



Universiteit  
Leiden  
The Netherlands

## **On the power efficiency, low latency, and quality of service in network-on-chip**

Wang, P.

### **Citation**

Wang, P. (2020, February 12). *On the power efficiency, low latency, and quality of service in network-on-chip*. Retrieved from <https://hdl.handle.net/1887/85165>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/85165>

**Note:** To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/85165> holds various files of this Leiden University dissertation.

**Author:** Wang, P.

**Title:** On the power efficiency, low latency, and quality of service in network-on-chip

**Issue Date:** 2020-02-12

## Chapter 4

# D-bypass Power Gating Approach

**Peng Wang**, Sobhan Niknam, Sheng Ma, Zhiying Wang, Todor Stefanov,  
"A Dynamic Bypass Approach to Realize Power Efficient Network-on-Chip"  
*in Proceedings of the 21st IEEE International Conference on High Performance Computing and Communications (HPCC-2019)*, Zhangjiajie, China, 2019.

---

**T**HIS chapter presents our dynamic bypass (D-bypass) power gating approach, which corresponds to **Contribution 2** introduced in Section 1.4, to further reduce the packet latency increase caused by power gating. This chapter is organized as follows. Section 4.1 highlights the advantages of bypass-based power gating approaches to overcome the drawbacks of power gating, that motivate the research and development of our D-bypass power gating approach. Section 4.2 gives a summary of the main contributions in this chapter. Then, Section 4.3 introduces the Node-Router Decoupling (NoRD) power gating approach which inspires our D-bypass power gating approach. It is followed by Section 4.4, which provides an overview of the related work. Section 4.5 elaborates our D-bypass structure and introduces the D-bypass power gating approach. Section 4.6 introduces the experimental setup and presents experimental results. Finally, a concluding discussion is given in Section 4.7.

### 4.1 Problem Statement

Conventional power gating approaches have two negative impacts on the NoC performance: 1) Wakeup delay, there is a notable wakeup delay (6-12 clock cycles) [CZPP15] before the powered-off routers are fully recharged to the active state. This wakeup delay blocks the packet transmission between routers and causes the packet latency to significantly increase; 2) Break even time (BET), the power gating process causes additional power consumption. Normally, we use breakeven time (BET) to measure

the idle time required to compensate the power overhead due to power gating. This implies that frequent power gating or power gating in a short time may cause more power consumption or inefficient power reduction.

Many approaches try to overcome the aforementioned drawbacks of power gating in different aspects. In order to reduce the negative impact of the wakeup delay, [MKWA08] and [CZPP15] switch on the routers ahead of packet transmission. Part of or the whole wakeup delay can be hidden, but these approaches have to power on the powered-off router every time when there is a packet going through the powered-off router, which may cause frequent power gating and results in more power consumption due to the frequent power gating. On the other hand, in order to avoid non-beneficial power gating caused by BET, many works [MKI<sup>+</sup>10, ZOG<sup>+</sup>15, WNWS17] adopt fine-grained power gating on router components, such as our duty buffer based (DB-based) power gating approach in Chapter 3. Instead of waking up the whole router, these approaches individually wake up part of the router components that are required to transfer packets and keep the rest of the router components powered off. In this way, some of the router components can have longer time to stay powered off. However, these approaches are at the expense of increasing the packet latency, because packets may experience more power gating processes over a routing path. In addition to the above mentioned approaches, bypass-based approaches such as in [CP12, BHW<sup>+</sup>17, ZL18] are more attractive and comprehensive to realize power efficient NoCs. This is because, by bypassing the powered-off routes along a routing path, packets do not need to be blocked and wait for the powered-off routers to be fully charged. Thus, the packet latency increase caused by the power gating is reduced. Furthermore, without frequent interruption of the sleeping state of the powered-off routers, routers have more idle time to stay powered-off and have less power consumption overhead caused by the power gating.

In [CP12], Chen proposes one feasible and applicable bypass-based NoC power gating approach called Node-Router Decoupling (NoRD). By using a bypass latch (in the network interface (NI)) in a downstream router as a transfer station, a packet can be ejected from the NoC to the network interface without the need of storing the packet into a powered-off router buffer. Then the packet can be re-injected (forwarded) to the next router without the need of going through the crossbar in the powered-off router. By repeatedly forwarding packets, the NoRD approach allows packets to go through the powered-off routers in any hop count. Meanwhile, as packets still go through powered-off routers, the conventional credit-based flow control is available to guarantee that there is no buffer overflow. Compared with other bypass-based NoCs [BHW<sup>+</sup>17], this feature greatly simplifies the flow control. However, NoRD does not support bypass in all directions, i.e., in a powered-off router, the bypass latch in a network interface can accept packets from only one specific upstream router and

forward packets to only one specific downstream router. As a consequence, when packets try to bypass the powered-off routers, there is only one available transmission direction and packets are forced to follow detour routing paths, not the shortest routing paths, which results in an inefficient packet transmission and poor scalability.

