

Shuffle-Exchange Mesh Topology for Networks-on-Chip

Reza Sabbaghi-Nadooshan¹, Mehdi Modarressi^{2,3}
and Hamid Sarbazi-Azad^{2,3}

¹*Islamic Azad University Central Tehran Branch, Tehran, Iran*

²*IPM School of computer science, Tehran, Iran*

³*Sharif University of Technology, Tehran, Iran*

1. Introduction

Network-on-Chip (NoC) is a promising communication paradigm for multiprocessor system-on-chips. This communication paradigm has been inspired from the packet-based communication networks and aims at overcoming the performance and scalability problems of the shared buses in multi-core SoCs (System on Chips) (Benini & Mecheli, 2002).

Although the concept of NoCs is inspired from the traditional interconnection networks, they have some special properties which are different from the traditional networks. Compared to traditional networks, power consumption is the first-order constraint in NoC design (Ogras et al., 2005). As a result, not only should the designer optimize the NoC for delay (for traditional networks), but also for power consumption.

The choice of network topology is an important issue in designing a NoC. Different NoC topologies can dramatically affect the network characteristics, such as average inter-IP distance, total wire length, and communication flow distributions. These characteristics, in turn, determine the power consumption and average packet latency of NoC architectures.

In general, the topologies proposed for NoCs can be classified into two major classes, namely regular tile-based and application-specific. Compared to regular tile-based topologies, application-specific topologies are customized to give a higher performance for a specific application. Moreover, if the sizes of the IP cores of a NoC vary significantly, regular tile-based topologies may impose a high area overhead. This area overhead can be compensated by some advantages of regular tile-based architectures. Regular NoC architectures provide standard structured interconnects which ensures well-controlled electrical parameters. Moreover, usual physical design problems like crosstalk, timing closure, and wire routing and architectural problems such as routing, switching strategies and network protocols can be designed and optimized for a regular NoC and be reused in several SoCs.

The mesh topology is the simplest and most popular topology for today's regular tile-based NoCs. On the other hand, the shuffle-exchange topology is a well-known network structure which was initially proposed by Stone (Stone, 1971) as an efficient topology for multicomputer interconnection networks. Several researchers have studied the topological

properties, routing algorithms, efficient VLSI layout and other aspects of shuffle-exchange networks (Steinberg & Rodeh, 1981; Sparso et al., 1991).

The fact that shuffle-exchange networks have smaller diameter than equal sized meshes motivates us to investigate them as the underlying topology for on-chip networks. In this chapter, we propose a 2D shuffle-exchange mesh (SEM) topology for NoC implementation. We compare the two most important NoC factors (latency and power) of the same sized mesh and SEM NoC architectures. To this end, we have implemented the networks in question in a NoC simulator. Using this simulator, a routing scheme for the SEM has been developed and the performance and power consumption of the two networks have been evaluated under similar working conditions. The simulation results show that the SEM, while having equal implementation cost, consumes lesser energy and exhibits higher performance compared to the traditional mesh network.

In this chapter, we will introduce the two-dimensional SEM topology, and develop a deadlock-free routing algorithm for it. We also compare the power consumption and network performance of equal sized SEM and mesh NoCs.

2. The 2D SEM topology

2.1 The structure

The traditional shuffle-exchange network (Figure 1 shows an 8-node shuffle exchange network) is first proposed in (Stone, 1971). This topology is one of the most popular interconnection architectures for multiprocessors and multicomputers due to its scalability and distributed self routing capability (Kim & Veidenbaum, 1995). Several researchers have studied the topological properties (Park & Agrawal, 1995; Pifarre et al., 1994) and efficient VLSI layout (Steinberg & Rodeh, 1981; Sparso et al., 1991) of the shuffle-exchange networks. In a shuffle-exchange network, each node is identified by a unique n -bit binary address, hence the network size (number of nodes), N , equals 2^n . Two nodes are connected to each other if either their addresses differ in the last bit or one is a one-bit cyclic shift of the other. To establish these connections, two operations namely, shuffle and exchange, are used. With shuffle and exchange operations, message is circulated among network nodes until it reaches the destination node.

