

An Empirical Investigation of Mesh and Torus NoC Topologies Under Different Routing Algorithms and Traffic Models

M. Mirza-Aghatabar⁺, S. Koochi⁺, S. Hessabi^{*}, M. Pedram[†]

⁺^{*}Sharif University of Technology, Tehran, Iran, [†]University of Southern California, CA, USA
⁺{Aghatabar, Koochi}@ce.sharif.edu, ^{*}Hessabi@sharif.edu, [†]Pedram@usc.edu

Abstract

NoC is an efficient on-chip communication architecture for SoC architectures. It enables integration of a large number of computational and storage blocks on a single chip. NoCs have tackled the SoCs disadvantages and are scalable. In this paper, we compare two popular NoC topologies, i.e., mesh and torus, in terms of different figures of merit e.g., latency, power consumption, and power/throughput ratio under different routing algorithms and two common traffic models, uniform and hotspot. To the best of our knowledge, this is the first effort in comparing mesh and torus topologies under different routing algorithms and traffic models with respect to their performance and power consumption.

Keywords: Network-on-Chip (NoC), Mesh, Torus, Routing Algorithm, Uniform, HotSpot.

1. Introduction

The number of transistors per chip will increase beyond billions, due to technology scaling to less than 50nm, at the end of this decade [1]. Therefore, we should apply new methods to manage this huge number of transistors on a chip. System-on-Chip (SoC) and Network-on-Chip (NoC) are two main implementation approaches. Nowadays, lots of products, such as cell phones and portable computers are implemented on a silicon chip [2]. However, SoCs have some disadvantages, such as: (1) non-reusability and (2) low scalability; Furthermore, due to ever increasing number of transistors, they will have: (3) complex design and (4) long time to market [3]. Traditionally, communication between processing elements was based on buses. However, for large multiprocessor SoCs with many processing elements, it is expected that the bus will become a bottleneck from a performance, scalability and power dissipation point of view [4][5]. Therefore, the idea of networks on chip,

which consists of a set of routers interconnected by links, has evolved [6].

NoC is an efficient on-chip communication architecture for SoC architectures. It enables integration of a large number of computational and storage blocks on a single chip. NoCs have tackled many disadvantages of the SoCs and are structured, reusable, scalable, and have high performance [4][7]. Lots of topologies have been proposed for NoCs so far, such as Mesh [9], Torus [8], Star [10], Octagon [11], SPIN [12]. Among these topologies, mesh topology has gained more consideration by designers due to its simplicity (cf. Figure 1(c)). The main problem with the mesh topology is its long diameter that has negative effect on communication latency. Torus topology was proposed to reduce the latency of mesh and keep its simplicity (cf. Figure 1(b)). The only difference between torus and mesh topology is that the switches on the edges are connected to the switches on the opposite edges through wrap-around channels. Every switch has five active ports: one is connected to the local resource while the others are connected to the closest neighboring switches. Although the torus architecture reduces the network diameter, the long wrap-around connections may result in excessive delay. However, this problem can be avoided by folding the torus, as illustrated in Figure 1(a) [13]. Due to importance of these two topologies, i.e., mesh and torus, we compare the performance and power consumption of these NoC topologies under different routing algorithms.

Routing algorithms can be classified according to their adaptivity. A routing algorithm can be either deterministic or adaptive. Deterministic routing algorithms always supply the same path between a given source-destination pair. Adaptive routing algorithms use information about network traffic and channel status to avoid congested or faulty regions of the network [9]. Adaptive algorithms may be implemented partially or fully.




Figure 1: NoC Architectures: (a) Folded Torus 4×4, (b) Torus 4×4, (c) Mesh 4×4

We used XY routing as an example of deterministic, Odd-Even and Negative First Turn models as examples of partially adaptive [14][9], and Duato as an example of fully adaptive routing algorithms [9]. These examples were chosen because they are all deadlock-free and incur minimal hardware cost.

The torus architecture needs at least 2 virtual channels (one is needed for the mesh architecture) to be deadlock-free [9] under deterministic and partially adaptive routing, and 3 virtual channels (2 for mesh) for fully adaptive routing. Therefore, we implement these topologies under 1, 2 and 3 virtual channels and compare their performance and power consumption under four routing algorithms and two (i.e., uniform and hotspot) traffic models.

