

EZ-Pass: An Energy & Performance-Efficient Power-gating Router Architecture for Scalable NoCs

Hao Zheng and Ahmed Louri

Abstract—With technology scaling into nanometer regime, static power is becoming the dominant factor in the overall power consumption of Network-on-Chips (NoCs). Static power can be reduced by powering off routers during consecutive idle time through power-gating techniques. However, power-gating techniques suffer from a large wake-up latency to wake up the powered-off routers. Recent research aims to improve the wake-up latency penalty by hiding it through early wake-up techniques. However, these techniques do not exploit the full advantage of power-gating due to the early wake-up. Consequently, they do not achieve significant power savings. In this paper, we propose an architecture called Easy Pass (EZ-Pass) router that remedies the large wake-up latency overheads while providing significant static power savings. The proposed architecture takes advantage of idle resources in the network interface to transmit packets without waking up the router. Additionally, the technique hides the wake-up latency by continuing to provide packet transmission during the wake-up phase. We use full system simulation to evaluate our EZ-Pass router on a 64-core NoC with a mesh topology using PARSEC benchmark suites. Our results show that the proposed router reduces static power by up to 31% and overall network latency by up to 32% as compared to early-wakeup optimized power-gating techniques.

Index Terms—Power-gating, Network-on-Chips, Energy-efficient

1 INTRODUCTION

NETWORK-ON-CHIPS (NoCs) have emerged as the standard communication fabric for connecting cores and memory modules on the chip. Current multi-core chips consist of hundreds of cores and future projections call for thousands of cores. However, today, NoCs consume a large portion (approximately 10%-36%) [1], [2], [3] of the entire chip's power budget. The problem will be further exacerbated by the continuous scaling of transistor feature size. This calls for innovative static power reduction techniques for future NoCs design.

Power-gating [5] is an effective technique that has been used to reduce static power by powering off the idle circuit blocks. The technique has been recently applied to NoC design [4], [6], [7], [8], [9]. There remains challenges to simultaneously maintaining performance (e.g. lower network latency) while reducing static power using power gating. Two of the challenges are (1) how to hide the large wake-up latency penalty and (2) how to extend the sleep time of the powered-off router while still providing adequate communication.

Powerpunch [10] attempts to improve the latency by leveraging the slack time. The slack time is the time that the network interface (NI) requires to packetize a flit. In Powerpunch, an early wake-up signal is sent from an active NI to the powered-off router while the NI is processing the packet. In doing so, the full wake-up latency is hidden. However, as the sleep time of the powered-off router is shortened, this negatively impacts the total power savings.

NoRD [4] provides a bypass ring network to bypass sleepy routers. However, such a technique has limited scalability due to the long latency of the ring topology. In [8], the authors exploit an adaptive routing algorithm to bypass powered-off routers, however, such an algorithm incurs a large latency penalty. In MP3 [11], the authors use a multi-stage intercon-

nection network, namely the Clos topology, to provide energy savings. However, the high-radix nature of the Clos network used is cost prohibitive.

In this paper, we propose an architecture to simultaneously tackle NoCs' energy consumption and performance. The main idea is inspired from the fact that during low traffic, it is more energy efficient to route packets through a simple switching technique rather than through a complex pipelined router. With low traffic, packets are separated in time and will not take advantage of the pipelined router; Consequently the pipelined router will go under-utilized while still consuming power.

The specific contributions of this paper are:

- (1) a low-cost energy-efficient router architecture called EZ-Pass for power-gating.
- (2) a flow control mechanism for both EZ-Pass and conventional routers.
- (3) a modified wake-up control policy for the proposed EZ-Pass router architecture.

2 BACKGROUND AND MOTIVATION

2.1 NoC Routers

Figure 1 illustrates the architecture of a 4-stage, five input ports wormhole NoC router and packet processing logic comprised of virtual channels (VCs) for storing arriving packets, Routing computation (RC) for calculating packet route, Virtual Channel Allocation (VA) for wormhole routing and flow control, and Switch Allocation (SA) for allocating an input port on internal crossbar.

In wormhole routing, a single packet is segmented into a single header flit, several body flits and a single tail flit. The route information of the header flit is read and computed by the control logic (RC, VA and SA) for storing and routing packets. As a result, a flit goes through a pipelined router in 4 stages, namely RC, VA, SA and switch traversing (ST). Moreover, as wormhole routing is a credit-based router, the information of credits is written into the Virtual Channel state tables.