These operators that are defined on an n -bit address pattern $(A_{n-1}A_{n-2} \dots A_1A_0)$ as follows:

Shuffle: $(A_{n-1}A_{n-2} \dots A_1A_0) = A_{n-2}A_{n-3} \dots A_1A_0A_{n-1}$

Exchange: $(A_{n-1}A_{n-2} \dots A_1A_0) = A_{n-1}A_{n-2} \dots A_1\bar{A}_0$

Each node generates two connections to other nodes via shuffle and exchange operations and accepts two connections from other nodes. Since these connections are unidirectional, the degree of the network is the same as the one-dimensional mesh (linear array). The diameter of a shuffle-exchange network with size N is $2 \times \log(N) - 1$ which is the minimum distance between nodes 0 and $2^n - 1$.

Some researchers, e.g. in (Padmanabhan, 1991), have proposed different flavors of shuffle-exchange network structures and corresponding routing algorithms to allow more flexible network sizes instead of a complete size of 2^n .

In this chapter we propose a two-dimensional shuffle-exchange network architecture for network-on-chips. The architecture of this network is depicted in Figure 2. In this network, the nodes in each row and column form a shuffle-exchange network.




Fig. 1. An 8-node shuffle-exchange network; the bold lines are generated by exchange operation and other lines are generated by shuffle operation (Dally & seitz, 1987).

In each direction, each node has two outgoing edges along which it can send data packets to other nodes and two incoming links in each dimension and thus, has 8 unidirectional links in two dimensions. Thus, the number of links per node in the 2D SEM is equal to that in a traditional mesh network (i.e., 4 bidirectional links). Since the node degree of a topology has an important contribution in (and usually acts as the dominant factor of) the network cost, the 2D SEM and mesh NoCs have almost the same cost.

However, the network diameter of the 2D SEM is smaller than the diameter of the equivalent mesh. More precisely, the diameters of a 2D SEM and a mesh are $4 \times \log(2N^{0.5}) - 2$ and $2(N^{0.5} - 1)$, respectively where N is the network size.




Fig. 2. A 2D SEM with 64 nodes

In shuffle-exchange networks, every link generated by an exchange operation has one corresponding link in the mesh network. However, the links generated by shuffle operations connect some non-adjacent nodes (in equivalent mesh) and reduce the distance between two end points of the network. Compared to a mesh, although establishing the shuffle links remove the link between some adjacent nodes (for example 2 to 1, 6 to 5 and 3 to 4 connections in Figure 1) and increases their distance by one hop, the distance between a larger number of nodes is decreased by one or multiple hops and this leads to a considerable reduction in average distance of the network.

Although the dominant factor of the network cost, the node degree, in 2D SEM and mesh networks are exactly the same, unlike the mesh topology the 2D SEM links do not always connect the adjacent nodes and hence, their lengths are not the same. This can lead to some variations in the delay and power of the network links and may also have link placement difficulty. The latter can be solved by a number of efficient VLSI layouts proposed for shuffle-exchange networks (Steinberg & Rodeh, 1981; Sparso et al., 1991). Moreover, since the operating frequency of a NoC is often determined by the router critical path, the long wires may not degrade the NoC speed. However, in the case of frequency degradation, the pipelined packet switching technique (Duato et al., 2002) which involves inserting some one-flit buffers for the links can solve the problem. The effect of longer links on power consumption has been considered in our simulation results (presented in the next section).

2.2 Routing algorithm

During past years, a number of routing algorithms have been developed for traditional shuffle-exchange networks. Dally (Dally & Seitz, 1987) presented a routing algorithm which routes the packets from the source node toward the destination by changing the address one bit at a time, starting from the most significant bit of the n -bit source address in a 2^n -node network. At the i -th step of the algorithm, the $(n-i)$ -th bit of the destination address is compared to the LSB of the current address. If these two bits are equal, the message is routed over the shuffle channel to keep the bit unchanged and rotate the address. Otherwise, the message is routed over the exchange channel to make the two bits identical and then over to exchange channel to rotate the address. This algorithm involves a maximum of $2n$ communication steps between adjacent nodes along the path from the source to the destination node. However, this algorithm can not always find the shortest path for some source and destination pairs (Dally & Seitz, 1987). In order to be deadlock-free, this algorithm requires n virtual channels per physical channel and the message uses the i -th virtual channel at the $(n-i)$ -th step. Since in this virtual channel selection scenario routing is performed in order of decreasing order of virtual channel number, the dependency graph of virtual channels is acyclic and the routing is deadlock-free (Dally & Seitz, 1987).

