Dynamic Voltage Scaling for Fully Asynchronous NoCs Using FIFO Threshold Levels

Abbas Rahimi, Mostafa E. Salehi, Siamak Mohammadi, Sied Mehdi Fakhraie School of Electrical and Computer Engineering University of Tehran, Tehran 14395-515, Iran

{ab.rahimi, smohammadi}@ece.ut.ac.ir {mersali, fakhraie}@ut.ac.ir

Abstract— In this paper, we propose a dynamic voltage scaling (DVS) policy for a fully asynchronous NoC suitable for low-power yet high-performance architectures. The DVS policy is a FIFO-adaptive DVS, which uses two FIFO threshold levels for decision. It judiciously adjusts switch voltage among only three voltage modes. The introduced architecture is simulated in 90nm CMOS technology with accurate Spice simulations. Experimental results show that the FIFO-adaptive DVS not only lowers the implementation cost, but also achieves another 31% energy-delay saving compared to the DVS policy based on link utilization, in a 90% saturated network.

I. INTRODUCTION

Technology scaling and the increasing device integration levels make power dissipation and on-chip communication as two major factors in the highperformance multiprocessor systems-on-chip (SoCs). Onchip communication is becoming increasingly important when SoCs grow in complexity and size [1]. Furthermore, power dissipation has emerged as the main design constraint in today complex SoCs, limiting performance, battery life and reliability. Network-on-chip (NoCs) [2], constitute a new design paradigm for scalable, highthroughput on-chip communication in SoCs with billions of transistors, and offers a perfect platform for power management. SPIN, a micro-network, attempts to solve the bandwidth bottleneck in SoCs interconnecting a large number of IP cores via NoCs [3][4]. A large number of researches on the synchronous NoC architectures have been performed such as AETHEREAL [5], XPIPES [6], and NOSTRUM [7].

Dynamic voltage and frequency scaling (DVFS) techniques, one of the most successful run-time techniques for improving power efficiency, are widely used for optimizing power in synchronous domains [8][9]. Most of these dynamic and static power saving techniques are related to scaling the voltage supply level which affects power consumption quadraticaly.

The dynamic power can be efficiently controlled by clock gating at both RT and architecture levels. On the other hand, the asynchronous logic scheme offers both RTL and architectural clock gating inherently without the need of any extra software [10]. Asynchronous circuits automatically switch to standby state when they are inactive, and have shown their interesting dynamic power savings, due to their unclocked nature [11]. As an alternative solution for NoC design, the MANGO clockless NoC [12] is one of the first asynchronous NoCs. ASPIN (asynchronous scalable packet-switching integrated network) [13] is another asynchronous micronetwork which is the asynchronous implementation of DSPIN (scalable distributed packet-switching integrated network) [14]. These two implementations are systematically compared in [15]. The other proposed asynchronous NoCs are QNOC [14], and ANOC [15].

Globally asynchronous locally synchronous (GALS) [16] paradigm merges the benefits of both synchronous and asynchronous designs, it is being widely investigated as a viable alternative to purely synchronous designs [17][18]. Better power efficiency is achieved in the GALS system, as it offers a natural way to operate each domain at different frequency and voltage, which facilitates the application of DVFS independently to different parts of circuit [19][20]. To enable GALS systems with multiple clock domains, including DVFS scaling per each synchronous module, the network should be implemented as an asynchronous circuit [21][22].

The rest of the paper is organized as follows: In Section 2, the most relevant recent researches are reviewed. In Section 3, the traffic model is described for a network containing a homogeneous 5x5 set of clusters. In Section 4, A FIFO-adaptive DVS policy is presented in details, including exploration of the threshold levels of FIFO and the three recommended voltage modes. Finally, Section 5 concludes the paper.

II. RELATED WORK

E. Beigné et al. [23] propose a dynamic voltage and frequency scaling policy for IP units integrated within a GALS NoC, but their policy is only applied for all IPs within an SoC, and ignores significant effects of the power dissipation of links and switches. Li Shang et al.

IFFF



[24] use a history-based dynamic voltage scaling policy for links, where the frequency and voltage of the links are dynamically adjusted to minimize power consumption. This work only targets the dynamic power optimization of the interconnection networks which realizes 4.6 times power savings on average at the expense of 15.2% increase in the average latency. S. E. Lee et al. [25] present a variable frequency link for a power-aware interconnection network, and apply a dynamic frequency scaling (DFS) policy which adjusts link frequency based on link utilization parameter.