• The authors are with the Department of Electrical and Computer Engineering, George Washington University, Washington, DC, 20052. E-mail: {haozheng,louri}@gwu.edu.
Manuscript submitted: 02-Nov-2017; Manuscript accepted: 13-Nov-2017;
Final manuscript received: 21-Nov-2017

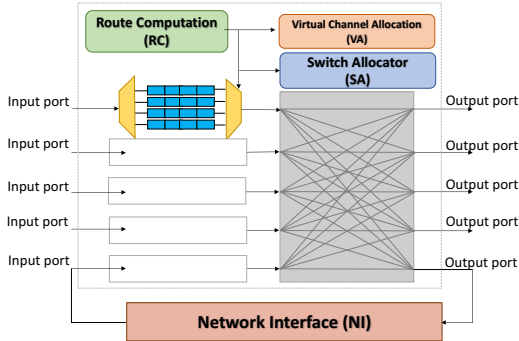


Fig. 1: NoC router architecture

NI provides the connectivity between router and higher-level protocols, and is responsible for encapsulating and sending/receiving flits to/from the network. In order to reduce router latency, prior research has proposed to utilize the NI to implement source and speculative routing [12] where the NI can perform the RC and VA stages instead of the corresponding router.

2.2 Power-gating of NoC Routers

Figure 2(a) depicts the use of power-gating to power off an idle circuit block. In the figure, the circuit block is controlled by a transistor, T1, acting as a switch. When the transistor is off, the circuit block is cut off from the power supply. T1 could be placed between Vdd and circuit block or between circuit block and ground.

Figure 2(b) shows how power gating technique is used for NoC routers. The figure shows two routers power-gated by two transistors and associated control blocks. In what follows, we use the following terminology:

- 1) **Cycles Between Consecutive Flits (CBCF):** This is the number of cycles between two unrelated flits arriving at a given router.
- 2) **Detection Time (DT):** This is the number of consecutive idle cycles detected by the router. This number is used to determine whether to power off the router or not. Prior work [5] has shown that 4 cycles is a reasonable detection time.
- 3) **Breakeven Time (BT):** This is the minimum number of consecutive cycles that a router stays in sleep mode to offset energy penalty caused by turning off the switch transistor T2. According to [6], the breakeven time should be at least 10 cycles.
- 4) **Beneficial Power-Gating (BPG) state:** This state is when the number of cycles between consecutive flits is larger than the breakeven time plus the detection time.
$$CBCF > DT + BT \quad (1)$$
- 5) **Unbeneficial Power-Gating (UPG) state:** This state is when the number of cycles between consecutive flits is smaller than the breakeven time plus the detection time which makes offsets the benefits of power gating.
$$CBCF \leq DT + BT \quad (2)$$

- 6) **Wake-up latency (WL):** The number of cycles required for a sleepy router (e.g. powered-off) to resume full activity (e.g. transitioning from powered-off to active). Prior research [4] has shown that this can be 8 cycles.

The number of cycles that a router stays in the sleep mode impacts the amount of energy savings using power-gating.




Fig. 2: (a) Power-gating technique and (b) its application to on-chip routers [4]

Therefore, the router power-gating is beneficial only when CBCF is at least as long as detection time plus breakeven time or at least 14 cycles.

2.3 Motivation

In order to understand the full benefits of the power-gating for NoCs, we studied application traffic behavior concentrating on PARSEC [13] benchmark suites. We divided the traffic into three categories based on the gaps in cycles between consecutive flits (CBCF factor). The first category, called *high traffic mode*, consists of flits with CBCF of less than 4 cycles. The second category, called *sporadic mode*, consists of flits with CBCF of 4-14 cycles, and the third category, called *low traffic mode*, consists of flits with CBCF of larger than 14 cycles. Since the power-gating detection time is 4 cycles, the high traffic mode is not suitable for powering off the routers. Traffic in this mode fully utilizes the router pipeline as shown as figure 3(a). In the sporadic mode, the powered-off router will be in UPG state due to the fact that CBCF is smaller than DT plus BT. In the low mode, powered-off routers will be in the BPG state due to the fact that CBCF is larger than DT plus BT. In our study, we found that 53% (see figure 4(a)) of overall traffic that qualified for power-gating is sporadic. Using a conventional pipelined router in sporadic mode is not energy-efficient since the pipeline is under-utilized as shown in figure 3(b). Further, as shown in figure 4(b), even the routers in BPG state will be reduced by another 23% due to rescheduling the next power-gating opportunity window (WL+DT).




