

A Multi-Vdd Dynamic Variable-Pipeline On-Chip Router for CMPs

Hiroki Matsutani¹, Yuto Hirata¹, Michihiro Koibuchi²,
Kimiyoshi Usami³, Hiroshi Nakamura⁴, and Hideharu Amano¹

¹Keio University
3-14-1 Hiyoshi, Kohoku-ku, Yokohama, Japan
{matutani,yuto,hunga}@am.ics.keio.ac.jp

²National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo, Japan
koibuchi@nii.ac.jp

³Shibaura Institute of Technology
3-7-5 Toyosu, Kohtoh-ku, Tokyo, Japan
usami@shibaura-it.ac.jp

⁴The University of Tokyo
7-3-1 Hongo, Bunkyo-ku, Tokyo, Japan
nakamura@hal.ipc.i.u-tokyo.ac.jp

Abstract—We propose a multi-voltage (multi-Vdd) variable pipeline router to reduce the power consumption of Network-on-Chips (NoCs) designed for chip multi-processors (CMPs). Our multi-Vdd variable pipeline router adjusts its pipeline depth (i.e., communication latency) and supply voltage level in response to the applied workload. Unlike dynamic voltage and frequency scaling (DVFS) routers, the operating frequency is the same for all routers throughout the CMP; thus, there is no need to synchronize neighboring routers working at different frequencies. In this paper, we implemented the multi-Vdd variable pipeline router, which selects two supply voltage levels and pipeline modes, using a 65nm CMOS process and evaluated it using a full-system CMP simulator. Evaluation results show that although the application performance degraded by 1.0% to 2.1%, the standby power of NoCs reduced by 10.4% to 44.4%.

I. INTRODUCTION

Recently, Network-on-Chips (NoCs) have been used in chip multi-processors (CMPs) [1][2][3][4][5] to connect a number of processors and cache memories on a single chip, instead of traditional bus-based on-chip interconnects that exhibit poor scalability. Fig. 1 illustrates an example of a 16-tile CMP. The chip is divided into sixteen tiles, each of which has a processor (CPU), private L1 data and instruction caches, and a unified L2 cache bank. The L2 cache banks are either private or shared by all tiles. These tiles are interconnected via on-chip routers, and a coherence protocol runs on the NoC.

NoCs are evaluated from various aspects, such as performance and cost. Power consumption, in particular, is becoming more and more important in almost all CMPs. Dynamic voltage and frequency scaling (DVFS) is a primary power saving technique that regulates the operating frequency and supply voltage in response to the applied load. It has been applied to various microprocessors and on-chip routers to reduce their power consumption [3][6].

However, there are certain difficulties with DVFS when there are multiple entities or power domains, in which the supply voltage and clock frequency are controlled individually, in a chip. That is, the operating frequencies of two neighboring power domains must be $1:k$ (k is a positive integer). Otherwise, an asynchronous communication protocol, which introduces a significant overhead, is required between them. Since an NoC typically has a strong communication locality in a chip, router-level fine-grained power management in response to the applied workload is preferred.

In this paper, we propose a low-power router architecture that dynamically adjusts its pipeline depth and supply voltage,

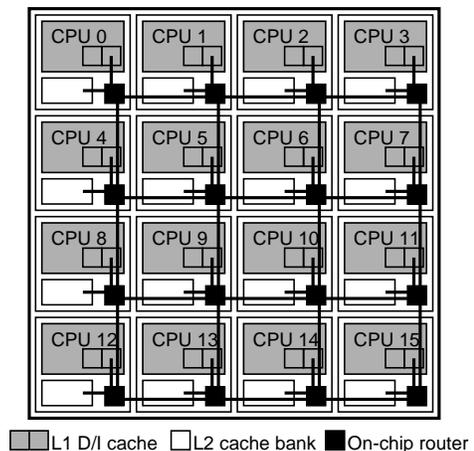


Fig. 1. Example of 16-tile CMP

instead of relying on traditional DVFS techniques that change the frequency of each router. The proposed router offers two pipeline structures: 2-cycle and 3-cycle modes. The 2-cycle mode provides low-latency and high-performance communications, while introducing a large critical path delay, which requires a high supply voltage to satisfy the timing constraints. On the other hand, the 3-cycle mode provides moderate performance, while dividing its critical path into multiple cycles; thus, the supply voltage can be reduced without violating the timing constraints. The 2-cycle mode is used for high workload, while the 3-cycle mode is for low workload.

The rest of this paper is organized as follows. Section II illustrates the multi-Vdd variable pipeline router architecture and its implementation using a 65nm CMOS process technology. Section III proposes the voltage and pipeline reconfiguration policies. Section IV evaluates the proposed router in terms of area, pipeline reconfiguration latency, and overhead energy. It also shows the system-level evaluation results from using a full-system CMP simulator. Section V surveys related work and Section VI concludes this paper.

