

ReDC: Reduced Deflection CHIPPER Router for Bufferless NoCs

Rose George Kunthara*, Rekha K James*, Simi Zerine Sreeba[†] and John Jose[‡]

* Division of Electronics, School of Engineering, CUSAT, Cochin, India

[†] Dept. of Electronics & Communication Engineering, Viswajyothi College of Engineering and Technology, Muvattupuzha, India

[‡] Dept. of Computer Science and Engineering, Indian Institute of Technology Guwahati, India

rosekunthara87@gmail.com, rekhajames@cusat.ac.in, simi.abie@gmail.com, johnjose@iitg.ac.in

Abstract—Network on Chip (NoC) is emerging as a promising design paradigm as an on-chip interconnect for multi-core architectures to overcome scalability and bottleneck issues of traditional bus-based and point-to-point communication architectures. Minimum packet latency, power and area with improved performance are the important characteristics that determine the efficiency of an NoC router. In this paper, we propose a bufferless deflection router, ReDC, which minimizes the deflection rate of flits by selecting a Permutation Deflection Network with an input combination that gives most number of productive output ports. Simulation results show that our proposed design improves the network saturation point, reduces the average flit latency and deflection rate without significant change in critical path delay when compared to existing state-of-the-art bufferless deflection routers.

Index Terms—Network-on-Chip, bufferless, deflection routing

I. INTRODUCTION

The advancements in IC technology have resulted in miniaturization due to the fall in transistor feature sizes to ultra-deep submicron levels. This has led to sharp increase in transistor densities resulting in the development of complex System on Chip (SoC). The traditional SoC is composed of IP (Intellectual Property) cores or predesigned functional blocks, which are interconnected using point to point intercommunication using dedicated wires or classical shared bus. But the shrinking technology resulted in an imbalance between on chip wire delay and gate delay leading to increased power consumption, on-chip synchronization errors, unpredictable delays, etc. Network on Chip (NoC) concept, which is a packet-switched network, was introduced to overcome scalability, predictability and bottleneck challenges faced by the traditional SoC architecture. NoC communication is gaining popularity due to its advantages like improved parallelism, scalability, simultaneous communication between multiple pairs of processing elements, inherent fault tolerance, improved load handling capability and modular topology to interconnect the processing cores [1], [2].

The regular tile-based NoC architecture consists of high-speed routers, Network Interfaces (NI) and communication links as its main components. Each tile can be a general purpose processor, a DSP processor, a memory subsystem, etc

and the communication between them is achieved by routing packets. A packet is the basic data unit in NoC and each packet is subdivided into flow control units called flits. In NoC, network traffic is initiated due to cache misses and coherence transactions. Router, which is the backbone of NoC, consists of five input and five output ports, one each for the north, south, east, west directions and one for local port connected to the processing core. The injection and ejection of flits to and from the network is done through local port [1], [3].

The conventional Virtual Channel (VC) based NoC has a set of buffers associated with each port of the router which increase their load handling capacity and throughput. But they consume significant portion of on-chip network power and area. Bufferless NoC routers, employing deflection routing have been proposed to overcome this rising power and area issues of the VC based router. In bufferless deflection router [4], [5], flits which do not get desired output link get deflected through a freely available output link. This increases latency of flits due to misrouting of flits. In this paper, we propose Reduced Deflection Chipper (ReDC), a modified version of CHIPPER [5], which uses a two stage router pipeline with single cycle latency at each stage, and is more energy efficient with reduced average latency and deflection rate while operating at almost the same speed as CHIPPER.

The remainder of this paper is organized as follows: An overview about the related work and the motivation behind the current work is present in Section II. Section III provides details about the new router architecture, ReDC. The experimental methodology followed is discussed in Section IV. Section V provides the results and analysis and finally Section VI concludes the paper.

II. RELATED WORK AND MOTIVATION

The scaling and applicability of the conventional input buffered virtual channel (VC) based router design is severely hampered by the presence of power hungry buffers [6], [7]. Even though VC routers eliminate unnecessary wastage of link bandwidth and deliver higher throughput, the buffers occupy significant area and consume large amount of static power when idle and dynamic power when active. Studies have shown that when the packet injection behaviour of real

workloads for mesh NoCs are considered, the input buffers in VC based NoC routers are overprovisioned [8], [9]. To improve the energy efficiency of the routers, researchers have come up with buffer-less and minimally buffered deflection routers to reduce power consumption and area.