The goal of this paper is thus to provide a detailed comparative evaluation of two prominent NoC topologies (Mesh and Torus) under different routing algorithms. In section II, we take a look at related work. Section III, describes the network architecture of our model and Section IV presents the evaluation metrics such as latency and power consumption. In Section V, experimental results are presented, and finally in Section VI, we conclude our work and give the summary.

2. Related Work

Pande and Grecu in [15] focused on performance evaluation of a set of recently proposed NoC architectures with realistic traffic models, using a deterministic routing without addressing the effect of different routing algorithms.

Li [16] proposed a deadlock-free routing algorithm used in the torus architecture, and analyzed its performance for different sizes of networks. The author also reported the network power consumptions of torus with cycle-breaking routing, and that of mesh with XY routing. The power analysis is based on switches with the decoupled admission (decouples the flit admission buffers from physical channels) and the mux-based crossbar. The power analysis results account for both the switch power and link power consumptions. In [17], the effect of traffic localization on energy dissipation in NoC-based interconnect was investigated and through system level simulation, the authors

showed that energy reductions of up to 50% can be achieved by exploiting locality in communication.

Chiu in [14] introduced a new turn model routing algorithm, named Odd-Even, for designing adaptive wormhole routing algorithms for meshes without virtual channels. The model restricts the locations, where some turns can be taken due to the state of the packet, to an even or an odd numbered column. This way, the proposed routing algorithm avoids deadlock. In our study, we have chosen the Odd-Even routing model among various turn models, because in comparison with previous methods, the degree of routing adaptiveness provided by this model is higher. The mesh network may also benefit from this feature (Odd-Even restriction of turns) in terms of improving the communication efficiency. Reference [14] showed that the even adaptiveness provided by the odd-even turn model makes message routing less vulnerable to non-uniform traffic models, such as hotspot traffic. In addition, adoption of this feature results in lower network performance fluctuation with respect to different traffic patterns.

None of these works has investigated the effect of routing algorithms on system power consumption. In this paper, we compare the torus and mesh topologies under different implementation and usage scenarios, (e.g., virtual channels, traffic models, and specially routing algorithms) in terms of their power dissipation and performance.

3. Network Architecture

The base component in the network architecture is a node or Intellectual Property (IP) block that consists of a Processing Element (PE) and a Router. In an IP block, the PE injects/ejects the generated/received packets based on a traffic model e.g., uniform, hotspot, first/second matrix transpose, etc. Routers receive packets on their input channels, and direct the packets toward their destinations according to destination addresses and routing algorithms by sending them to selected output channels. In this paper, we focus on unicast routing.

An internal structure of an IP block is shown in Figure 2. A router comprises of a number of different components such as Address Extractor, which receives the packets and determines their destination addresses and keeps the packets until it routes them; Multiplexer and De-Multiplexer which are used to manage virtual channel operations; Selector unit which selects the appropriate virtual channel; Crossbar switch which connects each input channel to each unoccupied output channel; and Reservator unit which implements the routing algorithm and controls the crossbar switch and

other related sub-modules. We adopt wormhole switching [9] in our router architecture.

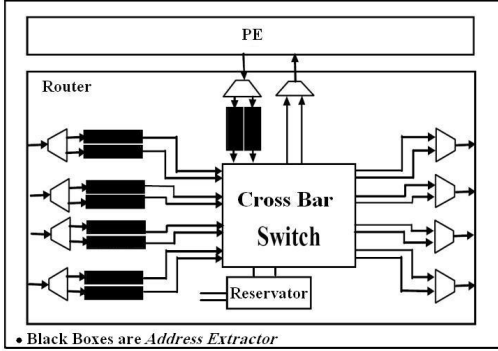


Figure 2: Hardware Implementation of an IP block [18]

4. Evaluation Metrics

Our evaluation metrics for comparing the mesh and torus architectures under different routing algorithms are latency, power consumption, and power/throughput ratio.

Throughput can be defined in a variety of different ways depending on the type of implementation.

Definition 4.1: Transport latency is defined as the time (in clock cycles) that elapses between the occurrence of a message header injection into the network at the source node, that include the queuing time in source, and the occurrence of the corresponding tail flit reception at the destination node [9].

