PREDICTION BASED MINIMALLY BUFFERED DEFLECTION ROUTER INTERCONNECT, IN NETWORK-ON-CHIP

[1]Kolpe Roshni Sadashiv
M.E, E&TC – Signal processing, Dr. D. Y. Patil School of engineering, Ambi, Pune
Email:[1]kolperoshni@gmail.com

Abstract— A conventional Network-on-Chip (NoC) router uses input buffers to store packets to improve performance, but consume significant power. While bufferless NoC design has shown reduction in area and power, and offers similar performance to conventional buffered designs for many workloads, than conventional buffered routers at high network load. This degradation is a significant problem for widespread adoption of bufferless NoCs.

In this work, we propose a new NOC router design called Minimally-Buffered Deflection (MinBD) Router combines deflection routing with a small buffer, which place some network traffic in this small buffer which would have been deflected otherwise. And using the prediction flow control technique for the side buffer. The side buffer would generate status signals, which are sent out to neighboring router switches, would help control the congestion in the network, by predicting the routing flow.

Key Words: Network-on-Chip, Minimally-Buffered Deflection (MinBD) Router, Prediction Buffer

I. INTRODUCTION

A network-on-chip is a component of current and future multi core and many core CMPs (Chip Multiprocessors) [10], and its design can be critical for system performance. As core counts rises, NoCs with designs such as 2Dmesh are expected to become more common to provide adequate performance scaling. Unfortunately, packet-switched NoCs are consuming significant power. In the Intel Terascale 80-core chip, 28% of chip power is consumed by the NoC [7]; for MIT RAW, 36% [35]; for the Intel 48-core SCC, 10% [3]. NoC energy efficiency is thus an important design goal [4], [5].

Bufferless yields simpler and more energy-efficient NoC designs: e.g., CHIPPER [12] reduces average network power by 54.9% in a 64-node system compared to a conventional buffered router. But, at high network traffic, deflection routing reduces performance and efficiency. This is because deflections occur more frequently when many flits contend in the network. Each deflection sends a flit further from its destination, causing unnecessary link and router traversals. Relative to a buffered network, a bufferless network with a high deflection rate wastes energy, and suffers worse congestion, because of these unproductive network hops. In contrast, a buffered router is able to hold flits (or packets) in its input buffers until the required output port is available, incurring no unnecessary hops. Thus, a buffered network can sustain higher performance at peak load. Our goal is to obtain the energy efficiency of the bufferless approach with the high performance of the buffered approach. One prior work, AFC (Adaptive Flow Control), proposes a hybrid design that switches each router between a conventional input-buffered mode and a bufferless deflection mode [8]. However, switching to a conventional buffered design at high load incurs the energy penalty for buffering

every flit: in other words, the efficiency gain over the baseline input-buffered router disappears once load rises past a threshold. AFC also requires the control logic for both forms of routing to be present at each network node, and requires power gating to turn off the input buffers and associated logic at low load. Ideally, a router would contain only a small amount of buffering, and would use this buffer space only for those flits that actually require it, rather than all flits that arrive. We propose minimally-buffered deflection routing (MinBD) as a new NoC router design that combines both bufferless and buffered paradigms in a more fine and efficient way. MinBD uses deflection routing, but also incorporates a small buffer and prediction buffer. The router always operates in a minimally-buffered deflection mode, and can buffer or deflect any given flit. When a flit first arrives, it does not enter a buffer, but travels straight to the routing logic. If two flits contend for the same output, the routing logic chooses one to deflect, as in a bufferless router. However, the router can choose to buffer up to one deflected flit per cycle rather than deflecting it. This fine-grained buffering-deflection hybrid approach significantly reduces deflection rate, and improves performance, as we show. It also incurs only a fraction of the energy cost of a conventional buffered router. Using prediction buffer and status signals, we can calculate or predict the total resource availability of the network. This information helps to avoid excess injection of packets in the network thus, reduces congestion and deflection rate. MinBD provides higher energy efficiency while also providing high performance, compared to a comprehensive set of baseline router designs.

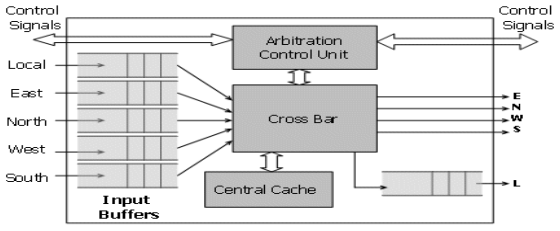

Fig. 1. Typical Router

II. BACKGROUND

On-chip networks most recently-proposed large-scale CMPs (chip multiprocessors) [3, 10, 12]. Such systems are cache-coherent shared memory multiprocessors. Interconnect has served as the substrate for large cache-coherent systems for some time, and the principles are the same in a chip multiprocessor: each core, slice of a shared cache, or memory controller is part of one “node” in the network, and network nodes exchange packets that request and respond with data in order to fulfill memory accesses. CMP NoCs are typically used to implement such a protocol between the cores, caches and memory controllers.