The bufferless deflection routers for mesh NoCs are first proposed in [10]. The routing mechanism employed in bufferless routers is either based on dropping and retransmission of packets [11] or on deflection of packets to an undesired port [4], [5], [8]. In the dropping mechanism models, high overheads are involved to coordinate the acknowledgements and retransmissions from the source. All the flits that arrive at the input ports pass through one of the output ports in a deflection router. When there are more than one incoming flit competing for the same output port, only one wins and gets the desired output port while the other flits that do not get the desired output port get deflected through available output ports. Thus the deflection rate of flits is high leading to an increased latency of flits and the network saturates earlier when compared to the conventional VC routers.

BLESS [4], a baseline bufferless router employs a sequential port prioritization using an age based flit ranking mechanism for output port selection. This causes increase in the router's critical path delay resulting in lowering of operating frequency of the NOC. To overcome this performance issue, CHIPPER [5] uses parallel port allocation scheme where golden packet mechanism is used for prioritizing the packets. CHIPPER has reduced pipeline stage delay compared to BLESS but at the expense of higher deflection rate as all the non-golden packets are assigned random ports. Both BLESS and CHIPPER has good performance under low and medium network traffic. WeDBless [12] is another bufferless deflection router which minimizes flit deflection rate and average latency by using output port allocation based on Deflection Weights (DW) and ranking the flits based on Weighted Deflection Count (WDC) inside the router.

A low power deflection routing method for bufferless NoC is proposed in [13], which uses a routing matrix for getting possible routing paths and the best route for each packet is then selected. At high network loads, there is increased deflection, power dissipation, delay and reduced network throughput due to frequent contentions in routers. To alleviate this problem, authors of [14] propose a lightweight link control mechanism to circumvent unnecessary network hops of deflected packets by allowing them to loop back to its current router, when possible, instead of being misrouted. Deflection Containment for bufferless NoC (DeC) [15] tries to overcome excessive flit deflections for power reduction and performance improvement by adding a link to each router for bridging subnetworks. A contending flit in one subnetwork is bypassed to another subnetwork instead of being misrouted, giving better path diversity and decreased number of deflections. SCEPTER architecture [17] is a high performance bufferless NoC that can dynamically set up single cycle virtual express paths across the chip to allow deflected packets to traverse along non-minimal paths with zero latency penalties.

In bufferless router designs, even though the power and area due to buffers is negligible, the deflection rate of flits is more and the network saturates much earlier compared to the conventional buffered routers. So to achieve better performance, the advantages of buffered and bufferless router designs are combined in minimally buffered routers by buffering a small fraction of the misrouted flits [8], [9], [16]. MinBD [8], which is a minimally buffered router, first employed a small side buffer to store a small number of misrouted flits in the router. DeBAR [9] is another minimally buffered router which uses a minimal central buffer pool to keep few of the deflected flits. It includes a hybrid flit ejection mechanism to provide the effect of dual ejection by using a single ejection port and flits are selectively buffered based on flit marking with better priority metrics. ADIEU [18] is an adaptive deflection router which incorporates dual injection and ejection units with minimal side buffering to improve overall performance of the system. A comparison of the buffered and bufferless design paradigms is done in [19] based on various design parameters.

CHIPPER is a bufferless deflection router where each incoming flit is moved to an output port in two cycles of operation. CHIPPER is superior to BLESS in terms of reduced critical path latency but the unnecessary flit deflections occur in CHIPPER due to random port allocation scheme which reduce their performance. In our proposed bufferless router design, ReDC, we employ two permutation deflection networks operating in parallel with same set of inputs but given in different order, for getting maximum number of productive ports and thus reducing flit deflection rate and dynamic power dissipation across NoC links.

III. REDC ARCHITECTURE

Figure 1 shows the architecture details of our proposed router, which is a modified version of the traditional two cycle deflection router. The details of various units are explained as follows: The input flits passing through the various units of router pipeline are carried by the four internal flit channels. The flits get stored in the corresponding pipeline registers at the end of each clock cycle. A flit is routed to a neighbouring router based on XY routing algorithm, which is a deterministic, minimal path routing algorithm, free of deadlock and livelock. In deflection routers, deadlock problem does not occur because cyclic dependency of resources will not happen [4] and livelock problem is resolved in CHIPPER using the golden flit priority mechanism [5]. The flits from neighbouring routers reach the pipeline register A at the beginning of each clock cycle.

