

The Impact of Routing Arbitration Mechanisms on 3D NoC Latency

Rodrigo Cataldo, Ramon Fernandes, Fernando Grandó, Thais Webber, Ana Benso, César Marcon

PUCRS - Pontifícia Universidade Católica do Rio Grande do Sul, Porto Alegre, Brazil – 90619-900
{cesar.marcon, ana.benso}@pucrs.br, {rodrigo.cataldo, ramon.fernandes, fernando.grando, thais.webber}@acad.pucrs.br

Abstract – This paper discusses the impact of routing arbitration mechanism on the packet latency for 3D NoC (Three-dimensional Network-on-Chip) architectures. We implemented several variations of Round-Robin mechanisms to explore how the arbitration efficiency affects the packet latency. The underlying objective is to discuss the compromise of increase router area and energy consumption through investing on a complex low-clock cycle router compared with a simpler but energy and area efficient router. The experimental setup is composed of two sizes of 3D mesh NoC, synthetic traffic pattern and several injection rates. Results demonstrate that the increase of arbitration latency does not affect proportionally the packet latency. In fact, for low traffic injection rates the arbitration algorithm does not influence the average of packet latency significantly, justifying the approach of employing less complex routers with less impact on area and energy consumption.

I. INTRODUCTION

Guiding packets through the network aiming to increase communication throughput and to reduce latency is a key characteristic of an efficient network architecture. Thus, the implementation of an efficient routing algorithm is an important design step in NoC architectures [1]. However, due to requirements as reduction of area and energy consumption, overly complex routing implementations that ensure greater throughput and reduced latency are not necessarily the best approach when designing a NoC, since it normally implies large spatial complexity that often translates into larger integrated circuit overhead. Aiming to overcome this problem, several 3D NoC architectures (e.g. [2][3][4][5]) employ deterministic routing algorithms, which implies less complex implementation, although efficient scheme implementation is still crucial for latency minimization. The amount of elementary operations of the routing algorithm (e.g. arbitration scheme) performed by the router to fulfill its design goal affects NoC latency and overall application performance.

Several works present different studies conducted on evaluating the arbitration mechanism efficiency, such as Wissem et al. [6] and Chan et al. [7], both proposing Quality-of-Service (QoS) approaches that avoid packet starvation and prioritize network traffic load, thus obtaining lower overall network latency under saturation traffic scenarios. The work of Kim et al. [8] propose a look-ahead arbitration scheme for reducing end-to-end communication latency, avoiding packet contention whenever possible by establishing a communication path from packet origin to its destination. It is noticeable that minimizing the router latency, sometimes increasing the capacity of taking parallel routing decision (e.g. parallel arbitration), implies some

degree of NoC latency minimization. It is also noticeable that there is a tradeoff between latency reduction and design complexity increase. Nevertheless, it is not obvious that a local router minimization influences directly and proportionally on the NoC latency minimization. To the best of our knowledge, we could not find a work that explores the effect of local router arbitration on a global NoC traffic; and this is the main purpose of this work.

This paper focuses on altering the routing arbitration mechanism targeting 3D mesh NoC by varying the amount of elementary operations that compose distinct steps in the Round-Robin packet switching algorithm, a topic not usually explored in literature where the best implementation of routing algorithms is usually taken for granted. Our goal, therefore, is to evaluate the performance impact of these changes on the latency estimations, based on the routing logic of a 3D mesh NoC. We focus on modeling the arbitration scheme by finite state machines, where each state represents a distinct step in the routing arbitration process; and state transitions are based on events of the routing process associated with control mechanisms relative to each router implementation.

This paper is organized as follows. Section II presents the Lasio 3D-NoC architecture used for the development of this work. Section III describes the evaluation criteria and test setup. Section IV presents the experimental results for different setup scenarios. Lastly, Section V presents the conclusions and further developments.

II. LASIO ARCHITECTURE

Lasio's architecture [9] is a 3D-Mesh NoC using TSV (Through Silicon Via) technology for interconnecting 2D mesh layers in the third dimension. Routing units are placed in a 3D mesh-based organization, with each router identified by its corresponding x , y and z coordinates. Packets are routed along the NoC through a deterministic XYZ routing algorithm. Each router connects directly a single Processing Element (PE) and each PE address is the same of the coordinates of the router it is connected. Figure 1 shows an example of a 4x4x4 Lasio highlighting the addressing of PEs according to x , y and z coordinates.

