Convex-Based DOR Routing for Virtualization of NoC

Guang Sun, Yuanyuan Zhang, Yong Li, Li Su, Depeng Jin, and Lieguang Zeng

Tsinghua University, Main Building 901c, Beijing, China
sung08@mails.tsinghua.edu.cn

Abstract. Network on Chip (NoC) is proposed as a promising intra-chip communication infrastructure. A simple and efficient routing scheme is important for large scale NoC to provide the required communication performance to applications with low area and power overheads. Although mesh is preferred for NoC, virtualization may lead to irregular topologies. In this paper, we propose a Convex-Based DOR (CBDOR) routing scheme for the convex topologies. We demonstrate the connectedness and deadlock-freedom of CBDOR. This routing mechanism relies only on two bits per switch. Simulation results show that the area overhead of CBDOR switch is just 2.2% higher than that of traditional DOR switch, with the added complexity negligible. Therefore, the simplicity in the routing mechanism and switch architecture makes CBDOR more practical and scalable when compared to LBDR and FDOR.

Keywords: Network on Chip (NoC); virtualization; irregular topology; Dimension-Order-Routing (DOR); connectedness; deadlock-free.

1 Introduction

With the development of the semiconductor technology, SoC (System-on-Chip) design is entering billion-transistor era [1], [2]. More and more IPs and smart interconnects are integrated in one chip. In order to alleviate the complex communication problems which arise as the number of on-chip components increases, Network on Chip (NoC) architecture has been recently proposed as a promising communication paradigm to replace global interconnects [3], [4]. NoC addresses the on-chip communication problem with a networking approach and provides notable improvements in terms of area, performance, scalability, reliability and flexibility over the traditional bus-based structures [5], [6], [7].

The topology and routing algorithm are the two important features that influence the network performance, cost and power consumption [8]. The Regular topology, especially the 2D mesh topology, becomes a kind of popular architecture for NoC design, for it is very simple and efficient from a layout perspective [9], [10]. For the traditional routing algorithms, logic-based routing algorithm (e.g. DOR) is preferred as it can reduce latency, power and area requirements [9]. However, most of them are used in regular topologies.

In order to fully exploit the increasing number of cores and get enough parallelism for applications, virtualization for multicore chips is becoming necessary [8], [10], [11], [12]. The virtualized NoC solution provides several advantages such as increasing resource utilization, reducing power consumption and increasing the yield of chips [11]. Although the concept of virtualization is not new (e.g. virtual memory, virtual machines), there are some challenges when applying it to NoC. A virtualized NoC may be viewed as a network that partitions itself into several different regions, with each region serving different applications and traffic flows concurrently [11]. In a virtualized NoC, the system should guarantee traffic isolation among regions. Therefore, virtualization for NoC may lead to irregular sub-networks within the original 2D mesh [9]. Fig. 1 shows an irregular topology resulting from virtualization of NoC.
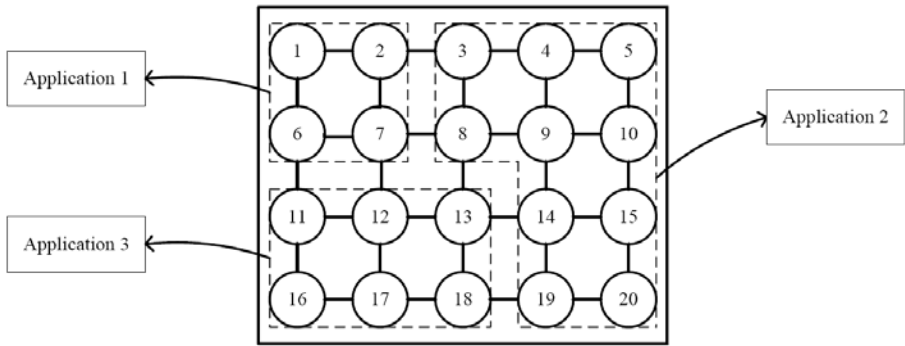