II. MULTI-VDD VARIABLE PIPELINE ROUTER

Fig. 2 illustrates the multi-Vdd variable-pipeline router architecture. This section introduces its design and implementation using a 65nm CMOS process.

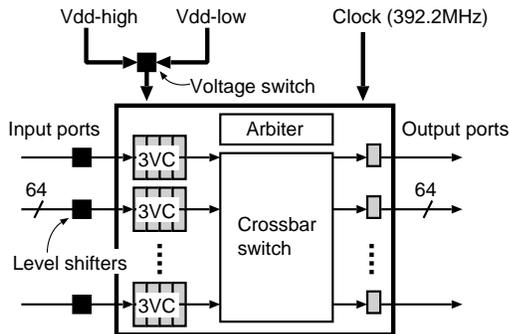


Fig. 2. Multi-Vdd variable pipeline router

TABLE I
SUPPORTED PIPELINE MODES

Mode	Pipeline	Vdd [V]	Freq [MHz]
2-cycle	[RC/VA/SA] [ST,LT]	1.20	392.2
3-cycle	[RC/VA/SA] [ST] [LT]	0.83	392.2

A. Baseline Router

As a baseline router architecture, we assume a 64-bit worm-hole router that has five physical channels. Each input physical channel has three virtual channels (VCs) each of which has a 4-flit input buffer, while each output physical channel has a single 1-flit output buffer. The flit width is 64-bit.

The packet processing in the router can be divided into the following four tasks: routing computation (RC), virtual-channel allocation (VA), switch allocation (SA), and switch traversal (ST). In addition, a single cycle is required to the link traversal (LT) between routers. Including the LT stage, we call this router a 5-cycle router in this paper.

B. Variable Pipeline Mechanism

The multi-Vdd variable pipeline router supports the 2-cycle and 3-cycle modes, as shown in Table I. The notation [X] denotes that task X is performed in a single cycle, [X,Y] denotes that tasks X and Y are performed sequentially within a clock cycle, and [X/Y] denotes a parallel execution.

In the 3-cycle mode, RC, VA, and SA operations are performed in the first cycle, since VA and SA can be performed in parallel ([VA/SA]) using a speculation technique [7] and RC and VA/SA are also performed in parallel ([RC/VA/SA]) using look-ahead routing [8]. ST and LT are performed in the second and third cycles, respectively. In the 2-cycle mode, on the other hand, ST is lumped into LT ([ST,LT]); thus the router pipeline depth is 1, while it exhibits a long critical path delay. These two modes can be switched in a single cycle if there are no flits in the router pipeline.

C. Pipeline Modes

The above-mentioned variable pipeline router was designed with Verilog HDL. Using a Fujitsu 65nm CMOS process, it was synthesized with Synopsys Design Compiler and placed and routed with Synopsys IC Compiler. From the timing analysis, the critical path delay of the 3-cycle mode is 1830.8 psec, while that of the 2-cycle mode is 2549.7 psec when the supply voltage is 1.20V.

The variable pipeline router does not change the operating frequency but changes the pipeline depth and supply voltage in response to the offered workload. Assuming that the variable pipeline router operates at 392.2MHz, the 2-cycle mode requires 1.20V while the 3-cycle mode requires only 0.83V to satisfy the

timing constraints, as shown in Table I. Here, 1.20V and 0.83V are denoted as Vdd-high and Vdd-low, respectively.

The voltage and pipeline reconfiguration is performed as follows.

- **3-cycle \mapsto 2-cycle:** Supply voltage is increased from Vdd-low to Vdd-high. Then pipeline mode is changed from 3-cycle to 2-cycle.
- **2-cycle \mapsto 3-cycle:** Pipeline mode is changed from 2-cycle to 3-cycle. Then supply voltage is reduced from Vdd-high to Vdd-low.

If the pipeline reconfiguration does not follow these rules, a timing violation may be introduced during the reconfiguration, resulting in bit errors. Overhead latency and energy for the reconfiguration is measured based on the circuit-level simulations described in Section IV-A2.

III. PIPELINE RECONFIGURATION

In this section, we first discuss how to reduce the power consumption of NoCs by using the multi-Vdd variable pipeline router. Then we propose a look-ahead power management method that can minimize both the performance penalty and standby power of NoCs.

A. Pipeline Reconfiguration Policy

Typically, DVFS provides lower performance with Vdd-low when the workload is low, while it provides higher performance when the workload is high. However, we believe this concept cannot be directly applied to on-chip routers for CMPs because of the following reasons.

- Low-power mode (i.e., low-performance mode) of on-chip routers increases the communication latency between processors and cache memories, which significantly increases the application execution time.
- Power consumption of processors is typically larger than that of on-chip routers; thus, increased application execution time by the low-power mode may adversely increase the total power consumption of the CMP.

Therefore, pipeline reconfiguration should be performed to reduce the power consumption of on-chip routers as long as the communication latency (i.e., application execution time) is not affected.

Consequently, the pipeline reconfiguration of the multi-Vdd variable pipeline router is performed as follows.

