration, pre-SSR transmission and arbitration, SSR transmission, and single-cycle multi-hop link traversal [15]. Similar to traditional NoC data transfers, in stage 1, the message competes and wins the local arbitration at the source node where the data transfer originates. At stage 2, using the multi-drop pre-SSR wires, a single-bit pre-cursor request is forwarded from the source node to all the intermediate routers that are within $VHPC_{max}$ *or* $PHPC_{max}$ hops. In the same stage, this pre-cursor request competes with other received pre-SSRs at the intermediate nodes. Upon successful pre-SSR arbitrations, a detailed bypass request is forwarded along the SSR wires during stage 3 from the source node to all the intermediate routers in the same dimension and are within $VHPC_{max}$ *or* $PHPC_{max}$ hops. Using the received SSR, intermediate routers establish the desired bypass connectivity at stage 4. This enables the data to traverse multiple intermediate hops in a single cycle.

Previous works [5] have already established the feasibility of planar multi-drop wires. However, no previous work has investigated the possibility of a multi-drop link spanning several vertical layers in a 3D NoC. To address this issue, we model a TSV-based vertical link that branches at each layer using the parameters in [16]. We design the TSV with driver and clock less buffers in 45nm and simulate using Cadence SPECTRE^{TM}. The propagation delay of the multi-drop vertical link spanning over four layers is 260ps. Hence, considering an arbitration delay of 90ps [5], the total time required for communication along a vertical link is approximately 350 ps. Therefore, it is possible to transfer data over multiple vertical layers within a single clock cycle considering a frequency of 2.75GHz. Thus, it is possible to extend the single cycle multi-drop SMART data transfer mechanism along the vertical links of the 3D NoC.

1) Advantages of Employing SMART Control in 3D NoC

For the planar links in SMART NoC, we define $PLPC_{max}$ as the maximum wire length traversed per cycle and is given by:

$$PLPC_{max} = PHPC_{max} \times Link\ length\ per\ hop. \quad (1)$$

For a SMART 2D NoC, employing a high $PLPC_{max}$ value can lower the average SMART hops. However, with increasing $PLPC_{max}$ values (or $PHPC_{max}$ values), the NoC clock frequency decreases linearly, ultimately degrading achievable NoC performance [5]. Compared to conventional planar NoC architectures, the 3D NoC reduces the average communication distance between cores by utilizing the vertical links as long-range shortcuts. Exploiting this advantage, we design high-performance SMART 3D NoC architectures with low $PLPC_{max}$ (and $PHPC_{max}$) values supporting high clock frequencies. For the same clock frequency we achieve comparatively higher $VHPC_{max}$ values, that in turn helps in improving the overall NoC performance. For example, considering a clock frequency of 2.75GHz, the planar metal wires only support single-cycle traversals across a link length of 4mm [5]. However, as mentioned above, the vertical links employing clockless buffers enable single-cycle traversals over four planar layers.

In order to handle the dense collective communication, a NoC operating at high clock frequency and enabling low network latency is desirable. Therefore, SMART 3D NoC is an efficient solution for on-chip collective communication.

B. High-Performance Routing for SMART 3D NoCs

A four-step private virtual tree (PVT) based broadcast distribution mechanism for SMART 2D NoCs was proposed in [6]. In NoCs with short $PLPC_{max}$ values, the PVT mechanism involves a high number of SMART hops. Therefore, PVT mechanism prefers NoCs with large $PLPC_{max}$ values and thus, operate at low clock frequencies. To overcome these issues, we

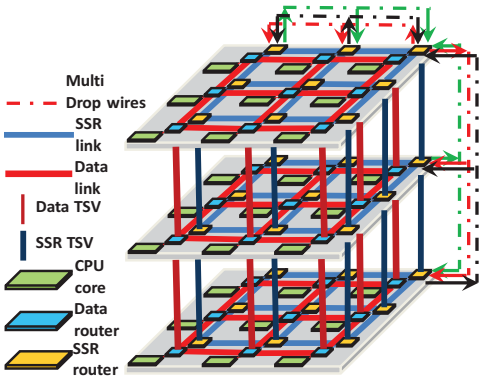

Fig. 1: Illustration of SMART 3D NoC. For better visibility, we only show the connectivity for a single row and column.

