

# A Novel Approach to Reduce Packet Latency Increase Caused by Power Gating in Network-on-Chip

Peng Wang<sup>\*†</sup>    Sobhan Niknam<sup>\*</sup>    Zhiying Wang<sup>†</sup>    Todor Stefanov<sup>\*</sup>

Leiden Institute of Advanced Computer Science, Leiden University, The Netherlands<sup>\*</sup>  
State Key Laboratory of High Performance Computing, National University of Defense Technology, China<sup>†</sup>  
{p.wang, s.niknam, t.p.stefanov}@liacs.leidenuniv.nl<sup>\*</sup>  
{pengwang, zywang}@nudt.edu.cn<sup>†</sup>

## ABSTRACT

The power gating technique is an effective way to reduce the high static power consumption in a Network-on-Chip (NoC). However, with notable wakeup delay, the power gating technique incurs significant packet latency increase. In this paper, we propose a novel Duty Buffer (DB) structure and an efficient DB-based power gating scheme to overcome this drawback. By keeping minimal number of DB active to replace any sleeping virtual channel in a router, our approach can efficiently reduce the packet latency increase along the whole routing path. Compared with a conventional five-stage pipeline router without power gating, our approach, with only one flit depth of the DB, increases the average packet latency by only 9.67%, which is much less than 57% and 21.75% latency increase in related approaches. With small hardware overhead, our approach can save on average 52.19% of the total power consumption in a NoC, which is comparable with 59.39% and 57.05% power savings in related approaches.

## CCS Concepts

•Hardware → Network on chip;

## Keywords

Network-on-Chip; Power Gating; Low latency

## 1. INTRODUCTION

A Network-on-Chip (NoC) communication infrastructure with its high scalability and low latency is widely applied in many-core systems [6]. However, the NoC accounts for a large part of the total power consumption in many-core systems. For example, the NoCs in Teraflop [10] and Scorpio [9] consume up to 28% and 19% of the total system power consumption, respectively. Due to the low average traffic load of real applications [7], most of the NoC power consumption

is contributed by the static power consumption of idle components. Even, under the minimal resource configuration, a NoC router static power consumption is still about 64% [4] of the total NoC power consumption. As a consequence, the high proportion of the static power consumption in a NoC makes it hard to meet the power budget for future many-core systems [3].

Power gating is a promising technique to reduce the static power consumption by powering off the idle components. However, there is a notable wakeup delay to power on the powered-off components during the wakeup process. When power gating is applied on a NoC, this wakeup delay, about 6-12 clock cycles [4], will interrupt the close cooperation between routers and will block the routing path for a while. As a consequence, the packet latency over the whole routing path dramatically increases and the NoC performance is degraded.

Several works [17, 13] try to reduce the packet latency increase caused by the power gating. As the drowsy SRAM has only 2 clock cycles wakeup delay, Zhan [17] uses the drowsy SRAM to build virtual channels (VCs) in a NoC router and implements a fine-grained power gating on virtual channels. In this way, the wakeup process becomes faster. By sending a wakeup signal ahead of the packet injection, Matsutani [13] switches on the power of the powered-off routers earlier, thereby hiding part of the wakeup delay. These approaches [17, 13] can reduce or hide part of the wakeup delay in a single wakeup process. But a packet may experience multiple wakeup processes and accumulate large delay along the routing path. In fact, as the average traffic load of real applications is low [7], there is high probability for packets to experience multiple wakeup processes. Furthermore, with more cores integrated on a chip in future many-core systems, this cumulative wakeup delay becomes much higher.

In order to further reduce the packet latency increase caused by the power gating, in this paper, we propose a novel and flexible hardware structure, called Duty Buffer (DB). Based on the DB structure, we propose a novel fine-grained DB-based power gating scheme to reduce the static power consumption of the VCs in routers. By using our DB to temporarily replace a powered-off VC and accept packets, an upstream router does not need to block packets while waiting for the VC in a downstream router to completely wake up. Thus, we can efficiently reduce the packet latency increase. Furthermore, as all VCs in the same input port of a router share the same DB, we can keep a minimal number of DBs powered on (on duty) and power off all of the VCs. In

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](http://permissions@acm.org).