Park (Park & Agrawal, 1995) improved Dally's routing (Dally & Seitz, 1987) using lower number of virtual channels per physical channel. They logically partition the network into several acyclic sub-networks and assign a rank to the sub-networks. Applying Dally's routing, the virtual channel number is increased only if the message enters a new partition with higher rank. As a result, the number of required virtual channels is reduced to $n - \lfloor (n-1)/2 \rfloor$.

Pifarre (Pifarre et al., 1994) introduced another deadlock-free routing algorithm for shuffle-exchange networks using only 4 virtual channels per physical channel regardless of the network size. However, in this algorithm, the maximum number of hops taken by a message increases from $2n$ (in Dally's algorithm (Dally & Seitz, 1987)) to $3n$. It first decomposes the network into some so called shuffle cycles by considering the network without exchange links. Note that every node in a shuffle cycle has the same number of 1s in its binary address which is defined as the level of a shuffle cycle. The routing algorithm involves two phases. In phase 1, at any step, a message stays in a shuffle cycle (if it is routed along a shuffle arc) or it is routed to a shuffle cycle of a higher level (if it is routed along a shuffle-exchange arc). In phase 2, the message is successively routed in shuffle cycles of decreasing levels. Consequently, every path has at most $3n$ steps: at most $2n$ shuffle steps and n exchange steps. The shuffle cycles can be made deadlock-free, in phase 1, by allocating two virtual channels. By allocating two more virtual channels for each shuffle arc, routing in shuffle cycles can be made deadlock-free, in phase 2.

For shuffle-exchange, we use a routing algorithm based on the algorithm proposed in (Pifarre et al., 1994). The algorithm decomposes the entire graph into several shuffle-cycles and constructs two increasing (in which the nodes are traversed in increasing number) and decreasing (in which the nodes are traversed in decreasing number) graphs as shown Figure 3. The algorithm involves two phases. The first phase, the increasing phase, visits the shuffle cycles in increasing order and the bit positions which are '0' in the source address and '1' in the destination address are changed to '1'. The other phase (the decreasing phase) visits the nodes in decreasing order in respect to their levels and bit positions which are '1' in the source address and '0' in the destination address are changed to '0'. We used the modified algorithm which removes the self loops and makes the path shorter.

As can be seen in Figure 3, the shuffle cycles in the increasing graph can be made deadlock-free by allocating two virtual channels which break the cycle. By allocating two more virtual channels for each shuffle cycle, routing in shuffle cycles can be made deadlock-free along the decreasing graph in phase 2, as well. Therefore, the network should have 4 virtual channels per physical channel to make our algorithm deadlock-free.

Now, after designing a routing scheme for the shuffle-exchange, we develop a deterministic and an adaptive routing mechanism for the 2D SEM. Like XY routing algorithm in mesh networks, the deterministic routing applies the above-mentioned routing mechanism in rows first in order to deliver the packet to the column at which the destination is located. Afterwards, the message is routed to the destination by applying the same routing algorithm in that column. Obviously, adding the second dimension in this routing scheme does not generate a cycle and is deadlock-free provided that the routing in each dimension is deadlock-free (Duato et al., 2002).

In the adaptive routing mechanism, on the other hand, all possible minimal paths between a source and a destination node are of potential use along the path depending on the traffic congestion and network conditions. Since each node is connected to the nodes in its row and column via a shuffle-exchange network, in each node, the routing algorithm routes the packets along one of the two networks based on the traffic congestion and resource availability. We avoid deadlocks using a deadlock-free routing methodology presented in (Duato, 1995) which divides the virtual channels into two adaptive and deterministic parts and uses the deterministic part upon message blockage in adaptive part.