Fig. 1. An irregular sub-network topology resulting from virtualization of NoC

According to the aforementioned motivations, it would be very necessary to develop an efficient routing algorithm for most practical irregular topologies. In this paper, we propose a Convex-Based DOR (CBDOR) routing scheme for most of the practical topologies that we might find in the near future in NoC. The CBDOR is a logic-based routing algorithm without the need of forwarding tables in the switches. Moreover, CBDOR is connected and deadlock-free. This routing scheme uses only two bits per switch to route in the convex topologies with minimal hops. Therefore, CBDOR is very simple. Simulation results show that the area overhead of CBDOR switch is just 2.2% higher than that of traditional DOR switch. Therefore, the simplicity in the routing mechanism and switch architecture makes CBDOR very practical and scalable.

The remainder of paper is organized as follows. Section 2 presents some related work. Then, Section 3 describes the CBDOR mechanism in detail. Section 4 analyses the connectedness and deadlock-freedom of CBDOR, and provides some evaluations. Finally, in Section 5, we conclude our work.

2   Related Work

Traditional DOR is very simple and efficient, but it can only be used in regular 2D mesh. Traditional routings used in irregular topology mainly include source-based routing and table-based routing [9]. In source-based routing, the entire path is stored in the packet header [9]. Because the packet header itself must be transmitted during routing, it consumes network bandwidth, especially when the routing path is very long or the packet has only a few bits. In table-based routing, a table at each switch stores the output port that must be used for each destination [9]. The main advantage of table-based routing is that it can be used in any topology and with any routing algorithm [9]. However, as the size of NoC increasing, the memory requirements for building such routing tables also increase, thus exhibiting longer access delay and consuming more areas and power overheads [8].

LBDR [9] is a logic-based routing without the need of forwarding tables in the switches. All the supported topologies share the same property: all the end-nodes can communicate with the rest of nodes through the minimal path defined in the original mesh topology, namely, convex property. However, LBDR need 12 bits per switch and some extra logic gates. Moreover, the routing computing is very complex.

FDOR [12] is also a logic-based routing, but the supported topologies only include the FDOR-topology, which must satisfy three requirements. First, it must consist of three meshes: the core, the X- flank and X+ flank. Second, the core mesh must have at least the same number of dimensions as any flank. Third, the border of flank is only a sub-border of the core mesh. Thus, the supported topologies in FDOR are just a very small part of the convex topologies.

In this paper, we propose CBDOR, which is a logic-based routing algorithm without the need of forwarding tables in the switches. The CBDOR uses only two bits per switch to routing in the convex topologies with minimal hops. So CBDOR is very simple and efficient. In addition, when the shape and the size of convex topology are varied, LBDR [9] need to regenerate the routing bits and run routing computing again, which is very complex. However, the CBDOR scheme directly routes according to the two connectivity bits per switch, without any more operations. Thus, this simplicity in the routing mechanism makes CBDOR more practical and scalable.

3   Proposed Routing Algorithm: CBDOR

This section describes the principles of CBDOR. Firstly, we define the topology conditions on which CBDOR can be used.

All the supported topologies in the CBDOR must be convex. It means that the topology region must contain all the line segments connecting any pair of points inside it. Intuitively, the supported topologies can not have holes inside the region. Therefore, in the CBDOR scheme, all the end nodes can communicate with the rest of nodes through a minimal path defined in the original 2D mesh topology (pictured in Fig.2 (a)). Fig. 2 (a-c) shows some examples of topologies supported by CBDOR. As