In our previous work [26], the energy/throughput trade-off was analyzed on a GALS NoC which is based on a fully asynchronous NoC, ASPIN [13], using sync-toasync and async-to-sync interfaces [27] to connect synchronous IP cores to the asynchronous network. Our experimental results show that although DFS techniques can improve power consumption in synchronous circuits, interval scaling and consequently throughput scaling in not recommended for energy saving in the fully asynchronous NoCs, and the best energy-delay (ED) [28] saving is achieved in high throughput regions. On the other hand, a DVS technique is able to save energy up to 40% at the expense of 13% throughput degradation, while a throughput scaling technique only achieves 0.6% energy saving with the same amount of throughput degradation. So these results, as the first study in this region, limit the range of throughput scaling and also limit voltage scaling between 1.0v to 0.75v as the optimum ranges for a DVS scheme.

In this paper we propose two dynamic voltage scaling policies for the GALS NoC architecture based on the previous optimum voltage scaling ranges. First, based on the link utilization parameter [25], a history-based [24] DVS policy is presented. Due to the limitations in implementing on-chip inductors [29], the proposed history-based DVS uses few number of voltage modes. Second, the FIFO-adaptive DVS policy is presented which uses only three voltage modes, and achieves considerable amount of energy saving at the expense of negligible throughput degradation.

III. TRAFFIC MODEL

To evaluate the GALS NoC performance, power, and saturation thresholds (the most important parameters), we have focused on a network containing a homogeneous 5x5 set of clusters. Details of the asynchronous routers and GALS NoC architectures are provided in [26]. Synchronous IP cores transmit and receive data to/from the asynchronous router through sync-to-async and asyncto-sync interfaces. The IPs which are connected to the local input port are used for generating traffic. Each IP consists of two parts, traffic generator (TG) and network analyzer (NA). The TG is connected to router's local input port and is used to model the uniform type of traffic [30], and inject packets to the network. The NA is connected to router's local output ports and consumes the generated traffic and check the delivery of packets. If too many IPs are generating traffic simultaneously, the network would be saturated. The saturation occurs when the traffic generated by each IP reaches a saturation threshold—that is, when the average packet latency rises exponentially to an infinite value. In our traffic model each TG generates 50 packets with the packet length of 8 flits and sends them to the other NAs. To account for network contention and to get a meaningful latency measurement, we have time-stamped the packets and posted them in FIFO buffers located in each TG.

We have measured the average packet latency as the time between the departure time in the source node and the arrival time in the destination node. The curve in Figure 1 depicts average packet latencies, at voltage 1.0v, versus the generated traffic by an IP. The network saturates in loads higher than 176 GFlits/s. In other words, if the IPs flit injection rate exceeds this rate the flits will not be delivered. Similarly, in 0.75v, the network saturates in loads higher than 144 GFlits/s.

According to our results, in 1.0v, the network is 100% saturated when the injection rate reaches 176 GFlits/s. Consequently the injection rates of 144 GFlits/s, 152 GFlits/s, and 160 GFlits/s saturate 82%, 86%, and 90% of the network respectively, and are used for our simulation results.



Figure 1. Packet latency versus different loads of the network in 1.0v.

IV. PROPOSED FIFO-ADAPTIVE DVS

In this new DVS policy the FIFO level is used to predict the upcoming workload instead of history-based [24] and link utilization parameter [25]. FIFO level is a good metric for knowing how many packets will traverse a switch and consequently set the voltage to the optimum value. To estimate the upcoming workload, we use the level of north, south, east, west, and local FIFOs, which is an indicator of traffic through a switch. Low FIFO occupancy level indicates low traffic intensity in a switch caused by light workloads in the incoming ports. Conversely, high FIFO occupancy level implies that higher voltages are required to pass the incoming flits to the destination ports.

To predict upcoming workloads in [25], the link utilization is used which is measured by sampling a link at a given time during a predefined period (T). Since this metric is evaluated based on fixed interval periods, it requires a clock for synchronization which is not perfectly matched with our fully asynchronous NoC. Furthermore, the link utilization requires a counter and other logics to count the number of input port requests. In addition to the link utilization component, the history-based method needs additional hardware to compute $\Psi(n)$ [31]. Therefore, in FIFO-adaptive DVS, we have monitored the traffic intensity by FIFO occupancy level to omit both hardware overhead cost and the clock signal from our asynchronous circuits.

The FIFO occupancy level of each port is an indicator of the traffic on that port- the FIFO depth is 8. Therefore, to filter out transient fluctuations from the input ports, we use the sum of FIFO occupancy levels of all input ports as the traffic indicator of each switch. To reduce the overhead of DC to DC convertors and on-chip inductors, the FIFO-adaptive DVS policy scales the operating voltage among the recommended voltage modes for all part of the switch, including five input ports and five output ports. This leads to a better decisions and also reduces the DC to DC convertors and other hardware components overhead and hence, facilitates its implementation.