Definition 4.2: In message passing systems, throughput (TP) may be defined as follows [15]:

$$TP = \frac{(total\ messages\ completed) \times (message\ length)}{(number\ of\ IP\ blocks) \times (total\ time)}$$

where *total messages completed* refers to the number of messages that successfully arrive at their destinations, *message length* is measured in flits, *number of IP blocks* is the number of functional IP blocks involved in the communication, and *total time* is the time (in clock cycles) that elapses between the occurrence of the first message generation and the last message reception. Thus, message throughput is measured as the fraction of the maximum load that the network is capable of physically handling. In this paper, we have: *total messages completed* = 8000, *message length* = 32, *number of IP blocks* = 16. The *total time* is a variable depending on the routing algorithms, traffic models, number of virtual channels, and NoC topology. We use the power/throughput ratio as a measure of the power-delay product (i.e., energy consumption) per message in the network.

Power consumption in NoCs consists of two components, the power consumed in routers, and the one associated with links:

$$P_{NoC} = P_{routers} + P_{links}$$

where $P_{routers}$ and P_{links} depend on the total capacitances and signal activity of the switch and each section of the interconnection, respectively. We calculated $P_{routers}$ using Power Compiler from Synopsys¹. The calculation accounts for both static and dynamic (switching and internal) power consumption. P_{links} is determined as follows:

$$P_{links} = \sum_{i=1}^{16} P_{link_i}$$

where

$$P_{link} = \alpha C_{wire} V_{DD}^2 f$$

We calculated the power dissipated in links of each router, and used UMC18 [19] which defines the V_{DD} for the 180nm technology as 1.98V. The clock frequency is set to 30MHz based on the critical path calculations. The switching activity (expected number of bit flippings on all nets in the NoC during one system clock cycle) of each link is extracted from *backsail* file, generated by Modelsim¹. The total capacitance of a wire can be approximated as:

$$C_{wire} = \frac{\epsilon}{d} WL$$

According to the International Technology Roadmap for Semiconductors [20] for the 180nm technology, the parameters are set to:

$$V_{DD} = 1.98\text{ v}, \epsilon = 3.5 \times 10^{-11}\text{ F/m}$$

$$d = 3.2 \times 10^{-4}\text{ m}, W = 3.2 \times 10^{-7}\text{ m}$$

We select the lengths of metal wire as 2mm for the mesh topology due to the size of local source, and 3mm for the torus topology, because the links in torus are longer than those in mesh. According to above data, $C_{wire} = 0.7\text{ pf}$ in mesh and 1.05 pf in torus.

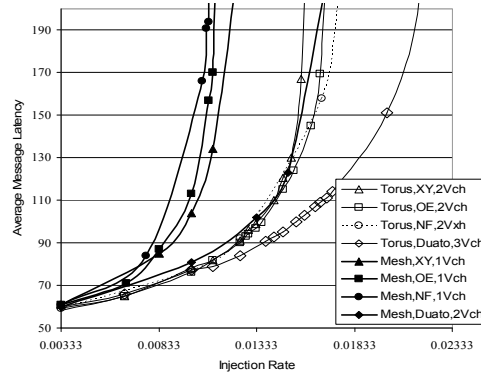

5. Experimental Results

In this section, we first compare the latency of mesh and torus topologies under different routing algorithms (XY, Odd-Even, Negative-First and Duato), different numbers of virtual channels, and two significant traffic models (Uniform and Hot Spot). Next we compare the power consumption and power/throughput of these topologies using the mentioned parameters.

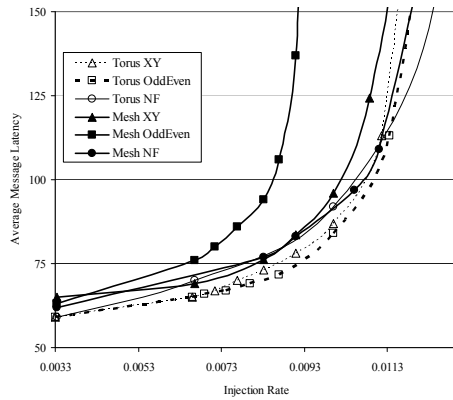
5.1. Latency

As mentioned before, we implemented the deadlock-free routing algorithms. We need at least *one*

