# Hybrid, Asymmetric and Reconfigurable Input Unit Designs for Energy-Efficient On-Chip Networks\*

Xiaoman LIU<sup>†</sup>, Yujie GAO<sup>†a)</sup>, Nonmembers, Yuan HE<sup>††,†b)</sup>, Member, Xiaohan YUE<sup>†</sup>, Haiyan JIANG<sup>†</sup>, and Xibo WANG<sup>†</sup>, Nonmembers

SUMMARY The complexity and scale of Networks-on-Chip (NoCs) are growing as more processing elements and memory devices are implemented on chips. However, under strict power budgets, it is also critical to lower the power consumption of NoCs for the sake of energy efficiency. In this paper, we therefore present three novel input unit designs for on-chip routers attempting to shrink their power consumption while still conserving the network performance. The key idea behind our designs is to organize buffers in the input units with characteristics of the network traffic in mind; as in our observations, only a small portion of the network traffic are long packets (composed of multiple flits), which means, it is fair to implement hybrid, asymmetric and reconfigurable buffers so that they are mainly targeting at short packets (only having a single flit), hence the smaller power consumption and area overhead. Evaluations show that our hybrid, asymmetric and reconfigurable input unit designs can achieve an average reduction of energy consumption per flit by 45%, 52.3% and 56.2% under 93.6% (for hybrid designs) and 66.3% (for asymmetric and reconfigurable designs) of the original router area, respectively. Meanwhile, we only observe minor degradation in network latency (ranging from 18.4% to 1.5%, on average) with our proposals.

key words: network-on-chip, router, input unit, network traffic, energy efficiency

## 1. Introduction

On-going process scaling has driven the rapid adoption of networks-on-chip (NoCs), which have also emerged to a critical part of the memory hierarchy in realizing scalable and efficient on-chip communications [1], [2]. Meanwhile, with increasing numbers of on-chip devices connected through NoCs, more routers or higher-radix routers are favored and they have higher demands for memory devices to implement input units in the routers, which makes the already tight power budget more of a concern. In practice, buffers (organized as virtual channels in input units) are the primary power consumers in NoCs, where they draw most of the network power (around two thirds) and this remains a critical issue for NoCs as technology continues to

Manuscript received November 11, 2022.

Manuscript revised February 26, 2023.

Manuscript publicized April 10, 2023.

<sup>†</sup>The authors are with Shenyang University of Technology, Liaoning, China.

<sup>††</sup>The author is with Keio University, Yokohama-shi, 223–8522 Japan.

\*This paper is extended from Y. Gao *et al.*, "Traffic-Aware Energy-Efficient Hybrid Input Buffer Design for On-Chip Routers," Proc. 15th IEEE International Symposium on Embedded Multicore/Many-core Systems-on-Chip, pp.375–381, Dec. 2022.

a) E-mail: ygao@smail.sut.edu.cn

scale [3], [4].

So far, many attempts to the power problem exist for NoCs. There are works focusing on cutting the power of routers with power gating [5]-[8] or dynamic voltage and frequency scaling [9], proportionally supplying power to the network based on the actual demand [10], completely eliminating routers (hence also buffers) through smart wiring [11] or improving the energy efficiency of networks through prediction [12], multicasting [13], traffic compression [14], pipeline bypassing [15] and hybrid flow control mechanisms [16]-[18]. When looking at the input units particularly, there are existing works having shared buffer design [19] or focusing on implementing energy-/area-efficient input units with other memory devices than the baseline static random-access memory (SRAM), such as planar embedded DRAM (eDRAM) [20] or spintronic domain-wall memory (DWM) [21]. Moreover, STT-MRAM (short for Spin Transfer Torque Magnetoresistive Random Access Memory) has also been used in NoC routers [22], [23]. The former work takes the advantage of STT-MRAM for its density (roughly four times the density of SRAM) to enlarge the buffer capacity and network throughput while hiding the long write latency and avoiding the high write energy of STT-MRAM devices through innovative control over buffer elements [22]. The latter proposes two hybrid input unit designs (hierarchical or banked) to mix SRAM and STT-MRAM for both energy- and area-efficiencies [23]. Despite being different from our purpose [22] or methods [23], these two works reveal an important fact, that is, for input units in NoC routers, direct replacements of SRAM with slower memory devices, such as STT-MRAM, is not an effective option. As a consequence, the bottom-line is, network latency should not be severely degraded while power consumption can still be suppressed [24].

Therefore, in this work, we present three novel input unit designs to help reduce the power consumption of NoC routers while conserving their performance. Firstly, with both SRAM and STT-MRAM devices, input units are redesigned so that they are operated under an important characteristic of on-chip traffic, the low frequency of long packets since most of the packets transmitted in NoCs are control messages (short packets without data payload). At the same time, it is also very rare to have multiple concurrent data messages (long packets with data payload) in one input unit at a time. This means the first buffer element in virtual channels (VCs) should always be kept to SRAM devices (low in

b) E-mail: isaacyhe@acm.org

DOI: 10.1587/transele.2022CTP0005

latency but leaky) while most of the other buffer elements in VCs may be replaced with slow but less leaky devices, such as STT-MRAM. Following this hybrid design, static power of the input unit can be dramatically reduced with buffer elements implemented in STT-MRAM while most packets are still stored in SRAM so that their latency and dynamic power are barely affected by STT-MRAM.