Fig. 3: (a) Pipelined router stages and (b) Unpipelined router stages




Fig. 4: (a) Sporadic mode traffic over all traffic that qualified for power-gating and (b) fraction of low mode traffic lost due to current power-gating technique

This traffic analysis study has inspired us to introduce a different router architecture for power-gating that mitigates the lack of power-saving in the UPG state. In what follows, we

introduce a new router architecture and a modified flow control and wake-up policies for the proposed scheme.

3 EZ-PASS ARCHITECTURE

3.1 EZ-Pass Router Architecture

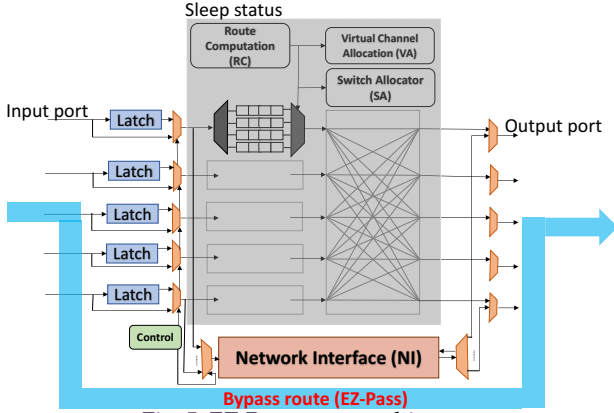


Fig. 5: EZ-Pass router architecture

Figure 5 shows the proposed EZ-Pass router architecture¹. It consists of a conventional router that is used for high traffic mode and EZ-Pass switch for handling sporadic and low traffic modes. This allows incoming flits to be routed without fully waking up the powered-off router. The EZ-Pass switch represents a by-pass [14], [15], [16] route and consists of single-flit latches, multiplexers (MUXs) and demultiplexers (DEMUXs). For example, when the router is powered-off, the incoming flits will be buffered into the single-flit latch. The EZ-Pass control logic routes the flit using a round robin scheme to the NI instead of the conventional router. The NI processes the incoming flit and switches it to the designated output port. The NI also records the VC information to be used later by the flow control policy.

As can be seen, the EZ-Pass route is much simpler and requires less power than the conventional pipelined router path.

3.2 Modified Flow Control

In wormhole routing, a flow control policy is needed to regulate communication between routers. Figure 6(a) shows a conventional pipelined router used for wormhole routing where there is a VC state table associated with each input port. The VC state table [12] contains Read pointer (RP), credits (CR), Output port (OP), Output VC (OVC) and Status. It, however, should be noted that when we use power-gating to power off a given router, the VC state table cannot be accessed which impacts the flow control mechanism.

In order to provide flow control for the proposed EZ-Pass router architecture and have the VC state table information available during power-off state, we modify the VC state information by unifying all VC state tables into a unified table as shown in figure 6(b) and move this information to NI. The unified VC state table is now accessed by both NI and the router. In case the router is powered-off, the NI can still access the unified VC state table for flow control purposes. We, therefore, add two more entries to the unified table namely (1) input port number (Port) and (2) downstream router status (S). The

1. The expression EZ-Pass is an electronic toll collection system in northeastern United States that vehicles can pass the toll station quickly without stopping at the toll booth.

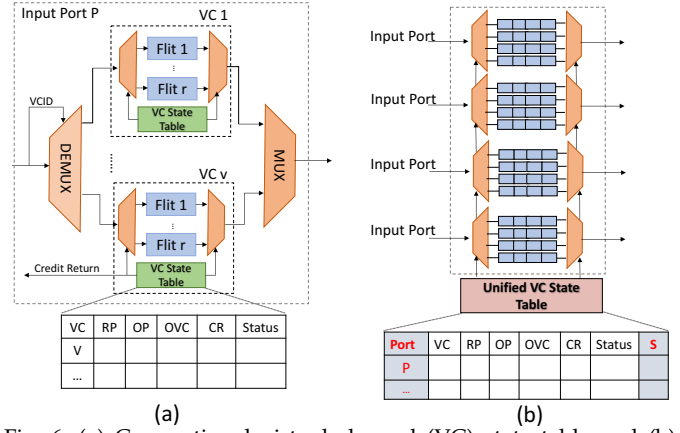


Fig. 6: (a) Conventional virtual channel (VC) state table and (b) unified virtual channel (VC) state table

input port number indicates the input port associated with the incoming flit, and therefore the router and NI can sufficiently identify the routing information. S indicates the power status of a downstream router. The current router can record the credit number (e.g. VC and latches) of its downstream router in the unified table.