Each router of Lasio is composed of six ports (i.e. *North*, *South*, *West*, *East*, *Top* and *Bottom*) for interconnecting with other NoC's routers and a dedicated port (i.e. *Local*) for connecting the local PE. Although port roles are distinct, all seven ports are structurally identical, offering bidirectional communication and configurable buffer depths.

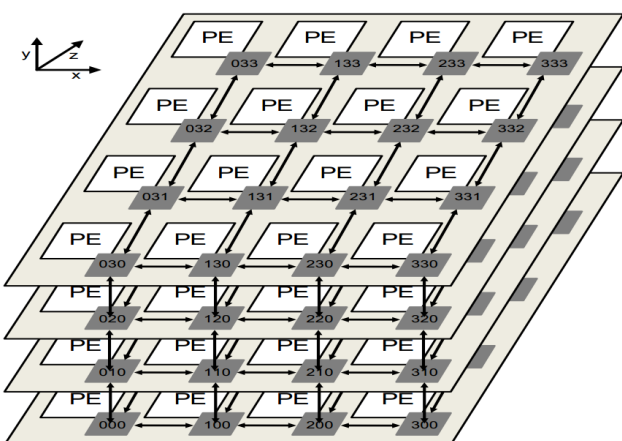


Figure 1 – Example of 4x4x4 Lasio NoC.

An in-house NoC generation tool support to implement Lasio with central or distributed packet arbitration. The distributed version encompasses one arbiter per output port, which increases its efficiency, but also increases area and energy consumption. The central version contains a single arbiter that coordinates incoming packet provided by all input ports, which uses a sequential approach (i.e. Round Robin). This work focus on the central packet arbitration, which is optimized in area and energy consumption.

Figure 2 shows a finite state machine (FSM) model of Lasio routers for the central packet arbitration.




Figure 2. The FSM that represents the arbitration scheme of a NoC's router.

Each state of Figure 2 represents a given phase in the routing process as follows:

- **S0 (Initializing state)** – the router passes only once by this state to perform initializing procedures (e.g., to set some register status), then FSM goes to state **S1** after a clock cycle;
- **S1 (Waiting state)** – the FSM remains waiting for incoming packets through any of the input ports on the router, or FSM deals with packets that remain waiting for the release of the destination port (i.e., the other packet that is using the same destination port release it). At that moment the FSM goes to state **S2**;
- **S2 (Verifying states)** – **S2** is a composition of two states implemented in two clock cycles that are responsible for verifying the packet destination address against the router

address and the corresponding destination port. If the destination port is free, the FSM finishes the arbitration and goes to **S3**. On the other hand, when destination port is busy, the FSM goes to **S1** for future arbitration (i.e. **reswitching**);

- **S3 (Ending state)** – **S3** is also implemented with two states. One is responsible for all flits delivering through the selected port, and the other one finishes the switching process by freeing the incoming data port. After that, the FSM goes to state **S1** to process further switching and routing requests.

This implementation of arbitration for XYZ routing algorithm takes in worst case 5 clock cycles to switch a requesting packet and 3 additional cycles for every reswitching due to network congestion (states **S1-S2**). Although low latency design is achieved with single-cycle switching implementations [10], it increases area overhead, and consequently energy consumption. The straightforward design of the abovementioned sequential logic aims to build a switching implementation without compromising area consumption.

III. EVALUATION CRITERIA

This section describes the criteria for performance evaluation of the NoC Lasio. As in many other studies of NoCs [11], packet latency is the foremost metric to this evaluation. Latency is influenced by a myriad of parameters. However, we chose to focus primarily on the efficiency of the routing unit represented as a FSM and the variation of injection rate of packets.

A. Latency Analysis

Packet latency is the specific metric chosen to evaluate the NoC performance, which comprises the interval of time between the moment a packet is inserted by the application and the moment this same packet is totally consumed by the target node.