NOCS '17, October 19–20, 2017, Seoul, Republic of Korea

© 2017 ACM. ISBN 978-1-4503-4984-0/17/10...\$15.00

DOI: <https://doi.org/10.1145/3130218.3130220>

this way, we use minimal static power consumption to keep a certain transmission ability, which is helpful to reduce the static power consumption. The specific novel contributions of this paper are the following:

- We propose a novel Duty Buffer structure, which can be used to replace any VC in a router. Taking the advantage of our novel DB, we propose a novel DB-based power gating scheme on VCs in a router. This scheme efficiently reduces the packet latency increase caused by the power gating. By keeping a minimal number of DBs powered on, our DB-based scheme also achieves a significant reduction of the static power consumption in a NoC.
- By experiments, we show that our DB-based approach can effectively reduce the packet latency increase caused by the power gating. Taking a conventional router without power gating as the baseline, our DB-based approach, with one flit depth of DB, increases the average packet latency only by 9.67%, which is much less than the 57% latency increase in [13] and 21.75% in [17]. With only 5.76% hardware overhead, our DB-based approach can save on average 52.19% of the total power consumption in a NoC, which is comparable with the 59.39% saving in [13] and 57.05% in [17].

The remainder of the paper is organized as follows: Section 2 gives an overview of the related work. Section 3 gives some background on the packet transmission process in a NoC. Section 4 elaborates on our DB structure, and the DB-based power gating scheme. Section 5 introduces the experimental setup and shows the results. Section 6 concludes this paper.

## 2. RELATED WORK

In this section, we discuss the related works on reducing the packet latency increase caused by power gating in a NoC.

As VCs are the main source of the static power consumption in a router, Zhan [17] applies the power gating on the VCs and uses the drowsy SRAM to build the VC because the wakeup delay of the drowsy SRAM is only two clock cycles, thus the wakeup process is much faster. However, this approach cannot completely remove the wakeup delay and packets accumulate large wakeup delay along the whole routing path. In contrast, our approach keeps a certain transmission ability of routers when the VCs are powered off. The packets are not blocked by the powered-off VCs in the routers. Thus, even though a packet experiences multiple wakeup processes along the routing path, the packet will not accumulate large wakeup delay. Therefore, our approach can reduce more efficiently the packet latency increase.

Based on the principle of the look-ahead routing, Matsutani [13] proposes a runtime power gating approach. By sending a wakeup signal into the next router ahead of the packet injection, this approach hides a few clock cycles in the wakeup process and wakes up the powered-off router earlier. However, the number of hidden clock cycles is determined by the number of the router pipeline stages, and is insufficient to cover the entire wakeup delay. By sending the wakeup signal to the rest of the routers along the routing path, Chen [4] improves the approach in [13]. In this way, the powered-off routers can be waked up much earlier and this approach almost achieves non-blocking power gating in deterministic

routing. However, this approach highly depends on the correct prediction of the routing path. If the routing path of a packet is adaptively changed, this approach may not find the correct routers to power on. Compared with [13, 4], the power gating mechanism in our approach does not need the extra wakeup or sleeping control signal between routers. By keeping a certain transmission ability of routers when the VCs are powered off, routers in our DB-based approach still can transfer packets during the wakeup process, which can be efficiently used to significantly reduce the overall packet latency increase in both deterministic routing and adaptive routing.

Kim in [12] proposed a Flexibuffer scheme to reduce the static power consumption of the VC buffers. Based on the buffer occupancy rate, the router predictively powers on the sleeping buffers in the VCs in advance. In this way, the negative impact of the power gating on the packet latency increase is reduced. However, the scheme in [12] can reduce the static power consumption, only when the VC size is large enough. This is because each VC in this approach has to keep at least  $b_{min} = \max(N_{wakeup}, N_{crt})$  ( $N_{wakeup}$  and  $N_{crt}$  are the number of clock cycles of the wakeup delay and credit round-trip delay, respectively) buffers powered on at run-time, which makes the Flexibuffer scheme not very efficient in reducing the static power consumption. Furthermore, considering the high wakeup delay, 6-12 clock cycles [4], for NoCs with small size of VCs, such as the NoCs in [9, 17, 4, 16, 8], the scheme in [12] cannot save any power consumption. Compared with the scheme in [12], the DB in our approach is shared with all VCs in the same input port, and the minimal DB size is one flit. In this way, our approach can be widely and efficiently used to reduce the static power consumption in NoCs with small and large size of VCs.

## 3. BACKGROUND

In order to understand our approach in Section 4, we provide some background information on the packet transmission process between routers in a NoC.

Figure 1(a) shows a 2-D mesh NoC with a virtual-channeled wormhole router. The router consists of input ports, a routing computation (RC) unit, a virtual channel allocation (VA) unit, a switch allocation (SA) unit, a crossbar, and output ports. A packet transmission needs close cooperation between routers. An input port receives packets from an upstream router or the network interface and stores them into the corresponding input VC. The RC unit determines the output port (transfer direction), which a packet should go through. After that, the VA unit selects an available input VC in the downstream router connected to the determined output port, and assigns it to the packet. The VC address will be put in the head flit of the packets. After the SA unit grants a packet to use the crossbar, the packet can be sent to the downstream router. When the head flit of the packet reaches the downstream router, the input channel decodes the VC address from the head flit and stores the head flit and the following flits into the corresponding input VC.

