Energy/Throughput Trade-off in a Fully Asynchronous NoC for GALS-Based MPSoC Architectures

A. Rahimi, M. E. Salehi, S. Mohammadi, S. M. Fakhraie, A. Azarpeyvand School of Electrical and Computer Engineering University of Tehran

Tehran 14395-515, Iran

{ab.rahimi, s.mohammadi}@ece.ut.ac.ir {mersali, fakhraie, azarpeyvand}@ut.ac.ir

Abstract—In this paper we evaluate the compromising effect of energy saving and throughput degradation on a fully asynchronous NoC architecture with regards to the dynamic voltage scaling guidelines. The investigated fully asynchronous NoC architecture is suitable for GALS-based MPSoCs architectures. The introduced architecture is simulated in 90nm CMOS technology with accurate Spice simulations, where the energy/throughput trade-off is reported and analyzed. Our results indicate that, although lower power may also be achieved by dynamic throughput scaling, this technique yields negligible energy saving for our asynchronous NoC. Therefore, we suggest a dynamic voltage scaling for this architecture which can save 40% energy at the expense of 13% throughput degradation.

I. INTRODUCTION

The increasing performance demands of recent embedded systems have complicated their architecture and design. These systems are currently composed of heterogeneous components, including: general purpose CPUs, custom IP blocks, DSP modules, and communication infrastructures. One method to manage design complexity and reduce time to market is to distinguish computation from communication [1]. This is possible through IP reuse and standard interfaces and system level communication modeling.

On-chip communication is becoming increasingly important when systems-on-a-chip (SoCs) grow in complexity and size. Point-to-point links, buses and networks on chip (NoCs) [2] are alternatives for SoCs communication infrastructure. Proposing NoCs for implementing communication in complex SoCs is justified by reusability, scalability, and energy efficiency properties displayed by these networks. The NoC approach was proposed as a promising solution to these complex on-chip communication problems [2][3]. For the NoC architecture, the chip is divided into a set of interconnected nodes. A router is embedded within each node connecting it to its neighboring nodes. As such, instead of routing design-specific global wires, the internode communication can be achieved by routing packets.

NoCs constitute a new design paradigm for scalable and high-throughput communication infrastructures in multiprocessor SoCs (MPSoCs) with billions of transistors. The NoC approach offers a perfect platform for implementing the globally asynchronous locally synchronous (GALS) paradigm [4], and makes clock distribution and timing closure problems more manageable. Given that for complex systems built at 65nm and below, it is almost impossible to move signals across the die in a single clock cycle or in a power efficient manner, it becomes obvious that a shift towards global on-chip asynchronous communication is needed. GALS approach divides the chip area into several independent subsystems, each clocked by a different clock signal. In addition, a GALS-based design style fits well with the concept of system-level power management schemes.

In recent years, GALS designs and dynamic voltage scaling (DVS) have emerged as some of the most popular approaches to address the ever increasing MPSoC energy consumption. Dynamic and adaptive techniques for voltage scaling are becoming essential because of the increasing gap between worst-case and average-case processing demands [5]. According to the quadratic dependency of energy to voltage, DVS is an effective controlling MPSoC technique for energy and performance. Magklis et al. present the effect of finegrain DVS on a clustered GALS microprocessor and propose two online algorithms for adjusting voltage and frequency of the front-end and back-end domains of a novel two-domain microprocessor [6]. An analytic scheme for dynamic voltage and frequency scaling in a multiple clock domain processor is also presented in [7].

The rest of the paper is organized as follows: In Section 2 the most relevant recent researches are reviewed. In Section 3 which is divided into three subsections the architecture of the fully asynchronous NoC along with the traffic model is described for a network containing a homogeneous 3x3 set of clusters, and energy dissipation of a single switch is analyzed. The experimental results and the related analysis for the presented architecture are discussed in Section 4. Finally, Section 5 concludes the paper.

II. RELATED WORK

A micro-network called SPIN [8][9] was the first published attempt to solve the bandwidth bottleneck, in MPSoCs interconnecting a large number of IP cores. After this, a large number of researches on NoC architectures have been performed which have synchronous architectures such as AETHEREAL [10], XPIPES [11], and NOSTRUM [12]. The proposed asynchronous NoCs are MANGO [13], QNOC [14], ANOC [15], and ASPIN [16].

Sheibanyrad et. al, provide a systematic comparison between performance parameters of two different implementations of the same micro-network architecture [17]. This NoC architecture has been designed to be used in shared memory MPSoCs based on GALS architecture. DSPIN implementation is multi-synchronous while ASPIN is implemented fully asynchronously. The packet latency of ASPIN is about 2.5 times smaller than DSPIN. On the other hand, during packet transmissions the energy consumption of ASPIN is higher than DSPIN, while in the idle state, due to clock power dissipation in synchronous designs, the energy consumption of DSPIN is three times higher than ASPIN.

The first study that targets dynamic power optimization of interconnection networks is presented in [18]. Authors propose a history-based DVS policy that judiciously adjusts link frequencies and voltages based on link utilization. Their approach realizes 4.6 times power savings on average. This is accompanied by a moderate impact on performance. GALS architecture and local clock generation motivates easy local DVFS. Although in [19], authors propose a dynamic voltage and frequency scaling for IP units integrated within a GALS NoC, their scheme ignores significant effects of energy dissipation in GALS NoC infrastructures.