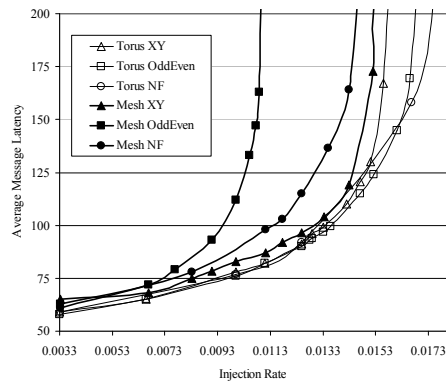
¹ Synopsys and Modelsim are registered trademarks of their respective owners



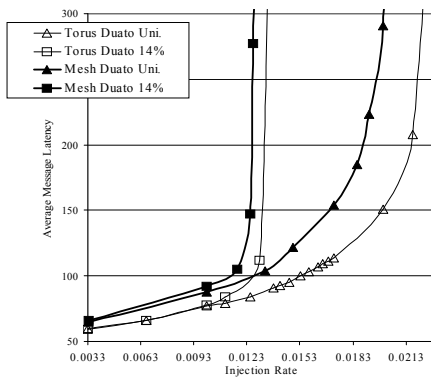
(a) Least Num. of Vchs to be deadlock-free-- Hot 14% (b) Least Num. of Vchs to be deadlock-free-- Uniform



(c) 2 Virtual Channels & Hot Spot 14%



(d) 2 Virtual Channels & Uniform



(e) 3 Virtual channels - Uniform & HotSpot

Figure 3: Latency versus injection rate and different virtual channels under uniform and hot 14% traffic distribution in torus and mesh topologies (Node 14 and 16 are hotspots in torus and mesh)

virtual channel for deterministic (XY) and partially adaptive (Odd-Even and Negative First) routing algorithms to be deadlock-free in mesh and two virtual channels in torus. In contrast, Duato fully adaptive routing algorithm needs two and three virtual channels,

respectively, for mesh and torus architectures to be deadlock-free [9]. First we compare the latency of these architectures when using the least number of virtual channels to remain free of deadlocks. Figure 3 shows the latency comparison of mesh and torus topologies under uniform (Figure 3(a)) and hot spot 14%² traffic models (Figure 3 (b)). We can see from these figures that the Duato routing algorithm in torus topology achieves the best latency, which is expected due to its wrap-around links and presence of three virtual channels.

Traffic distribution, which depends on traffic model and routing algorithm, has a direct impact on the network latency. Although partially adaptive routing algorithms have more adaptivity than the deterministic algorithm, they do not distribute the data traffic in a network any better than the deterministic one. For example, in Figure 3 (a,b,c,d), the XY routing results in lower latency than the Odd-Even routing in the mesh topology under both traffic models with one and two virtual channels. This comparison confirms that the XY routing results in a well-distributed traffic distribution due to its deterministic nature, i.e. XY