In order to avoid buffer overflow of input VCs, credit-based flow control is widely used. As shown in Figure 1(a), the output port contains a credit counter. Each credit stands for an unoccupied buffer space in the downstream router. When a flit of the packet leaves the router, the credit counter

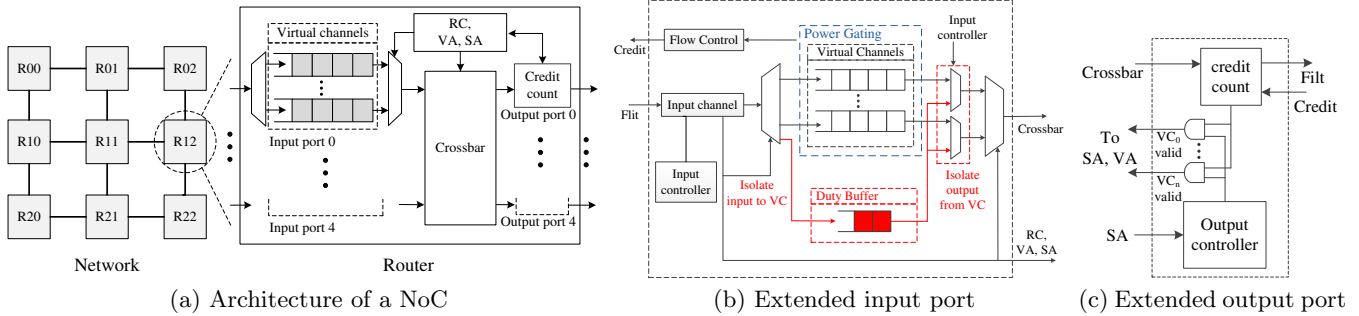


Figure 1: The NoC structure and extended input port and output port.

in the output port of the router decreases the credit to indicate that a buffer space in the downstream router is occupied. At the same time, the router will send a credit to its upstream router to indicate that the input VC in the router releases a buffer space.

The input VCs play a key role in the packet transmission process between routers. A larger number of VCs makes a contribution on eliminating the head of line (HOL) blocking and enhancing the NoC performance. By providing enough buffers, a deep VC also can cover the impact of the credit round-trip delay between routers. However, VCs cause a significant static power consumption, which accounts for over 80% [5] of the static power consumption of routers. So, reducing the static power consumption caused by VCs can efficiently decrease the power consumption of a NoC.

## 4. OUR DB-BASED APPROACH

The key idea of this paper is, by always having a small number of buffers powered on (on duty) in VCs, to keep a certain transmission ability of a router in the wakeup process when power gating is used. In this way, we can reduce the packet latency increase caused by the power gating. In order to achieve this goal, we have to overcome the following challenges.

- **Which VCs should be on duty?** Based on the discussion in Section 3, we know that the input VC address of a packet, which indicates where the packet will be stored, is determined in an upstream router. Therefore, a downstream router does not know when packets will arrive or which VCs will be occupied. As a consequence, it is unknown which VC should be on duty in a downstream router.
- **How to use as few buffers on duty as possible?** Keeping fewer buffers on duty in a VC is helpful to reduce the static power consumption in a NoC. However, the role of the VCs is not the same. Different classes of VCs are used to receive different packets. For example, the MESI [14] coherence protocol needs at least two different data VCs and one control VC to prevent a deadlock. If we want to keep some buffers in VCs on duty to reduce the packet latency increase, each VC class should have buffers on duty and ready to receive packets. As a consequence, we cannot efficiently reduce the static power consumption.

### 4.1 Input Port with Duty Buffer

In order to overcome the challenges, mentioned above, we propose the novel Duty Buffer structure shown in Figure 1(b) to extend input ports of a virtual-channeled worm-hole router. The following components are added to an input port: Input controller, Duty Buffer, and multiplexers. We apply power gating on the VCs as well. The Duty Buffer (DB) is a small buffer queue. The minimal size of the DB is one flit. It is always powered on and ready to receive packets at run-time.

For a downstream router, when a packet reaches an input port, but input VCs are powered-off or waking up, the input controller controls the input channel and demultiplexer to store the packet into the DB. The DB replaces the corresponding powered-off VC, which the packets should go into.

When the router tries to read the packets from the powered-off VC, the input controller controls the corresponding multiplexer to select reading the packets from the DB. By replacing the powered-off VC with DB, the router can transfer the packet as if this VC was powered on.

For an upstream router, as its corresponding downstream router can use the DB to replace the powered-off VCs to store packets, there is no need for the upstream router to block the packet and to wait for the VCs in the downstream router to completely wakeup. Thus, the upstream router still can send a packet to the downstream router, when the VCs in the downstream router are powered-off or waking up. As a result, the packet latency increase caused by the power gating on VCs is reduced.