## 4.2 Contributions

In order to overcome the aforementioned drawback, in this thesis, we propose a dynamic bypass (D-bypass) power gating approach. Based on a reservation mechanism to dynamically reserve a bypass latch in a powered-off router, the same bypass latch can be used by different upstream routers to dynamically build the bypass path. Thus, packets can bypass a powered-off router in any direction, which makes it possible for packets to always follow their shortest routing paths. Furthermore, as the reservation process is executed in parallel (overlaps) with the router pipeline, the timing overhead caused by the reservation process is minimized. The specific novel contributions of this work are summarized as follows:

- We extend the router structure to allow a bypass latch in a powered-off router to accept packets from any upstream router. Then, we propose a reservation mechanism to allow different upstream routes to share the same bypass latch at different times. In this way, the bypass path can be dynamically built based on the routing information of packets. Thus, when packets bypass the powered-off router, they can always follow the shortest routing paths.
- By experiments, we show that our D-bypass power gating approach can effectively reduce the power gating negative impacts on the performance and power consumption. Taking a conventional NoC without power gating as the baseline, our D-bypass power gating approach causes only 2.55% performance penalty, which is less than the 28.67% penalty in [MKWA08], 19.27% in [CP12], 7.24% in [WNWS17], and 5.69% in [ZL18]. Compared with a conventional NoC without power gating on real application workloads, our D-bypass power gating approach reduces on average 77.77% of the total power consumption of the conventional NoC, which is slightly better than 72.94%, 76.11%, 73.55% and 75.30% reductions in [MKWA08], [CP12], [WNWS17], and [ZL18], respectively. However, as a coarse-grained power gating approach, when the packet injection rate increases, most of the routers cannot be powered-off to reduce the power consumption. As a consequence, our D-bypass power gating approach is effective in reducing the power consumption only at low workloads.




Figure 4.1: Node-Router Decoupling.

### 4.3 Background

In order to better understand the contributions of this chapter, in this section, we briefly introduce the bypass-based power gating approach called Node-Router Decoupling (NoRD).

NoRD [CP12] introduces a feasible way to bypass the powered-off routers to transfer packets. As shown in Figure 4.1(b), two bypass paths are added in a router. When the router is powered-off, packets directly go through bypass path *A* in Figure 4.1(b) and are stored in the bypass latch shown in Figure 4.1(c). Then, packets go through bypass path *B* in Figure 4.1(b) to be forwarded to the next router. In this way, packets can go through the powered-off router and be forwarded to the next router. Furthermore, as the packets still go through the powered-off router, the conventional credit-based flow control still works to guarantee that there is no buffer overflow. However, constrained by the router structure, NoRD does not support bypassing of the powered-off router in all directions, i.e., in a powered-off router, each network interface can accept packets from only one specific upstream router and forward packets to only one specific downstream router. As shown in Figure 4.1(a) with the tick red arrow, in NoRD, a bypass ring is statically constructed to achieve full connectivity among routers. To bypass a powered-off router, packets have to go along the static bypass ring path. For example, as shown in Figure 4.1(a), *Router00* tries to send packets to *Router11*,

and its two downstream routers *Router01* and *Router10* are powered-off. *Router00* only can send packets to bypass *Router01*. However, as *Router01* only can forward packets along the bypass ring, packets are transferred to *Router02* in spite of the fact that there is only one hop from *Router01* to *Router11*. Then, after going through *Router02* and *Router12*, the packets reach the destination *Router11*. In this example, as NoRD only can forward packet to a special direction, packets have to be transferred in a detour/longer routing path, which undermines the transmission effectiveness. Furthermore, for a large size NoC, this static bypass ring is quite long, which extremely limits the scalability of NoRD.