Furthermore, under the same characteristics of on-chip traffic, it is only necessary to keep one complete VC for long packets while all other VCs can have most of the buffer elements (three fourths) removed but still working for short packets. Such asymmetric design can dramatically reduce the static and dynamic power of the input unit while not severely degrading the network performance. In consequence, when encountering long packets more frequently, a reconfigurable input unit design can be used where VCs dedicated for short packets can work together to form as one VC for long packets. The contributions of this work are summarized as follows:

- We have identified an important characteristic of traffic in modern NoCs, that is, a large portion of the network traffic are short packets.
- Based on this observation, we have implemented three novel input unit designs for on-chip routers which improve the energy efficiency dramatically, as in our evaluations.
- This simply tells us a fact that when characteristics of the network traffic is correctly considered, power consumption and transistor footprint of input unit designs can be easily reduced under small performance degradation.

The rest of this paper is organized as follows: Sect. 2 briefly reviews the baseline input unit design of NoC routers and motivates our work. In Sect. 3, our proposals and their qualitative discussions are introduced. Afterwards, the evaluation methodology and results are detailed in Sect. 4 and Sect. 5, respectively. Finally, Sect. 6 briefly discusses the related works while Sect. 7 concludes this work.

#### 2. Backgrounds and Motivations

In this section, we will cover some preliminaries of this work. These include the baseline input unit design of modern on-chip routers and how we are motivated by some important characteristics of the network traffic.

2.1 The Baseline Design of Input Units in On-Chip Routers

As shown in Fig. 1 (a), an on-chip router under VC flow control is mainly composed of five parts, the input unit, a route computation unit, a virtual channel allocator, a switch allocator and a crossbar switch. In general, these components are arranged to work in four pipeline stages (Fig. 1 (b)). When the head flit of a packet arrives at a router, it is fed to the route computation unit for finding where it is routed (RC) while also stored in one of the buffers arranged as VCs in an input unit. Afterwards, this flit is to be allocated a VC for the next hop (VA) and then a time slot of the crossbar switch (SA). Finally, this flit traverses the crossbar switch (ST) and leaves the router. If the packet is multiflit (therefore, a long packet), then a few body flits will also be allocated time slots of and traverse the crossbar switch. But RC and VA are not needed for them.

A baseline input unit in an on-chip router under VC flow control is presented in Fig. 1 (c) where buffers are arranged into VCs. Multiple VCs help alleviate head-of-line blocking and enlarge the router throughput. These VCs further belong to different virtual networks (VNs) which serve multiple classes of messages (such as requests, responses and acknowledgments) to avoid protocol-level deadlocks. Because of having such VNs and VCs, input units can consume a significant amount of power. In our evaluations, they are accounted for more than half of the static and over 80% of the dynamic power consumed by the on-chip network.

### 2.2 Motivations

When characterizing the traffic in NoCs<sup>†</sup> with certain applications from PARSEC [25], there are two interesting phenomena we found. Firstly, in Fig. 2 (a), we can see that short packets (having a header of 128 bits, therefore composed of 1 flit) dominate. This means, the first buffer elements of all VCs are more frequently-used than any other buffer elements. It also means, all other buffer elements are likely idle except the case when long packets (having a 128-bit header and a 512-bit payload, therefore composed of 5 flits) come to the routers. Secondly, when looking at the long packets (as in Fig. 2 (b)), it can be identified that when there is a long packet in an input unit, it is likely the only long packet in the input unit; that is, for majority of the time when long packets exist in an input unit, only one VC is utilized. These observations have motivated us in the following two ways:

- Since short packets dominate the network traffic, the first buffer elements in all VCs are more frequently used than other buffer elements so that replacing them with STT-MRAM will lead to much higher latency and energy. Therefore, the first buffer elements are better kept to SRAM devices if multiple memory devices are to be used for on-chip routers.
- When an input unit holds long packets, it is most likely the case that only one VC is fully active. This means, it is only necessary to implement one VC of any input units with fast but leaky SRAM devices to conserve the performance for long packets. Or in other words, it is only necessary to have one complete VC for an input unit while buffer elements other than the first one may even be removed.

<sup>&</sup>lt;sup>†</sup>Characteristics of the network traffic are collected using gem5 and GARNET under the same conditions as the "baseline system parameters" stated in Table 2.



Input Unit VN 2 All buffer elements VN 1 are implemented VN 0 with SRAM devices AC 0 S S S S

(a) A conventional on-chip router.

Fig. 1 Design of a conventional on-chip router, its pipeline workflow, and the design of its input unit.



Number of concurrent long packets



(a) Proportion of different types of packets in the network.

(b) Proportion of cycles with concurrent long packets in input units (when there is at least one long packet).

Fig. 2 Characteristics of the network traffic.