Since the DB can replace any powered-off VC, we do not need to determine which VC should be on duty. Furthermore, as all VCs in an input port share the same DB, we do not need to keep powered-on buffers for each class of VCs. Thus, we can keep as few buffers on duty in the DB as possible, and power off as many buffers in VCs as possible. In this way, we keep a certain transmission ability with minimal static power consumption.

### 4.2 Power Gating on VCs

In our extended input port, as shown in Figure 1(b), the input controller uses one switch to control the power of all VCs in the port. This is because powering on all VCs at the same time is beneficial to guarantee the NoC performance. That is, the traffic load in many-core systems is bursty [7], so routers tend to use a larger number of VCs in an input port at the same time. On the other hand, the average traffic load is very low [7]. This makes higher percentage of input VCs

idle during the applications execution. So, simultaneously powering off all VCs also can reduce significantly the static power consumption.

The rest of the components in a router, for instance, the RC unit, the VA unit, the SA unit, the crossbar and the output ports, are always powered on. In this way, if VCs in some input ports are powered off, the other input ports still can normally work. Packets will not be blocked in the router pipeline. Furthermore, compared with VCs, the power consumption in the rest of the router's components is much lower [5]. Always keeping them powered on, routers will not waste much static power consumption.

### 4.3 Power Gating Scheme

In order to correctly use our DB to reduce the latency increase during the waking up process, we have to achieve the following goals:

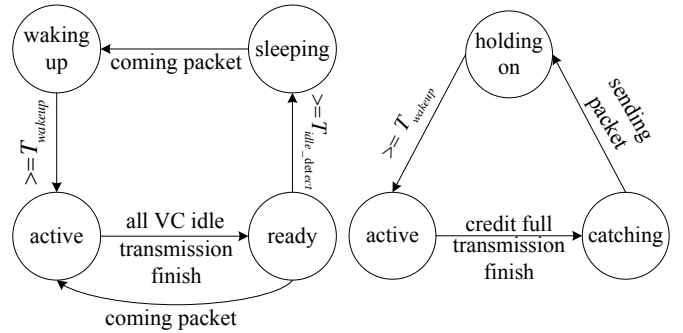
- **Keeping the flits order in a packet.** The flits of a packet may be separately stored in the DB and a VC when the state of VCs switches from waking up state to charging complete state. In order to keep the packet transmission correct, the transmission order of the flits in the same packet must not be changed.
- **No deadlock occurs.** In general, different class packets are allocated with different VCs. It is because that a network interface has to finish processing the reply packet first, then to deal with the new request packet. If a request packet and a replay packet are stored in the same buffer queue (FIFO), the request packet may be stored in the front of the reply packet and stalls the reply packet to access the network interface. As a consequence, the request packet and the replay packet block each other and the deadlock occurs [6]. Therefore, in the wakeup process, we have to guarantee that only packets with the same VC address enter the input port to prevent deadlock occurrence. This VC address should be determined by the head flit of the packet, which wakes up the powered-off VCs in the downstream router.

In order to achieve the above mentioned goals, the input controller and the output controller are added to a router, as shown in Figure 1(b) and Figure 1(c). In the following subsections, we introduce the working mechanisms of the input controller and the output controller.

#### 4.3.1 Input controller

In the extended input port, as shown in Figure 1(b), the input controller monitors the VC control to determine if it can switch off the power of VCs. The states of the input controller are shown in Figure 2(a).

In the **active** state, VCs in the input port are powered on. The input controller controls the input channel to inject packets into the corresponding VCs and keeps the DB idle. When all VCs and the DB are empty and the packet transmission is completed (the tail flit of the packet has left), the input port controller moves to the **ready** state and waits for  $T_{idle\_detect}$  clock cycles to switch off the power of VCs. We use the equation,  $T_{idle\_detect} = T_{credit\_delay} + T_{flit\_delay}$ , to compute the  $T_{idle\_detect}$ , where  $T_{credit\_delay}$  and  $T_{flit\_delay}$  are the credit and flit transmission delay between the neighbor routers, which are determined by design parameters.



(a) Input controller states (b) Output controller states

Figure 2: Controllers for input and output ports.

In the **ready** state, the VCs are still powered on and can be used to store packets. If there are incoming packets in the  $T_{idle\_detect}$  period, the packets will be stored to the corresponding VC and the input controller returns back to the **active** state. The **ready** state is used to avoid unnecessary power gating activities when the idle time of VCs is short.

In the **sleeping** state, the power of all VCs is switched off and VCs cannot be used to store packets. Once the head flit of a packet comes into the input port, it will be stored in the DB and the power of VCs will be switched on to charge the circuit. The input port goes into the **waking up** state.