- The 2-cycle mode with Vdd-high is used as often as possible for packet transfers in order to reduce the performance penalty.
- Otherwise, the 3-cycle mode with Vdd-low is used to minimize the standby power consumption of NoCs.

B. Look-Ahead Power Management

The voltage transition latency from Vdd-low to Vdd-high requires two cycles for 392.2MHz operation, as discussed in the circuit level evaluations in Section IV-A. To process packets with the low-latency 2-cycle mode, the voltage transition from Vdd-low to Vdd-high should be started in 2-cycle ahead before packets actually reach the router. To minimize the standby power of the router, the supply voltage should be reduced to Vdd-low after the packets leave.

For deterministic routing (e.g., XY routing), look-ahead routing computation [8] can be used to detect the packet arrivals two hops away [9]. That is, assuming a packet moves from Router A to Router C via Router B, Router C can detect the packet arrival when the routing computation at Router A is completed.

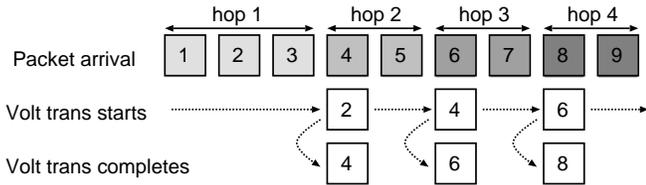


Fig. 3. Look-ahead based voltage transition

Only the first hop router cannot detect the packet arrivals preliminarily, because no prior routers that notify the packet arrival exist for the first hop. This is a serious problem in the power-gating router [9], since no packets cannot be forwarded before the wakeup procedure of the router components is completed, which significantly degrades application performance. In our multi-Vdd variable pipeline router, on the other hand, the performance overhead of the first hop problem is small. In fact, if the voltage reconfiguration from Vdd-low to Vdd-high is not completed before packet arrivals, the packet is processed with the 3-cycle mode; thus, the performance overhead is only a single cycle in this case. Consequently, packets are processed with the 3-cycle Vdd-low mode only in the first hop, while they are processed with the 2-cycle Vdd-high mode after the first hop.

Fig. 3 illustrates the look-ahead based voltage transition from Vdd-low to Vdd-high. The numbers in the boxes denote the clock cycle of an event, such as the packet arrival and voltage transition start and end times. For example, a packet reaches the first hop (hop-1) router at the first cycle (cycle-1). Similarly, it reaches hop-2, hop-3, and hop-4 routers in cycle-4, cycle-6, and cycle-8, respectively. As shown in this figure, only the hop-1 router forwards a packet with the 3-cycle mode and the others forward it with the 2-cycle mode. Since the look-ahead routing computation detects the next hop router, the low-to-high voltage transition of the hop-2 router is triggered at cycle-2 and is completed at cycle-4; thus, the hop-2 router can forward it with the 2-cycle mode. Similarly, the low-to-high voltage transition of the hop-3 router is triggered at cycle-4 and is completed at cycle-6, so it can forward the packet with the 2-cycle mode.

IV. EVALUATIONS

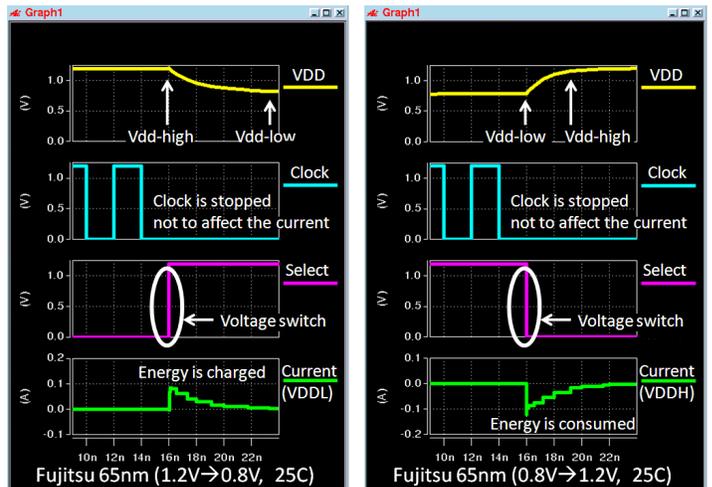
First, the proposed router is evaluated in terms of area, reconfiguration latency, and energy. Then, these circuit parameters are fed to a full-system CMP simulator, and the proposed and baseline routers are evaluated in terms of application performance and power consumption.

A. Circuit-Level Evaluations

The multi-Vdd variable pipeline router was implemented with a Fujitsu 65nm CMOS process, as described in Section II. From the GDS file of the router, its SPICE netlist is extracted using Cadence QRC Extraction. Then, the pipeline reconfiguration latency and energy of the router are measured using Synopsys HSPICE.