B. Application Scenario

Latency is strongly dependent of communication pattern between nodes in a NoC. However, when a deterministic synthetic traffic scenario is chosen as application, evaluation of latency can be made independent of the communication pattern, since the packet load can be uniformly distributed. For this reason, we selected the synthetic *All-to-all* traffic for evaluation in this paper. In this traffic scenario, all nodes send the same quantity of data in a deterministic way to all other nodes, except to itself. Firstly, every node sends simultaneously a packet to the first node (node with address 000). Then, in the same matter, each one of the node sends another packet to the next node (node 001), and so on.

C. Injection Rate Test Sets

The injection rate of packets is a major impact factor in NoC performance [12]. As the number of packets injected in the network increases, paths are increasingly blocked due to network congestion. Nevertheless, network congestion can be alleviated as the efficiency of the routing arbitration's unit improves. However, as the efficiency grows, so does the complexity, increasing area consumption. Aiming to evaluate this tradeoff, this work adopts a set of exponential

grown injections rates defined by 2^n , where n ranges from 0 to 6.

D. Experimental Setup

The setup parameters of Lasio NoC were defined as follows: 16-bit flit size, two depths of buffer (i.e., 8-flit and 16-flit) and two sizes of 3D mesh (i.e., 2x2x2 and 4x4x4). Additionally, we employed *all-to-all* traffic scenario (56 packets for 2x2x2 and 4,032 packets for 4x4x4) in order to inject packets in the network, according to the definition found on Section C. Figure 3 shows the test setup.

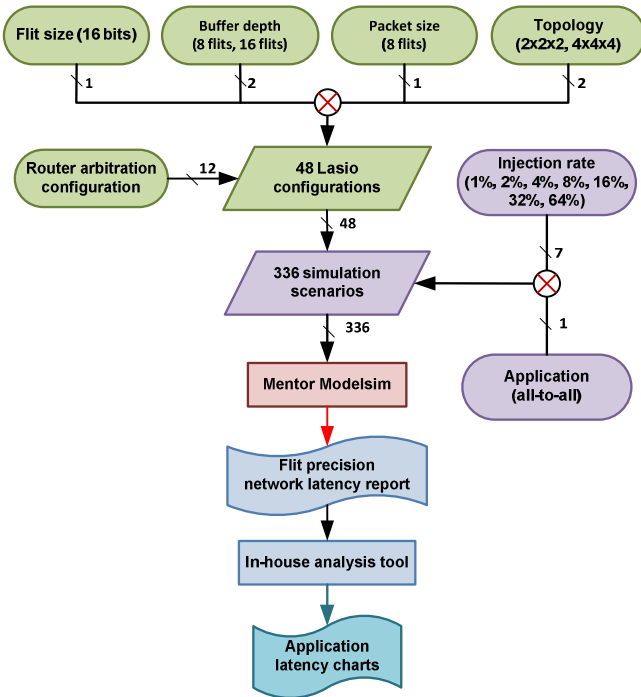


Figure 3. Composition and flow of the experimental setup.

For the evaluation of the impact of routing arbitration mechanism on latency, we construct two case scenarios. Firstly, additional states were inserted between states **S1** and **S2** (refer to Figure 2) composing scenario **#PSS** (*quantity of Packet Switching Stage*). This means that every single switching is delayed – including reswitching due to network congestion. The **#PFS** (*quantity of Packet Forwarding Stages*) is the second scenario resulting from inserting additional states after the switching process is completed (state **S3**). The **#PFS** scenario does not affect reswitching as scenario **#PSS**. The goal of multiple case scenarios is to evaluate the impact of sequential logic design on different locations of the routing arbitration’s state machine.

IV. EXPERIMENTAL RESULTS

This section presents the packet latencies evaluation according to the case scenarios from the routing arbitration’s state machine. In the next page, the Figure 4 shows the results obtained with scenario **#PFS**, whereas Figure 5 shows the results obtained with scenario **#PSS**. Additionally, both figures are divided in four charts aiming to compare the increase of NoC size (i.e. from 2x2x2 to 4x4x4) and buffer size (i.e. from 8 to 16).

