

An Asynchronous NoC Router Architecture Supporting Quality-of-Service

Youngbok Kim and Yong-Bin Kim
Department of Electrical and Computer Engineering
Northeastern University
Boston, MA 02115
Email: {youngbok, ybk}@ece.neu.edu

Abstract—This paper presents an asynchronous on-chip network router with Quality-of-Service (QoS) support. The router uses a virtual channel architecture with a priority-based scheduler to differentiate between multiple connections with various QoS requirements sharing the same physical channel. A gate-level prototype of the router has been built and its functionality and performance is evaluated. The simulations show that the router is capable of offering a high-level QoS within the capacity limitations of the network.

I. INTRODUCTION

A modern System-on-a-Chip (SoC) design represents a heterogeneous environment with various components interacting in many different ways (event-driven, data streaming, message passing, shared memory, etc.) [1]. Some of these have strict traffic requirements and constraints, and require guaranteed services such as minimum throughput and bounded communication latency. It is therefore essential for an interconnect to provide QoS capabilities in order to accommodate different components in the same network [2].

The work presented in this paper introduces a prototype architecture of an asynchronous on-chip network router with QoS support.

II. NOC ROUTER ARCHITECTURE

Instead of implementing a conventional input buffer organization where each input is associated with a FIFO queue, an input channel is associated with several lanes of small FIFO buffers in parallel (virtual channels). The buffers in each lane can be allocated independently of the buffers in any other lane. A blocked packet holds only a single lane idle and can be passed using any other lanes.

Each QoS connection is assigned to an individual virtual channel and best-effort traffic shares a single virtual channel. The network is therefore able to support $N-1$ QoS connections where N equals the number of virtual channels. When a new packet arrives at the router it is assigned a virtual channel at the appropriate output port according to the information saved in the header of the packet. This assignment persists until the last flit of the packet leaves the network node. If the particular virtual channel is already engaged the packet is blocked until the channel is released. The packet traverses the network following the same procedure at every node on its path until it reaches its destination node. This buffer organization provides

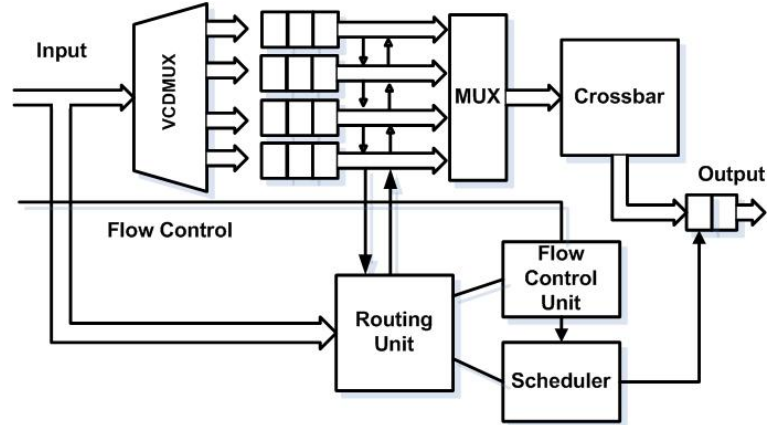


Fig. 1. Top level schematic of the router

means to establish a virtual path from a source node to a destination node and, consequently, to allocate buffer space for a particular connection using the path. As long as the network prevents the rest of the traffic from maliciously using a particular virtual channel, the connection will have the buffer resources available at any given time.

The scheduler allocates bandwidth on a per-flit basis according to the priority level of a virtual channel. This implies that a higher priority packet preempts transmission of the current packet with a lower priority. The flow control ensures that only virtual channels with free buffer space at the receiver compete for the physical channel.

A. Router Architecture

The router presented here has four bidirectional network ports to form a two-dimensional mesh network and a bidirectional service port to enable clients to inject and eject packets to/from the network. A top-level schematic of the implementation is shown in Fig. 1. The router consists of four main components: an input port controller, an output port controller, a switch and a routing unit. For clarity, the figure shows only one input port controller and one output port controller; there are actually 5 instances of each controller implemented in the design.