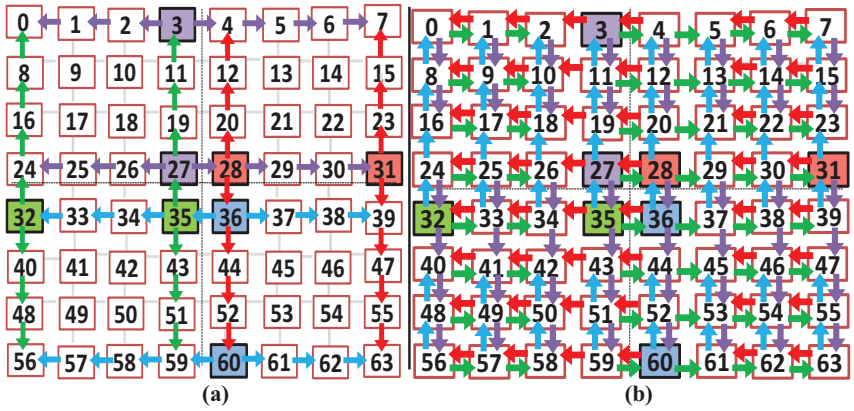

Fig 2: Illustration of planar HP-SMART routing. (a) Step 1 and (b) Step 2. Colored nodes are starting nodes. Each color represents the flow of a unique broadcast message.

introduce High-Performance SMART (HP-SMART) routing, a five-step collective communication mechanism for SMART 3D NoCs operating with high clock frequencies. During step 0 of HP-SMART, a broadcast message is first transferred from the source node to the two starting nodes lying at the center and the edge of the source planar layer. During this step, the message is treated as a normal unicast message and is dynamically arbitrated with other competing unicasts for link availability. Next, in steps 1-3, the message is passed on from the starting nodes to all the other nodes via an uncongested XYZ distribution tree, spanning across the three dimensions. Finally, during step 4, all the nodes absorb the received broadcast message. The salient features of the HP-SMART routing are:

1) HP-SMART: Planar Distribution

Figs. 2(a) and 2(b) elaborate our proposed HP-SMART broadcast distribution in a single planar die. In each layer, the number of incoming planar ports in the routers limits the number of simultaneous broadcast distributions that can be handled by the HP-SMART routing. Therefore, as we observe from Fig. 2, HP-SMART routing can distribute four simultaneous broadcasts in each planar layer without any link contention.

As we also note from Fig. 2, HP-SMART uses a quadrant-based approach in each layer to perform the four simultaneous broadcast distributions. Each quadrant has a set of two starting nodes, where one lies at the center of the die and the other lies along one of the edges (shown in Fig. 2). As an example, nodes 3 and 27 are the starting nodes for the upper left quadrant in Fig. 2(a). Similarly, considering all the four quadrants, eight nodes in each layer are marked as starting nodes. For arrangements where four uniform divisions cannot be made, the divisions are made such that the four center nodes lie in different quadrants.

As shown in Figs. 2(a) and (b), the broadcast distributions are performed along uncongested distribution trees originating from the starting nodes in each planar layer. Unlike the edge-to-edge transmissions in conventional SMART routing [5], HP-SMART follows center to edge X and Y distributions. Therefore, compared to conventional SMART, HP-SMART lowers the path lengths needed for broadcast distributions along X and Y dimensions to half. Hence, considering a planar die with $n \times n$ nodes, by employing a SMART 3D NoC with $PHPC_{max}=n/2$, each step of planar distribution can be performed in one cycle (a total of 2 cycles for complete planar distribution).

***Example*:** Here, we explain the XY planar distributions under the HP-SMART routing scheme through an example using Fig. 2. We consider an 8x8 arrangement of cores with $PHPC_{max}=4$. As shown in Fig. 2(a), during step-1 of the broadcast distribution, the message is forwarded from each starting node to a set of intermediate nodes. The intermediate nodes lie either along the same row or along the same column as that of the starting node. For instance, broadcasts originating from nodes 3 and 28 in Fig. 2(a), reach nodes 0-7 (along horizontal row) and 4-60 (along vertical column) respectively after completing step-1 distribution. Here, message starting from node 3 follows X-first while message starting from node 28, follows Y-first distribution. Similarly, step-1 planar distributions from other starting nodes follow either X-first or Y-first distribution. At the end of step-1 distribution, maximum of two broadcasts are accumulated at the local intermediate nodes. For example, node 7 accumulates two messages at the end of step-1, originating from nodes 3 and 31 respectively. Next, as shown in Fig. 2(b), step-2 planar distribution is initiated from the intermediate nodes to deliver the messages to the final destination nodes. At end of -step-2 distribution, all the nodes in the layer receive the four simultaneous broadcasts.

