



RETRACTED CHAPTER: Overview of Router Architecture in High Performance Computing

Dejun Shi¹, Xiaohu Han², Weijian Chen², and Hongliang Li²(✉)

¹ Information Engineering University, Zhengzhou 450000, China

² Jiangnan Institute of Computing Technology, Wuxi 214000, China
jurge@foxmail.com

Abstract. High Performance Computing (HPC) system has excellent computing power. The router is at the heart of a high-performance interconnection network, which is an important building block of HPC systems. Router architecture is classified into three types: crossbar, hierarchical, and network structures. The crossbar is the fairest, and the most basic component of the router, but it is also the most difficult to extend. Hierarchical structure using the hierarchical scheduling algorithm reduces arbitration complexity. According to the idea of network design, multiple crossbars are connected in a specific topology using integrated circuit technology implemented on a single circuit chip. This paper analyzes the characteristics of the three structures and related studies, which have guiding significance for the subsequent high-radix router design. This work also demonstrates the performance of various router architectures in terms of channel throughput and latency.

Keywords: crossbar · hierarchical structure · high-radix router · high performance computing · latency · network structure · throughput

1 Introduction

HPC is the system of integrating the computing power with a view to large problems in business, engineering, or science at surpassing higher performance than a typical computing system or workstation. HPC systems have excellent computing power, which is of great significance for high-speed computing fields such as high-speed numerical computing, complex system simulation, and massive data processing. The source of the powerful computing ability of HPC systems is the computing nodes. Thousands of computing nodes working in parallel provide the required computing performance for the entire system. A large number of computing nodes working in parallel and maximizing the computing efficiency of each computing node can't do without the high-performance interconnection network, which is another important component of an HPC system.

A high-performance interconnection network is an important part of an HPC system, which undertakes the function of high-bandwidth and low-latency data communication between the computing nodes. As high-performance computers gradually reach exascale or even beyond exascale, the need for a large-scale high-performance interconnection network is increasingly urgent.

The original version of this chapter was retracted: The retraction note to this chapter is available at https://doi.org/10.2991/978-94-6463-142-5_93

© The Author(s) 2023, corrected publication 2024

Y. Jiang et al. (Eds.): ICFIED 2023, AEBMR 237, pp. 493–506, 2023.

https://doi.org/10.2991/978-94-6463-142-5_57

The key component of the high-performance interconnection network is the router, which expands the node scale connected by a single router and plays an important role in the high-performance interconnection network. The router undertakes the function of data exchange in the high-performance interconnection network, and it is a key component of building a high-bandwidth and low-latency network. The core of the router is the crossbar, which can exchange input and output data completely unblocked, but the crossbar has the area and performance expansion problems. As the number of ports increases, the cost of crossbar becomes unacceptable, and hierarchical structure [1] and network structure [2] are successively emerging in the study of building high-radix routers.

The major contributions of this research work are to:

- Analyse the crossbar, hierarchical structure, and network structure of the router and summarize to guide the subsequent router design.
- Perform experimental study to demonstrate the performance of the crossbar, hierarchical structure, and network structure of the router architecture in terms of channel throughput and latency.

The rest of this article is structured as follows: Sect. 2 elucidates the state-of-the-art methods related to the crossbar router architecture. Section 3 analyzes the hierarchical router architectures. Section 4 investigates the network router architecture. Experimental results were presented in Sect. 5. Section 6 briefs about the recent studies related to the routing architecture. Section 7 presents a conclusion and future directions.

2 Crossbar

Crossbar is the core of the traditional router, which can exchange the data of all ports at the same time if there is no conflict among the target ports. However, there are obvious disadvantages to the crossbar. Crossbar implementation complexity is proportional to the number of ports, and as the number of ports increases, the throughput is limited. For the IQ (Input Queue) [3] router, the arriving packets are first stored in the input buffer, waiting to be transmitted through the crossbar. When an input buffer is managed as a single FIFO (First Input First Output) queue, a packet to an idle output port may be blocked, behind a packet that is waiting for a busy output port. Thus, the HoL (Head of Line) problem occurs. For uniform traffic, HoL limits the IQ router throughput to 58.6% [4].