Fig. 4 shows the measured waveforms during the voltage and pipeline reconfiguration. Fig. 4(a) shows the transition from Vdd-high (1.20V) to Vdd-low (0.8V), while Fig. 4(b) shows that from Vdd-low to Vdd-high. In each graph, the first (top) waveform shows the supply voltage (Vdd) of the router. The second and third waveforms show the clock signal and the select signal that selects the supply voltage level (0 for Vdd-high and 1 for Vdd-low). The fourth (bottom) waveform shows the current. To accurately measure the overhead current induced by the voltage switching, the clock is stopped during the simulation.

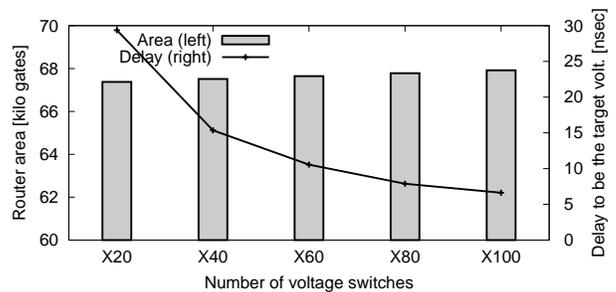
A certain latency is required so that Vdd reaches the target voltage level after the power source is switched (i.e., select signal



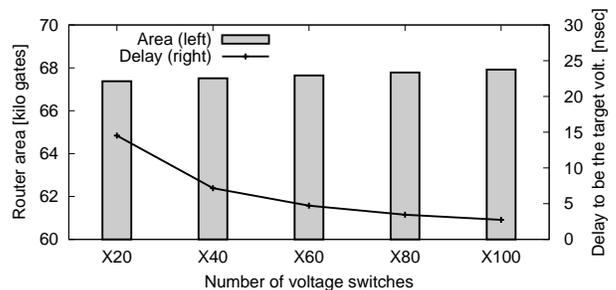
(a) Vdd-high to Vdd-low transition

(b) Vdd-low to Vdd-high transition

Fig. 4. Measured waveforms of Vdd, clock, select, and current of the proposed router



(a) Vdd-high to Vdd-low transition



(b) Vdd-low to Vdd-high transition

Fig. 5. Voltage transition latency vs. number of voltage switch cells (router hardware amount)

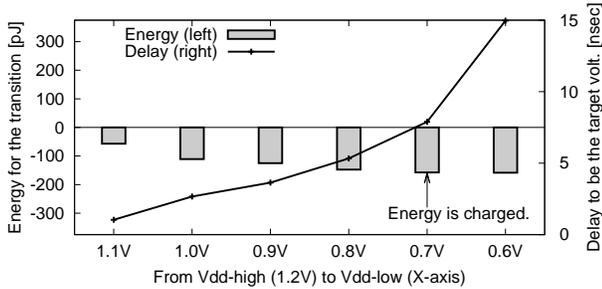
is changed), as shown in the graphs. This voltage-switching latency will affect the router performance. In particular, the transition latency from the 3-cycle mode with Vdd-low to the 2-cycle mode with Vdd-high should be minimized to reduce communication latency, which significantly affects application performance in the cases of shared memory CMPs. Fortunately, the latency for the low-to-high voltage transition is shorter than that for high-to-low, as shown in Fig. 4.

1) *Hardware Amount:* In the multi-Vdd variable pipeline router, voltage switch cells are inserted between the router and the power sources (i.e., Vdd-high and Vdd-low), as illustrated in Fig. 2. In addition, level shifter cells are inserted to all input (or output) ports of the router to convert the Vdd-low control signals to Vdd-high.

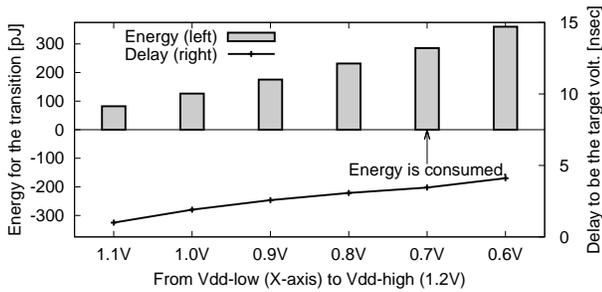
As the number (size) of voltage switch cells increases, the voltage-switching latency is shortened, while the area overhead of the voltage switch cells increases. Fig. 5 shows the voltage

TABLE II
BREAKDOWN OF ROUTER HARDWARE AMOUNT [KILO GATES]

	Router	Level shifter	Volt. switch	Total
Baseline	59.41	0.00	0.00	59.41
Proposed	63.13	4.11	0.54	67.78



(a) Vdd-high to Vdd-low transition



(b) Vdd-low to Vdd-high transition

Fig. 6. Voltage transition delay and energy vs. Vdd-low voltage

transition latency vs. number of voltage switch cells (total router area). Here, the voltage transition latency is defined as the required time for Vdd to reach $\pm 0.05V$ of the target voltage level after the voltage-select signal is changed. In Fig. 5, the x-axis shows the number of voltage switch cells ranging from 20 to 100 cells. A bar graph shows the gate count of the router (including the voltage switch and level shifter cells) [kilo gates], while a line graph shows the voltage transition latency [nsec]. As shown in this figure, more voltage switch cells reduce transition latency. However, the area overhead of the voltage switch cells is relatively small compared to the total router area. Thus, in this design, 80 voltage switch cells (X80) are inserted into each router by taking into account the latency and area overheads.