A. Ejection Unit and Injection Unit

The ejection unit and injection unit used in ReDC is the same as that of the CHIPPER. When there is an ejection flit or a flit destined to the local core in the current cycle, it is removed from the internal flit channel and send to the ejection port. The ejection unit supports only one ejection port per router. When a processing core wants to send a flit to another core, the core will be able to inject the flit as long as one of

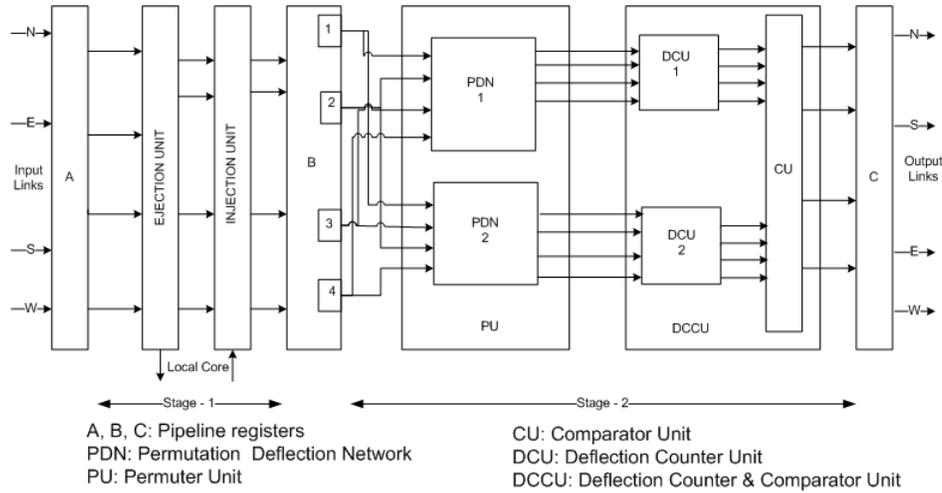


Fig. 1. Router pipeline architecture of ReDC

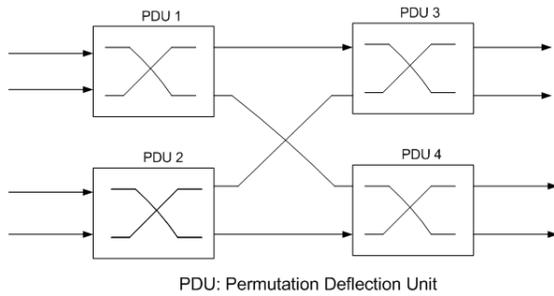


Fig. 2. Internal architecture of PDN.

its output links is free. As there is no buffer to hold the flit in the bufferless deflection router, the injection process can happen only when there is a free output link. The flit will remain queued at the processor core level in the absence of free output link.

B. Permuter Unit

When two or more flits request the same output port, deflection or hot-potato routing is used to resolve port contention problem because the flits need to go through the router pipeline without waiting or stalling as there are no intermediate buffering to hold the flits. The Permutation Deflection Network (PDN) used in ReDC for the parallel output port allocation is similar to the one used in CHIPPER and MinBD as shown in Figure 2 and the packets are prioritized using the golden packet scheme. PDN maps every input port of the router to every output port of the router efficiently. The PDN basically consists of two stages, with each stage employing two 2x2 crossbars. The priority levels and the required output port for the incoming flits decide the port allocation for each arbitration stage of the PDN. The highest priority flit will get the productive port and the other flits may or may not get a productive output port depending on the level of contention and port conflicts. In our proposed router, the Permuter Unit

(PU) consists of two such PDN networks but with the inputs connected in different order. In the first PDN network, North and East input channels are linked to PDU1 and the South and West input channels are connected to PDU2 whereas in PDN2, the North and West input channels are connected to PDU1 and the East and South input channels are linked to PDU2.

C. Deflection Counter and Comparator Unit

The outputs of both PDN networks go to the Deflection Counter and Comparator Unit (DCCU). DCCU contain two Deflection Counter Units (DCU) for counting the number of deflections from each PDN network. The Comparator Unit (CU) compares the number of deflections from each PDN network and selects the output of that PDN network which has less number of deflections and is passed to the next pipeline register, C. Reduced number of deflections denotes that there is decreased unproductive flit movement in the network leading to lower power consumption and reduced average flit latency.