Chen and Charol [5, 6] proposed a technique to reduce HoL in the IQ router by assuming that the first K ($K > 1$) packets in each FIFO queue can request to their respective output ports at the same time. Although issuing multiple requests simultaneously can improve throughput, it is more sensitive to traffic patterns because the first K packets of the FIFO queue in burst traffic may all go to the same output port.

Setting a separate queue for each output in the input buffer instead of just one queue completely eliminates HoL blocking. This is called VOQ (Virtual Output Queueing) and was proposed by Tamir et al. [7, 8]. HoL blocking in VOQ was eliminated because packets only queued behind those packets going to the same output port; and there were no packets pointing to different output ports before. When using VOQ, it has been demonstrated that the throughput of the input queue router can be increased from 58.6%

to 100% [9] for both uniform traffic and non-uniform traffic. In an IQ router with a VOQ queue, the number of VOQ queue is the square of the number of ports. When the number of ports is large, the large demand on the buffer of the VOQ makes the design difficult to implement, so it is difficult to extend to a high-radix router.

Opposite to the IQ router is the output OQ (Output Queued) router [10, 11], where the buffer is set only at the output port of the router. The packets arriving at each input port first enter the corresponding output buffer and then compete with other packets in the output buffer for the output port. Because the input packet is not blocked in the output buffer, i.e., with no HoL problem, the OQ router has high performance, but the output buffer of each port is required to store the packets from multiple input ports at the same time, which needs $N \times \text{speedup}$ and is costly and difficult to extend to more high-radix ports.

The IQ router could match the performance of the OQ router with a moderate speedup [12, 13]. Research has derived the CIOQ (Combined Input and Output Queued) hybrid router architecture from two opposite directions: either from pure IQ, where adding a moderate speedup factor can improve the delay-throughput characteristic; or from pure OQ, where adding input queues can provide a superior packet loss characteristic [14]. Adding an output buffer to an IQ router improves its performance to a certain extent, and the no packet loss characteristic required by high-performance interconnection networks can be implemented through the back pressure mechanism. The output buffer added here does not need to provide speedup capability. For the virtual channel router [15], the existence of the output buffer also decouples the crossbar allocation and physical channel arbitration. Without allocating the crossbar and arbitrating the output physical channel at the same time, when the crossbar allocation is complete, packets are forwarded to the corresponding virtual channel output buffer, and then each virtual channel multiplexes the single physical channel to simplify the design.

Jun et al. [16] proposed a scalable two-dimensional crossbar matrix switching architecture, which consists of multiple VOQ crossbars with an output buffer for each crossbar output, and proposed a hierarchical scheduling algorithm for crossbars that can provide 100% throughput in uniform traffic. The scalability of this architecture is mainly due to the fact that the hierarchical scheduling algorithm reduces the arbitration complexity, but the number of VOQ queues is proportional to the square of the number of ports, limiting the ability to make further expansion.

An $n \times n$ crossbar directly connects n inputs to n outputs, and each of the n inputs has connections to all n outputs, and each output selects one of the n inputs based on the selection signal. Figure 1 shows an implementation of the $n \times n$ crossbar [17].

The selection signal is the result of the input to the output scheduling arbitration. Allocation performs input-to-output matching, each input may request multiple outputs, and each output has requests from multiple inputs. The result of allocation is that at most one request is selected for each input, and each output receives at most one request. The allocation algorithm of crossbar should produce the maximal matching that meets the constraint, make as many inputs and output matches as possible, to improve the utilization rate of crossbar.

The OQ router has a crossbar and an output port buffer, which require the write speed of the output buffer to be N (number of ports) times as fast as the line speed, and are not

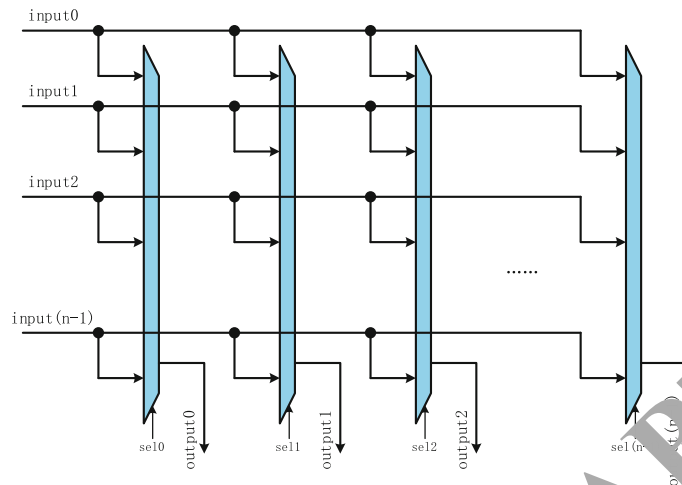


Fig. 1. $n \times n$ crossbar [self-drawn]




