

# A power efficient flit-admission scheme for wormhole-switched networks on chip

Zhonghai Lu, Li Tong, Bei Yin and Axel Jantsch  
Laboratory of Electronics and Computer Systems  
Royal Institute of Technology, Sweden  
{zhonghai,axel}@imit.kth.se, {lit,beiy}@kth.se

## Abstract

*Reducing power consumption is a main challenge when adopting a network as a global on-chip communication interconnect since the reduction in power dissipation should not at the expense of degrading the system performance. We investigate power in a wormhole-switched network with focus on the impact of flit-admission schemes, i.e., when and how the flits of packets are admitted into the network. We have proposed a novel flit-admission scheme that shows significant shrink of the switch complexity while maintaining equivalent network performance. This paper investigates its influence in network power involving both switches and links. We conduct experiments on a 2D mesh network. The results show that our flit-admission scheme achieves significant power and area reduction without performance penalty. To our knowledge, our work is the first study of power dissipation on flit admission schemes.*

**Keywords:** Power consumption, Network-on-Chip

## 1 Introduction

Power consumption is a key issue in network-on-chip (NoC) design. It is a critical criterion for most portable embedded systems-on-chip. In the nanometer regime, as the system capacity steadily increases, minimizing power consumption becomes more important in order to enhance system reliability and reduce packaging costs. Even from the performance perspective, as noted in [9], the communication bandwidth in future NoC architectures will probably be only limited by prohibitive levels of power consumption.

Network flow control governs how a packet is forwarded in a network, concerning shared resource allocation and contention handling. Wormhole switching is a network flow control mechanism that allocates buffers and physical channels (PCs) to flits instead of packets. A packet is split into one or more flits. A flit, the smallest unit on which flow control is performed, can advance once buffering in the next hop is available to hold the flit. This results in

that the flits of a packet are delivered in a pipeline fashion. For the same amount of storage, it achieves lower latency and greater throughput. However, wormhole switching uses channels inefficiently because a PC is held for the duration of a packet. If a packet is blocked, all PCs held by this packet are left idle. To mitigate this problem, wormhole switching adopts *virtual channels (lanes)* to make efficient use of the PCs [1]. Several parallel lanes, each of which is a flit buffer queue, share a PC. Therefore, if a packet is blocked, other packets can still traverse the PC via other lanes, leading to higher throughput. Because of these advantages, namely, *better performance, smaller buffering requirement* and *greater throughput*, wormhole switching with lanes is being advocated for on-chip networks [3, 13].

Flit admission in a wormhole-switched network concerns when and how the flits of packets are admitted or injected into the network. This is a critical issue since flits are the workload of switches, and also the source of network contentions for shared VCs and links. To achieve good network utilization and system throughput, flits should be admitted as fast as possible. However, the more flits admitted in the network, the higher contentions may occur, leading to performance degradation. We have looked into the problem and proposed a novel flit-admission scheme in [5] which shows a significant complexity reduction of the switch crossbar while maintaining equivalent performance. Since it couples a flit-admission queue with a PC in a *one-to-one* manner, this scheme is called *coupled flit admission*. In this paper, we present a study of power consumption on this scheme in comparison with a typical flit-admission scheme, called *decoupled flit admission* (see Section 3.1).

The remainder of the paper is organized as follows. Section 2 briefs the related work, highlighting the lack of research in optimizing network power. In Section 3, we describe the two flit-admission schemes, namely, the *decoupled admission* and *coupled admission*. We describe our power study in Section 4. Section 5 evaluates the flit-admission schemes under uniformly random traffic and locality traffic, i.e. traffic exploring the communication locality. Finally, we conclude the paper in Section 6.

## 2 Related work

There exists a variety of techniques to minimize power consumption of on-chip communication. A technical overview of these techniques operating at the circuit level, the architecture level, the network level and the system level is presented in [12]. However, few work has been performed for the power efficiency of on-chip networks. Power modeling and simulation were studied in [10, 15]. Shang et al. [14] proposed a dynamic voltage scaling technique for on-chip communication networks. In [4], Jantsch et al. analyzed the power consumption of link-level low power encoding techniques and end-to-end data protection in the Nostrum NoC [7]. The current research status strongly motivates to optimize the interconnection network power because it consumes a large fraction of the overall system power. For instance, the integrated router and links of the Alpha 21364 processor consumes about 20 percent of the total system power [8].