a counter example, Fig.2 (d) shows a topology with a hole inside the region. In this case, this topology is not convex. So the CBDOR can not be applied in this topology.
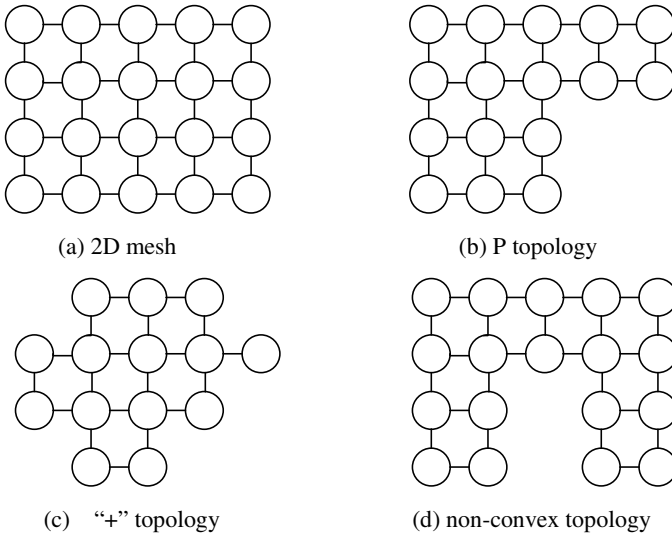

(a) 2D mesh   (b) P topology

(c)   "+" topology   (d) non-convex topology

Fig. 2. Examples of topologies (a-c) supported and (d) not supported by CBDOR

Implementation of CBDOR is very simple due to the fact that it is a pure logic-based routing scheme. According to the two connectivity bits per switch, the CBDOR can change the order in which the different dimensions are routed, so the routing logic in the CBDOR is slightly more complex than the one in the DOR. We verify this in the next section by synthesizing the two switches. Simulation results show that the added complexity of the CBDOR is negligible.

When irregular topologies are generated, the one connectivity bit per output port can be given by global manager (GM), which is responsible for system resource management [13]. The connectivity bits indicate whether a switch is connected with its neighbors [9]. For example, if the switch is not connected through the south port, C_s equals zero. Otherwise C_s equals one. Thus, the connectivity bits are C_s, C_n, C_e and C_w. In our scheme, we only use C_s and C_n to route. We assume that the X and Y coordinates of the final destination are stored in the message header (X_{des} and Y_{des}), and each switch knows it X and Y coordinates (through X_{cur} and Y_{cur} registers at each switch) [9]. The coordinates of the switch which is in the left and bottom corner of the 2D-mesh are smallest. Messages are routed with CBDOR scheme from current switch to destination switch, according to the offsets of coordinates and the two connectivity bits per switch. The CBDOR scheme is described in detail as below.

The pseudo-code of the CBDOR scheme:

```
if  Y_des <  Y_cur  and  C_s =1 then
    output_port= south;
elseif  Y_des >  Y_cur  and  C_n =1 then
    output_port= north;
elseif  X_des >  X_cur  then
    output_port= east;
elseif  X_des <  X_cur  then
    output_port= west;
else
    output_port= local_node;
end
```

4   Analyses and Evaluations

In this section, we analyze the connectedness among all the nodes and the deadlock-freedom in the CBDOR scheme. In order to evaluate the complexity of the CBDOR scheme, we provide the synthesis results for the switch architecture of the CBDOR and compared them with DOR, which is very simple in routing logic and economical in area overheads.

4.1   Connectedness

The connectedness of a routing scheme means that any pair of nodes in the supported topologies can communicate with each other by the routing scheme.

   All the supported topologies in the CBDOR must be convex. If the pair of nodes can be included in a rectangle topology, which is a part of original irregular topology, they can communicate with each other by the scheme just as DOR. Thus, in this case, the connectedness of the CBDOR is satisfied.

   If the pair of nodes can not be included in a rectangle topology, the connectedness of the CBDOR is satisfied through change the order in which the different dimensions are routed. For example, the destination switch is in the northeast of the source switch (as showed in Fig. 3). Before Y_{des} equals Y_{cur}, the CBDOR scheme in the intermediate switches can be divided into two cases. First, messages are routed in the north direction when C_n equals one. Second, messages are routed in the east direction when C_n equals zero. Because of the convex property of the supported topology, this scheme can make sure that the messages can be routed towards destination switch until Y_{des} equals Y_{cur}. After Y_{des} equals Y_{cur}, the messages are routed in east direction until X_{des} equals X_{cur}. Thus, in this case, the connectedness of the CBDOR is satisfied.