2) Uncongested 3D Broadcast Distribution

Considering a SMART 3D NoC with n vertical layers, a maximum of $4n$ broadcast transmissions can originate simultaneously (4 broadcasts from each layer). To establish a congestion free distribution of all $4n$ messages, a fixed Broadcast Distribution Interval (BI) is followed for steps 1-3[6]. In this fixed BI scheme, at specific cycles within the interval, a set of planar and vertical links are reserved and made available only for broadcast distributions originating from a particular starting node. In order to implement the BI scheme, it is important to determine the following three parameters:

a. The 3D (XYZ) distribution tree along which a broadcast is transmitted form the starting nodes to all the other nodes.
b. The instances within the broadcast interval at which a message will undergo planar and vertical transmissions along the distribution tree.
c. The minimum BI period to distribute all $4n$ broadcasts without any contention.

In the SMART 3D NoC, each message originating in a layer is first forwarded from the source node to the two local starting nodes in its own layer. Following this, three different methods are possible for a broadcast to achieve the XYZ distribution. In all the three methodologies, the possible turns in the routing are

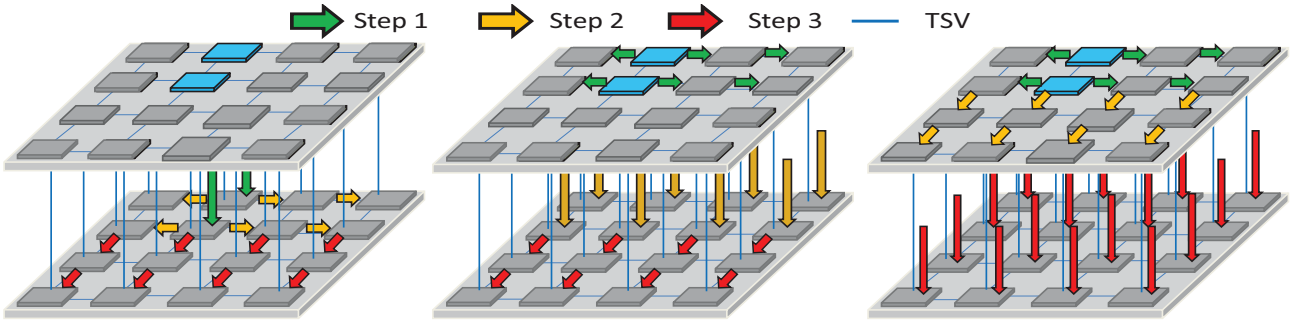

(a) *Vertical distribution first* (b) *Planar distribution first* (c) *Vertical distribution last*

Fig. 3: Illustration of the three XYZ-distribution methods using a 2-layered 32-core SMART 3D NoC employing $PHPC_{max}$=2, $VHPC_{max}$=1. Here we only show the distribution of one broadcast message in the bottom planar layer (message originates from the blue nodes in top layer). All other messages follow similar transmissions from their respective starting nodes.

restricted to ensure that no cyclic paths are present. Hence, the SMART 3D XYZ distribution schemes are deadlock-free. The three types of XYZ distribution are illustrated in Fig. 3.

a. *Vertical distribution first*: In this method, initially, the messages are forwarded from the starting nodes of the source layer to the starting nodes of the other layers using vertical links. Next, the messages undergo planar (X-Y) distribution in each layer as explained in section III.B.1.
b. *Planar distribution first*: Here, originating from the local starting nodes, the messages in the source layer first undergo step-1 planar distribution to reach a set of local intermediate nodes (shown in Fig. 2(a)). Next, the messages are forwarded from these local intermediate nodes to all other layers using the vertical links. Finally, in each layer, the messages undergo the step-2 planar distribution, starting from the intermediate nodes to all the final destination nodes following the paths in Fig. 2(b).
c. *Vertical distribution last:* Here, the messages first undergo a complete planar distribution in its source layer as explained in section III.B.1. The planar distribution accumulates four broadcast messages at each node in the source layer. Finally, using the vertical links, each node in the source layer forwards these accumulated messages to all the other nodes lying in the same vertical column.