To the best of our knowledge, our work is the first investigation of power consumption on flit admission schemes. While Shang et al. and Jantsch et al. examined dynamic power optimization from the network architecture level, our study targets the micro-architectural level.

## 3 The flit-admission schemes

### 3.1 The decoupled admission

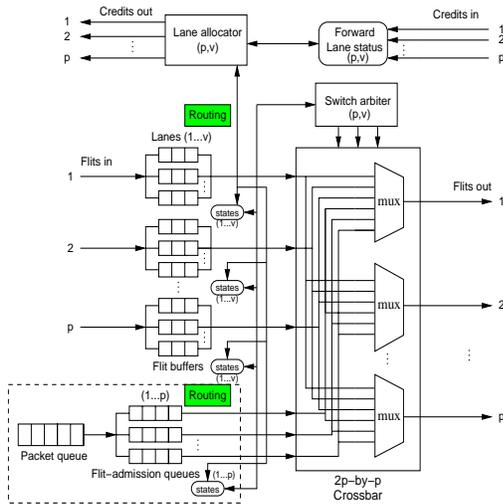


Figure 1. A canonical wormhole lane switch

Figure 1 illustrates a canonical wormhole switch architecture with virtual channels at input ports [1, 11, 13], which has  $p$  physical channels (PCs) and  $v$  lanes per PC. It employs credit-based flow control to coordinate packet delivery between switches. A network packet passes the switch

through four states: *routing*, *lane allocation*, *flit scheduling*, and *switch arbitration* [5]. The figure also shows a typical flit admission scheme which consists of  $p$  flit-admission queues with each connected to  $p$  multiplexers. Like a lane, a flit-admission queue is a FIFO storing the flits of packets. Initially, packets are stored in the packet queue. When a flit-admission queue is available, a packet is split into flits which are then put into the admission queue. This decomposition is called *flitization*<sup>1</sup>. A flit-admission queue transmits states similarly like a lane to inject flits into the network via the crossbar. By this scheme, flits from a flit-admission queue can be switched to any one of the  $p$  output PCs. As the flit-admission queues are decoupled from the PCs, the crossbar must be fully connected, resulting in a port size of  $2p \times p$ . We call this admission scheme *decoupled admission*.

### 3.2 The coupled admission

Figure 2 sketches the coupled flit-admission scheme in the switch model. Just like the decoupled admission, it uses  $p$  admission queues, but one queue is bound to one and only one multiplexer for a particular PC. Due to this coupling, flits from a flit-admission queue are dedicated to the PC. Consequently, an admission queue only needs to be connected to one multiplexer instead of  $p$  multiplexers. The size of the crossbar is sharply decreased from  $2p \times p$  to  $(p + 1) \times p$ , as shown in Figure 2. The number of control signals per multiplexer is reduced from  $\lceil \log(2p) \rceil$  to  $\lceil \log(p + 1) \rceil$  for any  $p > 1$ <sup>2</sup>.

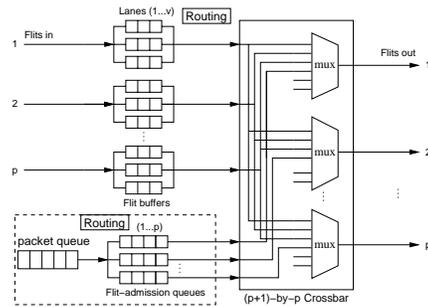


Figure 2. Sharing a  $(p+1)$ -by- $p$  crossbar

The implementation of the coupled admission incurs no additional hardware overhead. Specifically, the decoupled admission performs flitization before routing; the coupled admission does routing before flitization. By a routing algorithm, the output port of a packet is determined and then the flits of the packet are put into the flit-admission queue which aims to the output port of the switch.

<sup>1</sup>Flitization is named following packetization, i.e., encapsulate a message into one or more packets.

<sup>2</sup> $\lceil x \rceil$  is the ceiling function which returns the least integer that is not less than  $x$ .