Similarly, if the destination switch is in other position of the source switch, the connectedness of the CBDOR is also satisfied.
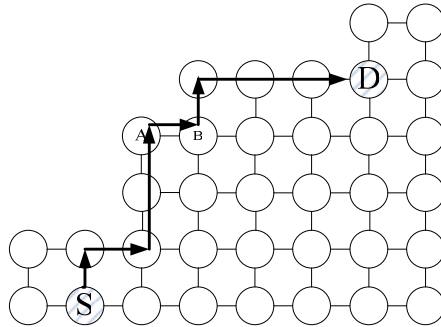

Fig. 3. The routing path from source switch to its northeast destination switch with the CBDOR scheme

From aforementioned analyses, we prove the connectedness of the CBDOR scheme. Moreover, because of the convex property of the supported topology, the CBDOR can route with minimal hops, which equals the Manhattan distance between the source and destination switches.

4.2   Deadlock-Freedom

A routing scheme for the supported topology is deadlock-free if there are no cycles in its channel dependency graph. Fig. 4 shows the channel dependency in YX routing. The turns marked with dotted lines are forbidden and the turns marked with solid lines are permissible.
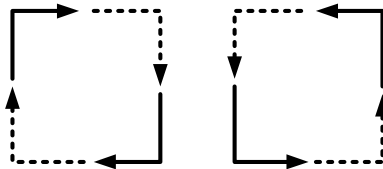

Fig. 4. The channel dependency in YX routing. The turns marked with dotted lines are forbidden and the turns marked with solid lines are permissible.

From Fig. 4, we can find WS, ES, EN and WN[1] turns are forbidden in YX routing. Although it is possible that WS, ES, EN and WN turns happen in the CBDOR routing, the CBDOR is deadlock-free. We take the EN turn as an example to demonstrate that there are no cycles in its channel dependency graph. If EN turn happens in a

---

[1] S, W, N and E represent south, west, north and east respectively.

switch, we can deduce that the connectivity bit C_n of its west neighbor switch equals zero. For example, in Fig.3, switch B may have EN turn, and the connectivity bit C_n of its west neighbor switch, namely the C_n of switch A, equals zero. This is to say, if a switch permits the EN turn, the north port of its west neighbor switch is unavailable. Because of the convex property of the supported topology, the EN turn can not form cycles in its channel dependency graph. Similarly, we can demonstrate that the WS, ES and WN turns in the CBDOR can not form cycles in its channel dependency graph. Therefore, the CBDOR scheme is deadlock-free.

4.3 Evaluation and Results

In this section, we implement the synthesis for the switch architecture of CBDOR, and compared them with tradition DOR.

In order to compare the area overheads of switch architectures using different routing schemes, VHDL language is used to design our performance simulation platform because it is believed that simulation platform designed by hardware description language is more similar to realistic on-chip network.

It is believed that the DOR scheme is very simple in routing logic and economical in area overhead. In order to evaluate the complexity of the CBDOR scheme, we synthesize the switch architectures for the CBDOR and traditional DOR with the VHDL-based platform. In the simulation, both schemes use wormhole switch technique. The synthesized results are given in Table 1. In Table 1, the unit of synthesized area is ALUT (adaptive look-up table). From Table 1, we can see that the area overhead of the CBDOR switch is just 2.2% higher than that of traditional DOR switch. Thus, the added complexity of the CBDOR is negligible. Regarding the network-throughput performance of the CBDOR, this scheme is slightly worse than DOR because it aggregates some traffic towards the edges of the topology.