#### 3. The Proposed Input Unit Designs for On-Chip Routers

In this section, we will cover how our hybrid, asymmetric and reconfigurable input units are designed and implemented. These are followed by discussions on the baseline and our proposed designs as well, to reveal their advantages and disadvantages in depth.

#### 3.1 The Hybrid Input Unit Designs

In this proposal, STT-MRAM devices are used to complement SRAM devices as buffer elements in the input unit. STT-MRAM is an emerging memory technology which stores data in a Magnetic Tunnel Junction (MTJ) [26], [27] and it is used to implement cache because of its relatively short access time and low power consumption [28]. The MTJ has two ferromagnetic layers (the free layer and the reference layer) which decides its change in resistance so that it can be used as a storage element.

STT-MRAM has the following advantages over

SRAM. Firstly, its access latency and energy consumption in reading can be close to SRAM since only a small voltage is needed to sense the difference of resistance at the MTJ. Second, its density is higher than that of SRAM as an STT-MRAM cell generally consists of one transistor and one MTJ (called 1T-1MTJ cell). Third, the above-mentioned organization also makes its leakage negligible when compared to SRAM. However, access latency and energy consumption in writing are typical concerns of STT-MRAM devices since a high voltage is needed to reverse or retain the direction of the free layer of the MTJ.

Following Sect. 2.2, we propose a hybrid buffer design to cut the static power consumption of routers in on-chip networks. As shown in Fig. 3 (a), this hybrid design is modified from the baseline one (Fig. 1 (c)) in the following aspects:

- 1. To retain the performance for short packets which dominate the network traffic, the first buffer elements of all VCs are kept to SRAM.
- 2. VC0 is also kept to SRAM to retain the performance for majority of the long packets.



(a) The hybrid input unit.

nit. (b) The asymmetric input unit.

(c) The reconfigurable input unit.

Fig. 3 Schematic view of different input unit designs ("S" denotes SRAM, while STT-MRAM is labeled as "M").

- 3. Other buffer elements than the first one in VC1, VC2, VC3 and VC4 are instead implemented with STT-MRAM devices so that the static power consumption can be cut.
- 4. For VC1, VC2, VC3 and VC4, we assume that the SRAM buffer elements and the STT-MRAM buffer elements are accessed through their own ports, respectively.

With the above architectural modifications, one variation of this hybrid design is, the order of the flits transmitted through a hybrid VC has two options. Firstly, if all flits of a long packet are stored to a hybrid VC under an in-order manner, then each flit has to wait for its predecessor to depart the VC. This means, the second to the fifth flits are slowed down despite the fact that the fifth one may actually be stored in the SRAM buffer element. Secondly, if out-of-order transfer of flits is allowed, then two flits in a long packet can be stored and departed from SRAM buffer elements immediately. This will at least speed up one more flit in a long packet, when compared to the in-order option. Therefore, for this hybrid proposal, we in fact have two designs, inorder and out-of-order, which decide the order of flits being stored and transmitted.

Moreover, this hybrid input unit design does not cause deadlocks since uneven access time at the buffers does not affect the VC allocation processes in the routers. VC allocation is carried out after the head flit of a packet, regardless of being short or long, arrives at a router and completes its route computation. This will help the packet get its VC at the next hop and it is exactly the same as conventional routers with baseline input units. However, this hybrid design may affect the releasing time of the VCs when encountering more than one long packet in the same input unit so that latency of some long packets is hindered because of being stored in the STT-MRAM buffers. Therefore, when being used, VCs with STT-MRAM buffer elements need more time to return to the available state.

#### 3.2 The Asymmetric Input Unit Design

Apart from using STT-MRAM to implement non-critical

buffer elements in the input units, we further propose an asymmetric input unit following Sect. 2.2. With this asymmetric design, we can cut the static power consumption of routers in on-chip networks by removing non-critical buffer elements from VC1 to VC4. As shown in Fig. 3 (b), this asymmetric design is modified from the baseline one (Fig. 1 (c)) in the following aspects:

- 1. To retain the performance for short packets which dominate the network traffic, the first buffer elements of all VCs are kept.
- 2. VC0 is also kept so that the performance for majority of the long packets does not degrade.
- 3. Other buffer elements than the first one in VC1, VC2, VC3 and VC4 are removed so that the static power consumption can be cut.
- 3.3 The Reconfigurable Input Unit Design

One step further from Sect. 3.2 is, if VC1, VC2, VC3 and VC4 can be used to store one more long packet, the on-chip router with such reconfigurable input unit will then be able to store two long packets at most, to deal with rare but unfavorable cases where more than one long packets come to a router. As shown in Fig. 3 (c), this reconfigurable design is modified from the baseline one (Fig. 1 (c)) in the following aspects:

- 1. To retain the performance for short packets which dominate the network traffic, the first buffer elements of all VCs are kept.
- VC0 is also kept in order to maintain the performance for majority of the long packets.
- 3. Other buffer elements than the first one in VC1, VC2, VC3 and VC4 are removed so that the static power consumption can be cut.
- 4. The four buffer elements from VC1, VC2, VC3 and VC4 can be reconfigured to work as a single VC, giving that all of them are empty when such reconfiguration happens.
- 5. To make the SRAM blocks inside an input unit capable of being configured between four short packet-oriented VCs and one long packet-oriented VC, it is required to

Input unit design	Baseline	Hybrid (in-order)	Hybrid (out-of-order)	Asymmetric	Reconfigurable
Throughput	+++	+++	+++	+	++
Latency	+++	+	++	++	+++
Static power	+	++	++	+++	+++
Dynamic power	+	+++	+++	++	++
Energy	+	++	++	+++	+++
Area	+	++	++	+++	+++

 Table 1
 Comparisons of input unit designs ("+" denotes "better" in performance-related metrics and "smaller" in power-, energy- and area-related metrics).