## 4 Network power consumption

### 4.1 Switch implementation

We have implemented the wormhole switch model including both flit-admission schemes in VHDL at the RT level. The switch has four input PCs, four output PCs, one packet admission channel and one packet ejection channel. The ejection model used in the design is the  $p$  sink model which differs from an ideal ejection model in that it uses  $p$  sinks to eject flits instead of  $p \cdot v$  sinks [6]. This switch model allows us to configure the switch parameters, such as the flit width  $W_{flit}$ , the number  $v$  of VCs per PC, the depth of a VC, and the depth of an admission queue etc. Each PC, called a *link*, consists of  $W_{flit}$  parallel wires for data flits,  $W_{credit}$  parallel wires for credits, and two additional wires functioning as enable/valid signals for flits and credits. The credit wire width  $W_{credit}$  depends on  $v$  by  $W_{credit} = \log v$ .

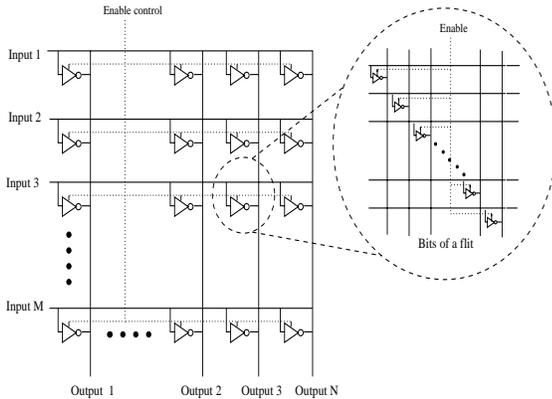


Figure 3. A MxN matrix crossbar

As discussed previously, the most area reduction of the coupling scheme is the crossbar. To explore the crossbar design space, we also examine the matrix crossbar shown in Figure 3, which is another common crossbar structure. In the matrix crossbar, the input ports and the output ports are connected by tri-state gates. The data of a flit from an input port propagate to the input ends of the tri-state gates in the same row; the open/close states of these gates determine which output ports receive the input data. We have also implemented the switch model with the tristate-gate crossbar. The two kinds of crossbar use the same arbiter which generates control signals for the multiplexer-based crossbar. Since the tristate-gate-based crossbar uses different control signals, we designed and included a decoder for the tristate-gate crossbar to realize the control signal conversion.

We synthesize the designs using Synopsys Design Compiler with medium mapping effort. The target technology is UMC18 (180nm). We do not synthesize storage elements,

i.e., buffers for lanes and packet/flit queues, and will neglect their power consumption in this study. The reason is that, since they consume a large fraction of silicon area, these buffers probably need to be customized as dedicated hardware FIFOs instead of RAM-based or register-based buffers [13]. In order to have a fair comparison, we do not place any constraints on timing and area for synthesis.

Table 1 shows the synthesized results for switch area in gates when  $W_{flit} = 32$  and  $v = 4$ . The routing algorithm implemented is dimension-ordered X-Y routing, which is deterministic. Particularly, it is simple and thus cheap to be implemented on silicon. We compare the decoupled-admission switch and coupled-admission switch using multiplexer-based crossbar (Mux-CB) and tristate-gate-based crossbar (Tristate-CB). As can be seen, due to the coupling scheme, the area reduction for the two types of crossbar and corresponding arbiter is about 41% and 15.5%, respectively. The reduction of switch area is up to 8%.

Implementation	Crossbar	Arbiter	Other logic	Total
Decoupled Mux-CB	2533	10279	37948	50760
Coupled Mux-CB	1479	8685	37422	47586
<i>reduction</i>	41.6%	15.5%	1.4%	6.3%
Decoupled Tristate-CB	5293	10279	37948	53520
Coupled Tristate-CB	3149	8685	37422	49256
<i>reduction</i>	40.5%	15.5%	1.4%	8%

Table 1. Comparison of area

Table 2 compares the critical path in ns and frequency in MHz of the four switch designs. The implementation of the switch model consists of a data path and a control path. It uses two clocks, one for the data path, the other for the control path. The control frequency runs two times as fast as the data frequency. The control path is the performance bottleneck. The link operating rate is equal to the data frequency. As can be seen, the coupling scheme does not degrade the implementation performance in both cases. It is worth mentioning that the switch designs are study implementations, which are not optimized. But they are sufficient to conduct power study for this paper, since we are concerned with the relative power instead of absolute power.