## 4.4 Related Work

A few approaches explore a bypass-based power gating NoC. Fly-over [BHW<sup>+</sup>17] switches off the power of an entire router (including output ports) and allows packets to bypass the powered-off routers, but Fly-over supports bypass in the horizontal ( $X + / X -$ ) and vertical ( $Y + / Y -$ ) directions. When a packet needs a router to change its transmission direction ( $X +$  to  $Y - / Y +$ ,  $X -$  to  $Y + / Y -$ ,  $Y +$  to  $X + / X -$ , and  $Y -$  to  $X + / X -$ ), this router must be woken up. Furthermore, as the output ports are powered off and all the credit information is lost, Fly-over has to utilize a complex flow control to recover the credit information when a powered-off router is powered on, which requires significant hardware overhead (a router needs 48 extra links to support this special flow control). Compared with Fly-over, Node-Router Decoupling (NoRD) [CP12] just uses the conventional credit-based flow to control the packet transmission. However, as we have introduced in Section 4.3, NoRD supports only one direction bypass in each powered-off router, which results in an inefficient packet transmission and poor scalability. Our D-bypass power gating approach also adopts the conventional credit-based flow that is similar to NoRD. However, in contrast to Fly-over [BHW<sup>+</sup>17] and NoRD [CP12], our D-bypass power gating approach is based on a reservation mechanism to dynamically build the bypass path, thus packets can bypass the powered-off routers in any direction and in any hop count. Furthermore, the reservation mechanism needs just 10 extra links for each router, which is much less than the 48 extra links in Fly-over [BHW<sup>+</sup>17]. With these aforementioned differences, our D-bypass power gating approach has better scalability than Fly-over [BHW<sup>+</sup>17] and has lower packet latency and less power consumption than NoRD [CP12].

EZ-bypass [ZL18] has similar bypass structure with our D-bypass power gating approach and allows packets to bypass the powered-off router in any direction. In EZ-bypass, each input port of a router needs one bypass latch to temporarily store packets. When a packet bypasses powered-off routers, this packet has to experience the multiple pipeline stages of routers to resolve the contention between packets that

may be in different input ports. However, in our D-bypass power gating approach, there is only one bypass latch in a router. Before using the bypass latch to go through the powered-off router, the upstream routers need to reserve this bypass latch first. In the process of the reservation, the contention between packets is resolved. In this way, when a packet is granted to use this bypass latch to go through the powered-off router, there are no other packets in the downstream powered-off router to contend with it and the router pipeline stages in the downstream powered-off router can be reduced to one stage, and some packet transmissions are accelerated. Furthermore, based on the number of reservation signals from the upstream routers, the powered-off router can detect the contention earlier. Thus, our D-bypass power gating approach can switch on the power of the powered-off router earlier than EZ-bypass.

## 4.5 D-bypass Approach

Fly-over [BHW<sup>+</sup>17] and NoRD [CP12] does not support bypassing in all directions. This limitation is mainly caused by the fact that the bypass latch cannot be shared by all upstream routers to forward packets. Therefore, in our D-bypass power gating approach, we first add one special hardware bypass structure in each router, which allows a bypass latch to accept packets from any of its upstream routers. Then, we propose a reservation mechanism to allow different upstream routers to use the same bypass latch at different times. By reserving the bypass latch at different times, the same bypass latch can be used to dynamically build the bypass paths from any upstream router to any downstream router. Consider the same example as described in Section 4.3, where a packet has to be sent from *Router00* to *Router11* and where *Router01* and *Router10* are powered off. Before packets are sent to the bypass latch in *Router01*, *Router00* reserves the bypass latch in *Router01*. Next the head flit of a packet is sent to the bypass latch in *Router01* and based on the routing information in the head flit, the bypass path is dynamically built from *Router01* to *Router11*, see Figure 4.2(a). Then, *Router01* can forward the packet to *Router11*. In this way, when packets go through the powered-off routers, they can always follow the shortest routing paths to their destinations.

### 4.5.1 Extended Router Structure

In this section, we introduce the extended router structure to support our D-bypass power gating approach. As shown in Figure 4.2(b)(c), and in contrast to NoRD [CP12], we remove the bypass latch from the NI and place it in the router, and put a NI controller (NI ctrlr) in the NI, which is used to reserve the bypass latch. In order to allow packets from all directions to skip the process of being stored in input buffers, thus






Figure 4.2: Extended router structure in D-bypass.

directly being stored in the bypass latch, we add a special hardware bypass structure to connect the input ports ( $X+$ ,  $X-$ ,  $Y+$ ,  $Y-$ , and output Inject of the NI) with the input multiplexer. We also add five multiplexers, one in each output port, and connect the bypass latch to these output multiplexers. Based on the above mentioned extension, without the need of the crossbar, the bypass latch can accept packets from all input directions and forward packets to any of the output directions. All multiplexers are controlled by the ctrlr unit.