A. Recommended Threshold Levels of the FIFO

The level of the proposed FIFO is monitored during a simulation with the load of 152GFlits/s (i.e. network saturated at 86%) versus different operating voltages and the results are summarized in Figure 2. As the results show, lower operating voltages lead to higher FIFO levels, and hence increase the probability of the network saturation. For example, when vdd is equal to 0.75v during 30% of the simulation time the FIFO contains 10 flits, while during 8% of the simulation time it contains the same amount of flits at vdd = 1.0v. This figure also shows a suitable range for the threshold levels of the FIFO between 10 and 24, because there is a very low probability that the number of flits in the FIFO exceeds 24. So this range can be used for selecting appropriate

voltage by the DVS policy.



Figure 2. Observed FIFO level during simulation versus different voltages.

To have optimum energy dissipation, the FIFOadaptive DVS algorithm should dynamically scale the operating voltage of the asynchronous switch between 0.75v and 1.0v [26], and improve the energy saving with negligible performance degradation. To reduce the number of voltage modes generated by regulators, the switch operating voltage should be selected among the three recommended voltage modes called high voltage (V_h) , medium voltage (V_m) , and low voltage (V_l) . Since we have proposed the FIFO occupancy level as the traffic intensity indicator, we need two FIFO levels called low threshold (Th_l) and high threshold (Th_h) to decide when to switch between the voltage modes. The decision of the FIFO-adaptive DVS is based on three simple assumptions:

If (FIFO_level < Th_l) set V_{switch} to V_l

If (
$$Th_l \leq FIFO$$
 level $\leq Th_h$) set V_{switch} to V_m

If ($Th_h \leq FIFO_level$) set V_{switch} to V_h

So, we have to find two suitable threshold levels for FIFO among the available range to have the best energy saving with least performance degradation. In [26], we have shown that throughput degradation does not improve the energy saving in asynchronous circuits. Therefore, we try to achieve the highest throughput with the least required voltage. High throughput equals low flit latency in each switch or in other words, low FIFO occupancy level. Therefore, we improve the throughput by minimizing the FIFO occupancy level and expect to have the best energy saving. The results will validate our assumption.

We have equaled V_h by the highest supported voltage (i.e. 1.0v), and V_1 by the lowest supported voltage (i.e 0.75v). For the sake of finding the threshold values, 0.85v is selected for V_m . Figure 4 shows the FIFO occupancy level during simulation for different sets of threshold values.



Figure 3. The threshold values (14,18) provides a lower occupancy FIFO relative to the threshold values.

These sets of threshold values are selected among the suitable range in Figure 2. As results show, when the (Th_I, Th_h) is equal to (14, 18), we have the lowest FIFO occupancy level and hence, the highest throughput. Therefore, we propose 14 and 18 as the optimum values for Th_I and Th_h , respectively. Figure 3 also shows the threshold values (14, 18) are highly effective in minimizing the FIFO occupancy level compared to (14, 22) and (10, 18).



Figure 4. FIFO occupancy level during simulation versus different values for (Th_1, Th_b) .

To validate the claim that higher throughput yields lower energy, we have calculated dynamic energy, throughput, and ED values versus different values for (Th_l, Th_h) in Figure 5 (a), (b), and (c), respectively. As shown in the figures, (14, 18) leads to the best results for all of these parameters.







Figure 5. Effects of (Th_l, Th_h) for different parameters of the system. a) Dynamic energy, b) Total energy, and c) ED values versus different configurations of (Th_l, Th_h) .

B. Three Recommended Voltage Modes

We will use (14, 18) as the optimum FIFO threshold levels in the rest of the paper. These threshold values are used to set the operating voltage to the three recommended values (i.e. V_h , V_m , V_l). V_h and V_l are set to 1.0v and 0.75v based on our observations in [26]. With 0.75v the ED is minimized as much as possible and 1.0v leads to the lowest packet latency and hence, the highest throughput when required. The next step is to find the suitable voltage value for V_m . The optimum V_m would be the value that leads to the lowest energy dissipation with the least throughput degradation. To find the optimum value for V_m , we have observed the effects of different voltage values for this parameter on total energy and average packet latency.

As shown in Figure 6, the minimum total energy dissipation and maximum packet latency are obtained with V_l , and the maximum total energy dissipation and minimum packet latency are provided with V_h . So we have tried to find a suitable V_m somewhere between these

two extremes, where both energy and packet latency are optimum. The middle of the curves should be a convergence point. As shown in Figure 6(a), this point should be around 0.86v, and Figure 6(b) proposes 0.81v for V_m. Therefore, we select different V_m in this range and observe their effect in energy dissipation and packet latency.