Implementation	Control path	Data path	Cont. Freq.
Decoupled Mux-CB	23.27	2.51	43
Coupled Mux-CB	18.72	2.49	53.4
Decoupled Tristate-CB	23.27	2.64	43
Coupled Tristate-CB	18.72	2.58	53.4

Table 2. Comparison of performance

With one of the switch designs, we can construct a network and then investigate its power consumption in switches and links. The switches are connected without any additional data link and physical layer logic. This implies that no encoding/decoding, multiplexing/de-multiplexing is

introduced between links, i.e., flits and credits are transmitted by parallel wires. Thus, the number  $N$  of wires per link equals to  $W_{flit} + W_{credit} + 2$ .

## 4.2 Switch and link power consumption

The power consumption of all design blocks (crossbar, routing logic, lane allocator, switch arbiter etc.) is analyzed using Synopsys Power Compiler with the UMC18 library. By the technology, the operating voltage is 1.98 volt. The switch operating frequency is conservatively set to be 20 MHz for both the multiplexer-based switch and the tristate-gate-based switch. We consider only switching power and neglect leakage power and short-circuit power. Besides, we do not include the power consumed by the clock tree.

We calculate the power consumption for a link wire (one bit) between the switches using

$$P_{wire} = \frac{1}{2} f C_{wire} V_{dd}^2 \cdot \alpha_{wire} \quad (1)$$

$$C_{wire} = \epsilon \frac{W_{net} L_R}{d} \quad (2)$$

where  $\alpha_{wire}$  is the switching probability of a wire which is either a flit data wire, a credit wire or a control wire;  $C_{wire}$  is the switching capacitance per wire;  $W_{net}$  and  $L_R$  are the width and length of a wire, respectively;  $d$  is the equivalent distance between wires and ground,  $V_{dd}$  is the supply voltage, and  $f$  is the operating frequency. Specifically, the wire width  $W_{net} = 3.2 \times 10^{-7} m$ , the distance  $d = 3.2 \times 10^{-8} m$ ; the permittivity of  $SiO_2$   $\epsilon = 3.5 \times 10^{-11} F/m$ . Assuming  $L_R$  is 4 mm, we have the link wire capacitance  $C_{wire} = 1.4 pF$ . Then, the power consumption of all the network links can be obtained by

$$P_{link} = \sum_{i=1}^M \sum_{j=1}^N P_{wire(i,j)}$$

where  $M$  is the number of network links;  $N$  is the number of wires per link.

## 5 Experimental results

### 5.1 Experiments setup

As regular low-dimension topologies are proposed for on-chip networks in order to simplify the control part of switches [2, 7], we construct four networks with 2D  $4 \times 4$  mesh topology using the four VHDL switch designs, namely, the *multiplexer-based switch with the decoupled admission*, the *multiplexer-based switch with the coupled admission*, the *tristate-gate-based switch with the decoupled admission* and the *tristate-gate-based switch with the*

*coupled admission*. For each network constructed, the switches operate synchronously. The X-Y routing guarantees deadlock-free for the mesh topology. Each switch has the same configuration parameters as follows: The number  $v$  of VCs per PC is chosen to be four, which is optimal for cost-performance tradeoff [1]; The depth of a lane is 2, which is minimal in order to pipeline flits; The depth of an admission queue is set to hold the flits of exactly one packet. All these buffer settings are intended to minimize the buffering cost. The flit width  $W_{flit}$  is set to 32 bits, thus the link width  $N$  is 36 ( $32 + \log 4 + 2 = 36$ ). For the  $4 \times 4$  mesh, the number  $M$  of links is 48.