It should be noted that employing the same type of transmission for all the broadcasts would create heavy link contention. For example, employing *Vertical distribution first* for all the *4n* broadcasts will lead to (a) accumulation of *n* messages at the starting nodes in each layer and (b) high utilization of planar links and low utilization of the vertical links. Compared to the *Vertical distribution first* method, *Planar distribution first* and *Vertical distribution last* methodologies better utilize the vertical links. Hence, to avoid the accumulation of the messages at the starting nodes, one has to selectively assign the broadcasts arising from specific layers to follow a combination of *Vertical distribution first, Planar distribution first* and *Vertical distribution last* methodologies. In this way, both the planar and vertical links are utilized more uniformly and multiple broadcasts can be transmitted at the same time, lowering overall BI period.

To find the three factors associated with the fixed BI scheme (the set of uncongested trees, instances at which transmissions take place and the BI period as listed above), we employ an exhaustive search mechanism. As explained above, there are three possible ways of distributing a broadcast arising from a particular layer. Hence, there exists 3^n ways of assigning them to *n* different layers. Since *n* is not large for a 3D NoC, an exhaustive search is easily feasible here. In this search, we assign one of the three different XYZ distribution methods for messages arising from each layer. Next, we identify time instances at which each broadcast can be distributed along the assigned paths without link contention to compute the BI period for the current XYZ distribution assignment. Then, by comparing the BI values of different XYZ distribution assignments, we can find the minimum BI value and the associated XYZ distribution instances for each message.

3) ACK aggregation and Unicast Communication

For collective communication messages, each destination node generates an acknowledgement (ACK) that must be delivered to the source node. We follow an ACK aggregation mechanism similar to [6], employing a set of dedicated multi-drop wires and ACK tables. Hence, the ACK messages retrace the path along which the initial broadcasts were distributed and are aggregated at intermediate and starting nodes in each layer. Since, a maximum of *4n* broadcasts can be distributed simultaneously, the total flit width for ACK multi-drop wires is given by $[\log_2(4n)+1]$ bits ($\log_2(4n)$ bits for message ID, 1 bit representing ACK/NACK). Unicast messages follow 1D SMART routing explained in [15].

Next, we incorporate the proposed HP-SMART broadcast distribution and ACK aggregation strategies in the SMART 3D NoC and undertake a detailed performance evaluation.

IV. Experimental Results

In this section, we evaluate and analyze the performance of the proposed SMART 3D NoC. We compare and contrast the performance of the SMART 3D NoC with respect to other state-of-the-art collective-communication aware architectures, viz., the hybrid partition based path multicast 3D NoC [14], tree multicast 3D NoC [8] and the SMART 2D NoC [5] [6].

A. Experimental Setup

We use Gem5 full system simulator [17]. We modify the Garnet network models within Gem5 to implement the different NoC architectures considered here. We choose a 64-core system comprised of x86 architecture based CPU cores. The size of each tile considered in this work is 2mm×2mm. Hence, we consider a die size of 16mm×16mm for the 2D NoC while the 3D NoCs employ four planar dies sized 8mm×8mm. The memory system

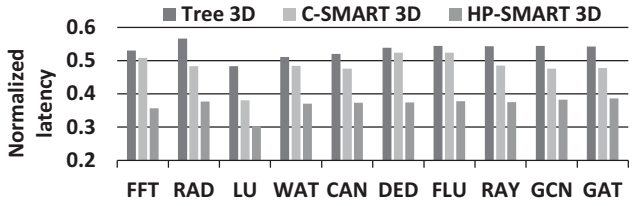

Fig. 4: Average message Latency of different 3D NoCs. Here, the values are normalized with respect to path multicast 3D NoC.

comprises of private 64KB L1 instruction and data caches, one shared 16MB L2 cache (256KB distributed L2 per core). All the NoCs operate with a clock frequency of 2.75GHz. For this clock frequency, SMART NoCs use $PHPC_{max}$=2 and $VHPC_{max}$= 3.

Following the XYZ distribution assignment methodology explained in Section III.B, for the considered 4-layer SMART 3D NoC, we employ *Vertical Transmission Last* for the broadcasts arising from top and bottom layers and *Vertical Transmission First* for the broadcasts arising from the middle layers. As explained in Section III.B, this routing mechanism balances planar and vertical link utilizations. Following this routing mechanism and using $PHPC_{max}$=2 and $VHPC_{max}$= 3, the minimum BI period obtained is 6 cycles.