In the **waking up** state, VCs are not stable and cannot accept packets. The incoming packets still are stored in the DB. After  $T_{wakeup}$  clock cycles, the VC charge is completed, and the state changes to **active** state and the VC can be used to store packets. The input controller stores the packets to the corresponding VC and stops injecting packets into the DB.

At the beginning of the **active** state, if there are any flits left in the DB, the input controller keeps replacing the output of the corresponding VC with the output of the DB until the DB is empty. In this way, the order of the flits in a packet will not be disturbed and the correct transmission is guaranteed.

#### 4.3.2 Output controller

As explained in the beginning of Section 4.3, to guarantee deadlock-free packet transmission during the wakeup process, an output port in a router has to send only packets with the same VC address to the downstream router. In order to achieve this, the output controller and some AND gates are added to the output port, as shown in Figure 1(c). In this extended output port, the output controller monitors the credit counter and the grant to use the crossbar from the SA, to identify the state of the input port in the downstream routers and controls the validation signals for the VA and SA, indicating which input VCs in the downstream router are available. The states of the output controller are shown in Figure 2(b).

In the **active** state, the credit counter is not full, or the packet transmission is not completed (the tail flit has not left). This indicates that the corresponding (the VCs) input port in the downstream router is powered on. So, the upstream router can normally send packets and receive credits.

When the credit counter is full and packet transmissions




Figure 3: The average packet latency.

Table 1: Parameters

topology	2D tours
router pipeline	5 stage
network size	$4 \times 4$
VC count	4 VCs per port
buffer depth	4 flits per VC
packet size	1 or 8 flits, 8B/flit
$T_{credit\_delay}$	1 clock cycles
$T_{flit\_delay}$	1 clock cycles
$T_{breakeven}$	10 clock cycles
$T_{wakeup}$	10 clock cycles

are completed, the output controller moves to the **catching** state. In the **catching** state, the output controller considers that the VCs of the corresponding input port in the downstream router are powered-off, even if, at this time, this input port may be in the **ready** state. The output port allows the VA to normally allocate VCs and SA to grant the use of the crossbar in the **catching** state.

Once a head flit of a packet is granted by the SA to use the crossbar, the output controller marks the allocated VC address and moves to **holding on** state. In the following  $T_{wakeup}$  clock cycles, only packets allocated to this marked VC address are allowed to use the crossbar and transferred to the downstream router. In the **holding on** state, the output controller assumes that all the flits sent are stored in the DB of the input port in the corresponding downstream router. The output controller guarantees that the used credits does not exceed the depth of the DB. In this way, the output controller can guarantee that there is no buffer overflow in the wakeup process. After  $T_{wakeup}$  clock cycles, the output controller moves to the **active** state. The VA and SA can normally allocate router resource for the packets.

## 5. EXPERIMENTAL RESULTS

In order to evaluate our approach in terms of performance and power consumption, we have implemented our approach on the cycle-accurate interconnection network simulator Booksim2.0 [11]. The parameters set in Booksim2.0 are shown in Table 1. Each input port of a router has four 4-flit depth VCs. The total number of buffers in a router is consistent with [8]. Based on the prior work in [8],  $T_{wakeup}$  is set to 10 clock cycles. To calculate the power consumption cost of power gating when the power of VCs is switched on and off, we set the  $T_{breakeven}$ , which is the number of sleep cycles required to compensate the overhead for charging VCs. The  $T_{breakeven}$  is 10 clock cycles and consistent

with the prior work in [17]

For comparison purpose, we have implemented the following schemes in Booksim2.0: (1) NO\_PG is the baseline NoC. It uses a conventional five-stage pipeline router without power gating; (2) LA\_PG is a NoC using the lookahead scheme [13] to hide 5 clock cycles of the wakeup delay; (3) DS\_PG is a NoC where the VCs of a router are implemented by the drowsy SRAM [17]. In the active state, the drowsy SRAM has the same power consumption as the conventional SRAM. When staying in the drowsy state, it only consumes about 10% of the SRAM static power consumption, and has 2 clock cycles wakeup delay; (4) DB\_PG is a NoC using our DB-based approach, presented in Section 4, with different depths of the DB.

### 5.1 Evaluation on Real Application Workloads

In order to evaluate our DB-based approach and compare it with the other approaches mentioned above, on real workloads of applications, we use six applications from the Splash2 benchmark suit: blockscholes, bodytrack, facesim, swaption, fluidanimate, and raytrace. We use Synfull [2] to capture the traffic behaviour of these applications. Synfull generates packets and feeds packets to Booksim 2.0 to evaluate the network performance. Based on the collected data from Booksim2.0, we use Dsent [15] to compute the power consumption under the 45nm technology and NoC frequency of 1GHz.