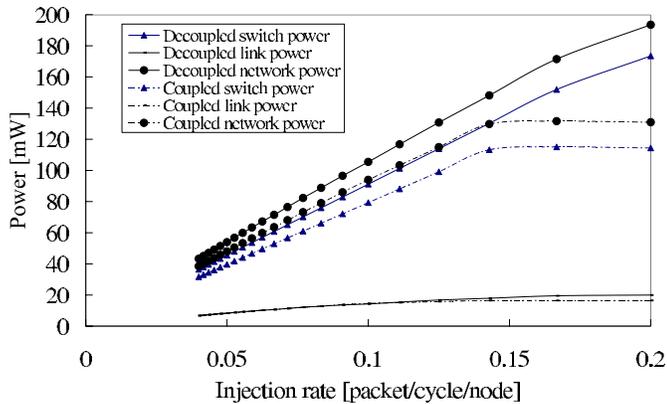
The purpose of our experiments is to compare the network power consumption between the switches using the decoupled admission and the switches using the coupled admission. We evaluate the networks with two types of traffic. One is random traffic which is distributed uniformly in the network. The other is *locality* traffic which is distributed in reversely proportional to distance. That is, if the source node is closer to the destination node, the communication probability is higher; if the distance is farther, the communication probability is lower. The locality traffic in fact explores the communication locality, which is a main optimization objective while mapping an application on a NoC in order to reduce latency and save power. Packets have a fixed length of four flits, with a head flit leading three data flits. They are injected into the network with a constant rate. The packet payload consists of random '0' and '1' bits.

We show the estimated power results in the following two subsections. We explain the text in the legends of the figures (Figure 4, 5, 6 and 7) as follows:

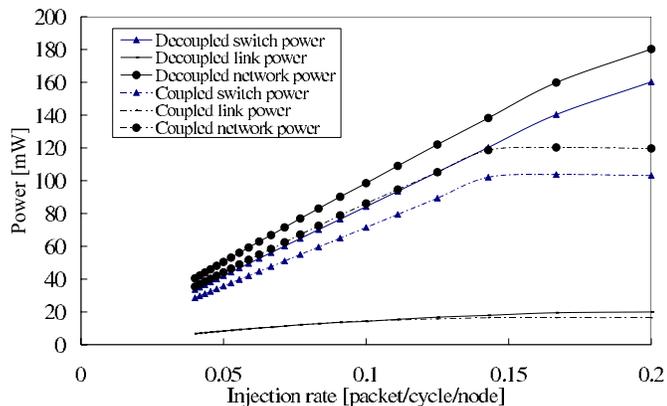
- Decoupled/Coupled switch power: Power consumed by all the switches in the network structured by switches using the decoupled/coupled admission.
- Decoupled/Coupled link power: Power consumed by all the links in the network structured by switches using the decoupled/coupled admission.
- Decoupled/Coupled network power: Power consumed by the network using the decoupled/coupled admission. It is the sum of the switch power and link power.

### 5.2 Power consumption with uniform traffic

Figure 4 shows the estimated power results for the networks with the multiplexer-based switches under uniform traffic. The power reduction due to the coupling scheme is 12.7% on average. The contributions are from both switches and links, with a reduction of 14.5% and 2.1% on average, respectively. The link power does not increase linearly with the injection rate. This is due to that the average link switching activity of the network does not go up with the injection rate in an exactly linear manner.



**Figure 4. Power of decoupled and coupled admission with mux-cb under uniform traffic**



**Figure 5. Power of decoupled and coupled admission with tristate-cb under uniform traffic**

The power dissipation of the networks with the tristate-gate-based switches under uniform traffic is shown in Figure 5. With the coupling scheme, the power saving on average is 14.3%; the switch power and link power are decreased by 16.5% and 2.1% on average, respectively.

In order to know the power distribution in a switch, we make a profile of the switch power at the injection rate of 0.2 packet/cycle/node. The power consumed by the logic blocks for the four designs is listed in Table 3, where SA denotes the average link switching activity. We can see that the crossbar consumes a significant portion of the switch power while the power consumed by the arbiter is almost negligible. The reduction in switch power is up to 35.7%.