3.3 Modified Wake-up Policy

In conventional power-gating, when a flit arrives at a powered-off router, the router is put in a wake-up state for all traffic modes and this incurs a wake-up latency of 8 cycles as stated in section 2.2. In the proposed architecture, and to save more energy, we only wake up the router for high traffic mode as follows:

When an incoming flit arrives while the router is powered off, the flit is passed to the NI for processing through the EZ-Pass route. It takes three cycles to process a flit through an EZ-Pass route (RC,VC,MUX). If during this time, at least three flits have arrived and have been buffered in the latches, we wake up the router and put it into the active state. Otherwise, we continue to process the flits through the EZ-Pass route.

4 EVALUATION

TABLE 1: Key Simulation Parameters

# of cores	64 on-chip, ALPHA, 2GHz
Router	4-stages
Private I/D L1	32KB, 2-way, LRU, 1-cycle latency
Shared L2 per bank	256KB, 16-way,LRU, 6-cycle latency
Cache block size	64 Bytes
Virtual channel	2 VCs/VN, 4-flit/VC
Protocol	MESI
Memory latency	128 cycles
Topology	Mesh

We evaluated the proposed architecture under full system simulation with the combined use of architecture-level and circuit-level simulators. The cycle-accurate gem5 simulator enhanced with GARNET was used for detailed timing simulation of the memory and on-chip network. We also used DSENT for the router area and power estimation using 45nm CMOS process and 0.8V operating voltage. A wake-up latency of 8 cycles is used assuming a 4ns wakeup delay, and we used 4 and 10 cycles for DT and BT, respectively. Table 1 lists the key parameters used in the evaluations. Full system simulation uses an 8×8 64-node mesh.

We analyzed our framework with PARSEC 2.0 benchmark suites. We compared with the following designs: (1) No-PG:

baseline design without power-gating; (2) Conv-PG [5]: conventional power-gating which has a 4-cycle consecutive idle-detection time; (3) Conv-OPT [6]: conventional power-gating with early wake-up to hide a portion of wake-up latency; (4) PowerPunch [10]: completely hide the wakeup latency; (5) EZ-Pass.

4.1 Network Latency Analysis

In figure 7, we plotted network latency. It can be seen that, as compared to CONV_PG, CONV_OPT and PowerPunch, the proposed architecture has a reduction of 52%, 32% and increase of 3% of overall network latency.

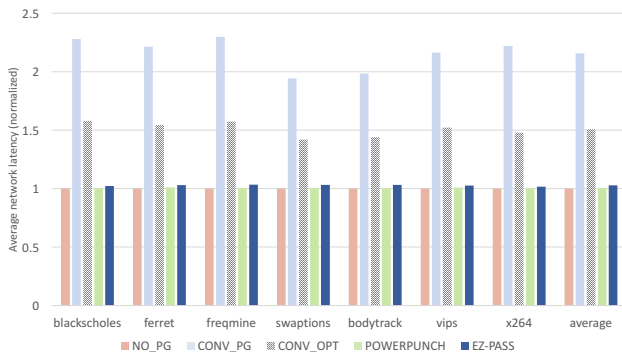


Fig. 7: Average network latency (normalized)

4.2 Power Analysis

In figure 8, we plotted the breakdown of router power. It can be seen that EZ-Pass has better power performance than all designs. Figure 9 shows that EZ-Pass has a static power reduction of 50%, 28% and 31% compared to CONV_PG, CONV_OPT, and PowerPunch, respectively.

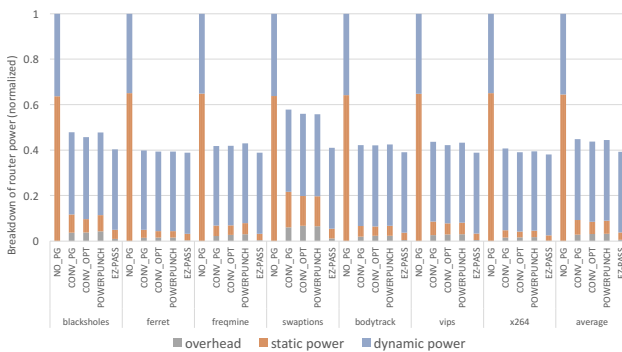


Fig. 8: Breakdown of router power (normalized)




Fig. 9: Router static power (normalized)

4.3 Area Analysis

We used DSENT with 45 nm technology parameters to estimate area overhead. EZ-Pass has a 4% area overhead compared to CONV_PG. Other approaches (CONV_OPT and PowerPunch) have not provided area estimation of their designs to compare against.