(a) Added VC status (red-colored) to enable reconfigurations.



(b) Modified VC allocator to enable reconfigurations.



have an additional set of states (a few extra bits in flipflops) added for the long packet mode. This is depicted in Fig. 4 (a).

- 6. Connections to the route computation unit are already sufficient to handle the reconfiguration as VC1~4 (as in the short packet mode) are already connected to the route computation unit. Therefore, when the head flit of a long packet is stored in any of them (after they are configured as one long packet-oriented VC), route computation can be carried out as usual.
- 7. Connections between the input units and the crossbar switch do not need to be altered since each input unit has a port connected to the crossbar switch and our reconfigurable design does not require any change on this.
- 8. Modifications on the VC allocator is necessary. In the

short packet mode, a VC allocator is needed to assign all five VCs to different packets, although there is only one VC for long packets. On the other hand, in the long packet mode, another VC allocator is needed assign the two full VCs to different packets. This is depicted in Fig. 4 (b).

9. We assume such reconfiguration can happen within a cycle.

Moreover, regarding the power and area overhead of the above modifications, the additional set of states and the added VC allocator are both negligible according to our evaluations.

#### 3.4 Qualitative Discussions

Following all the modifications in our proposals (Fig. 3), we can identify their consequences for on-chip routers in terms of performance, power, energy and area, as follows:

- Performance: In the hybrid designs, the number of buffers is not modified but 4 VCs are partially replaced with STT-MRAM, which is slower than SRAM. Hence, read and write to these VCs, especially with long packets, will be slower when compared to the baseline design. For the asymmetric design, when multiple long packets come to an input unit consecutively, only the first one could be stored in the complete VC while all others have to wait and this is going to congest the network. The reconfigurable design is better in the sense of possibly having two complete VCs, but when more than two long packets come to it consecutively, waiting and congestion will also occur. Fortunately, our novel designs follow characteristics of the network traffic, where multiple long packets coming to an input unit consecutively can be considered a rare case. Therefore, performance-wise, our proposals should be close to the baseline buffer design.
- *Static power:* With our proposals, 4 of 5 VCs are partially implemented with STT-MRAM or even removed. Therefore, they will consume smaller amount of static power than the baseline design.
- Dynamic power: With our proposals, replacing buffer elements in 4 VCs with STT-MRAM devices or removing them will reduce the amount of SRAM accesses; thus reducing their dynamic power.
- *Energy:* With similar performance and much lower power, our proposals are meant to consume less energy



Fig. 5 16-tile CMP connected through an on-chip network.

when compared to the baseline design.

• *Area:* Due to the higher density of STT-MRAM devices or smaller buffer size, the area consumption will also be cut with our proposals.

The discussions above are summarized in Table 1. In addition, further comparisons between our proposals are also stated in it.

#### 4. Evaluation Methodology

In this paper, various evaluations on performance and energy are carried out. To evaluate performance, we have modified gem5 [29] and GARNET [30] to provide cycle-accurate timing models of our proposals. In addition, evaluations on power and area are carried out through Mc-PAT [31]. In these evaluations, performance, power and area models of the STT-MRAM devices are extracted from NVSim [32].

As shown in Fig. 5, we assume a 16-tile mesh network with 128-bit links in evaluations. Each tile has an in-order processor core, a bank of L2 cache/a directory. These components are connected to a router individually. The entire network is set to have three virtual networks to support the MOESI directory coherence protocol which has three classes of traffic. We have found that, for this particular coherence protocol, both short and long packets exist in all three classes of traffic. In addition, each router with baseline input units has a maximum of six ports (therefore, six input units) and each input unit has five virtual channels while each virtual channel has four 128-bit buffer elements. Detailed evaluation conditions are summarized in Table 2.

In more details, the dynamic energy of STT-MRAM buffers is evaluated based on the read/write energy consumption we obtained from NVSim. In order to obtain the dynamic energy drawn by the STT-MRAM buffers, these numbers are multiplied by the number of read/write to them. More specifically, the high access energy of STT-MRAM shown in Table 2 is a consequence of the high voltage needed by STT-MRAM when reversing or retaining the direction of the free layer in the MTJ cell.

Our evaluations are based on eight synthetic traffic patterns (Table 3) and four applications (Table 4) from PAR-SEC [25]. For evaluations with synthetic traffic, we have injected long packets to the network; since all our propos-

Table 2Evaluation parameters.