#### 5.1.1 Effect on performance

Figure 3 shows the average packet latency for the six different realistic application workloads. The seventh set of bars in Figure 3 gives the average result over these six applications.

We use the average packet latency to measure the network performance. The LA\_PG hides 5 cycles of the wakeup delay by sending the wakeup signal in advance. However, compared with the baseline NO\_PG, it still incurs an average of 57% (about 16 cycles) packet latency increase throughout these six applications. The latency in the DS\_PG also increases by 21.75% (about 6 cycles) on average, even though the wakeup delay in the DS\_PG is only 2 clock cycles. These results indicate that packets experience over three wakeup processes on average and the cumulative wakeup delay is the main reason for the performance degradation.

In contrast, the average packet latency in our DB\_PG only increases by 9.67%, 5.67%, and 2.02% with 1, 2, and 3 flits depth of the Duty Buffer, respectively. Compared with the LA\_PG and the DS\_PG, our DB\_PG approach, with only one flit depth of the DB, achieves 47.33% and 12.08% less




Figure 4: The breakdown of the NoC power consumption.

Table 2: Area overhead

	NO_PG	DB_PG_1	DB_PG_2	DB_PG_3
input port ( $\times 5$ )	91914.734088	97289.549462	99265.1354	102494.10415
crossbar+VA+SA	6076.4996	6076.4996	6076.4996	6076.4996
output port ( $\times 5$ )	6595.446	7241.8025	7241.8025	7241.8025
total ( $um^2$ )	104586.679688	110607.851562	112583.437500	115812.406250

packet latency increase, respectively. This is because our DB can replace any sleeping VC in the wakeup process to store packets. The upstream routers do not need to block packets and wait for the VCs to completely waked up. Furthermore, powering on all of the VCs in an input port at the same times can effectively deal with the bursty traffic load of real applications.

The depth of our DB also has an obvious effect on the network performance, as shown in Figure 3. Since we use a five-stage pipelined router, there is a notable pipeline delay and credit round-trip delay in the router. In our DB\_PG with the smallest DB (DB\_PG\_1 in Figure 3), the packet transmission may be intermittently interrupted because of the credit round-trip delay between routers. In addition, as the DB only can replace one VC in a signal wakeup process, packets with different VC address may be blocked in the upstream routers, which also causes packet latency increase.

### 5.1.2 Effect on power consumption

Figure 4 shows the breakdown of the NoC power consumption across the six benchmarks and the seventh set of bars shows the average over these six benchmarks. The network power is decomposed into dynamic and static power consumption.

Compared with the baseline NO\_PG, the LA\_PG, and DS\_PG can reduce an average of 73.14%, and 68.83% of the static power consumption, respectively. The static power consumption in our DB\_PG with 1, 2, and 3 flits depth of the DBs is reduced by 64.11%, 58.49% and 53.63% on average, respectively. For the total power consumption, compared with the baseline NO\_PG, the LA\_PG, and DS\_PG reduce the total NoC power consumption by 59.39%, and 57.05%, respectively. Our DB\_PG with different DB depths reduces with 52.19%, 47.55%, and 45.14% the total power consumption, which is comparable to the LA\_PG and DS\_PG.

It is obvious that we can simply realize a non-block power gating scheme by combining the LA\_PG with the DS\_PG. However, in the sleeping state, the static power consumption of the drowse SRAM increases with the increase of the total number of buffers in a router. As a consequence, when a NoC has a large number of buffers, the drowse SRAMs still cause significant static power consumption in the sleeping

state. While, in our approach, the static power consumption in the sleeping state is only determined by the number of DBs, instead of the total number of buffers, which presents better scalability and is more suitable for a NoC with a larger number of buffers.

### 5.1.3 Effect on area overhead

In order to evaluate the area overhead of our DB\_PG scheme, we use Synopsys Design Compiler to synthesize the routers used in our experiments under 45nm NanGate Open Cell Library [1].

The area overhead of each component is shown in Table 2. Compared with the baseline router used in the NO\_PG, the routers used in our DB\_PG\_1, DB\_PG\_2 and DB\_PG\_3 cause about 5.76%, 7.65% and 10.73% area increases, respectively. Most of area overhead increase in our DB\_PG is contributed by the duty buffers and multiplexers in the input ports, while the area overhead of the input controller is very low. This is because the input controller is made up of simple logic circuits and a small number of registers, which does not cause large area overhead.

As our DB\_PG has no influence on the crossbar, VA and SA, there is no area overhead increase on these components. Compared the area overhead of the output ports in the baseline router, the output ports in our DB\_PG causes 9.8% area overhead increase. However, these area overhead increase takes neglectable percentage of the total router area, because the structure of the output port controller is simple logic circuits and only contains a small number of registers.