Table II lists the area breakdown of the baseline 3-cycle router and the proposed multi-Vdd variable pipeline router. 80 voltage switch cells (X80) are inserted to the proposed router. In addition, a level shifter cell is inserted into every input port of the proposed router. As a result, the proposed router is larger than the baseline router by 14.1%. Most of the area overhead comes from the level shifter cells and the control logic for dynamic reconfiguration of the router pipeline depth.

2) *Reconfiguration Latency*: Fig. 6 shows the voltage transition delay and energy vs. Vdd-low level. The line graph shows the high-to-low and low-to-high voltage transition latencies when the Vdd-low level ranges from 0.60V to 1.10V, while Vdd-high is fixed at 1.20V. When Vdd-low is 0.80V, the high-to-low latency is 5.3nsec, while low-to-high latency is 3.1nsec. That is, for the transition from the 3-cycle mode to 2-cycle mode, the router first increases the supply voltage and waits for 3.1nsec (2-cycle for 392.2MHz operation). After that, the pipeline reconfiguration from 3-cycle to 2-cycle is performed.

3) *Reconfiguration Overhead Energy*: A certain amount of dynamic energy is consumed when the supply voltage is changed

TABLE III
SIMULATION PARAMETERS (CMP)

Processor	UltraSPARC-III
L1 I-cache size	16 KB (line:64B)
L1 D-cache size	16 KB (line:64B)
# of processors	16
L1 cache response	1 cycle
L2 cache size	256 KB (assoc:4)
# of L2 cache banks	16
L2 cache response	6 cycles
Memory size	4 GB
Memory response	160 (± 2) cycles
# of memory ports	16

TABLE IV
SIMULATION PARAMETERS (NOC)

Topology	4 \times 4 mesh
Routing	dimension-order
# of VCs	3
Buffer size	4 flits
Flit size	64 bits
Control packet	1 flits
Data packet	9 flits

from Vdd-low to Vdd-high. In Fig. 6, the bar graphs show the overhead energy for the high-to-low and low-to-high voltage transitions. The energy consumption is completely different between the high-to-low and low-to-high transitions. In the low-to-high transition, current flows from the Vdd-high power source to the router; thus, overhead energy is consumed. In the high-to-low transition, on the other hand, current flows from the router to the Vdd-low power source. This current is charged at the capacitance of the Vdd-low power source, and it is consumed gradually by the router; thus, the overhead energy becomes negative. When Vdd-high and Vdd-low are 1.20V and 0.80V, respectively, the energy overhead for a low-to-high transition is 231.4pJ, while that for a high-to-low transition is -147.6pJ. That is, a pair of voltage transitions (i.e., low-to-high + high-to-low) consumes 83.8pJ on average.

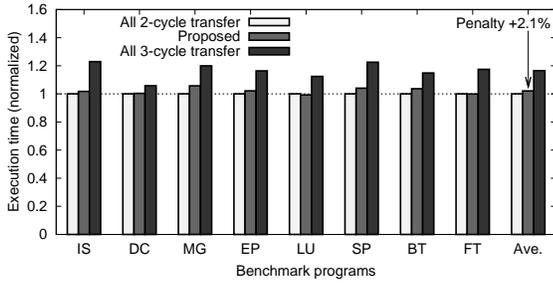
4) *Break-Even Time Analysis*: Energy reduction by remaining at the Vdd-low level should be larger than the energy overhead. Otherwise, total power consumption adversely increases due to the energy overhead. Here, break-even time (BET) is defined as the minimum time period that can compensate for the energy overhead (e.g., 83.8pJ) by remaining at the Vdd-low level. If the router remains at Vdd-low for a shorter time than BET, the energy overhead is larger than the benefit and total power consumption adversely increases.

The placed-and-routed design of the proposed router was simulated at 392.2MHz. From power analysis based on the switching activity information, the standby power consumption of the 2-cycle mode with Vdd-high is 2.78mW, while that of the 3-cycle mode with Vdd-low is 1.33mW since power consumption is proportional to the square of Vdd. That is, remaining in the 3-cycle mode with Vdd-low for 1sec can save 1.45mJ from the standby energy consumption. To compensate for the energy overhead of 83.8pJ, the router must remain in the 3-cycle mode at least 23 cycles. Any shorter stay at Vdd-low increases the total standby power consumption.