5 CONCLUSION

In this paper, we propose an EZ-Pass router to further reduce static power. Unlike previously proposed power-gating NoCs, the proposed architecture provides a simple by-pass routing mechanism to route messages during low traffic without completely waking up the powered-off router. This simple mechanism improves power savings and network latency. Our results show that overall network latency and static power can be reduced by up to 32% and 31% compared with early-wakeup optimized power-gating techniques, respectively. We note that EZ-PASS network latency shows a 3% increase over Powerpunch.

REFERENCES

- [1] Y. Hoskote, S. Vangala, A. Singha, N. Borkar, and S. Borkar, "A 5-ghz mesh interconnect for a teraflops processor," vol. 27, no. 5, pp. 51–61.
- [2] T. Mattson and et al., "The 48-core scc processor: The programmer's view," in *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, 2010, pp. 1–11.
- [3] G. Venkatesh, J. Sampson, N. Goulding, S. Garcia, V. Bryksin, J. Lugo-Martinez, S. Swanson, and M. B. Taylor, "Conservation cores: reducing the energy of mature computations," in *ACM SIGARCH Computer Architecture News*, vol. 38, no. 1. ACM, 2010, pp. 205–218.
- [4] L.Chen and T. Pinkston, "NoRD: Node-router decoupling for effective power-gating of on-chip routers," in *Intl. Symp. on Microarchitecture (MICRO)*, Feb. 2012.
- [5] Z. Hu, A. Buyuktosunoglu, V. Srinivasan, V. Zyuban, H. Jacobson, and P. Bose, "Microarchitectural techniques for power gating of execution units," in *Proceedings of the 2004 international symposium on Low power electronics and design*. ACM, 2004, pp. 32–37.
- [6] H. Matsutani, M. Koibuchi, D. Ikebuchi, K. Usami, H. Nakamura, and H. Amano, "Ultra fine-grained run-time nn/c gating of on-chip routers for cmps," in *Networks-on-Chip (NOCS), 2010 Fourth ACM/IEEE International Symposium on*. IEEE, 2010, pp. 61–68.
- [7] J. Zhan, J. Ouyang, F. Ge, J. Zhao, and Y. Xie, "Dimnoc: A dim silicon approach towards power-efficient on-chip network," in *Design Automation Conference (DAC), 2015 52nd ACM/EDAC/IEEE*. IEEE, 2015, pp. 1–6.
- [8] R. Parikh, R. Das, and V. Bertacco, "Power-aware nocs through routing and topology reconfiguration," in *51st Design Automation Conference (DCA)*, June 2014.
- [9] R. Das, S. Narayanasamy, S. K. Satpathy, and R. G. Dreslinski, "Catnap: Energy proportional multiple network-on-chip," in *Proceedings of ISCA-40*, 2013.
- [10] L. Chen, D. Zhu, M. Pedram, and T. M. Pinkston, "Power punch: Towards non-blocking power-gating of noc routers," in *High Performance Computer Architecture (HPCA), 2015 IEEE 21st International Symposium on*. IEEE, 2015, pp. 378–389.
- [11] L.Chen, L.Zhao, and T. M. Pinkston, "Mp3: Minimizing performance penalty for power-gating of clos network-on-chip," in *Intl. Symp. on High-Performance Computer Architecture (HPCA)*, Feb. 2014.
- [12] W. J. Dally and B. P. Towles, *Principles and practices of interconnection networks*. Elsevier, 2004.
- [13] C. Bienia and K. Li, "Parsec 2.0: A new benchmark suite for chip-multiprocessors."
- [14] A. Kumar, L.-s. Peh, P. Kundu, and N. K. Jha, "Express virtual channels: Towards the ideal interconnection fabric," in *Proceedings of ISCA-34*, 2007.
- [15] T. N. Jain, M. Ramakrishna, P. V. Gratz, A. Sprintson, and G. Choi, "Asynchronous bypass channels for multi-synchronous nocs: A router microarchitecture, topology, and routing algorithm," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 30, no. 11, pp. 1663–1676, 2011.
- [16] L. Xin and C.-s. Choy, "A low-latency noc router with lookahead bypass," in *Circuits and Systems (ISCAS), Proceedings of 2010 IEEE International Symposium on*. IEEE, 2010, pp. 3981–3984.