The results are presented in two types of charts. Firstly, for scenario **#PFS**, a relative comparison of average

latencies (in clock cycles) was used. Every dot in this chart represents the average increase of latency in contrast with the lowest average latency found in the set being examined. Secondly, for scenario **#PSS**, we plot an absolute comparison of average latencies (in clock cycles). Every dot in this last case represents an absolute average latency obtained through VHDL simulation. In both cases, six different types of arbitration’s unit were defined. For both topologies and scenarios, note that latency always increases with the addition of states at the routing arbitration’s unit.

In the scenario **#PSS**, packet latency increases largely (e.g. in Figure 5b, Buffer 8 and 16% of injection rate, increasing five states increases 193%, i.e. $st+=1$ is 400 clock cycles and $st+=6$ is 1170 clock cycles) taking into account all experiments the latency increases 160%, in average. As one can see, scenario **#PFS** is not affected as scenario **#PSS** is. The latency increases in worst case close to 80%, and taking into account all experiments the latency increases only 25%, in average, showing that, for the criteria been evaluated: (i) the efficiency of the routing arbitration’s unit affects considerably the performance of NoCs; (ii) for increasingly injection rates of packets, the efficiency of reswitching becomes more relevant. This is due to more occurrences of reswitching by the saturation of resources in the interconnection fabric.

Although both NoC sizes have similar results, there are discrepancies worth highlighting. For the application latency, 4x4x4 NoC topology and scenario **#PFS**, the addition of states has no significant impact on latency (i.e. below 10% of latency increase) up to 2% of injection rate. After that, the addition of three or more states culminates in latency increases of 20% and above. The highest increase is reached at 16% injection rate and then it decreases slowly. However, it always exceeds the base latency of the routing arbitration’s unit of Figure 2. Note that, for a 32% injection rate, the application creates and injects new packets every third clock cycle. In addition, the basic arbitration unit may take five clock cycles to switch between incoming ports. Therefore, it is safe to argue that from 32% and beyond of injection rate the network is overly saturated with traffic’s demand. Therefore, the increase of arbitration’s states at this point is not as damaging as before. This behavior is replicated on the 2x2x2 NoC architecture. However, the addition of states has no latency impact up to 4% of injection rate. This difference found is due to the lower demand of traffic by the reduced number of PEs using the network. Although the ratio of PEs is the same for both topologies, the XYZ routing algorithm does not take advantage of the full potential of routing links [12]. Hence, the increase on the absolute number of PEs results in increase of average packet latency.

The scenario **#PSS** effects on both switching and reswitching. States added at this stage of the FSM will always underperform in comparison with the previous scenario. For traffic that requires no reswitching (i.e., no network congestion occurs), an additional state in the scenario **#PSS** would result in the same delay of switching that an additional state in the scenario **#PFS**. For a traffic that requires reswitching, though, every reswitching in the scenario **#PSS** requires an additional cycle for every

additional state added at the arbitration's unit. This does not occur at the scenario #PFS.