Fig. 2. Crossbar adding buffer at crosspoint [self-drawn]

suitable for extending to larger high-radix routers. The IQ router has crossbar and input port buffers, requiring complex allocation algorithms to resolve conflicts between input and output ports. The architecture of a single input buffer queue limits the throughput performance due to HoL problems. The input buffer maintains a multi-queue architecture of one queue for each output, which, combined with iSLIP [1], DRRM [18] or other allocation algorithms, can achieve 100% throughput in uniform traffic. However, due to the limitation of the allocation time, VOQ routers are not applicable for high-radix routers with a large number of ports.

By adding buffers (Fig. 2) at the crosspoint of all inputs and outputs of the crossbar, the input packet is saved at the corresponding crosspoint buffer according to the outgoing output direction. The crosspoint buffers that hold packets for the same output compete for the output port, and the winner sends the packet to the output port. After adding the crosspoint buffer, the input and output are isolated by the buffer, the HoL problem disappears completely, and the throughput performance is fully scalable. However, at each crosspoint, a buffer is added, and the number of buffers is $O(n^2)$. When the number of ports increases, the presence of a large number of buffers makes such an architecture difficult to implement on a single chip.

3 Hierarchical Structure

Kim et al. [1] further proposed a hierarchical architecture (HC), constructed by dividing the router into multiple sub CIOQ crossbar switches. A k port router, consisting of $(k/p)^2$ $p \times p$ sub switches, each containing a $p \times p$ crossbar along with the input and output buffers. The input buffers of sub crossbar are called row buffers, and the output buffers of sub crossbar are called column buffers because they are located in the rows and columns of the sub crossbar array, respectively. The sub crossbar and the corresponding buffer and port logic are arranged in a block, called tiles, that are interconnected by the row bus and the dedicated wires in column direction. The hierarchical architecture router decomposes the router into input buffers, multiple CIOQ architecture sub crossbars, and global connections connecting these sub crossbars. Hierarchical routers in the architecture also benefit from logically separating control logic and reducing implementation logic complexity. Multi-port Binding Tile-based Router (MBTR) [20] is also a hierarchical architecture. The number of tiles in MBTR is less than the number of ports, and each tile contains multiple port logic and a sub switch, which is the result of a compromise of the implementation complexity, buffer requirements, and router performance.

Partitioned Crossbar Input Queued (PCIQ) architecture [23] groups the outputs, allows the use of simple arbiters and crossbars, and implements an efficient congestion management technique that eliminates HoL blocking. Hierarchical Asymmetric Crossbar (HAC) [23] consisting of asymmetric crossbars with input ports smaller than the output ports, HoL blocking has a very small or negligible impact on crossbar throughput. The HAC architecture forms a $N \times N$ high-radix router architecture through an N/m smaller asymmetric crossbar of $m \times N$, achieving high throughput without a special effort to eliminate HoL blocking. The PCIQ architecture, based on the CIOQ architecture, groups the output ports, increases the number of input buffer read ports to the number of output port groups, removes unnecessary output buffer, and through partition crossbars and congestion management technology, achieves high performance. The HAC architecture groups the input ports, which reduces the row buffer compared to the HC architecture, and through the asymmetric crossbar with naturally low HoL blocking, improves the overall exchange capability. In some ways, both the PCIQ and the HAC architectures are variants of the HC architecture.