IV. EXPERIMENTAL METHODOLOGY

We use the cycle accurate NoC simulator, Booksim [20] that models the traditional VC based router [1]. We make necessary modifications to model the two-cycle deflection router microarchitecture which is explained in CHIPPER [5]. Every flit that we consider has header information attached to it so as to facilitate independent routing of the flits within a packet which is the common standard in deflection routers. We use necessary reassembly mechanism for handling the out-of-order delivery of flits. The flit channel which is 140 bits wide contain: 128-bit data field and 12-bit header field. We make changes on this baseline deflection router simulator to model ReDC router and perform experimental analysis. All the evaluations that we conduct are using single flit packets.

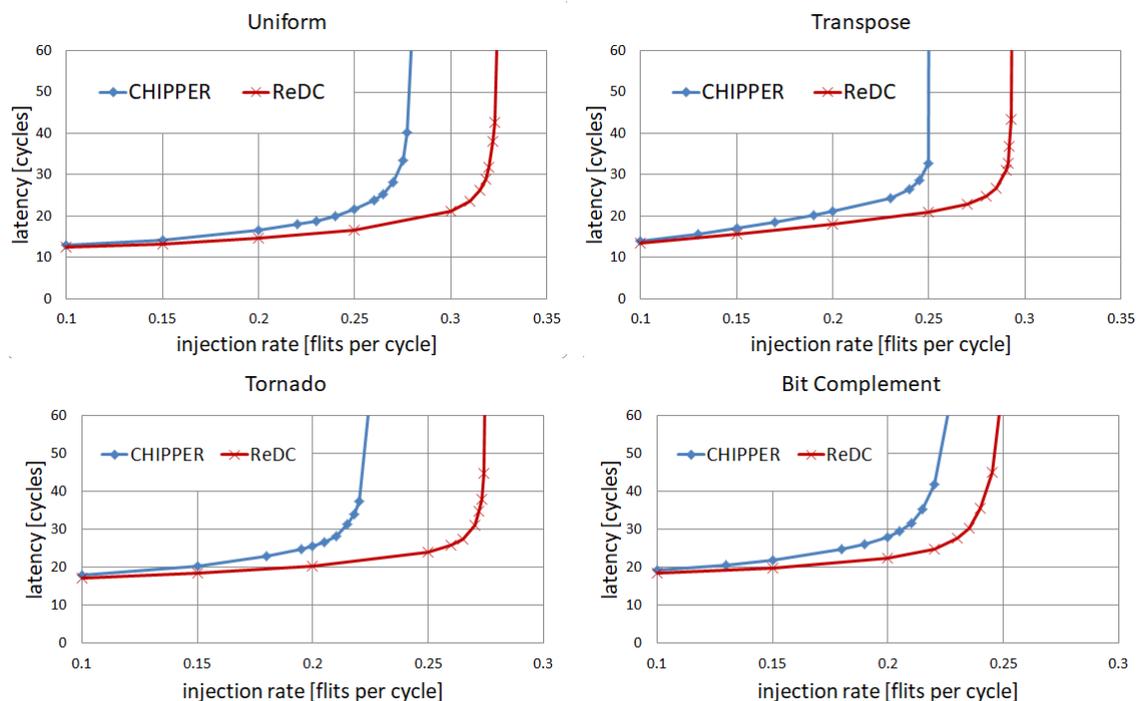


Fig. 3. Flit latency comparison under various synthetic traffic patterns in 8x8 mesh network.

TABLE I
PERCENTAGE OF DIFFERENT NETWORK INJECTION INTENSITY
APPLICATIONS IN VARIOUS BENCHMARK MIXES

Benchmark Mix	M1	M2	M3	M4	M5	M6	M7
% of Low	100	0	0	50	0	50	31
% of Medium	0	100	0	0	50	50	31
% of High	0	0	100	50	50	0	38

A. Synthetic Traffic

We analyse the performance of our design against CHIPPER for mesh topology using standard synthetic traffic patterns such as uniform, transpose, bit-complement, tornado, bit-reverse, shuffle and neighbor. Average flit latency and deflection rate values are collected for each of the traffic patterns after sufficient warm up time, with injection rate varied between zero and network saturation point.