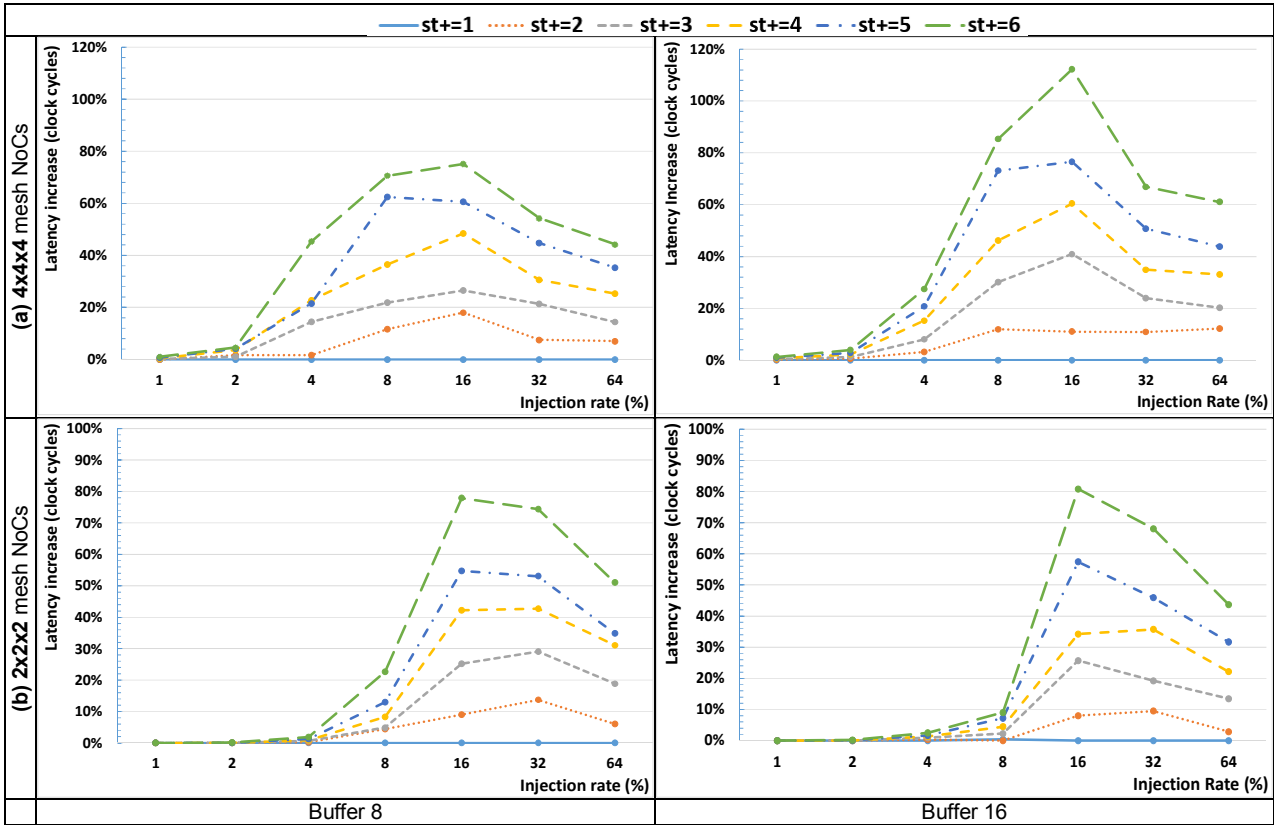


Figure 4. Applying additional switching states with scenario #PFS.