Table 1. Synthesized results of switches using the DOR and the CBDOR

	Area Overhead
DOR switch	6975 ALUT
CBDOR switch	7131 ALUT
Added	2.2%

5 Conclusion

In order to fully exploit the increasing number of cores and get enough parallelism for applications, it is necessary to allow for partitioning the whole interconnected network into several separate regions. However, virtualization for NoC may lead to irregular topologies. So we propose a Convex-Based DOR (CBDOR) routing scheme for most of the practical topologies that we might find in the near future in NoC. We demonstrate the connectedness and deadlock-freedom of CBDOR. This routing scheme is a pure logic-based routing and uses only two bits per switch to route in the convex topologies with minimal hops. Therefore, CBDOR is very simple. Simulation

results show that the area overhead of CBDOR switch is just 2.2% higher than that of traditional DOR switch, with the added complexity negligible. Therefore, the simplicity in the routing mechanism and switch architecture makes CBDOR more practical and scalable when compared to LBDR and FDOR.

References

1. Hu, J., Marculescu, R.: Energy-aware mapping for tile-based NoC architectures under performance constraints. In: ASP-DAC 2003, pp. 233–239. IEEE Press, Kitakyushu (2003)
2. Benini, L., Micheli, G.D.: Networks on chips: a new SoC paradigm. Computer 35, 70–78 (2002)
3. Marculescu, R., Ogras, U., Peh, L.S., Jerger, N., Hoskote, Y.: Outstanding research problems in NoC design: system, microarchitecture, and circuit perspectives. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 28, 3–21 (2009)
4. Kreutzand, M., Marcon, C.L., Carroand, N., Susin, A.: Energy and latency evaluation of NoC topologies. In: ISCAS 2005, pp. 5866–5869. IEEE Press, Kobe (2005)
5. Lee, H.G., Chang, N., Ogras, U.Y., Marculescu, R.: On-chip communication architecture exploration: a quantitative evaluation of point-to-point, bus, and network-on-chip approaches. ACM Transactions on Design Automation of Electronic Systems 12, 20–40 (2007)
6. Dally, W., Towles, B.: Route packets, not wires: On-chip interconnection networks. In: DAC 2001, pp. 684–689. IEEE Press, Las Vegas (2001)
7. Horowitz, M., Ho, R., Mai, K.: The future of wires. Proceedings of the IEEE 89, 490–504 (2001)
8. Mejia, A., Palesi, M., Flich, J., Kumar, S., Lopez, P., Holsmark, R., Duato, J.: Region-Based Routing: A Mechanism to Support Efficient Routing Algorithms in NoCs. IEEE Transactions on Very Large Scale Integration (VLSI) Systems 17, 356–369 (2009)
9. Flich, J., Rodrigo, S., Duato, J.: An efficient implementation of distributed routing algorithms for nocs. In: The Second ACM/IEEE International Symposium on Networks-on-Chip, pp. 87–96. IEEE Press, Newcastle (2008)
10. Gratz, P., Kim, C., McDonald, R., Keckler, S.W., Burger, D.C.: Implementation and evaluation of on-chip network architectures. In: International Conference on Computer Design 2006, pp. 477–484. IEEE Press, Las Vegas (2006)
11. Flich, J., Rodrigo, S., Duato, J., Sodring, T., Solheim, A.G., Skeie, T., Lysne, O.: On the Potential of NoC Virtualization for Multicore Chips. In: CISIS 2008, pp. 165–177. IEEE Press, Washington (2008)
12. Skeie, T., Sem-Jacobsen, F.O., Rodrigo, S., Flich, J., Bertozzi, D., Medardoni, S.: Flexible DOR routing for virtualization of multicore chips. In: International Symposium on System-on-Chip 2009, pp. 73–76. IEEE Press, Tampere (2009)
13. Chou, C.L., Ogras, U.Y., Marculescu, R.: Energy-and performance-aware incremental mapping for networks on chip with multiple voltage levels. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems 27, 1866–1879 (2008)