While many on-chip network designs have been proposed, Bufferless deflection routing was first proposed by Baran [2]. An early implementation of bufferless deflection routing for a NoC was BLESS [9], and CHIPPER [6] later provided a more efficient hardware implementation of the deflection routing and packet reassembly.

III. MOTIVATION

Previous NoC designs based on bufferless deflection routing [6], [9] were motivated largely by the observation that many NoCs in CMPs are over provisioned for the common-case network load. In this case, a bufferless network can attain nearly the same application performance while consuming less power, which yields higher energy efficiency.

For low-to-medium network load, a bufferless network has performance close to a conventional buffered network, because the deflection rate is low: thus, most flits take productive network hops on every cycle, just as in the buffered network. In addition, the bufferless router has significantly reduced power (hence improved energy efficiency), because the buffers in a conventional router consume significant power. However, as network load increases, the deflection rate in a bufferless deflection network also rises, because flits contend with each other more frequently. With a higher deflection rate, the dynamic power of a bufferless deflection network rises more quickly with load than dynamic power in an equivalent buffered network, because each deflection incurs some extra work. Hence, bufferless deflection networks lose their energy efficiency advantage at high load. Just as important, the high

deflection rate causes each flit to take a longer path to its destination, and this increased latency reduces the network throughput and system performance.

Overall, neither design obtains both good performance nor good energy efficiency at all loads. If the system usually experiences low-to-medium network load, then the bufferless design provides adequate performance with low power (hence high energy efficiency). But, if we use a conventional buffered design to obtain high performance, then energy efficiency is poor in the low-load case, and even buffer bypassing does not remove this overhead because buffers consume static power regardless of use. Finally, simply switching between these two extremes at a per-router granularity, as previously proposed [8], does not address the fundamental inefficiencies in the bufferless routing mode, but rather, uses input buffers for all incoming flits at a router when load is too high for the bufferless mode (hence retains the energy-inefficiency of buffered operation at high load). We now introduce our minimally-buffered deflection router which combines bufferless and buffered routing in a new way to reduce this overhead.

IV. MINBD: MINIMALLY-BUFFERED DEFLECTION ROUTER

The MinBD (minimally-buffered deflection) router is a new router design that combines bufferless deflection routing with a small buffer, which we call the "side buffer." We start by outlining the key principles we follow to reduce deflection caused inefficiency by using buffering:

1) When a flit would be deflected by a router, it is often better to buffer the flit and arbitrate again in a later cycle. Some buffering can avoid many deflections.

2) However, buffering every flit leads to unnecessary power overhead and buffer requirements, because many flits will be routed productively on the first try. The router should buffer a flit only if necessary.

3) Finally, when a flit arrives at its destination, it should be removed from the network (ejected) quickly, so that it does not continue to contend with other flits.

A. Basic High-Level Operation:

The MinBD router does not use input buffers, unlike conventional buffered routers. Instead, a flit that arrives at the router proceeds directly to the routing and arbitration logic. This logic performs deflection routing, so that when two flits contend for an output port, one of the flits is sent to another output instead. However, unlike a bufferless deflection router, the MinBD router can also buffer up to one flit per cycle in a single FIFO-queue side buffer.

The router examines all flits at the output of the deflection routing logic, and if any are deflected, one of the deflected flits is removed from the router pipeline and buffered (as long as the buffer is not full). From the side buffer, flits are re injected into the network by the router, in the same way that new traffic is injected. Thus, some flits that would have been deflected in a bufferless deflection router are removed from the network temporarily into this side buffer, and given a second chance to arbitrate for a productive router output when re-injected. This reduces the network's deflection rate (hence improves performance and energy efficiency) while buffering only a fraction of traffic.

B. Using a Small Buffer to Reduce Deflections

The key problem addressed by MinBD is deflection inefficiency at high load: in other words, when the network is highly utilized, contention between flits occurs often, and many flits will be deflected. We observe that adding a small buffer to a deflection router can reduce deflection rate, because the router can choose to buffer rather than deflect a flit when its output port is taken by another flit. Then, at a later time when output ports may be available, the buffered flit can re-try arbitration. Thus, to reduce deflection rate, MinBD adds a "side buffer" that buffers only some flits that otherwise would be deflected.