Fig. 3. The logical partitioning of a shuffle-exchange network of size 8

3. Comparison results

In order to compare the energy dissipation and performance of the 2D SEM with the mesh, we have used a modified version of the Popnet NoC simulator (Popnet, 2007). The simulator can simulate and calculate the performance measures of NoCs under different traffic patterns and supports virtual channel-based wormhole switching. It also includes the Orion power library (Wang et al., 2002) that can calculate the energy dissipated in the NoC under simulation. For our experiments, we set the network link width to 32 bits (flit size = phit size = 32 bits). The power is calculated based on a NoC with 180 nm technology whose routers operate at 250 MHz.

The simulation results is obtained for an 8×8 mesh interconnection network with XY routing algorithm and an 8×8 2D SEM using the routing algorithms described in the previous section. The message length is assumed to be 32 and 64 flits and 4 and 6 virtual channels per physical channel are used. Messages are generated according to a Poisson distribution with rate λ , and the destinations of the messages are uniformly selected from the network nodes.

In Figure 4, the average message latency is plotted as a function of message generation rate at each node for the mesh and 2D SEM networks using deterministic routing (which involves 4 virtual channels) for two different message sizes. As can be seen in the figure, the 2D SEM has smaller average message latency with respect to the equivalent mesh network. The reason is that the average inter-node distance of the 2D SEM network is lower than the equivalent mesh network.

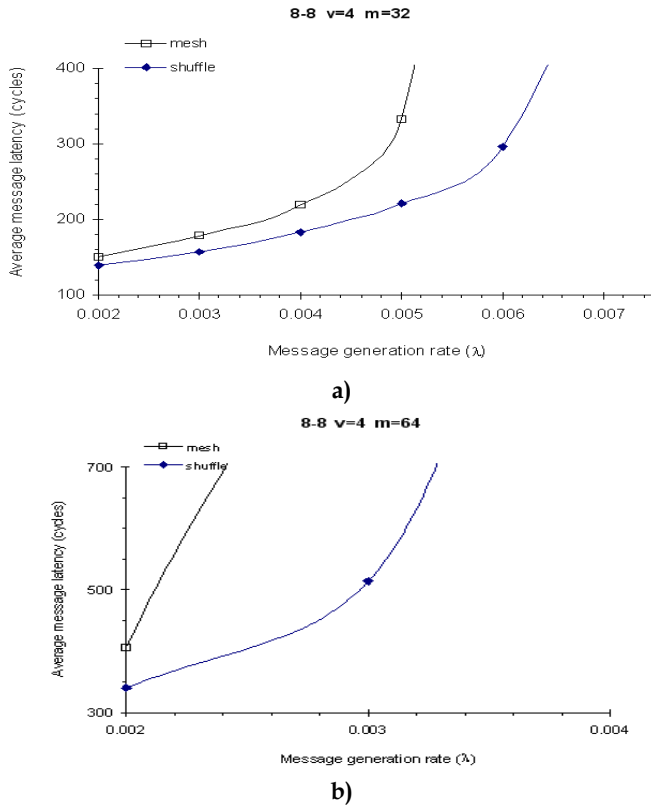


Fig 4. The average message latency of deterministic routing in the 64-node 2D SEM and mesh networks using 4 virtual channels per physical channel with message length a) 32 flits and b) 64 flits.

Figure 5 compares the latency results of adaptive and deterministic routing schemes in a 2D SEM. In order to conduct a fair comparison, both routing algorithms use 6 virtual channels per physical channel (deterministic routing algorithm employs 6 virtual channels per physical channel while adaptive routing algorithm divides the virtual channels into 2-virtual channel adaptive and 4-virtual channel deterministic parts). It can be seen that the adaptive routing algorithm has improved the average message latency compared to the deterministic routing. The improvement is more significant in high-traffic regions where adaptive resolves contentions more effectively.