B. Real Workloads

We analyse the performance of ReDC in comparison with CHIPPER using real application mixes from SPEC CPU2006 benchmark application suite [22]. Using Multi2Sim [21] simulator, a 64-core multiprocessor system is modelled where each core consists of an out-of-order x86 processing unit with 64KB, 4-way set-associative L1 cache and 512KB, 16-way set-associative shared distributed L2 cache. Each core runs one of the applications from the SPEC CPU2006 benchmark suite. Based on the misses per kilo instructions (MPKI) values, the benchmark applications are classified into Low, Medium and High MPKI. We generate 7 multiprogrammed workload mixes

by combining the various applications from the benchmark suite as given in Table 1. The network traffic generated by these workload are fed to the NoC simulator to simulate the network operations for the comparison of CHIPPER and ReDC routers.

V. RESULTS AND ANALYSIS

We compare the average deflection rate and average flit latencies of CHIPPER and ReDC to analyse the effect of our design in reducing the deflections. The network size during simulation is assumed to be of 8x8 mesh. The performance parameters such as average flit latency, average deflection rate and network saturation point are analysed in depth.

A. Effect on Average Flit Latency

The flit latency is the time instant when the flit is first created to the time instant when the flit is ejected to the destination node, including the queuing time at the source. The average flit latency should be minimal for better performance since higher latency value for the flits in the network increase the stall time of the applications leading to their throttling and poor application performance. Lower and wider latency curve indicates that the performance of the router is better. Figure 3 shows the flit latency comparison of CHIPPER and our proposed design under some typical synthetic traffic patterns in 8x8 mesh network. ReDC reduces the average flit latency by 23% compared to CHIPPER for uniform-random traffic. For non-uniform traffic patterns like bit-complement and transpose, there is an average flit latency reduction of 13% and 14% respectively.

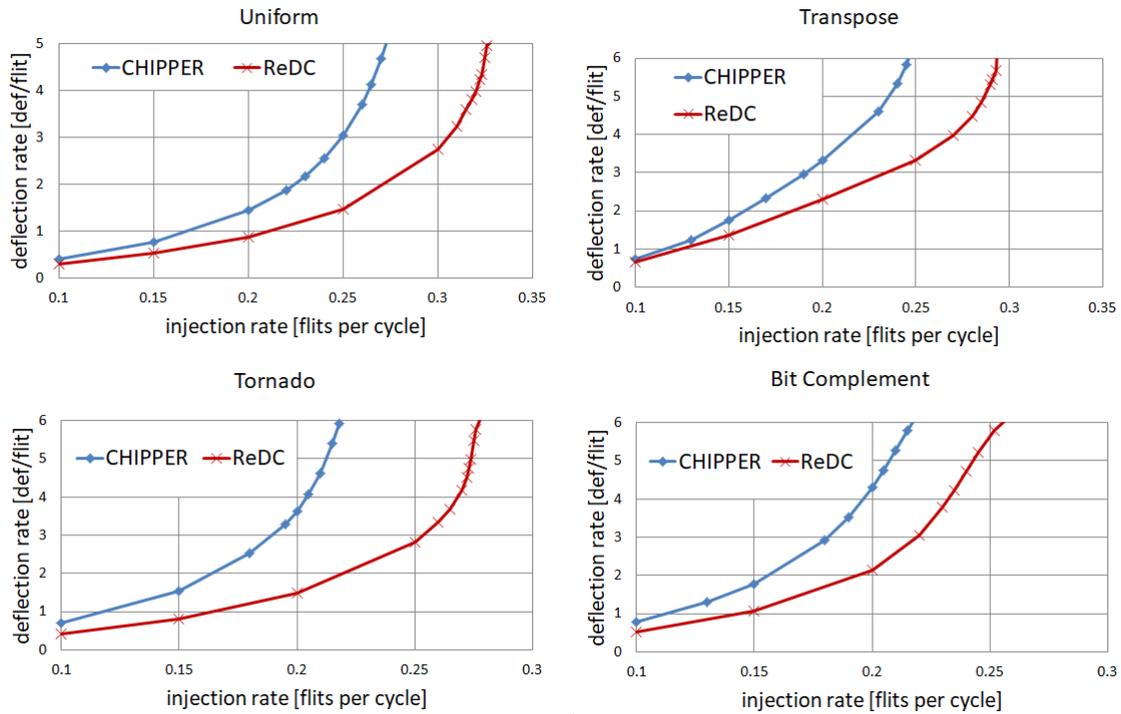


Fig. 4. Flit deflection rate comparison under various synthetic traffic patterns in 8x8 mesh network.