For performance evaluations, we consider the on-chip traffic induced by MOESI Hammer cache coherence mechanism. The area overhead introduced by Hammer is less than that of the conventional directory protocol and hence, Hammer is more desirable for large-scale manycore platforms. Therefore, designing an efficient NoC for Hammer is a highly relevant research topic. Irrespective of the application under execution, Hammer coherence induces high volumes of collective communication. Here, we use a set of ten applications across the SPLASH-2 [18], PARSEC [19] benchmark suites and Grappolo [20]. Grappolo is used to identify the naturally existing communities in real-world graphs. The applications chosen are: FFT, RADIX(RAD), LU, WATER-SPATIAL (WAT), CANNEAL (CAN), DEDUP (DED), FLUIDANIMATE (FLU), RAYTRACE (RAY) and GRAPPOLO with COND (GCN) and ASTRO (GAT) datasets.

B. Performance evaluation of Proposed SMART 3D NoC

In this section, we compare the network latency of the NoC architectures considered here. In addition, to highlight the performance benefits achieved due to the proposed routing, we consider two different flavors of SMART 3D NoC, a) incorporating the conventional edge-to-edge routing [6] (labeled as C-SMART 3D NoC) and b) incorporating the proposed center-to-edge high performance routing (labeled as HP-SMART 3D NoC). For C-SMART 3D NoC, to achieve 2-cycle end-to-end planar distributions, we employ $PHPC_{max}$=3 (thus, $PLPC_{max}$=6mm) resulting in a NoC clock frequency of 1.83GHz. In Fig. 4, we compare the performance of HP-SMART and C-SMART 3D NoCs with respect to the existing state-of-the art 3D collective communication aware NoCs. For path multicast, at each router along the path, the header flit needs to be read and repackaged with the updated list of destinations. Hence, we use a 3-cycle router for the path multicast NoC following [4]. For the tree multicast 3D NoC, we use the router microarchitecture as discussed in [8].

All the values in Fig. 4 have been normalized with respect to the network latency of the path multicast based 3D NoC (Normalized path multicast latency=1, not shown for better visibility). We observe that the path multicast based 3D NoC has the highest latency among all the NoCs considered due to the high hop counts and intra-router delays associated with it [4]. On the other hand, tree multicast based 3D NoC employs an advanced router microarchitecture lowering the per-hop latency. Also, tree multicast lowers the average hop count needed for broadcast distribution compared to the path-based scheme [4]. Hence, tree multicast based 3D NoC performs significantly better than the path multicast NoC.

Both the SMART architectures perform better than the conventional path and tree routing schemes as they employ single-cycle multi-hop traversals. We also note from Fig. 4 that on average C-SMART 3D NoC performs 52% and 10% better than Path and Tree multicast based NoCs respectively, despite running at 33% lower clock frequency. Compared to the path and tree multicast- based 3D NoCs, our proposed HP-SMART 3D NoC achieves a performance improvement of 65% and 31% respectively, on average for the considered applications. This demonstrates the advantage of employing SMART 3D data transfer mechanism in handling dense collective communication traffic. Fig. 4 also highlights the performance improvement achieved solely due to our proposed HP-SMART routing in a 3D NoC. Compared to the C-SMART 3D NoC, the proposed HP- SMART 3D NoC shows on an average 24% lower latency. Since HP-SMART follows a center-to-edge routing scheme, the path length to be covered during planar distribution is reduced by half when compared to the C-SMART 3D NoC. Since both implementations are based on 3D NoCs, the improvement gains are solely due to the proposed HP-SMART routing mechanism.

In Fig. 5, we compare the HP-SMART 3D NoC with the HP-SMART 2D NoC (SMART 2D NoC with HP-SMART routing scheme). Both NoCs operate at 2.75GHz (*i.e.* $PLPC_{max}$ = 4mm). We observe that the proposed HP-SMART 3D NoC reduces latency on an average by 28% compared to its 2D counterpart. The latency improvement of the SMART 3D NoC is not only due to the reduction in physical path length but also due to the reduction in the queueing latency.

To corroborate this, in Fig. 6, we compare the number of broadcast messages waiting to access network resources (queuing latency) in SMART 2D and SMART 3D NoCs. When M number of broadcasts are generated in the same quadrant

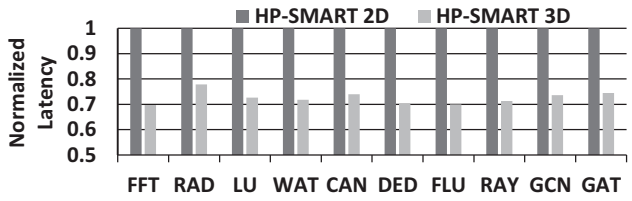

Fig. 5: Average message latency of HP-SMART 2D and 3D NoCs. The values are normalized with respect to HP-SMART 2D NoC.