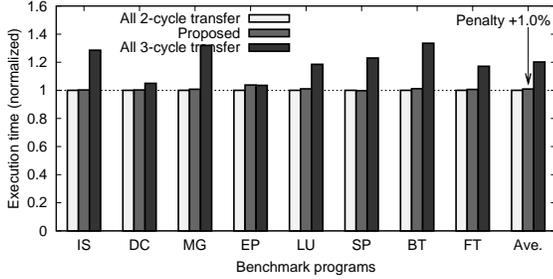
B. System-Level Evaluations

The obtained circuit parameters discussed in the previous subsections were fed to a full-system CMP simulator to evaluate the application execution time and the standby power reduction by taking into account the energy overhead.

1) *Simulation Environments*: An NoC used in the 16-tile CMP illustrated in Fig. 1 was simulated. Regarding the L2 cache organization, the multi-Vdd variable pipeline router architecture was applied to the following two CMP configurations.



(a) 16-tile CMP (shared L2 cache)



(b) 16-tile CMP (private L2 cache)

Fig. 7. Execution time of NPB applications (1.0 denotes execution time with ideal 2-cycle routers)

- **Shared L2 cache banks:** L2 cache banks in all tiles are shared by all tiles. They form a single shared L2 cache in a chip.
- **Private L2 cache banks:** Each L2 cache bank is used as a private L2 cache inside the same tile. It cannot be accessed by the other tiles.

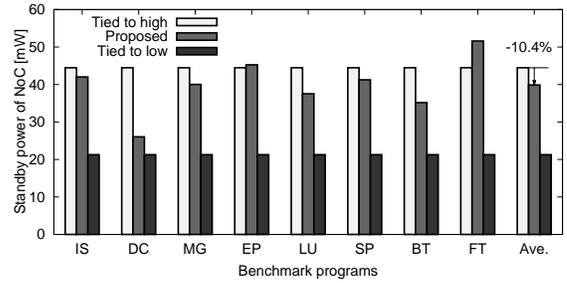
A directory-based coherence protocol is used to maintain the cache coherence. To avoid the end-to-end protocol (i.e., request and reply) deadlocks, three virtual channels (VCs) are used for different message classes. The traffic amount on the shared L2 CMP is larger than that on the private L2 CMP, since L2 cache banks in all tiles are shared by all tiles via NoC in the case of the shared L2 CMP.

To simulate the above-mentioned CMP, we use a full-system multi-processor simulator: GEMS [10] and Wind River Simics [11]. We modified a detailed network model of GEMS, called Garnet [12], to accurately simulate the proposed multi-Vdd variable pipeline router. Table III lists the processor and memory system parameters, and Table IV lists the on-chip router and network parameters. To clearly show the worst-case performance degradation induced by the voltage and pipeline reconfiguration, we assume relatively rich main memory bandwidth, namely one one memory port for each tile.

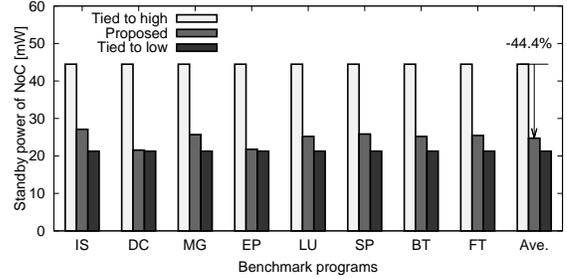
To evaluate the application performance on CMPs with the proposed router, we use eight parallel programs (IS, DC, MG, EP, LU, SP, BT, FT) from the OpenMP implementation of NAS Parallel Benchmarks (NPB) [13]. Sun Solaris 9 operating system is running on the CMPs. These benchmark programs were compiled by Sun Studio 12 and are executed on Solaris 9. The number of threads was set to sixteen for the 16-tile CMP.

2) *Application Execution Time:* The multi-Vdd variable pipeline router is applied to the above-mentioned NoC of the CMPs, and the eight NPB programs are performed on the NoC. The following three router modes are compared in terms of the execution time of these applications.

- **All 2-cycle transfer:** The router mode is fixed in the 2-cycle Vdd-high mode (high-performance but high-power).



(a) 16-tile CMP (shared L2 cache)



(b) 16-tile CMP (private L2 cache)

Fig. 8. Standby power of NoC (reconfiguration power is included)

- **Proposed:** The router mode is changed to the 2-cycle Vdd-high mode whenever a packet approaches the router. Otherwise, it remains at the 3-cycle Vdd-low mode (high-performance and low-power).
- **All 3-cycle transfer:** The router mode is fixed at the 3-cycle Vdd-low mode (low-power but low-performance).

Fig. 7 shows the application execution times of these router modes. Fig. 7(a) shows the results on the shared L2 CMP, and Fig. 7(b) shows those on the private one. The results are normalized so that the execution time of the All 2-cycle transfer mode is 1.0. Although the proposed power management method degrades the application performance by 1.0%-2.1% compared to the All 2-cycle transfer mode that cannot reduce the power, performance degradation is very small. This slight performance degradation comes from the 3-cycle transfers at the first hop of the packet transfers, because the look-ahead mode control cannot detect packet arrival and forces the 3-cycle transfer only at the first hop, as mentioned in Section III-B. The next section explains how much standby power can be reduced with the proposed power management compared to the All 2-cycle transfer mode at Vdd-high.