Although the hierarchical router architecture simplifies router routing and allocation logic, the number of internal cross wires remains proportional to the square of the number of ports, and each sub crossbar contains input and output buffers, resulting in a large

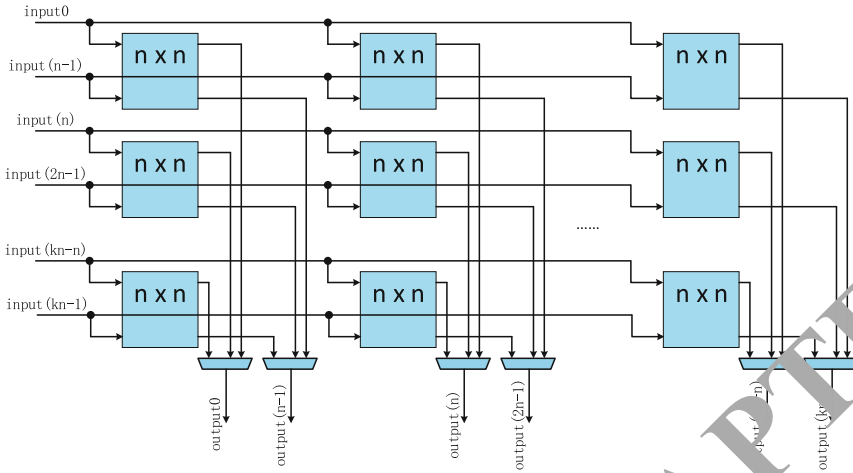


Fig. 3. $kn \times kn$ crossbar decomposed into $k^2 n \times n$ sub crossbars [re-drawn]

number of buffers in the router and occupying a large number of chip area, which reduces reliability and limits scalability.

The hierarchical architecture can construct a $kn \times kn$ crossbar through k^2 smaller crossbars of $n \times n$. Figure 3 is the diagram of $kn \times kn$ crossbars decomposed into $k^2 n \times n$ sub crossbars. The $k^2 n \times n$ sub crossbars are arranged into a matrix of k rows and k columns. Each crossbar of one row receives the same n inputs, and each output of the sub crossbar competes with the output of the $kn \times kn$ crossbars and the corresponding output of the other sub crossbars in the same column. From the logical structure, the $kn \times kn$ hierarchical crossbar composed of $k^2 n \times n$ sub crossbars are no different from a single $kn \times kn$ crossbar, but after adding buffers at the input and output ports of the sub crossbars (Fig. 4), smaller sub crossbars reduce the allocation complexity. The allocation time makes it easy to meet the design requirement, so hierarchical architecture makes it easier to build a high-radix router.

After adding buffers at the input and output ports of the sub crossbars, the input and output of the $kn \times kn$ crossbars and the sub crossbars are isolated. There are three stages from the input to the output: 1) the input of the $kn \times kn$ crossbar to the input buffer of the sub crossbar; 2) the input buffer of the sub crossbar to the output buffer of the sub crossbar; 3) the output buffer of the sub crossbar to the output port of the $kn \times kn$ crossbar. Stage 1): Determine which sub-crossbar's input buffer to use based on the output direction of the input packet. Stage 2) perform the allocation of the $n \times n$ crossbar. Stage 3, the output of the $kn \times kn$ crossbar competes in the same column with the output of the k sub crossbar. Each stage has less selection or arbitration complexity and can be completed within a specified time period.

The number of buffers in the hierarchical architecture is $O(nk^2)$, which is proportional to the square of the number of sub crossbars. To implement the high-radix router, the appropriate size and number of sub crossbars should be selected, so that the router can achieve the optimal performance, power, and area.

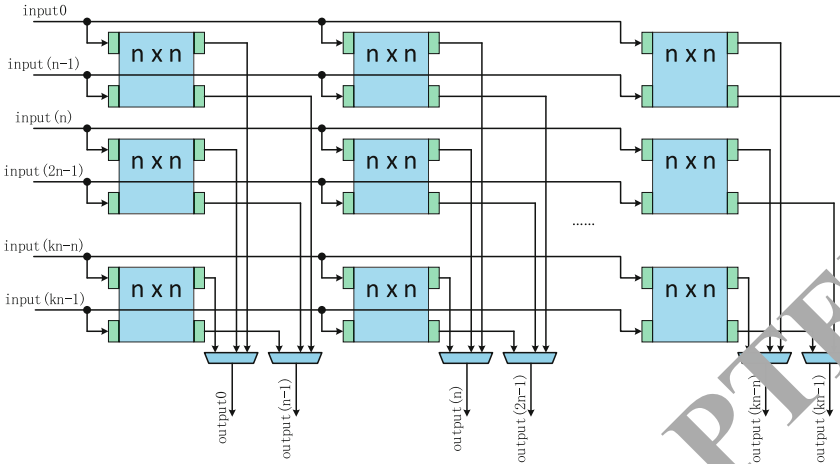


Fig. 4. Adding buffers at the input and output ports [see above]