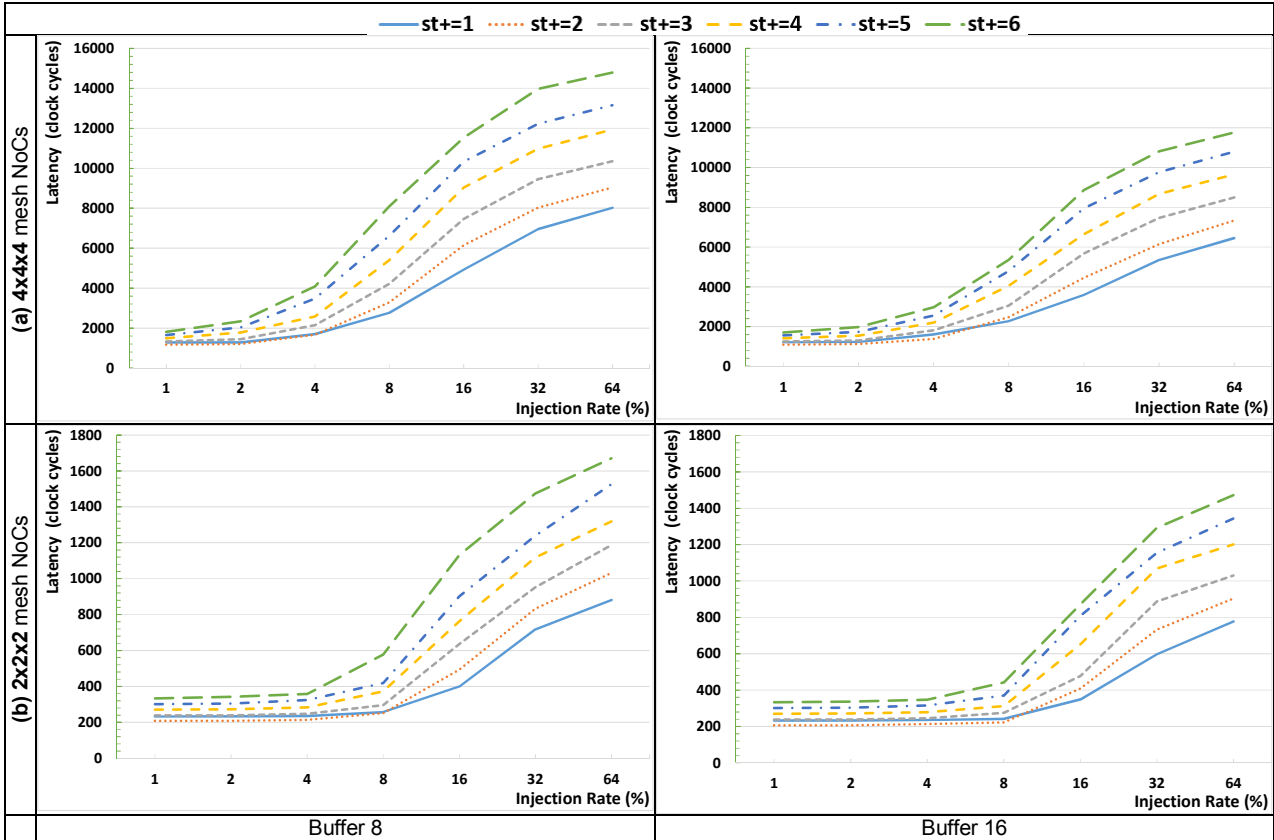


Figure 5. Applying additional reswitching states with scenario #PSS.

There is one caveat at this logic rationale. Since the base arbitration's unit takes three clock cycles for reswitching, it is possible that an extra clock cycle at this unit result in faster transmission. Take the following scenario: let be five consecutive clock cycles ($ck1, ck2, ck3, ck4, ck5$) and two routers that need to switch a packet to the same destination. The routers have a different arbitration's unit. $R1$ has the basic arbitration's unit and $R2$ has an additional state between $S1$ and $S2$ (Figure 2). Now, the destination router is blocked in cycles $ck1$ through $ck4$, and is free otherwise. Both routers will try, and they will fail to switch in $ck1$ due to the condition described earlier. Then, $R1$ will need three cycles to retry, and $R2$ will need four cycles to do the same. In $ck4$ the destination router is still blocked, and $R1$ will need another three cycles to retry. However, $R2$ will try and succeed in switching his packet at $ck5$. This scenario shows that it is possible for an additional state in the scenario #PSS to transmit faster than the basic arbitration's unit. Nevertheless, this is a peculiar and probably rare case in a traffic containing thousands of switching events.

The packet latency in scenario #PSS always increases in comparison with the base arbitration's unit, as shown in Figure 5. On a $2 \times 2 \times 2$ mesh NoC, an approximately fixed variation of latency occurs for an injection rate of 1% through 4%. This means that there is a very low percentage of reswitching occurring. Over these injection rates, latency increases dramatically as the performance is undermined by the increased delay in both switching and reswitching. A $4 \times 4 \times 4$ mesh NoC has similar results than the ones achieve with $2 \times 2 \times 2$ mesh NoC. However, the highest average latency is reached at 8% of injection rate. As explained earlier, this is expected due to the use of XYZ routing algorithm.

The results show that the increase in arbitration latency using sequential logic design does not affect proportionally the packet latency. In most cases, latency increase is kept under 100% even when the time of packet switching is doubled. Furthermore, the packet latency of reswitching logic becomes increasingly more relevant as the network's resources becomes scarce.

These results point out that although the optimization of arbitration and switching routing mechanisms implies the minimization of global NoC latency, on average, there are other aspects that compromises this gain, implying a tradeoff between the costs of minimizing router latency and the increase of its complexity that spend more energy and area.