3) *Standby Power Reduction:* The proposed variable pipeline router is operated in the 2-cycle Vdd-high mode for packet transfers to prevent performance degradation, while it remains at the 3-cycle Vdd-low mode to minimize power consumption when no packets are processed. Thus, the packet processing power is constant since packet processing is performed with the 2-cycle Vdd-high mode except the first hop. In this section, therefore, we focus on standby power reduction including the overhead energy.

As mentioned in Section IV-A3, a pair of voltage transitions (i.e., low-to-high + high-to-low) consumes 83.8pJ on average. The standby power consumption of the 2-cycle Vdd-high mode is 2.78mW, while that of the 3-cycle Vdd-low mode is 1.33mV. To evaluate the standby-power including overhead energy, all reconfiguration events are recorded during the full-system simulation of each application. By multiplying the frequency of the mode reconfiguration by its energy overhead, we estimated the reconfiguration power in addition to standby power consumption.

Fig. 8 shows the standby power of the NoCs. In the graphs, “Tied to Vdd-high” corresponds to “All 2-cycle transfer”, and “Tied to Vdd-low” corresponds to “All 3-cycle transfer”. The standby power with the proposed power management includes the reconfiguration power for each application. Fig. 8(a) shows the results on the shared L2 CMP, and Fig. 8(b) shows those on the private one. The traffic amount in the private L2 case is smaller than that in the shared one; thus, the reconfigurations in the private one are infrequent compared to the shared one. As a result, in the private L2 case, the standby power is reduced by 44.4% on average even when the energy overhead is included. In the shared L2 case, the standby power is reduced by 10.4%. Applications with high traffic volume (e.g., FT) introduce frequent reconfigurations compared to BET, and thus it sometimes overwhelms the benefit of dynamic voltage reduction. A more sophisticated power management policy that takes into account BET may be able to prevent unnecessary reconfigurations, while router design complexity is increased.

V. RELATED WORK

DVFS has been applied to various microprocessors and on-chip routers [3][6]. Regarding the pipeline stage optimization, [14] proposes the time-stealing technique for on-chip networks. It improves the router operating frequency by exploiting the timing imbalance between router pipeline stages. That is, a router can operate at a higher frequency which is determined by average delay of all pipeline stages, not the slowest pipeline stage. In addition, some techniques that optimize the pipeline structure in response to the workload have been developed for microprocessors [15] and on-chip routers [16]. For example, router pipeline structure, supply voltage, and operating frequency are changed in response to the workload in [16].

Since a NoC typically has a strong communication locality in a chip, router-level fine-grained power management in response to the applied workload is more efficient. However, these previous approaches are not suited to fine-grained power management of NoCs, because these approaches adjust the operating frequency and supply voltage for each power domain. In this case, the operating frequencies of two neighboring power domains must be the same or divisible by another. Otherwise, an asynchronous communication protocol is required. Therefore, unlike these previous approaches, our multi-Vdd variable pipeline router dynamically adjusts its pipeline depth and supply voltage, instead of relying on traditional DVFS techniques that change the frequency of each router.

Runtime power gating [9] is another approach to reduce the power consumption of routers. However, it suffers a wakeup latency to activate the sleeping components; thus, a sophisticated wakeup mechanism is required to mitigate the wakeup latency.

VI. CONCLUSIONS

In this paper, the multi-Vdd variable pipeline router was designed and implemented with a 65nm CMOS process. We evaluated it in terms of area overhead, latency, and energy overhead for pipeline reconfiguration. We also estimated BET of a mode reconfiguration required to compensate for the reconfiguration energy overhead. As a result, area overhead for the pipeline reconfiguration controller, voltage switch cells, and level shifter cells is 14.1%. Voltage transition latency is 3.1nsec and 5.3nsec for low-to-high and high-to-low transitions, respectively. A pair of voltage transitions consumes 83.8pJ. BET to compensate for this energy overhead is 23 cycles for 392.2MHz operation.

We also proposed a look-ahead based power management that detects packet arrivals and pre-configures itself to the 2-cycle Vdd-high mode for minimizing both performance penalty

and power consumption. The proposed power management was applied to the multi-Vdd variable pipeline router and evaluated using real 16 threads parallel applications on a full-system CMP simulator. The simulation results showed that, although the application performance is slightly affected (1.0% to 2.1%) compared to the All 2-cycle transfer mode at Vdd-high, the standby power is decreased significantly (10.4% to 44.4%) compared to the Tied to Vdd-high mode. For applications with high traffic volume, however, the reconfiguration power overwhelms the power reduction by the low-power mode. Exploring more sophisticated power management policies that take into account BET while keeping router design simple is our future work.

ACKNOWLEDGEMENTS

This research was performed by the authors for STARC as part of the Japanese Ministry of Economy, Trade and Industry sponsored “Next-Generation Circuit Architecture Technical Development” program. The authors thank to VLSI Design and Education Center and Japan Science and Technology Agency CREST for their support. This work was also supported by Grant-in-Aid for Research Activity Start-up #23800053.