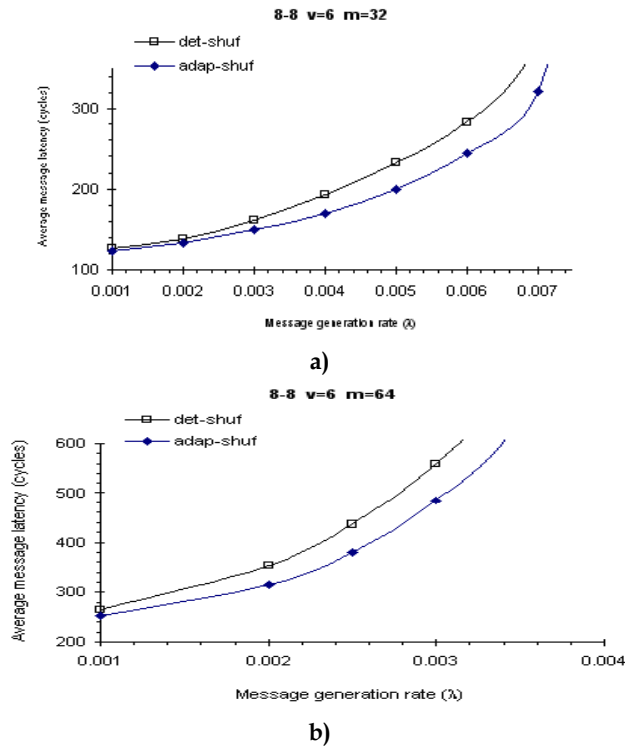


Fig. 5. The average message latency of the deterministic and adaptive routing algorithms in a 64-node 2D SEM using 6 virtual channels per physical channel with a) 32-flit messages and b) 64-flit messages.

As mentioned before, the effect of wire lengths in power consumption is considered in the calculation of consumed power by Orion. Based on the core size information presented in (Mullins et al., 2006), we set the side size of the cores of our simulated 8×8 NoCs to 2 mm. The length of the shuffle wires in the 2D SEM is set based on the number of cores they pass. Figure 6 displays the power consumption of the mesh and 2D SEM networks using deterministic routing scheme in the scenario used in figure 4. As can be seen in the figure, the proposed 2D SEM topology can effectively reduce the power consumption of the NoC. The main source of this reduction is the long wires which bypass some nodes and hence save the power which is consumed in intermediate routers in an equivalent mesh topology. Note that when the mesh network reaches to its saturation region, the 2D SEM network still can handle the traffic and thus the saturation rate for the 2D SEM is higher than that in the mesh. The extra messages communicated in the network have increased the total power consumption in the 2D SEM after the saturation rate of the mesh network. This is of course natural to have more energy consumed for higher traffic crates.

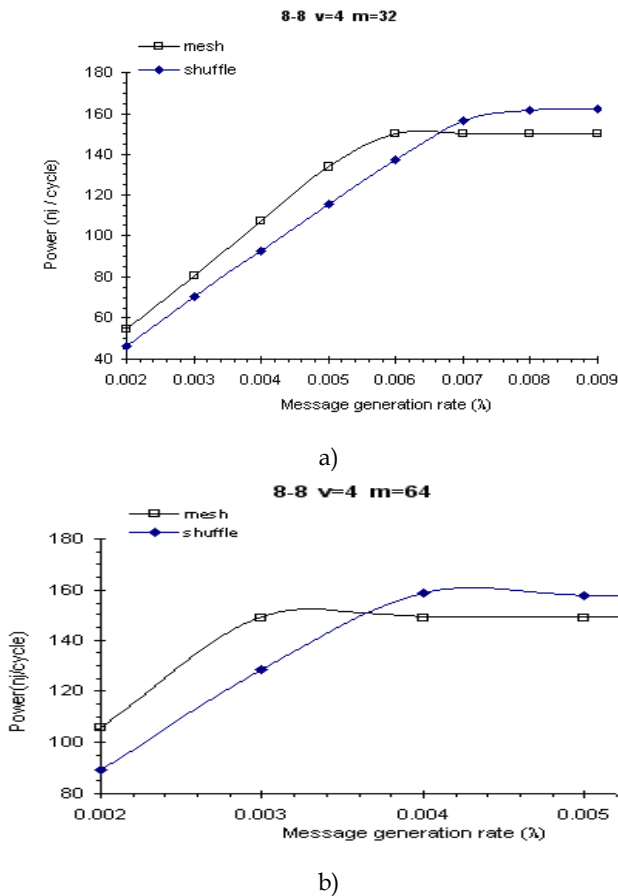


Fig. 6. The power consumption of 64-node mesh and 2D SEM NoCs using deterministic routing and 4 virtual channels per physical channel with a) 32-flit and b) 64-flit messages.