4 Network Structure

Ahn et al. [2] proposed the idea of designing the router as a network, connecting multiple small switches into a network through on-chip interconnect wires, integrated into a single chip to form a high-radix router with more ports. On chip switches can be connected as folded Clos [17], 2D torus networks, and 2D HyperX [24] topologies. Ahn et al. [2] proposes a bilateral butterfly topology that has fewer sub crossbars and half of the global wires than the folded Clos topology, while achieving lower load latency and equivalent saturation throughput.

Using multistage switches is also a natural way to implement high-radix routers according to the network as a router idea, and three-stage Clos network switches are attractive due to their modularity and scalability. A Three-stage Clos network switch [25] with buffered intermediate stage uses buffers only in the input and output phases, and the middle stage is a buffer-free architecture, known as the MSM (memory-space-memory) Clos network switch [26]. The packets in MSM can reach the output port from the input port through any intermediate stage switch, providing multiple paths, and the packets are order-preserving because there is no buffer in the second stage. The routing of multistage switches prefers buffer less intermediate stage switches to enable routing packets independently without disorder, although this requires matching algorithms to achieve efficient scheduling and routing allocation. Chrysos et al. [27] build a 136-port high-radix router SCOC (Scalable Clos On-Chip) through a three-stage Clos network, with no buffer inside the Clos network. The SCOC uses the large number of connections available on the chip, and the cheap on-chip speedup, to compensate for its inefficient scheduling problems. The combination of packet level multipaths and speedup gives SCOC significant performance. Although it is indistinguishable from crossbars in terms of performance, the SCOC configuration is not a strictly non-blocking architecture, and different end-to-end paths in routing paths may conflict, necessitating internal-provided acceleration to maintain high performance.




Fig. 5. $kn \times kn$ switch built with stage Clos network [see redrawn]

Using the on-chip network constructed by the direct or indirect network as a router architecture, is also a common way to build a high-radix router [2]. On-chip networks can use the Clos network [17], 2D torus, 2D HyperX [24] etc. Figure 5 constructs the router architecture of $kn \times kn$ with three-stage Clos networks. The first and third stages of the Clos network contain $k \times n$ crossbars, with n inputs per crossbar in the first stage, and n outputs for each crossbar in the third stage. The second stage of the Clos network contains $n \times k$ crossbars, and k inputs of each $k \times k$ crossbar connect to the corresponding outputs of k crossbars in the first stage, and the k outputs connect to the corresponding inputs of k crossbars in the third stage. The configuration of the three-stage Clos network is represented by a triple (m, n, r) , where m is the number of intermediate crossbars, n is the number of inputs (outputs) of the first stage (third stage) crossbars, and r represents the number of the first stage (third stage) crossbars. Figure 5 is the (n, n, k) Clos network, and the number of the second stage crossbars is equal to the input of the first stage crossbar, so the architecture is non-interfering [17]. Ahn et al. [21] also introduce high-radix routers built with other balanced networks, and they are compared with hierarchical routers in terms of router area, power consumption, and performance. Although the overall performance of Clos network structure routers is worse than that of hierarchical structure routers, they have a smaller area and lower power consumption.

Chrysos et al. [27] built a 136-port router SCOC through a three-stage Clos network, with the input and output buffers at the first and third stages respectively, and no buffer in the middle stages to make SCOC comparable to crossbar from a performance perspective through reasonable routing scheduling and some speedup. Generally, parallel applications running on HPC systems faces an end-to-end latency issue. Several studies have stated that end-to-end latency can be lowered by using randomly connected inter-switch networks. Kawano et al. [29] proposed the Layout-Oriented Routing with

Entries for Neighbors (LOREN) distributed routing mechanism for irregular networks with limited link length. The experimental results show that the proposed work achieves low system latency.