REFERENCES

- [1] T. W. Ainsworth and T. M. Pinkston, “Characterizing the Cell EIB On-Chip Network,” *IEEE Micro*, vol. 27, no. 5, pp. 6–14, Sep. 2007.
- [2] B. M. Beckmann and D. A. Wood, “Managing Wire Delay in Large Chip-Multiprocessor Caches,” in *Proceedings of the International Symposium on Microarchitecture (MICRO’04)*, Dec. 2004, pp. 319–330.
- [3] J. Howard *et al.*, “A 48-Core 1A-32 Message-Passing Processor with DVFS in 45nm CMOS,” in *Proceedings of the International Solid-State Circuits Conference (ISSCC’10)*, Feb. 2010, pp. 108–109.
- [4] A. S. Leon, K. W. Tam, J. L. Shin, D. Weisner, and F. Schumacher, “A Power-Efficient High-Throughput 32-Thread SPARC Processor,” *IEEE Journal of Solid-State Circuits*, vol. 42, no. 1, pp. 7–16, Jan. 2007.
- [5] D. Wentzlaff, P. Griffin, H. Hoffmann, L. Bao, B. Edwards, C. Ramey, M. Mattina, C.-C. Miao, John F. Brown III, and A. Agarwal, “On-Chip Interconnection Architecture of the Tile Processor,” *IEEE Micro*, vol. 27, no. 5, pp. 15–31, Sep. 2007.
- [6] E. Beigne, F. Clermidy, H. Lhermet, S. Miermont, Y. Thonnart, X.-T. Tran, A. Valentian, D. Varreau, P. Vivet, X. Popon, and H. Lebreron, “An Asynchronous Power Aware and Adaptive NoC Based Circuit,” *IEEE Journal of Solid-State Circuits*, vol. 44, no. 4, pp. 1167–1177, Apr. 2009.
- [7] L.-S. Peh and W. J. Dally, “A Delay Model and Speculative Architecture for Pipelined Routers,” in *Proceedings of the International Symposium on High-Performance Computer Architecture (HPCA’01)*, Jan. 2001, pp. 255–266.
- [8] M. Galles, “Scalable Pipelined Interconnect for Distributed Endpoint Routing: The SGI SPIDER Chip,” in *Proceedings of the IEEE Symposium on High-Performance Interconnects (HOTI’96)*, Aug. 1996, pp. 141–146.
- [9] H. Matsutani, M. Koibuchi, D. Ikebuchi, K. Usami, H. Nakamura, and H. Amano, “Performance, Area, and Power Evaluations of Ultrafine-Grained Run-Time Power-Gating Routers for CMPs,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits*, vol. 30, no. 4, pp. 520–533, Apr. 2011.
- [10] M. M. K. Martin, D. J. Sorin, B. M. Beckmann, M. R. Marty, M. Xu, A. R. Alameldeen, K. E. Moore, M. D. Hill, and D. A. Wood, “Multifacet General Execution-driven Multiprocessor Simulator (GEMS) Toolset,” *ACM SIGARCH Computer Architecture News (CAN’05)*, vol. 33, no. 4, pp. 92–99, Nov. 2005.
- [11] P. S. Magnusson, M. Christensson, J. Eskilson, D. Forsgren, G. Hallberg, J. Hogberg, F. Larsson, A. Moestedt, and B. Werner, “Simics: A Full System Simulation Platform,” *IEEE Computer*, vol. 35, no. 2, pp. 50–58, Feb. 2002.
- [12] N. Agarwal, L.-S. Peh, and N. Jha, “Garnet: A Detailed Interconnection Network Model inside a Full-system Simulation Framework,” Princeton University, Tech. Rep. CE-P08-001, 2008.
- [13] H. Jin, M. Frumkin, and J. Yan, “The OpenMP Implementation of NAS Parallel Benchmarks and Its Performane,” in *NAS Technical Report NAS-99-011*, Oct. 1999.
- [14] A. K. Mishra, R. Das, S. Eeachempati, R. Iyer, N. Vijaykrishnan, and C. R. Das, “A Case for Dynamic Frequency Tuning in On-Chip Networks,” in *Proceedings of the International Symposium on Microarchitecture (MICRO’09)*, Dec. 2009.
- [15] H. Shimada, H. Ando, and T. Shimada, “Pipeline Stage Unification: Low-Energy Consumption Technique for Future Mobile Processors,” in *Proceedings of International Symposium on Low Power Electronics and Design (ISLPED’03)*, Aug. 2003, pp. 326–329.
- [16] Y. Hirata, H. Matsutani, M. Koibuchi, and H. Amano, “A Variable-pipeline On-chip Router Optimized to Traffic Pattern,” in *Proceedings of the International Workshop on Network on Chip Architectures (NoCArc’10)*, Dec. 2010, pp. 57–62.