When multiple upstream routers or the NI need the bypass latch to forward packets, since there is only one bypass latch, as shown in Figure 4.2(b), the bypass latch cannot simultaneously forward packets coming from multiple upstream routers and the NI. However, it is possible for multiple upstream routers and the NI to share the same bypass latch by using it at different points in time. To achieve such sharing, we have devised a reservation mechanism and its hardware support. As shown in Figure 4.2(b), the handshaking control signals, i.e., the incoming signals ( $IC$ ) and reservation success signals ( $RS$ ), are added between routers. The indexes *up* and *down* in Figure 4.2(b) and Figure 4.2(c) are used to distinguish which router (upstream and downstream) the signals are connected to. The  $IC$  signals are also used in NoRD. In an upstream router, the  $IC$  signal is asserted to inform a downstream router that a packet is coming.




Figure 4.3: Example of the reservation process.

Besides the aforementioned  $IC$  signal functionality in NoRD, the important role of the  $IC$  signal in our D-bypass power gating approach is to reserve the bypass latch in the powered-off router. When an upstream router tries to send packets to a powered-off router, instead of asserting the  $WU_{down}$  signal, it asserts the  $IC_{down}$  signal to reserve the bypass latch in the powered-off downstream router. When the ctrlr unit in the powered-off downstream router detects this  $IC$  signal (for this downstream router, it is  $IC_{up}$ ), the ctrlr unit marks the bypass latch as reserved and does not allow other upstream routers to use it. Meanwhile, the downstream router asserts the  $RS_{up}$  to inform the upstream router that it gets the right to use this bypass latch to forward packets. Once the upstream router receives this  $RS$  signal (for this upstream router, it is  $RS_{down}$ ), it can send packets to that powered-off router. As our D-bypass router can forward packets to any output direction, when the packet is stored in the bypass latch, the ctrlr unit can, based on the routing information in the packet, forward the packet along its shortest routing path. In this way, according to the requirement of the packet transmission, the bypass path in a powered-off router can be dynamically built. When the upstream router finishes the packet transmission, it clears the  $IC_{down}$  signal. Then, the powered-off downstream router releases the reservation of the bypass latch and allows other upstream routers to reserve it.

Based on the aforementioned reservation mechanism, at different times, the bypass latch in a powered-off router can be used by different upstream routers and the bypass path can be dynamically built to forward packets along their shortest routing path.

#### 4.5.2 An Example of the Reservation Process

In order to show the details of our reservation mechanism, we use the example in Figure 4.3 to illustrate the reservation process in our D-bypass power gating approach. We assume a four-stage pipeline router, which consists of route computation (RC), virtual channel allocation (VA), switch allocation (SA), and switch traversal (ST). The link traversal (LT) takes one more clock cycle. *Router A* tries to send packets to *Router B*, but *Router B* is powered-off. The reservation process is shown in Figure 4.3.

In Cycle 0, *Router A* executes the RC stage for a packet and is aware that the

packet should go to *RouterB*. So, *RouterA* asserts the  $IC_{down}$  to reserve the bypass latch in *routerB*.

In Cycle 1, *RouterA* executes the VA stage for packets. Meanwhile, the ctrlr unit in *RouterB* receives the  $IC$  signal (for *RouterB*, it is  $IC_{up}$ ), sets the input multiplexer to select the corresponding input port, marks the bypass latch as reserved, and asserts the corresponding  $RS_{up}$  signal to acknowledge that *RouterA* can forward packets through *RouterB*. If there are multiple  $IC_{up}$  signals simultaneously received to reserve the bypass latch, the ctrlr unit utilizes a round robin arbitration to grant the bypass latch to one of the upstream routers asserted these  $IC$ s.

In Cycle 2, *RouterA* executes the SA stage. As the  $RS$  (for *RouterA*, it is  $RS_{down}$ ) signal has arrived at this moment, *RouterA* gets the right to forward packets to *RouterB*. The head flit of one packet is granted to go to *RouterB*. The rest of the flits are blocked at the SA stage until that *RouterA* receives the credit from *RouterB* or *RouterB* is powered on.

In Cycle 3, in the ST stage of *RouterA*, the head flit of the packet is sent to the crossbar. Then, in Cycle 4, in the LT stage of *RouterA*, the head flit is sent to *RouterB*.