Figure 6. Effects of voltage modes on: a) Total energy, b) Average packet latency.

Figure 7 shows the effects of different V_m values on dynamic/leakage/total energy, ED, and power consumption of the NoC. As shown, $V_m = 0.82v$ leads to the lowest dynamic energy, leakage energy, total energy, power, and ED (normalize values). Thus, we have selected voltage modes $V_1 = 0.75v$, $V_m = 0.82v$, $V_h = 1.0v$.



C. Comparison of FIFO-adaptive and Link-Utilization-Based DVS

So far we have specified the FIFO threshold levels and

the voltage modes as well. At the end, we have compared the two DVS algorithms: FIFO-adaptive DVS algorithm, and history-based DVS policy which uses the link utilization as the traffic intensity indicator for asynchronous NoCs [32]. The ED saving results are shown in Figure 8 for different loads in comparison to the system in which the voltage is fixed at 1.0v. The FIFOadaptive DVS has not only lower implementation cost, but also surpasses the DVS based on link utilization in ED saving for different loads. It achieves more than 31% and 29% ED savings compared to the DVS based on link utilization in 90% and 86% saturated networks, respectively.



Figure 8. Comparison of ED savings in FIFO-adaptive DVS and DVS based on link utilization for different loads.

V. CONCLUSION

In this paper we exploited a fully asynchronous NoC architecture for GALS-based MPSoC architectures and proposed a DVS scheme for low-power and low-energy applications. To evaluate the GALS NoC energy, power, and performance, we introduced a traffic model and found the related saturation thresholds in different voltage link utilization indicator and the modes. The recommended voltage scaling regions were then introduced and a history-based DVS algorithm based on link utilization were proposed accordingly. We also augmented the DVS algorithm with FIFO and explored the effective threshold levels for FIFO. Then a FIFOadaptive DVS algorithm were proposed, which uses the FIFO level as the traffic intensity indicator and scales the operating voltage to three recommended optimum voltage modes. The FIFO-adaptive DVS has not only lower cost of implementation, but also achieves better ED saving compared to the link-utilization-based DVS in saturated networks.

REFERENCES

- W. J. Dally and B. Towles, "Route packets, not wires: Onchip interconnection networks," *Proc. of DAC*, Jun. 2001, pp. 684–689.
- [2] Benini, L. and De Micheli,G. "Networks on chips: a new SoC paradigm". *IEEE Comp.*, v.35(1), 2002, pp. 70-78.
- [3] P. Guerrier, A. Greiner. "A generic architecture for on chip packet-switched interconnections," *Proc. of DATE 2000*, pp. 250-256.
- [4] A. Adriahantenaina and A. Greiner, "Micro-network for SoC: implementation of a 32-port SPIN network," *Proc. of DATE 2003*.
- [5] J. Dielissen, A. Rădulescu, K. Goossens, and E. Rijpkema,"Concepts and implementation of the philips network-on-chip", *IP-SOC 2003*.
- [6] M. Dall'Osso, G. Biccari, L. Giovannini, D. Bertozzi, and L. Benini, "Xpipes: a latency insensitive parameterized network-on-chip architecture for multi-Processor SoCs," *Proc. of the 21st ICCD*, 2003, pp. 536-542.
- [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," *Proc. of DATE 2004.*
- [8] C. Xian, Y. H. Lu, and Z. Li, "Dynamic voltage scaling for multitasking real-time systems with uncertain execution time," *IEEE Trans. on computer-aided design of integrated circuits and systems*, vol. 27, no. 8, august 2008, pp. 1467-1488.
- [9] U. Y. Ogras, R. Marculescu, D. Marculescu, and E. G. Jung, "Design and management of voltage-frequency island partitioned networks-on-chip," *IEEE Trans. on Very Large Scale Integration Systems*, vol. 17, no. 3, pp. 330-341, March 2009.
- [10] M. Es Salhiene, L. Fesquet, and M. Renaudin, "Dynamic Voltage Scheduling for Real Time Asynchronous Systems", *Proc. of PATMOS*'2002, 2002.
- [11] Van Gageldonk H., Van Berkel K., Peeters A., Baumann D., Gloor D., and Stegmann G., "An asynchronous lowpower 80C51 microcontroller," *Proc. of ASYNC*'98, 1998, pp. 96-107.
- [12] T. Bjerregaard and J. Sparsø, "A router architecture for connection-oriented service guarantees in the MANGO clockless Network-on-Chip," *Proc. of DATE 2005*, pp. 1226–1231.
- [13] A. Sheibanyrad, A. Greiner, and I. Miro-Panades, "Multisynchronous and fully asynchronous NoCs for GALS architectures," *IEEE Design & Test*, vol. 25, Issue 6, November 2008, pp. 572-580.
- [14] I. Miro Panades, A. Greiner, and A. Sheibanyrad, A Low Cost Networkon-Chip with Guaranteed Service Well Suited to the GALS Approach, *Nano-Net 2006*, 2006.
- [15] A. Sheibanyrad, I. Miro-Panades, and A. Greiner, "Systematic comparison between the asynchronous and the multi-synchronous implementations of a network on chip architecture," *Proc. of DATE 2007*, pp. 1090-1095.
- [16] D. M. Chapiro, "Globally asynchronous locally synchronous systems," PhD thesis, Stanford University, 1984.
- [17] A. Iyer and D. Marculescu, "Power and performance evaluation of globally asynchronous locally synchronous processors," *Proc. ISCA*, 2002, pp. 652-661.