Baseline system parameters			
Number of cores:	16		
Topology:	$4 \times 4$ mesh		
Processor:	3 GHz, In-order		
L1 I/D cache:	32 KB per Processor, 2-way set associative,		
	2 cycles per access		
L2 cache:	256 KB per Bank, 8-way set associative,		
	20 cycles per access		
Cache line:	64 Bytes		
Main memory:	8 GB, 180 cycles per access		
Coherence protocol:	MOESI, Directory		
Link:	128-bit, 1 cycle traversal		
Packet:	128-bit control (short), 640-bit data (long)		
Router:	3 GHz, 4-stage pipeline,		
	103.15 mW static power,		
	1534.26 mW peak dynamic power		
Virtual channel:	5 per Virtual network		
Virtual network:	3 per Physical link		
Routing algorithm:	X-Y routing		
Process technology:	22 nm		
Vdd:	1 V		
Parameters for STT-M	RAM devices in the hybrid input unit design		
Capacity:	64 bytes (same as a VC)		
Access latency:	5 cycles per read, 31 cycles per write		
Access energy:	31.64 pJ per read, 128.8 pJ per write		
Leakage power:	534.52 uW		
Parameters for the asy	mmetric and reconfigurable input unit designs		
Router:	3 GHz, 4-stage pipeline,		
	42.63 mW static power		
	405.48 mW peak dynamic power		

 Table 3
 Synthetic traffic employed in evaluations.

Traffic patterns:	bit complement, bit reverse, bit rotation, neighbor, shuffle, tornado, transpose, and uniform random
Packet sizes:	5-flit

 Table 4
 Application traffic employed in evaluations.

Applications:	blackscholes, canneal, ferret, and x264
Input set sizes:	small and medium for <i>blackscholes</i> , <i>canneal</i> , and <i>ferret</i> ; test and small for <i>x264</i>

als work as good as the baseline input unit design. On the other hand, long packets are used to illustrate the worst case latency of different input unit designs. Moreover, there are also two input set sizes used for the application benchmark programs, as stated in (Table 4).

#### 5. Evaluation Results

In this section, we present our evaluation results and discussions. Evaluation results are carried out with both synthetic traffic pattern (Sect. 5.1) and benchmark programs from PARSEC (Sect. 5.2). In addition, we also present the area consumption of routers with our proposals.

5.1 Results under Synthetic Traffic Patterns

Evaluation results with synthetic traffic patterns are pre-



Fig. 6 Network latency per flit (cycles) versus injection rates (flits/node/cycle) under synthetic traffic patterns.



sented in Fig. 6. All these figures show network latency per flit versus injection rates. With 1-flit packets, our proposals work exactly the same as the baseline input unit so we only present results evaluated with 5-flit packets, which represent long packets, to reflect the worst case performance of different input unit designs.

Firstly, we can see that the reconfigurable input unit performs better than other proposals and they it is the most similar one to the baseline design, especially when the injection rate is low. This is because of the fact that it can be reconfigured to function as two complete SRAM-based VCs. However, when the injection rate gets higher, our reconfigurable design exacerbates faster than the baseline design. This is reasonable as five complete VCs have much better throughput than two of them.

Secondly, the asymmetric and the out-of-order hybrid designs perform very similar. When the injection rate is low, the asymmetric design is slightly better than the out-of-order hybrid design. On the other hand, the out-of-order hybrid design saturates slower than the asymmetric design since its size of buffers is larger than the asymmetric design.

Thirdly, the in-order hybrid design performs the worst within all counterparts. This simply means that it is not a good idea to utilize slower memory devices too often.

Fourthly, when the injection rate is very low (at the beginning of the curves), all our proposals can approach the baseline design in terms of per-flit latency. This means, when the network is not busy, removing some buffer elements or replacing them with slower memory devices does not affect the network performance much.

Fifthly, within all traffic patterns, we can observe that "bit reverse" and "transpose" are the most stringent ones as latency under these two traffic patterns climbs the fastest with increasing injection rate. On the other hand, latency increases much more slowly under "neighbour" and "tornado".

### 5.2 Results with PARSEC Applications

Evaluation results on network latency per flit are shown in Fig. 7. For the hybrid designs, although our proposal uses slower memory technology, it achieves similar latency to the baseline design. We can observe that our out-of-order hybrid proposal is roughly 1 cycle slower than the baseline design in "x264" workloads with an average slowdown of 13.9%. This has proved that our traffic-aware philosophy works well. It can also be found that the in-order hybrid design is the slowest within all counterparts, this is due to the extra flit buffered in the STT-MRAM devices. For the asymmetric and the reconfigurable designs, it is clear that they are simply faster than the hybrid ones, with the reconfigurable design only slowing flits down by 1.5% on average.