In Cycle 5, *RouterB* stores the head flit in the bypass latch. As no other packets can enter *RouterB*, there is no need to execute the VA, SA, and ST stages, so the pipeline stages are reduced to one stage, i.e., Forward Packet (FP). In the FP stage, according to the routing information in the head flit, the ctrlr unit builds the bypass path for the packet, i.e., the ctrlr unit determines the output port and selects an available VC for the packet, then sets the corresponding output multiplexer to forward the head flit and the rest of flits of the packet to the downstream router of *RouterB* (if *RouterB* is the destination router, the packet will be directly ejected to the NI). In this way, the bypass path can be dynamically built. Furthermore, if there are multiple packets transfers through *RouterB* at different times, different bypass paths can be dynamically built for each packet.

It should be noted that the  $IC_{down}$  signal from *RouterB* to a downstream router of *RouterB* is also asserted in this clock cycle. If the downstream router of *RouterB* is also powered off, the head flit is blocked at the FP stage until *RouterB* gets the  $RS_{down}$  signal from its downstream router. In this way, the packet can bypass multiple powered-off routers. When one flit leaves *RouterB*, one credit is sent to *RouterA*.

In Cycle 6, *RouterA* gets the credit to send another flit. In our example, the packet has two flits, so, the packet transmission is finished in this clock cycle and the  $IC_{down}$  signal is de-asserted.

In Cycle 7, *RouterA* executes the ST stage for the last flit. *RouterB* is aware that the  $IC$  coming from *RouterA* signal (it is  $IC_{up}$  for *RouterB*) is de-asserted and de-asserts the corresponding  $RS_{up}$  signal.

After experiencing the LT stage in Cycle 8, the last flit arrives in *RouterB*. In Cycle 9, the last flit is forwarded to the downstream router of *RouterB*. The ctrlr unit in *RouterB* releases the reservation of the bypass latch and allows other upstream routers to reserve the bypass latch.

Based on the reservation process exemplified above, the bypass latch in the powered-off routers can be used by all upstream routers and the NI to forward packets to any direction at different times. By reserving multiple bypass latches in different routers, packets can bypass multiple powered-off routers along their routing path. Furthermore, as shown in this example, the reservation process is executed in parallel (overlaps) with the router pipeline. Thus, the timing overhead of the reservation process is minimized.

### 4.5.3 Power Gating Conditions

In this section, we introduce the conditions which drive the ctrlr unit in Figure 4.2(b) to control the power supply of a router.

#### Powering off a router

When there is no packet left in a router, and the ICs and WUs signals from all its upstream routers are de-asserted, the router goes into the idle state and the PG signals are asserted to all upstream routers, but at this moment, the power supply is not cut off yet. After waiting  $T_{idle\_detect}$  clock cycles, the ctrlr unit asserts the sleep signal (Figure 4.2(b)) and cuts off the power supply. If there is any IC or WU signals asserted during  $T_{idle\_detect}$ , the ctrlr unit immediately de-asserts the PG signals. By waiting  $T_{idle\_detect}$  clock cycles to cut off the power supply, we can avoid non-beneficial power gating caused by short idle time of routers, which causes frequent power gating and additional power consumption.

#### Powering on a router

To keep good NoC performance, the routers should be powered on at the right moment to deal with high traffic workloads. In our D-bypass power gating approach, we use two metrics to determine when a router should be powered on.

- $N_{IC}$  is the number of  $IC_{up}$ s simultaneously received by a powered-off router. In a powered-off router, when  $N_{IC}$  exceeds a threshold  $th_{IC}$ , the powered-off router is woken up. In this situation, the condition of powering on a router is triggered by the  $IC_{up}$  signals. As an  $IC_{up}$  signal is sent ahead of a packet transmission, part of the wakeup delay is hidden. Furthermore, during the time of

charging the powered-off router, one of the upstream routers can forward packets through the powered-off router. Thus, the packet latency increase caused by the wakeup delay is reduced.

- $N_{IVC}$  is the number of input VCs, in one upstream router, contending for the same downstream router to forward packets.  $N_{IVC}$  indicates the workload of an upstream router. As there is only one bypass latch in a router, our D-bypass power gating approach has significant credit round-trip delay, which blocks a packet transmission to wait for credits. Powering on the downstream routers can reduce this impact. In an upstream router, when  $N_{IVC}$  to a powered-off downstream router exceeds a threshold  $th_{IVC}$ , the corresponding WU signal is asserted to wakeup the downstream router. During the time of waiting the downstream router to fully charge, the upstream router can forward packets through the bypass latch of the downstream router, so the impact of the wakeup delay is also reduced.