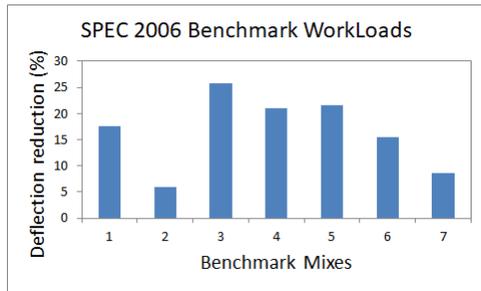


Fig. 5. Percentage reduction in deflection rate for real applications.

B. Effect on Deflection Rate

Deflection rate is computed as the average number of deflections per injected flit. Lower values for the deflection rate denotes that the unproductive flit movement in the network is decreased leading to reduction of dynamic power dissipation. When the injection rate increases, the deflection rate will also increase due to more port contention. Figure 4 shows that ReDC has less deflection rate compared to CHIPPER since we have selected that PDN network which gives more number of productive ports and thus less number of deflections. The average deflection rate reduces by 32% for uniform-random traffic pattern. For non-uniform traffic distributions like transpose and bit-complement there is a reduction of 22% and 35% respectively for ReDC compared to CHIPPER. From Figure 5, we can see that our proposed design outperforms CHIPPER in the deflection rate analysis of the real applications. There

is no significant change in throughput for both CHIPPER and ReDC.

C. Effect on Saturation Injection Rate

Injection rate indicates the number of flits that are being injected into the network per router per cycle. Saturation point denotes the injection rate at which the average flit latency reaches five times the value that of the zero load latency. As the injection rate approaches saturation, the average latency increases exponentially due to flooding of the network with flits. A router which has high saturation injection rate indicates that it has better load handling capacity. From Figure 3, it is quite obvious that ReDC has an improved saturation injection rate compared to CHIPPER for synthetic traffic patterns for an 8x8 mesh network. For our proposed router, the network saturation point extends by 20%, 17% and 15% for uniform-random, transpose and bit-complement traffic distributions.

D. Effect on Dynamic Power Consumption across NoC links

Orion 2.0 tool [23] is used for estimation of dynamic power consumption through NoC links for various injection rates and load rate. At pre-saturation level (injection rate of 0.1 flit per cycle approximately), dynamic power dissipation decrease by 32%, 22%, and 40% respectively for uniform, transpose and bit-complement traffic patterns. The dynamic power dissipation across links reduces by 51%, 30% and 50% for uniform, transpose and bit-complement traffic distributions at saturation area (injection rate of 0.2 flit per cycle approximately). By merit of our proposed routing scheme, deflections are reduced which leads to decreased activity on the links and subsequent

reduction in dynamic power consumption. The reduction in power consumption is more significant in the case of high network load indicating that ReDC has better load handling capacity.

E. Hardware Overhead

We implement Verilog HDL models of CHIPPER and ReDC router and synthesize using Xilinx ISE Design Suite 14.2 to compute the router pipeline latencies. Router delay is defined as the total time taken by a flit to move from the router input port to the router output port. Since similar functional units are used in first stage of CHIPPER and our proposed router, both router architectures have the same delay in their first stage of router pipeline. CHIPPER consists of a PDN in its output stage whereas ReDC include two PDN in parallel followed by DCCU, incurring an additional delay of 22% in its output stage. But router pipeline frequency of both routers remains the same as the pipeline latency of first stage dominates over second stage. Even though hardware area overhead of proposed design is 46% more than that of CHIPPER, there is significant reduction in deflection rate and dynamic power consumption across NoC links.

VI. CONCLUSION

The efficiency and performance of an NoC depend on the design of high performance and efficient routers. In this paper, we proposed ReDC, a bufferless deflection router that uses two PDN with the same set of inputs, but given in different order, for getting the maximum productive output ports. The proposed design gives better performance by selecting the PDN which gives more number of productive output ports so that the deflection rate of the flits are reduced and the average flit latency is minimal for the same critical path latency when compared to the state-of-the-art bufferless deflection router.