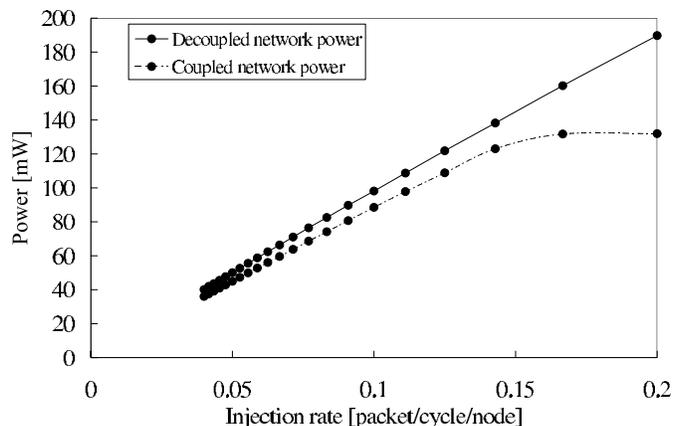
The reasons for the power saving with the coupling scheme are two-fold. First, as the number of input ports of the crossbar is shrunk from  $2p$  to  $p + 1$ , the gate count and switching capacitance are decreased. Since the crossbar is very power-consuming, this complexity reduction leads to

Implementation	CB	Arbiter	Other	Total	SA
Decoupled Mux-CB	19.38	0.08	34.71	54.17	21.1%
Coupled Mux-CB	8.63	0.11	26.99	35.73	17.5%
<i>reduction</i>	55%	-37%	22%	34%	17.1%
Decoupled Tristate-CB	16.89	0.08	33.07	50.04	21.1%
Coupled Tristate-CB	5.22	0.12	26.81	32.15	17.5%
<i>reduction</i>	69%	-50%	19%	35.7%	17.1%

**Table 3. Power distribution in a switch**

less power consumption in the switches. Second, with the coupled admission, the average link switching activity is reduced when the packet injection rate is higher. At a higher injection rate, the coupling scheme admits flits into the network slower than the decoupled admission due to head-of-line blocking [5]. Therefore, the switching activities inside switches and link switching activity become less, contributing to switch power and link power reduction, respectively.

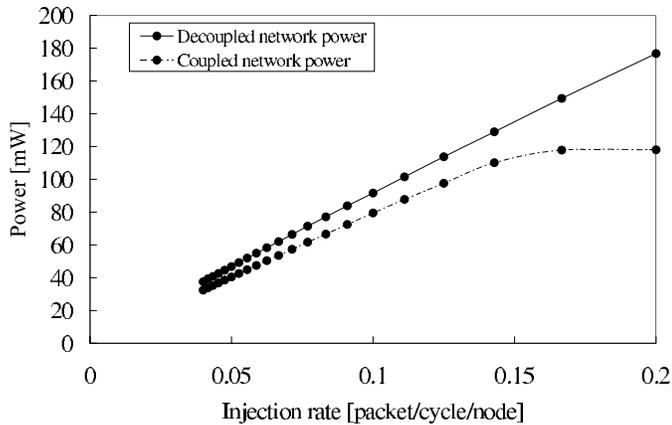
### 5.3 Power consumption with locality traffic



**Figure 6. Power of decoupled and coupled admission with mux-cb under locality traffic**

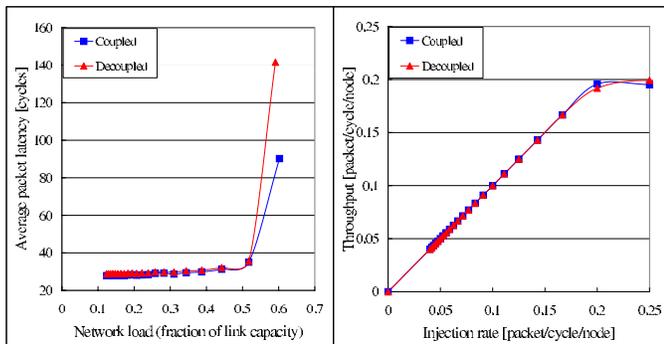
We depict the power consumption of the networks with the multiplexer-based switches and tristate-gate-based switches under locality traffic in Figure 6 and Figure 7, respectively. Thanks to the coupling scheme, in the first case, the power reduction ranges from 10.1% to 30.5%, 11.5% on average; in the latter case, the power decreases in the range [13.5%, 33.2%], 14.9% on average.

For the uniform traffic, the coupling scheme does not exhibit performance degradation with the single-cycle switch model and switch configuration parameters [5]. For the locality traffic, we draw the network performance simulated with the RTL mux-based switches in Figure 8. Each switch uses 16 flit-admission queues in order to alleviate the impact of multi-cycle delivery on the flit-admission speed. It



**Figure 7. Power of decoupled and coupled admission with tristate-cb under locality traffic**

takes five data cycles for a head flit and three data cycles for other flits to pass through a switch. The decoupled admission uses one more cycle for routing than the coupled one. The average packet latency is calculated from the instant the packet is split into flits to that the packet is assembled after receiving all the flits and ejected from the network. Packet source queuing time is not included. The network load is the average percentage of active links.