It is clear that there is a risk of deadlock when multiple upstream routers need the same powered-off router to transfer packets, but the powered-off router may be continuously occupied by a router and the other routers cannot get a chance to send packets. In order to avoid this deadlock problem, we set the threshold  $th_{IC} = 1$ . On the other hand, in order to avoid performance penalties as much as possible, we aggressively set the threshold  $th_{IVC} = 1$ , which implies that when multiple packets are sent simultaneously to the same powered-off router, the powered-off router should be powered on. The low  $th_{IC}$  and  $th_{IVC}$  may tend to trigger more often the condition of powering on a router, which may cause frequent power gating on a router. However, considering the low average injection rate in real applications, there is still high probability of transferring packets through powered-off routers without frequently triggering the condition for powering on a router.

## 4.6 Experimental Results

In order to evaluate our approach in terms of performance and power consumption, we have implemented our approach using a full-system simulator called Agate [CZPP16]. Agate is based on the widely used full-system simulator GEM5 [BBB<sup>+</sup>11], and Agate supports the simulation of the key items in NoC power gating techniques. The NoC model and power model used in Agate are based on Garnet [AKPJ09] and Dsent [SCK<sup>+</sup>12], respectively. The key parameters used in our experiments are shown in Table 4.1. We choose a four-stage pipeline router. The number of VCs and the buffer size of control VCs and data VCs are set based on the related works [CZPP15] and [CP12]. For simplicity, we use a X-Y deterministic routing algorithm in our D-bypass power

Table 4.1: Parameters.

Network topology	$8 \times 8$ mesh
Router	4-stage pipeline
Virtual channel	2 VCs/VN, 3 VNs
Input buffer size	1-flit/ ctrl VC, 5-flit / data VC
Routing algorithm	X-Y, Adaptive
Link bandwidth	128 bits/cycle
Wakeup delay	8 clock cycles
Break even time	10 clock cycles
Private I/D L1\$	32 KB
Shared L2 per bank	256 KB
Cache block size	16 Bytes
Coherence protocol	Two-level MESI
Memory controllers	4, located one at each corner

gating approach and other related approaches, but for the NoRD approach, we have implemented the special adaptive routing algorithm required by NoRD [CP12] to fairly compare with the NoRD approach. The value of the wakeup delay and break even time (BET) are according to the related works [CZPP15] and [CP12]. As there are additional components added in our D-bypass router and the routers in related approaches, in order to evaluate the power consumption of these components, we use Dsent [SCK<sup>+</sup>12] to estimate the power consumption of the major components, such as the buffers and multiplexers, to make the experimental results more accurate.

For comparison purpose, we have implemented the following power gating approaches: (1) NO\_PG: the baseline NoC without power gating; (2) Conv\_PG: conventional power-gating NoC, which is deeply optimized by sending WU (Look ahead [MKWA08]) and de-asserting PG signals [CZPP16] in advance, thus 6 clock cycles of the wakeup delay are hidden in our experiments; (3) NoRD\_PG [CP12]: the power gating NoC with the NoRD approach; (4) DB\_PG [WNWS17]: our DB-based power gating approach introduced in Chapter 3. In each input port of a router, a one-flit size duty buffer is added to implement the DB-based power gating approach; (5) EZ\_bypass [ZL18]: the power gating NoC with the EZ-bypass approach in which the bypass structure is similar to our approach; (6) D-bypass: the NoC with our D-bypass power gating approach introduced in Section 4.5.

#### 4.6.1 Evaluation on Synthetic Workloads

In order to explore the behavior of our D-bypass power gating approach under a wider range of packet injection rates, in this section, we evaluate the performance of our D-bypass power gating approach under synthetic traffic patterns. We select three syn-




Figure 4.4: Packet latency across different injection rates.

thetic traffic patterns: 1) Uniform random: packets' destinations are randomly selected; 2) Bit-complement: packets from source router  $(x, y)$  are sent to destination router  $(N-x, N-y)$ ,  $N$  is the number of routers in the  $X$  and  $Y$  dimensions of a NoC; 3) Transpose: packets from source router  $(x, y)$  are sent to destination router  $(y, x)$ ;

### Effect on NoC Network Latency

As shown in Figure 4.4(a) and Figure 4.4(b), when the injection rate is around 0.001 packets/node/cycle, our D-bypass has higher average packet latency than DB\_PG and EZ\_PG, but lower than Conv\_PG and NoRD\_PG. This is because in our D-bypass approach, multiple packets cannot simultaneously bypass the same powered-off routers at the same time, and some packets are blocked due to power gating. However, compared with Conv\_PG, there are significant number of packets that can bypass the powered-off routers. On the other hand, when the packet bypasses the powered-off router, the powered-off router pipeline stages are reduced to one stage and some packets' transmissions can be accelerated. Thus, in Figure 4.4(c), our D-bypass has the lowest packet latency among all the approaches.