Ahn et al. [2] use the influence of the traffic characteristics of the system network on the router's internal network design. When the system network topology is folded Clos, the optimization design of the router's internal network is a bilateral butterfly network with fewer crossbars and global wires compared with the Clos network, further reducing the area and power consumption. The topology selection space in network architecture is vast, and theoretically, the system-level used network with the appropriate topological parameters can be integrated on a single chip. Under a certain process level and packaging and assembly conditions, increasing the scale of a single router chip as much as possible can significantly reduce the construction cost of the entire system network.

Camacho and Flich [28] proposed a topology for homogeneous parallel-concentrated-mesh (HPC-Mesh) to improve fault tolerance and save energy. This network uses the injection algorithm to enable connectivity to all the networks with the help of network interface. The topology is adjusted dynamically according to the characteristics of the network, which is the main advantage of this study. This demonstrates how we can improve performance for high traffic rates while decreasing power consumption by just using a portion of the network for lower traffic rates. The suggested network can adapt to the volume of traffic with the aid of the injection algorithm.

Modern research based on topology design concentrates on lowering the diameter of networks. Topologies in the low-diameter network category, such as Jellyfish, significantly reduce power consumption, cost, and latency. But these kinds of networks offer shorter path lengths than other well-known topologies, like torus. Hence, multi-path routing schemes are not used in these network topologies. Besta et al. [30] studied the state-of-the-art protocols and architectures to analyse the minimal and non-minimal paths. This work also discusses the scope of developing high performance routing with multipath in data centres and supercomputers.

Essentially, inside network architecture router is a small network, which needs to consider the routing and scheduling problems, although the three-stage Clos network architecture shown by Fig. 5 is non-interfering, it is not non-blocking. An existing connection in Fig. 5 may block a potential match between a pair of idle inputs and outputs, which is very different from a crossbar or hierarchical architecture, and it is also an important cause of performance loss.

5 Experiments

The evaluation of communication performance for the crossbar, hierarchical structure, and network structure of the router architecture was performed pursuant to numerous simulation experiments conducted with the help of the self-developed computer program. The self-developed computer program is called simwork. Simwork is able to simulate the performance of multiple router structures and network topologies in different traffic with periodic accuracy.

The efficiency of router architectures such as crossbar, network, and hierarchical structures has been evaluated by assessing the communication performance of a single




Fig. 6. Comparison of average channel throughput for crossbar, network and hierarchical structure [self-drawn]




Fig. 7. Comparison of latency for crossbar, network and hierarchical structure by varying the offered traffic [self-drawn]

router by using channel throughput and latency against the offered traffic. Here, the radix of the router is fixed at 64, and each port is connected to a node. Figure 6 shows the relation between the routing architectures, such as crossbars, networks, and hierarchical structures, and the maximal throughput. It is observed that hierarchical structure achieves a balanced and high channel throughput when compared to other alternatives such as crossbars and network structure.

Figure 7 depicts the impact of flow control and router architectures on the communication performance. In fact, the relation between latency and offered load is demonstrated in Fig. 7 when the size of the packet is set to 128 flits. The average delay of the hierarchical structure is considerably low when compared to other structures.

6 Discussion

When performing meticulous data interchange workloads, the performance of a parallel computing system will be severely affected if the network is unable to leverage the growing volume of information produced by the processing elements. When executing latency-sensitive jobs, latency will be just as important as throughput. Puente [31] proposed a complete router architecture for parallel computing systems intended to enhance network throughput while preserving a low node pass time, accordingly satisfying the necessities of the future multiprocessor systems. In this study, Cache-coherent, non-uniform memory access (CC-NUMA) type of multiprocessors is concentrated. This is widely used architecture in the HPC field owing to its decent scalability and simple programming.

Following Moore's law, high performance computing systems have seen their computational power multiplied over the years, lately reaching exascale. Expeditious computing provides substantial challenges for both the design of network topologies and processor system architectures. Several improvements in routing and topology have been made to meet the design goals of high-performance computing. Santhosh [29] made an attempt to investigate the modern topologies and optimize routing methods that make efficient use of networks for the purpose of offering low latency and high throughput. This work makes use of fat-trees to devise effective routing schemes. Experimental evaluations show significant improvements in this approach.