² 14% of messages are sent to Hot Node

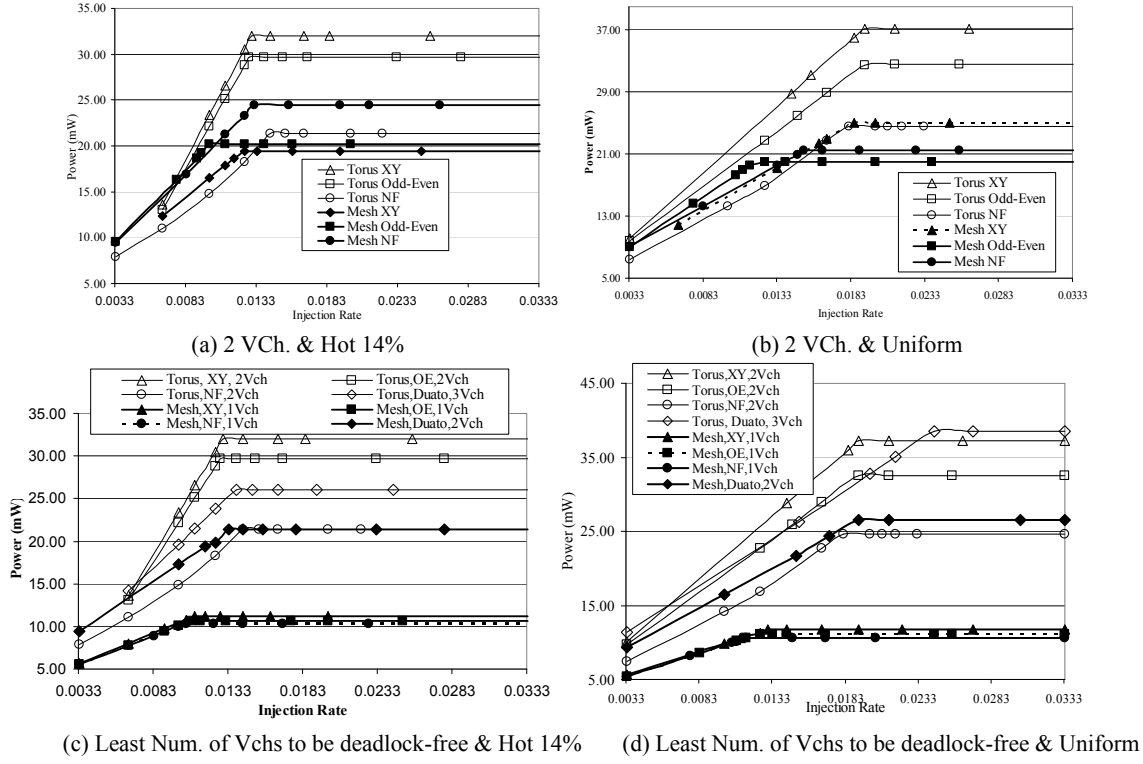


Figure 4: Power dissipation diagrams in torus and mesh topologies

routing distributes traffic based on dimension order, whereas traffic distribution in the Odd-Even routing is not as well-distributed (especially in mesh) and depends on the addresses of the source-destination nodes. This fact is also mentioned in [14]. Figure 3(c) shows that the Negative First routing has lower latency than XY routing under hot spot traffic model while Figure 3(d) shows its worse latency under uniform traffic model.

Figure 3(c,d) show that in a torus topology, partially adaptive routing algorithms have a little better latency than XY routing algorithm, but at all they are the same to some extent. For example the Odd-Even routing algorithm under all previous conditions has lower latency than the XY routing in torus topology. This higher performance of the Odd-Even routing

algorithm, which is not considerable, stems from its higher adaptivity, which in turn results from its better utilization of the wrap-around links compared to the XY routing algorithm. Therefore, the presence of wrap-around links in torus make this topology more efficient compared to mesh due to the better traffic distribution. As we can see in Figure 3(c,d), the torus architecture under hot spot and uniform traffic models with partially adaptive, fully adaptive and deterministic routing has better performance than the mesh architecture with equal number of virtual channels.

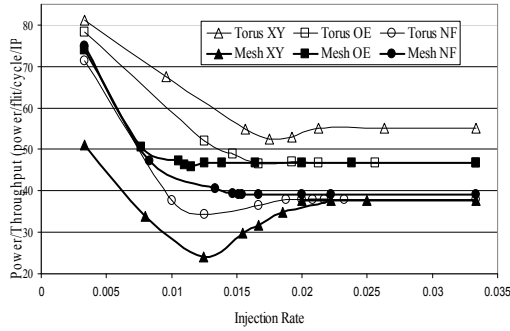
The same set of conclusions hold with respect to the throughput performance of the mesh and torus topologies. Results are not shown for brevity.

5.2. Power Dissipation and Power/throughput


Table 1: comparison of mesh and torus topologies

		Mesh/Torus						Selection			
		HotSpot			Uniform			HotSpot		Uniform	
Routing Alg.	#ViCh in Mesh	PP ¹	P/T ²	S ³	PP	P/T	S	Mesh	Torus	Mesh	Torus
XY	2	0.60	0.4	0.962	0.67	0.68	0.967	*		*	
NF	2	1.14	0.57	0.828	0.87	1.03	0.859		*		*
OE	2	0.68	0.41	0.78	0.61	1.04	0.635	*			*
Duato	3	0.89	0.86	0.96	0.72	1.18	0.854	*			*