- *Switching*

Switching is closely related to the internal flow-control of a network and has a great impact on the amount of buffering required in individual network nodes. Buffer space in a NoC(Network On Chip) directly impacts the silicon area overhead of the network and must therefore be kept to a minimum. Wormhole switching [3] reduces the theoretical minimum buffer space to one flit per virtual channel. However in this case the input buffers are able to store up to three flits for two reasons: (1) to decouple the input link from the switch to introduce more concurrency to the design, and (2) to close down the flow-control loop between two neighboring network nodes in order for a single connection to be able to use 100.

- *Packer organization*

Determining the right packet size is crucial to make optimum use of the network resources. The optimum size depends highly on the characteristics of the application. If a message has to be split into too many packets, the overhead of disassembling and reassembling them might be too high. On the other hand, if the packet is too large it might block other traffic affecting the performance of the system. A variable-packet-length organization is proposed in this work in order to improve the flexibility of the network. This way, it can be decided dynamically how to split a message into packets in order to achieve the best performance.

- *Flow-control*

The router employs a credit based flow-control mechanism to prevent data being sent to a buffer. Each virtual channel has a separate credit based counter which is decremented when a request is forwarded to the scheduler. If the counter is zero the request is blocked until new credits are received from the receiving node.

- *Switch*

The router employs a 5x5 multiplexed crossbar switch. Based on a restriction that packets are not allowed to be sent back to the source node, the crossbar is only partly connected to minimize the silicon area. In synchronous network router there is usually a single control unit which schedules packets through the crossbar. The controller has a global knowledge of all inputs, and it is able to optimize the sequence in which packets traverse the crossbar to achieve optimal throughput and prevent contention between the virtual channels sharing the same input port. In asynchronous networks this is rather impractical because it would require synchronization between all of the inputs. Therefore, each output of the switch has a separate controller.

- *Scheduling*

The scheduler uses a low latency asynchronous arbiter with a fixed priority algorithm [4]. The function of the

output buffers in Fig. 3 is to decouple an arbitration cycle from an output transaction cycle. This way the system is capable of pipelined operation performing arbitration for the next flit while transmitting the current flit over a network link. Furthermore, if the arbitration is faster than the output transaction cycle the system can allocate more than 50 times the output bandwidth to a single contender [4].

- *Routing*

A dimensional ordered routing algorithm has been implemented in the design because it offers a deadlock-free and livelock-free operation with deterministic behavior and it is relatively simple to implement in hardware.

III. EVALUATION

The asynchronous NoC architecture has been intensively validated by simulations using traffic generators. This traffic generation through a simple network allowed us to validate all the functionalities of the asynchronous node : send/accept protocol, packet management, and virtual channel arbitration policy. The 5 inputs/5 outputs asynchronous node has been designed. The node cycle time is approximately equivalent to both high and low priority channels. This cycle time brings a network link throughput of about 250 MHz for a 32-bit data flit. This will provide 5 Gbytes/sec total throughput when all inputs/ outputs of the node run concurrently without any packet contention. Regarding the latency, the node latency is about 2ns for high priority virtual channel and 2.5ns for low priority virtual channel.

IV. CONCLUSIONS

A new Asynchronous NoC architecture that provides Quality-of-Service(QoS) is proposed. The router employs virtual channels and a priority based algorithm to differentiate between the high priority traffic and low priority traffic. The presented router is capable of offering a high-level QoS within the capacity limit of the network

REFERENCES

- [1] M. Sgmi, et al., "Addressing the System-on-a-Chip Interconnected Wavelength Communication-Based Design", DAC 2001, pp. 667- 672, June 2001.
- [2] K. Clwssem, et al., "Network on Silicon: Combining Best Effort and Teed Service", In Proc. of DATE, pp. 423-425, Mmcb 2002.
- [3] W.J. Dally and C. L. Seitz, "The Torus Routing Chip", Distributed computing, vol. 1, pp. 187-196, 1986.
- [4] T. Felicijan, et al., "An Asynchronous Low Latency Arbiter for QoS Applications", In Proc of ICM 03, pp 123-126, December 2003.