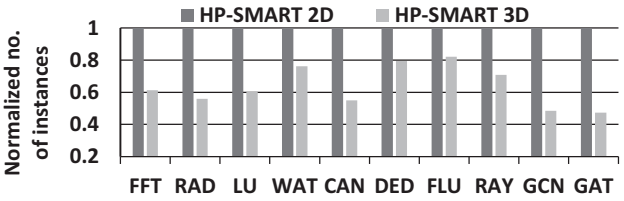

Fig. 6: Number of instances(normalized), when multiple messages were generated in the same quadrant within a single BI period.

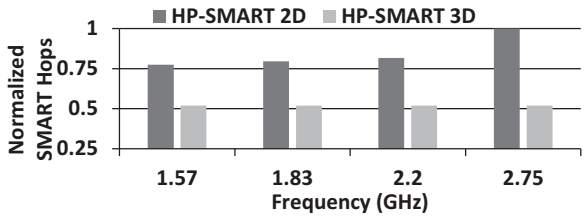

Fig. 7: Number of SMART hops in a 256-core system, normalized with respect to HP-SMART 2D NoC operating with *freq*=2.75GHz

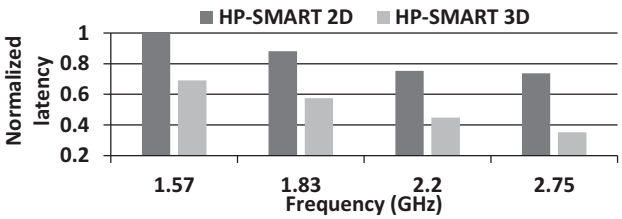

Fig. 8: Zero load latency for different frequencies in a 256-core system, normalized with respect to a 2D NoC with *freq*=1.57GHz

within one BI period, *M-1* messages are forced to wait for the upcoming BIs. This leads to queueing latency at the starting nodes. From Fig. 6 it is evident that the number of cases where multiple broadcasts have been sent to the same starting nodes within a single BI period is lowered by 37% on an average in HP-SMART 3D NoC compared to its 2D counterpart. This reduces the average queueing latency for the messages, leading to better performance for HP-SMART 3D NoC. To explain further, the HP-SMART 2D NoC only allows four simultaneous transmissions and hence, only one broadcast can be initiated in a single BI period from each quadrant. On the other hand, as explained in Section III. B, the proposed HP-SMART 3D NoC can handle *4n* broadcast messages within one BI period. As a result, messages experience lower queueing latency in HP-SMART 3D NoC compared to its 2D counterpart.

C. Frequency Scalability of SMART 3D NoCs

To compare the frequency scalability of the SMART 2D and 3D NoCs (4 layers), we choose a larger system size (256-cores, tile size 1mm×1mm). We assume core sizes to shrink to enable higher degree of integration. A large system size allows us to consider a wider range of $PLPC_{max}$ (varying from 4mm to 7mm) and in turn a larger range of clock frequencies. We show the average SMART hops and network latency variations for different frequencies (hence, different $PLPC_{max}$) values in Figs. 7 and 8 respectively. Here both 2D and 3D architectures employ HP-SMART routing. As seen from Fig.7, for the SMART 2D NoC, the average SMART hop increases with increasing NoC frequencies due to the reduction in $PLPC_{max}$. On the other hand, the average SMART hops for the HP-SMART 3D NoC at different frequencies considered here remain the same since an $PLPC_{max}$ = 4mm is sufficient to cover the entire die along one dimension, allowing the data to be transferred in one cycle. Hence, as evident from Fig. 8, SMART 3D NoC achieves higher latency gains for higher clock frequencies. This analysis demonstrates that the SMART 3D NoC is more scalable with increasing clock frequencies than its 2D counterpart.