1: Peak Power 2: Power/Throughput 3: Saturation Point
 *: means better selection between torus and mesh (under equal conditions)



(a) Power/Throughput- 2 VCh. & Hot Spot 14%



(b) Power/Throughput - 2 VCh. & Uniform



(c) Power/Throughput -3 VCh -Uniform & Hot 14%

Figure 5: Power/Throughput in torus and

Power consumption in NoCs depends on the total capacitances and signal activities of the switch and each section of the interconnect wire. The latency of different routing algorithms, and their traffic distribution capability, in mesh topology (with one virtual channel and two for Duato) is as follows:

$$Perf(Duato) > Perf(XY) > Perf(Odd-Even) \text{ and } Perf(Negative First)$$

Traffic distribution usually has a direct effect on the NoC power consumption i.e., better traffic distribution leads to less blocking time and more message transformation in a cycle which results in higher link activities and so more power consumption in network. As stated earlier, routing algorithms and traffic models determine the traffic distribution. The power dissipated in torus and mesh topologies under different conditions such as routing algorithms, virtual channels, and traffic models are reported in Figure 4. Power consumptions in mesh topology with 1 virtual channel (and 2 for Duato), which is shown in Figure 4(c,d) follows the partial order equation given above. This is also evident from Figure 4(a,b) where the power dissipated in mesh with XY routing is higher than that in mesh with Odd-Even routing. This observation is compatible with our previous statement, that traffic distribution is better in mesh with XY routing, which subsequently leads to higher power consumption. In other words, the cost one pays for lower latency with XY routing in mesh is

higher power consumption. In different applications one can thus choose one or the other of these routings depending on the desired latency versus power dissipation trade off point.

Next we evaluate the power consumption of routing algorithms in the torus topology. We have seen that routing algorithms with higher adaptivity benefit more from the wrap-around links of the torus topology, which in turn leads to a better traffic distribution. We shown that partially adaptive and XY routing algorithms have a similar latency in torus topology and XY is a little worse. At the same time, XY routing also results in the highest power consumption in the torus topology as seen in Figure 4(a,b). The reason is that, the traffic in each dimension is very high, which in turn results in high overall power consumption. We thus conclude that the XY routing is not suitable for torus architectures because it doesn't have a good latency and also has the highest power consumption compared to the adaptive routing algorithms.

Let's take a closer look at Figure 4(b). The power dissipated in torus or mesh topologies under uniform traffic is more than that of the hot spot traffic model. This means that although in a topology under hot spot traffic, the peak power of the hot node is more than the peak power of a node under uniform traffic, the total power dissipated in the uniform traffic is higher than that of the hot spot traffic because in the former case, one encounters lower message blocking in the network and better traffic distribution which results in higher total switching activity in the NoC. Figure 4(c,d) report the power consumption in the mesh and torus topologies with the least number of virtual channels to be deadlock-free under four different routing algorithms and two traffic models. We observe that the highest power consumption under the uniform traffic model is associated with the Dutao fully adaptive routing algorithm with 3 virtual channels. Having more virtual channels along with the effect of routing algorithm on traffic distribution are necessary

conditions for higher power consumption because they result in more multiplexing and switching in each channel as seen in Figure 4(c). However, these conditions are not sufficient to ensure higher power consumption. For example, Figure 4(c) shows that power consumption of XY and Odd-Even routing algorithms under hot spot traffic model with 2 virtual channels in torus topology are worse than that of the Duato routing with 3 virtual channels. This means that traffic models along with routing algorithms and number of virtual channels have an effect on the NoC power consumption. As another example, Figure 4 (d) shows that power consumption of the XY routing under uniform traffic with 2 virtual channels in torus topology is less than that of the XY routing under uniform traffic with 3 virtual channels. We conclude that when a routing algorithm leads to a better traffic distribution, such as XY routing in mesh topology, power consumption increases. At the same time, higher adaptivity of the routing algorithm results in lower power consumption under equal conditions (number of virtual channels, traffic model and topology).

The least power consumption is associated with the mesh topology with one virtual channel. Power consumed in a torus topology is always more than power consumed in a mesh topology with equal conditions.