## 5.2 Evaluation on Synthetic Workloads

In order to further explore our DB\_PG behaviour under a wider range of packet injection rates, in this section, we evaluate the performance of NO\_PG, LA\_PG, DS\_PG and our DB\_PG under synthetic traffic patterns. Booksim2.0 provides abundant synthetic traffic patterns. We select four synthetic traffic patterns: 1) uniform random: packets' destinations are randomly selected; 2) transpose: packets from source node  $(x, y)$  are sent to destination node  $(y, x)$ ; 3) bit-comp: packets from  $(x, y)$  are sent to  $(N - x, N - y)$ ,  $N$  is the number of nodes in the  $X$  and  $Y$  dimensions of a NoC; 4) tornado: packets from  $(x, y)$  are sent to  $(x + \frac{N}{2} - 1, y)$ .

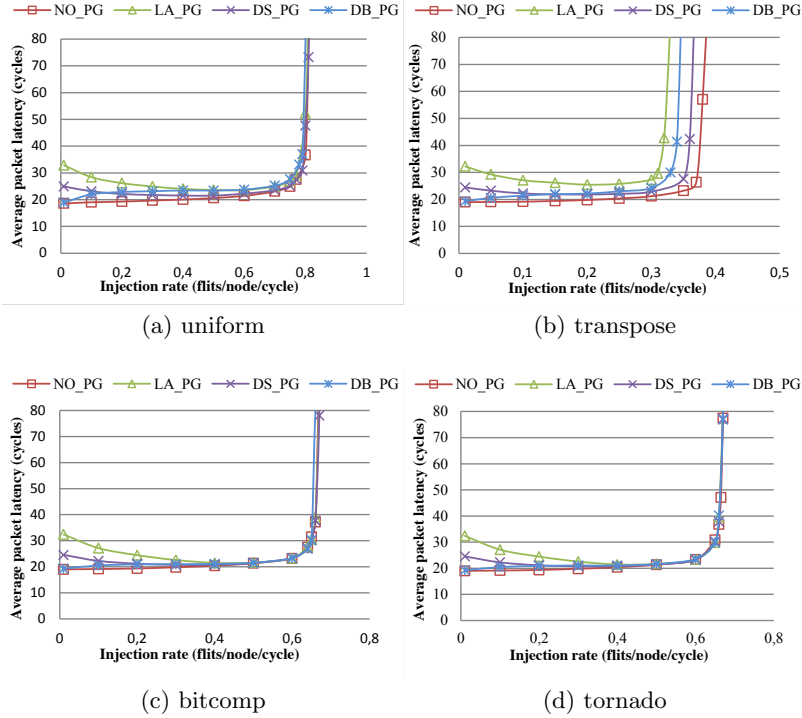


Figure 5: Packet latency across full range of workloads.

Figure 5 shows the average packet latency curves with different injection rates under the different traffic patterns. Since DB\_PG with different depths of the DB has similar trend in terms of the average packet latency curve, in order to clearly show the experimental results, in Figure 5, we only show our DB\_PG with DB of depth 1.

The zero-load latency is an important performance indicator for a NoC. As shown in Figure 5, when the injection rate is almost zero (the injection is about 0.01), where the packet latency is close to the zero-load latency, our DB\_PG has less packet latency than LA\_PG and DS\_PG. This is because, when the injection rate is low, most of the input ports of routers are idle. A packet experiences multiple wakeup processes along the routing path. As a consequence, a packet in LA\_PG and DS\_PG accumulates large wakeup delay. In contrast, by keeping a certain transmission ability in the wakeup process, our proposed DB\_PG can avoid the situation where a packet accumulates too much latency.

With the injection rate increasing from 0 to about 0.2, as shown in Figure 5, our DB\_PG still has less packet latency increase than LA\_PG and DS\_PG. However, the packet latency in our DB\_PG slowly increases, while the packet latency in LA\_PG and DS\_PG decreases. This is because, with the injection rate increasing, more and more input ports of routers are always busy and cannot be powered off. The probability of accumulating a larger wakeup delay becomes lower. On the other hand, multiple packets may compete for a transfer to a downstream router in the same wakeup process, but the DB only can replace one VC in a signal wakeup process in order to avoid deadlock occurrence. As a consequence, in the wakeup process, some of the packets are blocked in our DB\_PG.

When the injection rate further increases and is close

to the saturation injection rate where the packet latency sharply increases, as shown in Figure 5(a)(b), the packet latency in the LA\_PG and DB\_PG is higher than that in DS\_PG, because, in this situation, the wakeup delay in a single wakeup process becomes the main reason for the packet latency increase. The LA\_PG and DB\_PG have higher wakeup delay than DS\_PG. However, because of the DB structure, our DB\_PG outperforms LA\_PG.