The area estimation is done based on the hybrid synthesis-analytical area models presented in (Mullins et al. , 2006; Kim et al., 2006; Kim et al. 2008). In these papers, the area of the router building blocks is calculated in 90nm standard cell ASIC technology and then analytically combined to estimate the router total area. Table 1 outlines the parameters. The analytical area models for NoC and its components are displayed in Table 2. The area of a router is estimated based on the area of the input buffers, network interface queues, and crossbar switch, since the router area is dominated by these components.

The area overhead due to the additional inter-router wires is analyzed by calculating the number channels in a mesh-based NoC. A $n \times n$ mesh has $2 \times n \times (n-1)$ channels. The 2D SEM has the same channels as mesh with longer wires. In the analysis, the lengths of packetization and depacketization queues are considered as large as 64 flits.

In Table 3, the area overhead of 2D SEM NoC is calculated for 8×8 network size in a 32-bit wide system. The results show that, in an 8×8 mesh, the total area of the 2mm links and the

routers are 0.0633 mm² and 0.1089 mm², respectively. Based on these area estimations, the area of the network part of the 2D SEM network shows a 27% increase compared to a simple 2D mesh with equal size. Considering 2mm×2mm processing elements, the increase in the entire chip area is less than 2%. Obviously, by increasing the buffer sizes, the network node/configuration switch area increases, leading to much reduction in the area overhead of the proposed architecture.

Parameter	Symbol
Flit Size	F
Buffer Depth	B
No. of Virtual channels	V
Buffer area (0.00002 mm ² /bit (Kim et al., 2008))	B _{area}
Wire pitch (0.00024 mm (ITRS, 2007))	W _{pitch}
No. of Ports	P
Network Size	N (= n×n)
Packetization queue capacity	PQ
Depacketization queue capacity	DQ
Channel Area (0.00099 mm ² /bit/mm (Mullinset al. , 2006))	W _{area}
Channel Length (2mm)	L
No. of Channels	N _{channel}

Table 1. Parameters

	Symbol	Model
Crossbar	RCX _{area}	W _{pitch} ² ×P×P×F ²
Buffer (per port)	RBF _{area}	B _{area} ×F×V×B
Router	R _{area}	RCX _{area} +P×RBF _{area}
Network Adaptor	NA _{area}	PQ×B _{area} +DQ×B _{area}
Channel	CH _{area}	F×W _{area} ×L×N _{channel}
NoC Area	NoC _{area}	n ² ×(R _{area} +NA _{area})+CH _{area}

Table 2. Area analytical model

Network	Link Area	Router Area	Increase percent to mesh	increase percent in the entire chip
mesh	.06338	.1089	0	0
2D SEM	.0905	.1089	27.69	1.91

Table 3. 2D SEM area overhead

4. Conclusion

The mesh topology has been used in a variety of interconnection network applications especially for NoC designs due to its desirable properties in VLSI implementation. In this chapter, we proposed a new topology based on the shuffle-exchange topology, the 2D

shuffle-exchange mesh (2D SEM), and conducted latency and power consumption comparative simulation experiments for the proposed topology and mesh network. Simulation results showed that the 2D SEM can improve the latency of the network especially for high traffic loads. The power consumption in the 2D SEM is also shown to be less than that of the equivalent mesh network.

We also analyzed the effects of the various wire lengths in the implementation of the 2D SEM. Finding an optimal mapping scheme for the 2D SEM NoCs and also a VLSI layout based on the design considerations in deep sub-micron era is the future work in this line.