REFERENCES

- [1] W. Dally and B. Towles, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, USA, 2004.
- [2] William Dally, "Route packets, not wires: On-Chip interconnection networks", in *Design Automation Conference (DAC-01)*, pages 684-689, New York, ACM Press, June 2001.
- [3] W. Dally, "Virtual-channel flow control," *IEEE Transactions on Parallel and Distributed Systems*, vol. 3, no. 2, pp. 194-205, 1992.
- [4] T. Moscibroda and O. Mutlu, "A case for bufferless routing in on-chip networks," in *ISCA*, pp. 196-207, 2009.
- [5] C. Fallin et al., "CHIPPER: A low complexity bufferless deflection router," in *HPCA*, pp. 144-155, 2011.
- [6] Y. Hoskote et al., "A 5-GHz mesh interconnect for a teraflops processor," *IEEE Micro*, vol. 27, no. 5, pp. 51-61, 2007.
- [7] M. B. Taylor et al., "Evaluation of the raw microprocessor: An exposed-wire-delay architecture for ILP and streams," in *ISCA*, 2004.
- [8] C. Fallin et al., "MinBD: Minimally-buffered deflection routing for energy-efficient interconnect," in *NOCS*, pp. 1-10, 2012.
- [9] J. Jose et al., "DeBAR: Deflection based adaptive router with minimal buffering," in *DATE*, pp. 1583-1588, 2013.
- [10] E. Nilsson et al., "Load distribution with the proximity congestion awareness in a network-on-chip," in *DATE*, pp. 1126-1127, 2003.
- [11] M. Hayenga et al., "SCARAB: A single cycle adaptive routing and bufferless network," in *MICRO*, pp. 244-254, 2009.
- [12] S. Z. Sleeba, J. Jose and M. G. Mini, "WeDBless: Weighted Deflection Bufferless Router for Mesh NoCs," in *Proceedings of the 24th Great Lakes Symposium on VLSI*, ACM, pp. 77-78, 2014.

- [13] Chung-Kai Hsu, Kun-Lin Tsai, Jing-Fu Jheng, Shanq-Jang Ruan and Chung-An Shen, "A low power detection routing method for bufferless NoC," in *International Symposium on Quality Electronic Design (ISQED)*, pp. 364-367, 2013.
- [14] Igor Z. Stojanovic, Milica D. Jovanovic and Goran Lj. Djordjevic, "Dual-mode inter-router communication channel for deflection-routed networks-on-chip," in *The Journal of Supercomputing*, vol. 71, no. 7, pp. 2597-2613, 2015.
- [15] Xi-Yue Xiang and Nian-Feng Tzeng, "Deflection Containment for Bufferless Network-on-Chips," in *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 113-122, 2016.
- [16] B. Nayak et al., "SLIDER: Smart Late Injection Deflection Router for Mesh NoCs," in *International Conference on Computer Design (ICCD)*, pp. 377-383, 2013.
- [17] B. K. Daya, L. Peh, and A. P. Chandrakasan, "Towards High-Performance Bufferless NoCs with SCEPTER," in *IEEE Computer Architecture Letters*, vol. 15, no. 1, pp. 62-65, 2016.
- [18] J. Jose and A. Das, "An Adaptive Deflection Router with Dual Injection and Ejection Units for Mesh NoCs," in *31st International Conference on VLSI Design (VLSID)*, pp. 374-379, 2018.
- [19] Z. Lu et al., "Evaluation of on-chip networks using deflection routing," in *GLSVLSI06*, pp. 296-301, 2006.
- [20] Nan Jiang, Daniel U. Becker, George Michelogiannakis, James Balfour, Brian Towles, John Kim and William J. Dally, "A Detailed and Flexible Cycle-Accurate Network-on-Chip Simulator," in *IEEE International Symposium on Performance Analysis of Systems and Software*, 2013.
- [21] R. Ubal et al., "Multi2sim: A simulation framework to evaluate multicore-multithreaded processors," in *SBAC-PAD*, pp. 62-68, 2007.
- [22] "SPEC2006 CPU benchmark suite," <http://www.spec.org>.
- [23] A. B. Kahng et al., "ORION 2.0: A Fast and Accurate NoC Power and Area Model for Early-Stage Design Space Exploration," in *Design, Automation Test in Europe (DATE)*, pp. 423-428, 2009.