Results presented in Fig. 8 reflect the network energy consumption (including energy consumed through both



Fig. 8 Energy consumption per flit (solid indicates dynamic power).



static and dynamic power) of different input unit designs. Because of utilizing STT-MRAM devices or having less buffer elements, our proposals outperform the baseline design dramatically, with an average energy reduction per flit of up to 56.2% (under the reconfigurable design). This means, by considering the characteristics of network traffic and properly utilizing the buffer elements, it is possible to achieve much better energy efficiency. Another observation is, the asymmetric and reconfigurable designs consume less static energy than the hybrid designs but more dynamic energy. This is due to the fact that they are solely implemented with power-hungry SRAM buffers but their sizes of buffers are smaller.

Moreover, with different input set sizes, we do not observe any significant difference in network latency, thus it seems that all input sets we evaluated do not overload the evaluated NoC. On the other hand, we observe slightly higher energy consumption with larger input sets for "canneal", "ferret", and "x264", especially in dynamic energy. This is simply due to the fact that larger inputs sets incur higher amount of network traffic.

#### 5.3 Router Area with Different Input Unit Designs

As shown in Fig. 9, another benefit of using STT-MRAM or simply removing buffer elements is the reduction of area. After replacing VCs with STT-MRAM devices, on-chip routers with our proposal consume 93.6% of the area of conventional routers. Moreover, the asymmetric and reconfigurable designs are the best (66.3% of a conventional router) in this evaluation since their buffers are removed.

A drawback of our evaluations is, we do not model the hardware changes in our proposals that are utilized to enable out-of-order flit transfer and the reconfiguration of buffers. These changes can affect the energy and area consumption but fortunately not the network performance and we assume flit re-ordering and buffer reconfiguration can be easily carried out. For their energy and area overhead, we plan to address them in our future work at the circuit level.

#### 6. Related Works

Many optimization techniques were proposed in the past to help suppress the power consumption and improve the energy efficiency of NoCs. For example, there were ideas focusing on directly reducing the power of routers through power gating [5]–[8] and dynamic voltage and frequency scaling [9]. Also, there were works which proportionally supply power to the network based on the actual demand [10] or even completely eliminate routers (hence also buffers) through smart wiring [11]. Speculative routing was also brought up through ideas like prediction [12], multicasting [13], and pipeline bypassing [15]. They were used to raise the energy efficiency through improved performance. In addition, traffic compression [14] and hybrid flow control mechanisms [16]–[18] were also used for similar purposes.

For input units that are common in most of the router designs, studies targeting at lower power consumption and higher energy efficiency can generally be categorized into two groups. Firstly, shared buffer design was proposed to utilize the imbalance of different VCs in the input unit to save both power and area [19]. Secondly, there were also works focusing on implementing energy-/area-efficient input units with other memory devices and a typical example was STT-MRAM. Jang et al. made use of the density advantage of STT-MRAM in order to enlarge the buffer capacity and network throughput [22]. They also tried to hide the long write latency and avoid the high write energy of STT-MRAM devices through innovative control over the usage of different buffer elements. Zhan et al. proposed two hybrid input unit designs (hierarchical and banked) to mix SRAM and STT-MRAM for both energy- and area-efficiencies [23].

In more details, the proposal from Jang *et al.* is different from our work in the purpose. With larger total buffer capacity, it should outperform the designs from us and Zhan *et al.* in network throughput but would definitely consume larger power and area. Moreover, the two hybrid input unit designs from Zhan *et al.* should have similar power and area consumption to our hybrid design if the number of buffer elements and the process technology are the same; since all of them have similar ratios of STT-MRAM buffer elements over SRAM buffer elements. On the other hand, our asymmetric and reconfigurable designs should consume less power and smaller area as we simply removed some SRAM buffer elements in them.

#### 7. Conclusions

Input units in routers are an important aspect of NoCs since they determine both power consumption and throughput of the network. In this paper, we have proposed three novel input unit designs. With the hybrid design, two memory devices, SRAM and STT-MRAM have been mixed in the router. On the other hand, with the asymmetric and reconfigurable designs, half of the buffer elements have been simply removed.

These proposals have allowed the frequency of long packets to be utilized to optimize the accesses to buffers in the input units and this has further resulted in a significant cut to the energy consumption of flits in the network. On the other hand, also benefiting from the frequency of long packets, our hybrid, asymmetric and reconfigurable buffer designs have incurred very little negative effect on the latency despite using a slow but less leaky memory technology or having much less buffering spaces. Such effectiveness has proved that our proposals are more future-proof as the importance of energy efficiency escalates.

#### Acknowledgments

First and foremost, we would like to sincerely thank the editor and reviewers for their help and valuable comments on our submission. In addition, we want to express our gratitude to Hiroshi Nakamura, Shinobu Miwa and Takashi Nakada for their insightful advice at the early stage of this work. We also want to thank Jinyu Jiao for his help in the evaluations. This work was supported, in part, by JSPS KAKENHI with Grants JP20K23315 and JP22K11958, and by the Okawa Foundation for Information and Telecommunications under Grant 21-04.