Of late, the development of interconnects to facilitate low-latency and high-bandwidth is crucial in the direction of HPC and data center systems. To resolve this problem, Teh et al. [30] investigated reconfigurable network architectures that can adapt to traffic patterns at runtime using optical circuit switching. The study of how network performance, cost, scalability, and power consumption differ based on optical circuit switch (OCS) placement in the physical topology is presented. In this work, architectures such as ToR-reconfigurable networks (TRNs) and pod-reconfigurable networks (PRNs) are utilized to improve performance.

7 Conclusion

This paper introduces the related research work on router architecture and then analyses several typical architectures of routers, including the traditional crossbar for a small number of ports and the hierarchical architecture and network architecture for high-radix. The crossbar architecture is of the best fairness and is the most basic component of the router, but it is also the most difficult to extend due to the limitation of arbitration time and HoL problems, so the single crossbar is not suitable to extend to a high-radix router with more ports. The hierarchical architecture decomposes the crossbar into a two-dimensional crossbar array architecture, and each sub-crossbar in the array only needs exchange of partial input for partial output, alleviating the HoL problem, and the hierarchical architecture uses a hierarchical scheduling algorithm to reduce the allocation complexity. According to the idea of network design, multiple crossbars are connected with a specific topology, using integrated circuit technology to implement it on a single circuit chip. As long as the internal network is load-balanced, the network

architecture router can achieve full bandwidth. However, the network architecture of routers is generally not strictly non-blocking, and the established connections will affect the later potential input and output matching, even if the input and output are both idle, resulting in latency and a decrease in throughput performance. An experimental result show that Hierarchical structure outperforms well in terms of average channel throughput and latency when compared to the crossbar and network router architecture.

The non-blocking router architecture based on the Clos network will be proposed in the following work, combining the benefits of high-performance hierarchical structure and low-cost network structure. The Clos network replaces the sub crossbar to construct the hierarchical high-radix router and realizes a non-blocking high-radix router with a large number of ports.

References

1. J. Kim, W. Dally, B. Towles and A. K. Gupta, "Microarchitecture of a high radix router," 32nd International Symposium on Computer Architecture (ISCA'05), pp. 420-424, 2005.
2. J. Ahn, S. Choo and J. Kim, "Network within a network approach to create a scalable high-radix router microarchitecture," IEEE International Symposium on High-Performance Comp Architecture, pp. 1-12, 2012.
3. N. McKeown, "The iSLIP scheduling algorithm for input-queued switches," IEEE/ACM Trans. Netw, vol. 7, pp. 188-201, 1999.
4. M. Karol, M. Hluchyj and S. P. Morgan, "Input Versus Output Queueing on a Space-Division Packet Switch," IEEE Trans. Commun, vol. 35, pp. 1347-1356, 1987.
5. M. Chen and N. D. Georganas, "A fast algorithm for multi-channel/port traffic scheduling," in Proc. IEEE Supercom/ICC'94, pp. 96-100, 1994.
6. M. Karol and M. Hluchyj, "Queueing in high-performance packet switching," IEEE J. Select. Areas Commun, vol. 6, pp. 1587-1597, 1988.
7. Y. Tamir, and G. L. Frazier, "High-performance multi-queue buffers for VLSI communications switches," The 15th Annual International Symposium on Computer Architecture, pp. 343-354, 1988.
8. Y. Tamir, and G. L. Frazier, "Dynamically-Allocated Multi-Queue Buffers for VLSI Communication Switches," IEEE Trans. Computers vol. 41, pp.725-737, 1992.
9. N. McKeown, A. Mekittikul, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switches," IEEE Trans. Commun., vol. 47, pp. 1260-1267, 1999.
10. H. Ahmadi, W. L. Denzel, C.A. Murphy and E. Port, "A high-performance switch fabric for integrated circuit and packet switching," Seventh Annual Joint Conference of the IEEE Computer and Communications Societies. Networks: Evolution or Revolution? pp. 9-18, 1988.
11. T. Suzuki, H. Nagano, T. Suzuki, T. Takeuchi and S. Iwasaki, "Output-buffer switch architecture for asynchronous transfer mode," Journal of Communications and Networks, vol. 2, pp. 269-276, 1989.
12. B. Prabhakar, and N. McKeown, "On the speedup required for combined input- and output-queued switching," Autom, vol. 35, pp. 1909-1920, 1999.
13. S.T. Chuang, A. Goel, N. McKeown and B. Prabhakar, "Matching output queueing with a combined input/output-queued switch," IEEE J. Sel. Areas Commun, vol. 17, pp. 1030-1039, 1999.
14. C. Minkenberg, T. Engbersen, "A combined input and output queued packet switched system based on PRIZMA switch on a chip technology," IEEE Communications Magazine, vol. 38, pp. 70-77, 2000.