**Figure 8. Performance under locality traffic**

As can be observed, the network employing the coupling scheme achieves equivalent performance in terms of latency and throughput, compared with the decoupled admission.

## 6 Conclusion

We have presented a study of power consumption on the novel flit-admission scheme that couples a flit-admission queue with an output physical channel. Our experiments show that this scheme results in both area reduction and power saving. Although our discussions are equally applicable to macro wormhole-switched networks in parallel

computing, the experiments were designed for a NoC that employs a low-dimension topology, deterministic routing, and smaller buffering cost.

From both network performance and power perspective, we believe this coupled admission may serve as a promising flit-admission technique for on-chip network design. Our future work is to optimize network power combining flit admission with flit ejection. As discussed in [6], exploring the flit-ejection design space shows the potential to simplify the switch design with limited performance penalty.

## References

- [1] W. J. Dally. Virtual-channel flow control. *IEEE Transactions on Parallel and Distributed Systems*, 3(2):194–204, March 1992.
- [2] W. J. Dally and B. Towles. Route packets, not wires: On-chip interconnection networks. In *DAC*, 2001.
- [3] J. Hu and R. Marculescu. Exploiting the routing flexibility for energy/performance aware mapping of regular noc architectures. In *Proceedings of Design Automation and Test in Europe*, 2003.
- [4] A. Jantsch, R. Lauter, and A. Vitkovski. Power analysis of link level and end-to-end data protection on networks on chip. In *Proceedings of the IEEE International Symposium on Circuits and Systems*, 2005.
- [5] Z. Lu and A. Jantsch. Flit admission in on-chip wormhole-switched networks with virtual channels. In *Proceedings of International Symposium on System-on-Chip*, November 2004.
- [6] Z. Lu and A. Jantsch. Flit ejection in on-chip wormhole-switched networks with virtual channels. In *Proceedings of the IEEE Norchip Conference*, November 2004.
- [7] M. Millberg, E. Nilsson, R. Thid, and A. Jantsch. Guaranteed bandwidth using looped containers in temporally disjoint networks within the Nostrum network on chip. In *Proceedings of the Design Automation and Test Europe Conference (DATE)*, 2004.
- [8] S. Mukherjee, P. Bannon, S. Lang, and A. S. D. Webb. The alpha 21364 network architecture. *IEEE Micro*, 22(1), Jan.-Feb. 2002.
- [9] D. Pamunuwa, J. Öberg, L.-R. Zheng, M. Millberg, A. Jantsch, and H. Tenhunen. A study on the implementation of 2D mesh based networks on chip in the nanoregime. *Integration - The VLSI Journal*, 38(2):3–17, October 2004.
- [10] C. Patel, S. Chai, S. Yalamanchili, and D. Schimmel. Power constrained design of multiprocessor interconnection networks. In *Proceedings of International Conference on Computer Design*, pages 408–416, October 1997.
- [11] L. S. Peh and W. J. Dally. A delay model for router microarchitectures. *IEEE Micro*, pages 26–34, Jan.-Feb. 2001.
- [12] V. Raghunathan, M. B. Srivastava, and R. K. Gupta. A survey of techniques for energy efficient on-chip communication. In *Proceedings of Design Automation Conference*, June 2003.
- [13] E. Rijpkema et al. Trade offs in the design of a router with both guaranteed and best-effort services for networks on chip. In *Proceedings of Design Automation and Test Conference in Europe*, Mar. 2003.
- [14] L. Shang, L.-S. Peh, and N. K. Jha. Dynamic voltage scaling with links for power optimization of interconnection networks. In *Proceedings of the 9th International Symposium on High-Performance Computer Architecture (HPCA)*, January 2003.
- [15] H.-S. Wang, X. Zhu, L.-S. Peh, and S. Malik. Orion: A power-performance simulator for interconnection networks. In *Proceedings of the 35th International Symposium on Microarchitecture (MICRO)*, November 2002.