#### References

- W. Dally and B. Towles, Principles and Practices of Interconnection Networks, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Dec. 2003.
- [2] Y.J. Yoon, N. Concer, M. Petracca, and L. Carloni, "Virtual channels vs. multiple physical networks: A comparative analysis," Proc. 47th Annual Design Automation Conference, pp.162–165, June 2010.
- [3] A.N. Udipi, N. Muralimanohar, and R. Balasubramonian, "Towards scalable, energy-efficient, bus-based on-chip networks," Proc. IEEE 16th International Symposium on High-Performance Computer Architecture, pp.1–12, Jan. 2010.
- [4] A. Banerjee, R. Mullins, and S. Moore, "A power and energy exploration of network-on-chip architectures," Peoc. 1st International Symposium on Networks-on-Chip, pp.163–172, May 2007.
- [5] H. Matsutani, M. Koibuchi, D. Wang, and H. Amano, "Runtime power gating of on-chip routers using look-ahead routing," Proc. 2008 Asia and South Pacific Design Automation Conference, pp.55–60, Jan. 2008.
- [6] H. Matsutani, M. Koibuchi, D. Ikebuchi, K. Usami, H. Nakamura, and H. Amano, "Ultra fine-grained run-time power gating of on-chip routers for CMPs," Proc. 4th ACM/IEEE International Symposium on Networks-on-Chip, pp.61–68, May 2010.

- [7] L. Chen and T.M. Pinkston, "NoRD: Node-router decoupling for effective power-gating of on-chip routers," Proc. 45th Annual IEEE/ACM International Symposium on Microarchitecture, pp.270–281, Dec. 2012.
- [8] L. Chen, D. Zhu, M. Pedram, and T.M. Pinkston, "Power punch: Towards non-blocking power-gating of NoC routers," Proc. 21st IEEE International Symposium on High Performance Computer Architecture, pp.378–389, Feb. 2015.
- [9] X. Chen, Z. Xu, H. Kim, P. Gratz, J. Hu, M. Kishinevsky, and U. Ogras, "In-network monitoring and control policy for DVFS of CMP networks-on-chip and last level caches," Proc. 6th IEEE/ACM Sixth International Symposium on Networks-on-Chip, pp.43–50, May 2012.
- [10] R. Das, S. Narayanasamy, S.K. Satpathy, and R.G. Dreslinski, "Catnap: Energy proportional multiple network-on-chip," Proc. 40th Annual International Symposium on Computer Architecture, pp.320–331, June 2013.
- [11] F. Alazemi, A. AziziMazreah, B. Bose, and L. Chen, "Routerless network-on-chip," Proc. 24th IEEE International Symposium on High Performance Computer Architecture, pp.492–503, Feb. 2018.
- [12] H. Matsutani, M. Koibuchi, H. Amano, and T. Yoshinaga, "Prediction router: yet another low latency on-chip router architecture," Proc. 15th IEEE International Symposium on High Performance Computer Architecture, pp.367–378, Feb. 2009.
- [13] Y. He, H. Sasaki, S. Miwa, and H. Nakamura, "McRouter: Multicast within a router for high performance network-on-chips," Proc. 22nd International Conference on Parallel Architectures and Compilation Techniques, pp.319–330, Sept. 2013.
- [14] R. Das, A.K. Mishra, C. Nicopoulos, D. Park, V. Narayanan, R. Iyer, M.S. Yousif, and C.R. Das, "Performance and power optimization through data compression in network-on-chip architectures," Proc. 14th IEEE International Symposium on High Performance Computer Architecture, pp.215–225, Feb. 2008.
- [15] A. Ejaz, V. Papaefstathiou, and I. Sourdis, "FreewayNoC: A DDR NoC with pipeline bypassing," Proc. 13th IEEE/ACM International Symposium on Networks-on-Chip, pp.1–8, Oct. 2018.
- [16] N.D.E. Jerger, L.-S. Peh, and M.H. Lipasti, "Circuit-switched coherence," Proc. 2nd ACM/IEEE International Symposium on Networks-on-Chip, pp.193–202, April 2008.
- [17] A.K. Lusala and J.-D. Legat, "Combining SDM-based circuit switching with packet switching in a router for on-chip networks," International Journal of Reconfigurable Computing, vol.2012, pp.1–16, Sept. 2012.
- [18] J. Jiao, Y. He, T. Cao, and M. Kondo, "Enabling circuit-switching in modern on-chip networks," Microprocessors and Microsystems, vol.95, 104712, 2022.
- [19] H. Farrokhbakht, H. Kao, and N.E. Jerger, "UBERNoC: Unified buffer power-efficient router for network-on-chip," Proc. 13th IEEE/ACM International Symposium on Networks-on-Chip, pp.1–8, Oct. 2019.
- [20] C. Li and P. Ampadu, "A compact low-power eDRAM-based NoC buffer," Proc. 2015 IEEE/ACM International Symposium on Low Power Electronics and Design, pp.116–121, July 2015.
- [21] D. Kline, H. Xu, R. Melhem, and A.K. Jones, "Domain-wall memory buffer for low-energy NoCs," Proc. 52nd Annual Design Automation Conference, June 2015.
- [22] H. Jang, B.S. An, N. Kulkarni, K.H. Yum, and E.J. Kim, "A hybrid buffer design with STT-MRAM for on-chip interconnects," Proc. 6th IEEE/ACM International Symposium on Networks-on-Chip, pp.193–200, May 2012.
- [23] J. Zhan, J. Ouyang, F. Ge, J. Zhao, and Y. Xie, "DimNoC: A dim silicon approach towards power-efficient on-chip network," Proc. 52nd Annual Design Automation Conference, June 2015.
- [24] H. Matsutani, M. Koibuchi, D. Wang, and H. Amano, "Adding slowsilent virtual channels for low-power on-chip networks," Proc. 2nd ACM/IEEE International Symposium on Networks-on-Chip, pp.23–32, April 2008.