15. W.J. Dally, "Virtual-Channel Flow Control," IEEE Trans. Parallel and Distributed Systems, vol. 3, no. 2, pp. 194-205, 1992.
16. J. Jun, S. Byun, B. Ahn, S. Y. Nam, and D. Sung. Two-Dimensional Crossbar Matrix Switch Architecture. In Asia Pacific Conf. on Communications, pp. 411–415, 2002.
17. W. Dally, and B. Towles. "Principles and Practices of Interconnection Networks," Morgan Kaufmann Publication, 2004.
18. Y. Li, S. Panwar and H. J. Chao. "On the performance of a dual round-robin switch," in Proceedings of Twentieth Annual Joint Conference of the IEEE Computer and Communications Society, vol. 3, pp. 1688-1697, 2001.
19. M. Y. Teh, Z. Wu, M. Glick, S. Rumley, M. Ghobadi, and K. Bergman, "Performance trade-off in reconfigurable networks for HPC," J. Opt. Commun. Netw. 14, pp. 454-468, 2012.
20. Y. Dai, K. Wang, G. Qu, L. Xiao, D. Dong, and X. Qi, "A Scalable and Resilient Micro-architecture Based on Multiport Binding for High-Radix Router Design," 2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS), pp. 429–438, 2017.
21. M. Santhosh, "Routing Algorithms for the Emerging Topologies in HPC and Data Center Networks", Electronic Theses, Treatises and Dissertations, Florida State University Libraries 2013.
22. G. Mora, J. Flich, J. Duato, P. López, E. Baydal and O. Lyne, "Towards an efficient switch architecture for high-radix switches," 2006 Symposium on Architecture For Networking And Communications Systems, pp. 11–20, 2006.
23. K. Wang, F. Ming, and S-Q. Chen, "Design of a Tile-based High-Radix Switch with High Throughput," vol. 17, pp. 277–285, 2011.
24. J. Ahn, L. Nathan, A. Binkert, M. Davis, M. McLaughlin and R. Schreiber. "HyperX: topology, routing, and packaging of efficient large scale networks." Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, pp. 1–11, 2009.
25. F. Chiussi, G. Kneuer and V. Kumar, "Low cost scalable switching solutions for broadband networking: the ATLANTA architecture and chipset," IEEE Communications Magazine, vol. 35, pp. 44-53, 1997.
26. H. Chao, J. Z. Jing and S. Liew. "Matching algorithms for three-stage buffer less Clos network switches," IEEE Commun. Mag. vol. 41, pp. 46–54, 2003.
27. N. Chrysos, C. Minlenberg, M. Rudquist, C. Basso and B. Vanderpool, "SCOC: High-radix switches made of bufferless clos networks." in Proceedings of 21st International Symposium on High Performance Computer Architecture (HPCA), pp. 402–414, 2015.
28. J. Camacho and J. Flich, "HPC-Mesh: A Homogeneous Parallel Concentrated Mesh for Fault-Tolerance and Energy Savings," 2011 ACM/IEEE Seventh Symposium on Architectures for Networking and Communications Systems, pp. 69–80, 2011.
29. F. Kawanishi, H. Nakahara, I. Fujiwara, H. Matsutani, M. Koibuchi, and H. Amano, "A Layout-Oriented Routing Method for Low-Latency HPC Networks," IEICE Transactions on Information and Systems, vol. E100.D, no. 12, pp. 2796–2807, 2017.
30. M. Besta et al., "High-Performance Routing With Multipathing and Path Diversity in Ethernet and HPC Networks," in IEEE Transactions on Parallel and Distributed Systems, vol. 32, no. 4, pp. 943-959, 2021.
31. V. Puente, J. A. Gregorio, R. Beivide and C. Izu, "On the design of a high-performance adaptive router for CC-NUMA multiprocessors," IEEE Transactions on Parallel and Distributed Systems, vol. 14, no. 5, pp. 487-501, 2003.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