It has been shown that power is strongly correlated to throughput in an NoC because the throughput determines the switching activity. So in general it is unfair to compare different architectures in terms of their power efficiency without also reporting their throughputs. It is thus more desirable to examine the power/throughput ratio (which is the same metric as the often-quoted power-delay product) of competing architectures. Figure 5. provides the power/throughput diagrams of mesh and torus topologies under different conditions such as routing algorithms. Clearly, a routing algorithm with lower power/throughput ratio is more sought-after. As we mentioned before, XY is not a fitting routing algorithm for torus topology. This claim can also be deduced from results reported in Figure 5(a,b), where it is seen that the XY routing has the largest power/throughput ratio under hot spot and uniform traffic models in torus topology among all examined routing algorithms. This arises from the fact that the XY routing algorithm profits the wrap-around links in torus topology lower than adaptive routing algorithms, due to its deterministic nature. Surprisingly, however, the XY routing in the mesh topology has the best power/throughput. This fact means that XY routing outperforms adaptive routing in mesh topology due to elimination of wrap-around links.

Table 1 reports a summary comparison of mesh and torus topologies under different conditions and network parameters. We report about three figures of merit in this table: PP, P/T and S which mean the ratio of peak power, power/throughput and injection rate at the brink of saturation in mesh and torus topologies, respectively. A value lower than unity for each of these figures of merit (PP, P/T and S) means lower power consumption, better power/throughput and smaller brink of saturation in mesh compared to torus, and vice versa.

Our selection mechanism, which is shown in the last four columns, is based on majority voting between these three factors (PP, P/T and S). For example, we conclude that under the uniform traffic model, torus topology is a better selection than mesh whereas under the hot spot traffic model, mesh topology is superior. Our proposed selection is a simple one which can be changed by NoC designers based on their target performance levels.

We also used PP, P/T and S factors as three different functions to select the best routing algorithm for each topology under uniform and hotspot 14% traffic models with the least number of virtual channels to be deadlock-free. The selection results are shown in Table 2. For example with P/T selection function Negative First is the best routing algorithm in torus topology under both uniform and hot spot 14% traffic models, and in mesh topology XY is the best one under hot spot 14% and Duato is the best one under uniform traffic model.

Table 2: best routing algorithm selection

Selection function	Mesh		Torus	
	HS_14% ⁺	Uniform	HS_14%	Uniform
PP	NF [*]	NF	NF	NF
P/T	XY	Duato	NF	NF
S	Duato	Duato	Duato	Duato

*: Negative First +: HotSpot 14%

In all, when latency is a constraining criterion, it is better to use the torus topology and when power consumption is a constraining criterion, it is better to use the mesh topology.

6. Summary and Conclusions

NoCs have tackled the SoCs disadvantages. Mesh and torus are two well-known and universal topologies among many presented NoC topologies.

We carried out detailed comparisons of mesh and torus topologies for different figures of merit such as latency, power consumption and power/throughput under equal conditions such as: routing algorithm, traffic model, number of virtual channels, etc, torus

always has better latency than mesh. However the cost we pay for this improvement is higher power consumption. We showed that the more adaptiveness a routing algorithm has, the lower power consumption in the torus topology is. So the XY routing as a deterministic routing is not a suitable routing algorithm for torus topology. Also we showed that the XY routing under hot spot traffic model is a good routing for the mesh topology.

Routing algorithms, traffic models and number of virtual channels have a direct effect on power consumption and give rise to interesting trade-offs. We showed that adaptive routing algorithms effectively utilize the wrap-around links in torus topology. We also showed that the XY routing algorithm has the largest power/throughput value in torus topology while, interestingly, it is the best routing algorithm for the mesh topology.

We proposed a simple selection function (with majority voting between three figures of merit: PP, P/T and S) between mesh and torus topologies. With our selection function torus was a better topology than mesh under both uniform and hot spot 14% traffic models.

Finally, we selected the best routing algorithm with the least number of virtual channels for each routing to be deadlock-free, in mesh and torus topologies under uniform and hotspot 14% traffic models with respect to PP, P/T and S factors as three different selection functions. Our overall conclusion is that when latency (power consumption) is a constraining criterion, it is better to use the torus (mesh) topology.