V. Conclusion

Designing efficient collective communication infrastructure for 3D NoCs is a challenge that has not been sufficiently addressed yet. Conventional SMART 2D NoC provides efficient collective communication support. However, the conventional SMART 2D NoC has a limited frequency scalability. To address these challenges and limitations, we have proposed a High-Performance 3D NoC (HP-SMART 3D NoC). HP-SMART 3D NoC not only outperforms existing 3D collective communication mechanisms but also performs significantly better than the conventional SMART implementation employing 2D NoCs. Our experimental results show that the proposed HP-SMART 3D NoC achieves 28% lower network latency that its 2D counterpart for the applications considered in this work. Therefore, HP-SMART 3D NoC is an excellent choice for efficient collective communication.

References

[1] B.S. Feero and P.P. Pande, "Networks-on-Chip in a Three-Dimensional Environment: A Performance Evaluation," *IEEE Trans. on Computer*, 53(1), pp. 32-45, 2008.

[2] P. Zhou, et al., "Application-specific 3D network-on-chip design using simulated allocation", In *Proc.of ASP-DAC*, pp. 517–522, 2010.

[3] P. Conway and B. Hughes, "The AMD opteron northbridge architecture," in IEEE Micro, vol. 27, no. 2, pp. 10–21, 2007.

[4] K. Duraisamy, et al., "Multicast-Aware High-Performance Wireless Network-on-Chip Architectures," *IEEE TVLSI*, 25(3), pp. 1126-39, 2017.

[5] T. Krishna et al., "Breaking the on-chip latency barrier using SMART," in *Proc. of HPCA*, pp. 378-389, 2013.

[6] T. Krishna and L. S. Peh, "Single-cycle collective communication over a shared network fabric," in *Proc. of NOCS*, Ferrara, pp. 1-8, 2014.

[7] K. Duraisamy and P. P. Pande, "Performance evaluation and design trade-offs for wireless-enabled SMART NoC," in *DATE*, pp. 1360-1365, 2017.

[8] T. Krishna, L.-S. Peh, B. Beckmann, and S. K. Reinhardt, "Towards the ideal on-chip fabric for 1-to-many and many-to-1 communication," in *Proc. of MICRO*, pp. 71–82, 2011.

[9] A. Karkar *et al*., "Hybrid wire-surface wave architecture for one-to-many communication in networks-on-chip," in *Proc. of DATE*, pp. 1–4, 2014.

[10] M. Kamali, et al.,, "Formal Modeling of Multicast Communication in 3D NoCs," in *Euromicro Conf. on Digital System Design*, pp. 634-642, 2011.

[11] W. Wang *et a*l., "Design and realization of 3D NOC multicast router base on multicast rotational routing arithmetic," in *ICCCNT*, pp. 1-6, 2014.

[12] S. R. Moosavi, et al. "Enhancing Performance of 3D Interconnection Networks using Efficient Multicast Communication Protocol," in *Euromicro Intl. Conf. on Parallel, Distributed, and Network-Based Processing*, pp. 294-301, 2013.

[13] X. Wang *et a*l., "Low latency and energy efficient multicasting schemes for 3D NoC-based SoCs," in *IEEE/IFIP Intl. Conf. on VLSI and System-on-Chip*, pp. 337-342, 2011.

[14] M. Ebrahimi, et al., "Partitioning methods for unicast/multicast traffic in 3D NoC architecture," in *IEEE Symp. on Design and Diagnostics of Electronic Circuits and Systems*, pp. 127-132, 2010.

[15] X. Chen and N. K. Jha, "Reducing Wire and Energy Overheads of the SMART NoC Using a Setup Request Network," in *IEEE Trans. on Very Large Scale Integration Systems*, 24 (10), pp. 3013-3026, 2016.

[16] P. Batra et al., "Three-dimensional wafer stacking using Cu TSV integrated with 45nm high performance SOI-CMOS embedded DRAM technology," in *Jnl. of Low Power Elec. and Appln*. 4(2), pp.77-89.

[17] N. Binkert, et al. "The GEM5 Simulator", in *ACM SIGARCH Computer Architecture News*, 39(2):1-7.

[18] S.C. Woo, et. al., "SPLASH-2 Programs: Characterization and Methodological Considerations," In *Proc. of ISCA- 95*, pp. 24-36.

[19] C. Bienia, "Benchmarking Modern Multiprocessors," Ph.D. dissertation, Dept. Comp., Sci., Princeton Univ., NJ, 2011.

[20] H. Lu, M. Halappanavar and A. Kalyanaraman, "Parallel heuristics for scalable community detection," in *Parallel Computing*. 47, 19-37, 2015.