With the injection rate increasing up to the saturation injection rate (around 0.13 packets/node/cycle in uniform random, 0.07 packets/node/cycle in bit-complement, 0.05 packets/node/cycle in transpose), the curve of the average packet latency in our D-bypass approach slowly drops, and it is lower than the curve of Conv\_PG and NoRD\_PG, and gradually gets close to the curve of NO\_PG. This indicates that our D-bypass approach can more efficiently deal with high bursty traffic workloads than Conv\_PG and NoRD\_PG, which meets requirements of real applications where traffic workloads are bursty.

The saturation injection rate is also an important parameter to evaluate the NoC performance. A NoC with higher saturation injection rate can achieve higher throughput. As shown in Figure 4.4, our D-bypass approach has the same saturation injection rate as the baseline NO\_PG, but NoRD\_PG and DB\_PG have lower saturation in-




Figure 4.5: Power consumption across different injection rates.

jection rate. This is because, at the saturation injection rate, all routers are powered on and our D-bypass approach works the same as NO\_PG. However, the routers in NoRD\_PG are not as efficient as the routers in NO\_PG. This is because NoRD\_PG needs VCs to support its special adaptive routing along the bypass ring. As a consequence, NoRD\_PG cannot fully utilize VCs to achieve the same saturation injection rate as NO\_PG. Therefore, compared with the bypass-based power gating scheme NoRD\_PG, our D-bypass approach can achieve higher throughput.

### Effect on NoC Power Consumption

As shown in Figure 4.5, when the packet injection rate is 0.001 packets/node/cycle, our D-bypass approach has the lowest power consumption. This is because, at such low injection rate, our D-bypass approach can transfer packets through the powered-off routers without the need of powering them on. Thus, our D-bypass approach can reduce more the power consumption compared to Conv\_PG. Furthermore, compared with DB\_PG and EZ-bypass, we need less hardware to implement our D-bypass approach. It means that our D-bypass approach causes less extra power consumption. Thus, when most of the routers are powered-off in a NoC, our D-bypass approach consumes less power than DB\_PG and EZ-bypass. In addition, compared with NoRD\_PG, our D-bypass transfers packets through the powered-off routers along the shortest routing path, which is more efficient in transferring packets and helpful to reduce the power consumption.

However, when the injection rate increases, the power consumption in our D-bypass approach increases and reaches the power consumption of NO\_PG, which is the same for Conv\_PG, NoRD\_PG, and EZ\_PG. This is a common drawback for all coarse-grained power gating approaches, because power gating is applied on the granularity of a router. When the injection rate increases, more routers become busy and cannot be powered off. As a consequence, compared with DB\_PG, which is a fine-grained power gating approach and is effective in reducing the power consumption






Figure 4.6: Execution time.

under a wider range of packet injection rates, our D-bypass approach can efficiently reduce the power consumption only at low packet injection rates.

## 4.6.2 Evaluation on Real Application Workloads

In this section, we use real application workloads to compare the approaches in terms of the application performance, the NoC average packet latency, and the NoC power consumption. To do so, we use nine applications from the Parsec [BKSL08] benchmark suite.

### Effect on Application Performance

Figure 4.6 shows the execution time of the nine applications, which is normalized to the baseline NO\_PG, and the tenth set of bars in Figure 4.6 gives the average results over these nine applications. Our D-bypass approach causes less performance penalty (execution time increase) than the related approaches. Compared with the baseline NO\_PG, our D-bypass causes an average of 2.55% performance penalty, which is less than the 28.67% performance penalty in Conv\_PG, 19.27% in NoRD\_PG, 7.24% in DB\_PG, and 5.69% in EZ\_bypass. In the ferret benchmark, our D-bypass has its largest performance penalty of 6.03%, and Conv\_PG, NoRD\_PG, DB\_PG, and EZ\_bypass have also their largest performance penalty of 47.39%, 37.18%, 21.22%, and 19.51%, respectively.




Figure 4.7: Average packet latency.