5. References

- <http://www.princeton.edu/~lshang/popnet.html>, August 2007.
- Benini, L. & Micheli GD. (2002). Networks on Chip: A New Paradigm for Systems on Chip Design, *Design Automation and Test in Europe (DATE)*, pp. 418-419.
- Dally, WJ. & Seitz, C. (1987). Deadlock-free Message Routing in Multiprocessor Interconnection Networks, *IEEE Trans. on Computers*, Vol. 36, No. 5, pp. 547-553.
- Duato, J. (1995). A Necessary and Sufficient Condition for Deadlock-free Adaptive Routing in Wormhole Networks, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 6, No. 10, pp. 1055-1067.
- Duato, J.; Yalamanchili, S. & Ni, L. (2002). *Interconnection Networks: An Engineering Approach*, Morgan Kaufmann Publishers.
- ITRS. (2006). International technology roadmap for semiconductors. *Tech. rep.*, International Technology Roadmap for Semiconductors.
- Kim, M.; Kim, D. & Sobelman, E. (2006). NoC link analysis under power and performance constraints, *IEEE International Symposium on Circuits and Systems (ISCAS)*, Greece.
- Kim, MM.; Davis, JD.; Oskin, M & Austin, T. (2008). Polymorphic on-Chip Networks, *International Symposium on Computer Architecture (ISCA)*, pp. 101 -112.
- Kim, S. & Veidenbaum, AV. (1995). On Shortest Path Routing in Single Stage Shuffle-Exchange Networks, *In Proc. 7th Annual ACM Symposium on Parallel Algorithms and Architectures*, pp. 298-307.
- Mullins, R.; West, A. & Moore, S. (2006). The Design and Implementation of a Low-Latency On-Chip Network, *Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 164-169.
- Ogras, UY.; HU, J. & Marculescu, R. (2005). Key Research Problems in NoC Design: A Holistic Perspective, *CODES+ISSS*, Jersey City, NJ, pp. 69-74.
- Padmanabhan, K. (1991). Design and Analysis of Even-Sized Binary Shuffle-Exchange Networks for Multiprocessors, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 2, No. 4, pp. 385-397.
- Park, H.; Agrawal, DP. (1995). Efficient Deadlock-free Wormhole Routing in Shuffle-based Networks, *7th IEEE Symposium on Parallel and Distributed Processing*, pp. 92-99.
- Pifarré, GD. et al. (1994). Fully Adaptive Minimal Deadlock-Free Packet Routing in Hypercubes, Meshes, and other Networks: Algorithms and Simulations, *IEEE transaction on Parallel and Distributed Systems*, Vol. 4, pp. 247-263.
- Sparso, J. et al. (1991). An Area-efficient Topology for VLSI Implementation of Viterbi decoders and Other Shuffle-Exchange type Structures, *IEEE journal of solid-state circuits*, Vol. 24, No. 2, pp.90-97.

- Steinberg, D. & Rodeh, M. (1981). A Layout for the Shuffle-Exchange Network with $O(N^2/\log^3/2N)$ Area, *IEEE Trans. On Computers*, Vol. C-30, No. 12, pp. 971-982.
- Stone, H. (1971). Parallel Processing With Perfect Shuffle, *IEEE Trans. on Computers*, Vol. 20, pp. 153-161.
- Wang, H.; Zhu, X.; Peh, L. & Malik, S. (2002). Orion: A Power-Performance Simulator for Interconnection Networks, *35th International Symposium on Microarchitecture (MICRO)*, Turkey, pp. 294-305.



Parallel and Distributed Computing

Edited by Alberto Ros

ISBN 978-953-307-057-5

Hard cover, 290 pages

Publisher InTech

Published online 01, January, 2010

Published in print edition January, 2010

The 14 chapters presented in this book cover a wide variety of representative works ranging from hardware design to application development. Particularly, the topics that are addressed are programmable and reconfigurable devices and systems, dependability of GPUs (General Purpose Units), network topologies, cache coherence protocols, resource allocation, scheduling algorithms, peertopeer networks, largescale network simulation, and parallel routines and algorithms. In this way, the articles included in this book constitute an excellent reference for engineers and researchers who have particular interests in each of these topics in parallel and distributed computing.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Reza Sabbaghi-Nadooshan, Mehdi Modarressi and Hamid Sarbazi-Azad (2010). Shuffle-Exchange Mesh Topology for Networks-on-Chip, Parallel and Distributed Computing, Alberto Ros (Ed.), ISBN: 978-953-307-057-5, InTech, Available from: <http://www.intechopen.com/books/parallel-and-distributed-computing/shuffle-exchange-mesh-topology-for-networks-on-chip>

INTECH

open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2010 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](#), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.