- [18] G. P. Semeraro et al., "Hiding synchronization delays in GALS processor microarchitecture," *Proc. of ASYNC*, 2004, pp. 159-169.
- [19] G. Magklis, P. Chaparro, J. Gonzalez, and A. Gonzalez, "Independent front-end and back-end dynamic voltage scaling for a GALS microarchitecture," *Proc. of ISLPED* '06, October 2006, pp. 49-54.
- [20] G. Semeraro et al., "Energy-efficient processor design using multiple clock domains with dynamic voltage and frequency scaling," *Proc. of ISHPC*, 2002, pp. 29-40.
- [21] A. Lines, "Nexus: an asynchronous crossbar interconnect for synchronous system-on-chip designs," *Proc. of the 11th Symposium on High Performance Interconnects*, 2003, pp. 2–9.
- [22] R. Dobkin, V. Vishnyakov, E. Friedman, and R. Ginosar, "An asynchronous router for multiple service levels networks on chip," *Proc. of ASYNC*, 2005, pp. 44–53.
- [23] E. Beigné, F. Clermidy, S. Miermont, and P. Vivet, "Dynamic voltage and frequency scaling architecture for Units integration within a GALS NoC," *Proc. of the Second ACM/IEEE International Symposium on Networkson-Chip*, 2008, pp. 129-138.
- [24] L. Shang, L.-S. Peh, and N.K. Jha, "Dynamic voltage scaling with links for power optimization of interconnection networks," *Proc. of the 9th International Symposium on High-Performance Computer Architecture*, 2003, pp. 91–102.
- [25] S. E. Lee and N. Bagherzadeh, "A variable frequency link for a power-aware network-on-chip (NoC)," *Integration*, *the VLSI Journal*, v.42 n.4, September 2009, pp.479-485.
- [26] A. Rahimi, M. E. Salehi, S. Mohammadi, S. M. Fakhraie, and A. Azarpeyvand, "Energy/throughput trade-off in a fully asynchronous NoC for GALS-based MPSoC architectures," Proc. of 5th International Conference on Design & Technology of Integrated Systems in Nanoscale era (DTIS), 2010.
- [27] A. Sheibanyrad and A. Greiner, "Two efficient synchronous asynchronous converters well-suited for network on chip in GALS architectures," *Proc. Integrated Circuit and System Design. Power and Timing Modeling, Optimization and Simulation (PATMOS 06)*, LNCS 4148, Springer Berlin, 2006, pp. 191-202.
- [28] M. Pedram and J. M. Rabaey, "Power aware design methodologies," Kluwer: Academic, 2002.
- [29] B. Gorjiara, N. Bagherzadeh, P. Chou, "An efficient voltage scaling algorithm for complex SoCs with few number of voltage modes," *Proc. of ISLPED*, 2004, pp. 381-386.
- [30] S. Koohi, M. Mirza-Aghatabar, S. Hessabi, and M. Pedram, "High-Level Modeling Approach for Analyzing the Effects of Traffic Models on Power and Throughput in Mesh-Based NoCs," *Proc. of 21st International Conference on VLSI Design*, 2008, pp. 415-420
- [31] V.Soteriou, N.Eisley, and L.-S.Peh, "Software-directed power-aware interconnection networks," ACM Trans. on Architecture and Code Optimization (TACO), vol. 4 no. 5, 2007.
- [32] A. Rahimi, M. E. Salehi, M. Fattah, and S. Mohammadi, "History-Based Dynamic Voltage Scaling with Few Number of Voltage Modes for GALS NoC", Proc. of 5th IEEE International Conference on Future Information Technology (DATICS-FutureTech), 2010.