### Effect on NoC Network Latency

Figure 4.7 shows the average network latency across the nine applications. Our D-bypass approach can efficiently reduce the network latency increase caused by power gating. Compared with NO\_PG across the applications, the average network latency in our D-bypass approach slightly increases, but is much lower than Conv\_PG and NoRD\_PG. This is because our D-bypass approach can dynamically build the bypass path and allow packets to bypass the powered-off router in all directions. Thus, packets can go along the shortest routing paths to bypass the powered-off routers, and are not blocked due to the power gating processes.

In most of the applications, our D-bypass approach has slightly lower average network latency than DB\_PG and EZ\_bypass. This is because DB\_PG is a fine-grained power gating approach and causes more power gating processes. Compared with EZ\_bypass, our D\_bypass is based on a reservation mechanism which can power on the powered-off router earlier when multiple upstream routers need the same powered-off router to forwards packets. However, in the benchmarks ferret, fluidanimate, swaptions, and x264, our D-bypass approach has slightly higher average network latency than EZ\_bypass, because each input port in EZ\_bypass has a bypass latch to hold one flit of a packet, whereas in our D-bypass approach, all input ports in a router have to share one bypass latch to forward packets, which may result in more contention and blocking of some packet transmissions. However, in our D-bypass, as only one packet is allowed to go through a powered-off router at a time, the router pipeline stages can be reduced to one stage when packets bypass the powered-off routers. Thus, some packet transmissions are accelerated and our D-bypass approach has lower application execution time than EZ\_bypass in ferret and swaptions, in spite of the fact that




Figure 4.8: Breakdown of the NoC power consumption.

our D-bypass approach has slightly higher average packet latency than EZ\_bypass.

### Effect on NoC Power Consumption

Figure 4.8 shows the breakdown of the NoC power consumption across the nine applications and the tenth set of bars shows the average over these nine applications. The NoC power consumption is broken down into three parts: the extra power consumption caused by the power gating (PG\_overhead) and the dynamic/static power consumption of routers (dynamic/static).

As can be seen in Figure 4.8, our D-bypass approach reduces slightly more the power consumption than the related approaches. Compared with NO\_PG, our D-bypass reduces on average 77.77% of the total NoC power consumption, which is slightly better than 72.94% in Conv\_PG, 76.11% in NoRD\_PG, 73.55% in DB\_PG, and 75.30% in EZ\_bypass. This is because, for real application workloads, the traffic is busy for very short periods of time, thus the average packet injection rate is low for a long period of time. Therefore, all of these power gating approaches can power off routers for a long time to reduce the static power consumption. In addition, our D-bypass approach can transfer packets through the powered-off routers without waking them up. Thus, our D-bypass can power off the routers for even longer time and it can reduce more the router static power consumption and PG\_overhead compared to Conv\_PG. Even though NoRD\_PG is also a bypass-based power gating approach, it does not support bypass in all directions and forces packets to go along the bypass ring. Packets have to go through more routers, which may cause more power gating processes. As a consequence, NoRD\_PG consumes slightly more router static power and PG\_overhead than our D\_bypass. Furthermore, in order to transfer packets through the powered-off routers or the powered-off input ports, our D\_bypass, EZ-bypass, and DB\_PG need to always keep some components powered on, that always consume static power. However, compared with DB\_PG and EZ-bypass, our D\_bypass needs

to keep fewer components always powered-on. Therefore, our D\_bypass is more efficient to reduce the static power consumption of the routers.

## 4.7 Discussion

In this chapter, we propose a dynamic bypass (D-bypass) power gating approach to allow packets to bypass powered-off routers in any hop count and in any direction. Based on a reservation mechanism, all the upstream routers can share the same bypass latch to dynamically build the bypass path for different packets. In this way, packets can be transferred along their shortest routing paths. With small hardware overhead, our D-bypass approach can efficiently reduce the power consumption and has less performance penalty.

Even though our D-bypass power gating approach allows packets to bypass the powered-off routers in any direction, the efficiency of the bypass path is limited by the single bypass latch in a router. The packets maybe frequently blocked to wait for the free bypass latch. As a result, in some applications, there is still significant packet latency increase in our D-bypass power gating approach. Furthermore, like most of the course-grained power gating approaches, our D-bypass power gating approach cannot fully utilize the idle time of each component in a router. When the traffic workload is high, most of the routers in a NoC become busy and cannot be powered off to reduce the static power consumption. As a consequence, our D-bypass power gating approach is effective in reducing the power consumption only at low traffic workloads.