Based on the results of the synthetic traffic patterns, when the injection is below 0.2, our DB\_PG can efficiently prevent a packet accumulating large wakeup delay along the routing path and achieves better performance than LA\_PG and DS\_PG. However, when the injection rate is close to the saturation injection rate, in some traffic patterns, our DB\_PG causes a little bit performance loss in terms of the average packet latency. Considering the injection rate of the real application [7] is much lower than 0.2, our DB\_PG can be widely used to reduce the packet latency increase caused by the power gating.

## 6. CONCLUSIONS

In this paper, we proposed a novel Duty Buffer structure and DB-based power gating scheme to reduce the packet latency increase caused by power gating in a NoC. As the DB can replace any VC in the input port, the input port can power off all VCs without the need to consider which VC will be used in the future. In this way, we can keep minimal number of buffers “on duty” to reduce the packet latency increase caused by the power gating, and power off most of the VCs to reduce the static power consumption. The experimental results show that our approach outperforms the lookahead and drowsy SRAM approaches. With a small

amount of additional hardware overhead, our DB-based approach can efficiently reduce the static power consumption, which is comparable with the lookahead and drowsy SRAM approaches.

## 7. ACKNOWLEDGEMENT

This work is partially supported by National Natural Science Foundation of China (No. 61672526, 61572508).

## 8. REFERENCES

- [1] Nangate 45nm library. <http://www.nangate.com/>.
- [2] M. Badr and N. E. Jerger. Synfull: synthetic traffic models capturing cache coherent behaviour. In *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*, pages 109–120. IEEE, 2014.
- [3] S. Borkar. Thousand core chips: a technology perspective. In *Proceedings of the 44th annual Design Automation Conference*, pages 746–749. ACM, 2007.
- [4] L. Chen, D. Zhu, M. Pedram, and T. M. Pinkston. Power punch: Towards non-blocking power-gating of noc routers. In *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, pages 378–389. IEEE, 2015.
- [5] H.-Y. Cheng, J. Zhan, J. Zhao, Y. Xie, J. Sampson, and M. J. Irwin. Core vs. uncore: the heart of darkness. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2015.
- [6] W. J. Dally and B. P. Towles. *Principles and practices of interconnection networks*. Elsevier, 2004.
- [7] R. Das, O. Mutlu, T. Moscibroda, and C. R. Das. Application-aware prioritization mechanisms for on-chip networks. In *Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 280–291. ACM, 2009.
- [8] R. Das, S. Narayanasamy, S. K. Satpathy, and R. G. Dreslinski. Catnap: energy proportional multiple network-on-chip. In *ACM SIGARCH Computer Architecture News*, volume 41, pages 320–331. ACM, 2013.
- [9] B. K. Daya, C.-H. O. Chen, S. Subramanian, W.-C. Kwon, S. Park, T. Krishna, J. Holt, A. P. Chandrakasan, and L.-S. Peh. Scorpio: a 36-core research chip demonstrating snoopy coherence on a scalable mesh noc with in-network ordering. *ACM SIGARCH Computer Architecture News*, 42(3):25–36, 2014.
- [10] Y. Hoskote, S. Vangal, A. Singh, N. Borkar, and S. Borkar. A 5-ghz mesh interconnect for a teraflops processor. *IEEE Micro*, 27(5):51–61, 2007.
- [11] N. Jiang, J. Balfour, D. U. Becker, B. Towles, W. J. Dally, G. Michelogiannakis, and J. Kim. A detailed and flexible cycle-accurate network-on-chip simulator. In *Performance Analysis of Systems and Software (ISPASS), 2013 IEEE International Symposium on*, pages 86–96. IEEE, 2013.
- [12] G. Kim, J. Kim, and S. Yoo. Flexibuffer: Reducing leakage power in on-chip network routers. In *Design Automation Conference (DAC), 2011 48th ACM/EDAC/IEEE*, pages 936–941. IEEE, 2011.
- [13] H. Matsutani, M. Koibuchi, D. Wang, and H. Amano. Run-time power gating of on-chip routers using look-ahead routing. In *Proceedings of the 2008 Asia and South Pacific Design Automation Conference*, pages 55–60. IEEE Computer Society Press, 2008.
- [14] M. S. Papamarcos and J. H. Patel. A low-overhead coherence solution for multiprocessors with private cache memories. *ACM SIGARCH Computer Architecture News*, 12(3):348–354, 1984.
- [15] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic. Dsent-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*, pages 201–210. IEEE, 2012.
- [16] Y. Yao and Z. Lu. Dvfs for nocs in cmps: A thread voting approach. In *High Performance Computer Architecture (HPCA), 2016 IEEE International Symposium on*, pages 309–320. IEEE, 2016.
- [17] J. Zhan, J. Ouyang, F. Ge, J. Zhao, and Y. Xie. Dimnoc: A dim silicon approach towards power-efficient on-chip network. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2015.