The side-buffer interfaces to the router pipeline in both the inject/eject and permute stages. First, it removes some deflected flits from the pipeline after the permutation network assigns output ports, because the router only knows at this point which flits are deflected. Second, it eventually re-injects the flits that it buffers in the eject/inject stage, using a second instance of the injector that

is placed before the ordinary injector (to give priority network access to buffered flits over new traffic). Finally, a "redirection" block is placed in the pipeline preceding the re-injection block in order to provide livelock-free delivery for buffered flits.

C. Injection and Ejection

A flit must enter and leave the network at some point. To allow traffic to enter (inject) and leave (eject), the MinBD router contains inject and eject blocks in its first pipeline stage. When a set of flits arrive on router inputs, these flits first pass through the ejection logic. This logic examines the destination of each flit, and if a flit is addressed to the local router, it is removed from the router pipeline and sent to the local network node.3 If more than one locally-addressed flit is present, the ejector picks one, according to the same priority scheme used by routing arbitration. Flits from the side buffer are re-injected before new traffic is injected into the network. However, note that there is no guarantee that a free slot will be available for an injection in any given cycle.

D. Prediction based flow control in MinBD

We would further improve upon the MinBD design, using the prediction flow control technique for the side buffer. The side buffer would generate status signals, which are sent out to neighbouring router switches.

We assume that there is an external control mechanism, which accumulates such status signals from neighbouring router switches, which helps in its decision making process of routing appropriate flit/s to the appropriate router switch. This mechanism would help control the congestion in the network, by predicting the routing flow.

V. IP CORE FPGA DESIGN IMPLEMENTATION METHODOLOGY

a. Goal: To develop a synthesizable IP core in HDL, for the input block (receiving packets) of the routing logic.

b. HDL:

i. Design Description using Verilog-HDL

c. Verification Scheme:

i. Functional Simulation using Xilinx ISE Simulator

ii. Test bench for top level design using Verilog-HDL

iii. Stimulus for testing IP Core

iv. Simulation results / Timing diagram Analysis

d. FPGA Implementation:

i. Design to be synthesized using Xilinx ISE Software.

ii. Generate a design bit stream for FPGA device.

iii. Demo the download of design bit stream on FPGA device.

iv. No further verification on hardware / FPGA Kit

v. FPGA kit will NOT be provided

vi. Target Technology: Xilinx Spartan / Virtex Device

VI. CONCLUSION

MinBD combines deflection routing with a small buffer, such that some network traffic that would have been deflected is placed in the buffer instead. By using the buffer for only a fraction of network traffic, MinBD makes more efficient use of a given buffer size than a conventional input-buffered router. Its average network power is also greatly reduced: relative to an input-buffered router, buffer power is much lower, because buffers are smaller. Relative to a bufferless deflection router, dynamic power is lower, because deflection rate is reduced with the small buffer. And using prediction scheme injection rate of the packet also controlled, to improve performance.

REFERENCES

[1] P. Abad et al. Rotary router: an efficient architecture for cmp Interconnection networks. ISCA-34, 2007.

[2] P. Baran. On distributed communications networks. IEEE Trans. on Comm., 1964.

[3] S. Borkar. NoCs: What's the point? NSF Workshop on Emerging Tech. for Interconnects (WETI), Feb. 2012.

[4] S. Borkar. Thousand core chips: a technology perspective. DAC – 44, 2007.

[5] S. Borkar. Future of interconnect fabric: a contrarian view. SLIP ’10, 2010.

[6] C. Fallin, C. Craik, and O. Mutlu. CHIPPER: A low-complexity bufferless deflection router. HPCA-17, 2011.

[7] Y. Hoskote et al. A 5-GHz mesh interconnect for a teraflops processor. IEEE Micro, 2007.

[8] S. A. R. Jafri et al. Adaptive flow control for robust performance and energy. MICRO-43, 2010.

[9] T. Moscibroda and O. Mutlu. A case for bufferless routing in on-chip networks. ISCA-36, 2009.

[10] A. Snavely and D. M. Tullsen. Symbiotic jobscheduling for a simultaneous multithreaded processor. ASPLOS-9, 2000.

[11]M. Taylor, J. Kim, J. Miller, and D. Wentzlaff. The raw Microprocessor: A computational fabric for software circuits and general-purpose programs. IEEE Micro, Mar 2002.

[12]H. Wang, L. Peh, and S. Malik. Power-driven design of router micro architectures in on-chip networks. MICRO-36, 2003.