10. References

- [1] A. Allen, D. Edenfeld, W.H. Joyner, A.B. Kahng, M. Rodgers, and Y. Zorian, "2001 Technology Roadmap for Semiconductors," IEEE Computer, pp. 42-53, Jan. 2002.
- [2] E. Nilsson, "Design and Implementation of a Hot-Potato Switch in a Network-on-Chip," M.S. thesis, Royal Institute of Technology, Stockholm, Sweden, Jun. 2002.
- [3] M. Dall'Osso, G. Biccari, L. Giovannini, D. Bertozzi, and L. Bennini, "Xpipe: a Latency Insensitive Parameterized Network-on-Chip Architecture for Multi-Processor SoCs," ICCD03, pp.536-539, San Joes, CA, USA, Oct. 2003.
- [4] L. Benini and G. de Micheli, "Networks-on-Chip: A new Paradigm for System on Chip Design," Design Automation and Test in Europe, IEEE computer, vol. 35, no. 1, pp. 70-78, January 2002.
- [5] Pascal T Wolkotte, Gerard J M Smit, Nikolay Kavaldijev, Jens E Becker, Jurgen Becker, "Energy Model of Networks-on-Chip and a Bus," Accepted for the International Symposium on System-on-Chip (SoC 2005) Tampere, Finland, November, 2005.
- [6] S. Bhat, "Energy Models for Network-On-Chip Components," M.S. thesis, Royal Institute of Technology, Eindhoven, Netherlands, Dec. 2005.
- [7] F. Moraes, N. Calazans, A. Mello, L. Moller, L. Ost, "HERMES: an Infrastructure for Low Area Overhead Packet-Switching Networks on Chip," Integration, the VLSI Journal, vol. 38, pp. 69-93, 2004.
- [8] W.J. Dally and B. Towles, "Route Packets, Not Wires: On-Chip Interconnection Networks," Proc. Design Automation Conf. (DAC), pp. 683-689, 2001.
- [9] J. Duato, S. Yalamanchili, and L. Ni, "Interconnection Networks—An Engineering Approach," Morgan Kaufmann, 2002.
- [10] S. B. Akers and B. Krishnamurthy, "A Group-Theoretic Model for Symmetric Interconnection Networks," IEEE Transactions on Computers, vol. C-38, no. 4, pp. 555–566, April 1989.
- [11] F. Karim, A. Nguyen, and S. Dey, "An Interconnect Architecture for Networking Systems on Chip," IEEE Micro, vol. 22, no. 5, pp 36–45, September/October 2002.
- [12] Adriahtenaina, H. Charlery, A. Greiner, L. Mortiez, and C. A. Zeferino, "Spin: A scalable, Packet Switched, on Chip Micro-Network," In DATE 03 Embedded Software Forum, pages 70–73, 2003.
- [13] W.J. Dally and C.L. Seitz, "The Torus Routing Chip," Technical Report 5208:TR: 86, Computer Science Dept., California Inst. of Technology, pp. 1-19, 1986.
- [14] G. M. Chiu, "The Odd-Even Turn Model for Adaptive Routing," IEEE Transactions on Parallel and Distributed Systems, vol. 11, pp. 729-38, 2000.
- [15] Partha Pratim Pande, Cristian Grecu, Michael Jones, André Ivanov, Resve A. Saleh, "Performance Evaluation and Design Trade-Offs for Network-on-Chip Interconnect Architectures," IEEE Trans. Computers, vol. 54, no. 8, pp 1025-1040, 2005.
- [16] T. Li, "Estimation of Power Consumption in Wormhole Routed Networks on Chip," Stockholm, Sweden, May. 2005.
- [17] Partha Pratim Pande, Cristian Grecu, Michael Jones, André Ivanov, Res Saleh: "Effect of Traffic Localization on Energy Dissipation in NoC-based Interconnect," ISCAS (2), pp 1774-1777, 2005.
- [18] D. Rahmati, A. E. Kiasari, S. Hessabi and H. Sarbazi-Azad, "A Performance and Power Analysis of WK-Recursive and Mesh Networks for Network-on-Chips," Proceedings of the 24th International Conference on Computer Design (ICCD), 2006
- [19] United Microelectronics Corporation (UMC), <http://www.umc.com/>, Dec, 2006.
- [20] International Technology Roadmap for Semiconductor (ITRS), <http://www.itrs.net/>, Dec, 2006.