V. CONCLUSIONS

Low latency design is a necessary trend to pursue, especially for designs of 3D NoCs. However, sophisticated switching/arbitration implementations increase the consumption of area and energy. We discussed a specific analysis on the arbitration scheme, which is a potential bottleneck on intrachip communication. The packet latency is dependent on the number of elementary operations composing the arbitration logic. Our purpose was to quantify

this dependence, highlighting the tradeoffs on low latency design and area overhead. Our experiments showed that the specific modifications on the routing arbitration's unit affect considerably the performance of NoC. Experimental results also demonstrate that for exponential packet's injection rates, the efficiency of reswitching becomes more relevant for estimating the latency on applications. Finally, experimental results indicate that the minimization of router's arbitration mechanism does not significantly influence on the overall NoC latency. Therefore, the latency gains may not compensate the consumption of area and energy of the router, which is required to achieve this optimization.

ACKNOWLEDGMENT

This work is partially funded by FAPERGS under grants PqG 12/1777-4 and Docfix SPI n.2843-25.51/12-3 and CNPq process number 132778/2014-9.

REFERENCES

- [1] A. Ahmed, A. Abdallah. **Low-overhead Routing Algorithm for 3D Network-on-Chip**. *International Conference on Networking and Computing (ICNC)*, pp.23-32, 2012.
- [2] Y. Yaoyao, L. Duan, J. Xu, J. Ouyang, M. Kwai Hung, Y. Xie. **3D optical networks-on-chip (NoC) for multiprocessor systems-on-chip (MPSoC)**. *IEEE International Conference on 3D System Integration (3DIC)*, pp.1-6, 2009.
- [3] A.-M. Rahmani, P. Liljeberg, J. Plosila, H. Tenhunen. **Developing a power-efficient and low-cost 3D NoC using smart GALS-based vertical channels**. *Journal of Computer and System Sciences*, vol.79, n.4, pp.440-456, Jun. 2013.
- [4] M. Ahmed, R. Kumar. **Parameterized path-based, randomized, oblivious, minimal routing in 3D mesh NoC**. *IEEE Region 10 Conference (TENCON)*, pp.1-6, 2012.
- [5] C. Marcon, T. Webber, R. Fernandes, R. Cataldo, F. Grando, L. Poehls, A. Benso. **Tiny - optimised 3D mesh NoC for area and latency minimisation**. *Electronics Letters*, vol.50, n.3, pp.165-166, 2014.
- [6] C. Wissem, B. Attia, A. Noureddine, A. Zitouni, R. Tourki. **Quality of Service Network on Chip based on a new priority arbitration mechanism**. *International Conference on Microelectronics (ICM)*, pp.1-6, 2011.
- [7] C.-H. Chan; K.-L. Tsai; F. Lai; S.-H. Tsai. **A priority based output arbiter for NoC router**. *IEEE International Symposium on Circuits and Systems (ISCAS)*, pp.1928-1931, 2011.
- [8] K. Kim; S.-J. Lee, K. Lee; H.-J. Yoo. **An arbitration look-ahead scheme for reducing end-to-end latency in networks on chip**. *IEEE International Symposium on Circuits and Systems (ISCAS)*, vol.3, pp.2357-2360, 2005.
- [9] Y. Ghidini, M. Moreira, L. Brahm, T. Webber, N. Calazans, C. Marcon. **Lasio 3D NoC vertical links serialization: Evaluation of latency and buffer occupancy**. *Symposium on Integrated Circuits and Systems Design (SBCCI)*, pp.1-6, 2013.
- [10] A. Kumar, P. Kundu, A. Singh, L.-S. Peh, N. Jha. **A 4.6Tbits/s 3.6GHz Single-cycle NoC Router with a Novel Switch Allocator in 65nm CMOS**. *International Conference on Computer Design (ICCD)*, pp. 63-70, 2007.
- [11] E. Salminen, A. Kulmana, T. Hämäläinen. **On network-on-chip comparison**. *Euromicro Conference on Digital System Design Architectures, Methods and Tools (DSD)*, pp.503-510, 2007.
- [12] M. Haghi, M. Rani. **Evaluation of Effect of Packet Injection Rate and Routing Algorithm on Network-on-Chip Performance**. *International Journal of Innovative Research in Science, Engineering and Technology (IJIRSET)*, vol. 3, n. 2, pp. 9589-9597, 2014.