- [25] C. Bienia, S. Kumar, and K. Li, "PARSEC vs. SPLASH-2: a quantitative comparison of two multithreaded benchmark suites on chipmultiprocessors," Proc. 2008 IEEE International Symposium on Workload Characterization, pp.47–56, Sept. 2008.
- [26] X. Wu, J. Li, L. Zhang, E. Speight, R. Rajamony, and Y. Xie, "Hybrid cache architecture with disparate memory technologies," Proc. 36th Annual International Symposium on Computer Architecture, pp.34–45, 2009.
- [27] D. Apalkov, A. Khvalkovskiy, S. Watts, V. Nikitin, X. Tang, D. Lottis, K. Moon, X. Luo, E. Chen, A. Ong, A. Driskill-Smith, and M. Krounbi, "Spin-transfer torque magnetic random access memory (STT-MRAM)," ACM Journal on Emerging Technologies in Computing Systems, vol.9, no.2, pp.1–35, May 2013.
- [28] C.W. Smullen, V. Mohan, A. Nigam, S. Gurumurthi, and M.R. Stan, "Relaxing non-volatility for fast and energy-efficient STT-RAM caches," Proc. 17th IEEE International Symposium on High Performance Computer Architecture, pp.50–61, Feb. 2011.
- [29] N. Binkert, B. Beckmann, G. Black, S.K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D.R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M.D. Hill, and D.A. Wood, "The gem5 simulator," ACM SIGARCH Computer Architecture News, vol.39, no.2, pp.1–7, Aug. 2011.
- [30] N. Agarwal, T. Krishna, L.-S. Peh, and N.K. Jha, "GARNET: a detailed on-chip network model inside a full-system simulator," Proc. 2009 IEEE International Symposium on Performance Analysis of Systems and Software, pp.33–42, April 2009.
- [31] S. Li, J.H. Ahn, R.D. Strong, J.B. Brockman, D.M. Tullsen, and N.P. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," Proc. 42nd Annual IEEE/ACM International Symposium on Microarchitecture, pp.469–480, Dec. 2009.
- [32] X. Dong, C. Xu, Y. Xie, and N.P. Jouppi, "NVSim: A circuitlevel performance, energy, and area model for emerging nonvolatile memory," IEEE Trans. Comput.-Aided Design Integr. Circuits Syst., vol.31, no.7, pp.994–1007, 2012.



Yuan He is a lecturer (fixed-term) in the Faculty of Science and Technology, Keio University, Japan and an adjunct associate professor at the School of Information Science and Engineering, Shenyang University of Technology, China. He received his Ph.D. from The University of Tokyo in 2014, M.E. with First Class Honours and B.Sc. from the University of Auckland in 2009 and 2005, respectively. His research interests include computer architecture, domain-specific accelerations and in-

memory processing. He is a member of the ACM, CCF, IEEE and IEICE.



Xiaohan Yue is an Associative Professor in the School of Information Science and Engineering at Shenyang University of Technology. He received B.E. in electronic information and technology, M.E. and Ph.D. in computer application technology from the Northeastern University, China, in 2005, 2009, and 2013, respectively. His primary research interests are trusted computing, wireless network security, security and privacy in ubiquitous computing, and cryptography. He is a member of the CCF and IEEE.



Haiyan Jiang received her Master of Engineering in computer science and technology from the School of Information Science and Engineering at Shenyang University of Technology in June 2022. Her research interests include embedded systems and network security.



Xiaoman Liu received his Bachelor of Engineering in computer science and technology from the School of Mathematics and Information Science at Anshan Normal University in June 2019. He joined the School of Information Science and Engineering at Shenyang University of Technology as a postgraduate student working on his master degree from September 2020. His research interests include computer architecture and machine learning.



Xibo Wang is a Professor and the Dean of the School of Information Science and Engineering at Shenyang University of Technology. He received his Ph.D. in computer software and theory from the Northeastern University in China. He was a senior visiting scholar at the Allen University of Technology in Germany. He was awarded with the Hundred Talents Project of Liaoning Province. His main research interests include computer detection and control, management information system design, real-

time systems and embedded software. He is a senior member of the CCF.



Yujie Gao received her Master of Engineering in software engineering from the School of Software at Shenyang University of Technology in June 2022. Before joining Shenyang University of Technology, she obtained her Bachelor of Engineering in computer science and technology from the School of Information Engineering, Shenyang University. Her research interests include the programming and design of embedded systems.