In this paper we analyze the compromising effect of energy saving and throughput degradation on a fully asynchronous NoC architecture with regards to dynamic voltage scaling guidelines. The fully asynchronous NoC architecture and dynamic voltage scaling guidelines are suitable for GALS-based MPSoC architectures used in low-power yet high-performance applications. The presented fully asynchronous NoC is based on the work done in [16][17] but in a more energy-efficient manner.

III. ASYNCHRONOUS NOC ARCHITECTURE

In MPSoC designs, a fundamental challenge is the capability of operating under totally independent timing assumptions for each subsystem. Such a multi-clock-domain synchronous system contains several subsystems with independent clocks. A network with a fully asynchronous design, which does not involve the issue of synchronization, is a natural approach to construct GALS architectures. An asynchronous NoC limits the synchronization failure to the network interfaces, where the synchronous/asynchronous data has to enter into the asynchronous/synchronous regions.

Most NoCs have a 2D mesh topology for simplifying physical implementation. The routers are distributed in each subsystem and are connected to their north, south, east, and west neighbors via global asynchronous interconnects.

A. Asynchronous NoC Switch

As mentioned above, the asynchronous NoC architecture introduced in this paper is a fully asynchronous NoC based on [16]. This architecture exploits two special FIFOs for connecting synchronous IP cores to the asynchronous network. Synchronous IP cores transmit and receive data to/from the router through syncto-async and async-to-sync interfaces presented in [20]. Both interfaces are high-throughput, low-latency converters that can be used to convert synchronous communication protocol to asynchronous one and vice versa. Sync-to-async FIFO (SA_FIFO) and async-to-sync FIFO (AS_FIFO) are shown in Figure 1 in black and dashed arrows, respectively.

Each switch contains four (input/output) ports that connect router to the neighbors and the local port is used for connecting router to the IP cores via converters (Figure 2). In the asynchronous NoC implementation, the wires between each input and output port is double railed and the communication is based on the four-phase protocol. The asynchronous NoC uses the distributed X-First algorithm for routing the packets between different ports. The X-First algorithm which has been implemented in the input ports of the router, guarantees the in-order-delivery for the network. With this algorithm, packets are first routed on the X direction and then on the Y direction. As a result, there is no need to connect the north and south input ports to the west and east output ports. When an input port of router receives the header of a packet, the destination address is analyzed and a flit is forwarded to the corresponding output ports. Output port continuously sends flits to the neighbor's router as long as there is enough space in the FIFO of neighbor's input port (the FIFO depth is 8).



Figure 1.Aasynchronous NoC with SA_FIFO and AS_FIFO.



Figure 2. Asynchronous NoC infrastructure.

To avoid starvation, when there are simultaneous requests for the same output port, the Round-Robin algorithm is used for scheduling the requests. Deadlockfree and starvation-free benefits are provided for NoC by using X-First routing and Round-Robin scheduling algorithms.

B. Energy Dissipation of a Single Switch

To evaluate the energy dissipation, latency, and throughput, we have modeled the NoC switch and its corresponding links in gate-level using VERILOG hardware description language (HDL). Fundamental asynchronous components such as C-gate, latch, AND, OR, NOT, MUTEX, etc. are simulated with a 90nm Spice model using Spice simulations. Physical parameters such as average power and propagation delay are then extracted from the simulation results and back annotated to the gate-level HDL modules. Spice simulations are performed in a wide range of voltage values, then energy dissipation and latency are measured in different operating voltages by gate-level simulation.

Figure 3 shows the relation of normalized energy, power, and latency (average values per flit) of a single switch based on voltage variations. According to this figure, higher voltages lead to higher energy, higher power, and also lower latency. This figure presents the high performance and low energy extremes and indicates the compromising effect of voltage on energy and latency. However, this figure does not propose an optimum voltage, in which the highest energy saving is achieved with the lowest latency penalty.

The important measurements for presenting the interaction between delay and energy are energy-delay product (ED) and energy-delay-delay product (ED²) metrics [21]. These metrics would give an idea about the optimum voltage value. Figure 4 shows the normalized

energy-delay saving in voltage variations. According to this figure, 25% ED saving is achieved by reducing voltage from 1.0v to 0.8v, while only 10% ED saving is further achieved by reducing voltage from 0.8v down to 0.5v. Therefore, Figure 4 can locate a voltage value range as the optimum working area which has the most energy savings with the lowest latency degradations.



Figure 3. Relation of normalized energy, power, and latency according to voltage variations.



Figure 4. Energy-delay saving according to voltage variations.

In addition to achieving considerable amount of ED saving in 1.0v to 0.8v, Figure 5, shows that the lowest ED^2 is also obtained in this voltage range. Although lower voltages lead to lower energy, Figure 5 shows that in voltages lower than 0.7v, ED^2 is highly raised which yields very poor performance.



Figure 5. Energy-delay-delay according to voltage variations.

Figure 6 specifies the power/energy saving and latency degradation trade-off in different voltage values. According to this figure, reducing voltage from 1.0v to 0.8v leads to 48% energy saving versus 21% latency degradation and scaling voltage from 1.0v to 0.6v yields 65% energy saving against 99% latency degradation. Another important observation is that in scaling voltage from 1.0v to 0.7v, energy savings are higher than latency degradations; therefore, we introduce this region as the low-energy region. In other words, for voltages below 0.7v, little energy saving is gained versus large latency degradation.



Figure 6. Trade-off between power/energy saving versus latency degradation relative to the values of nominal working point (voltage=1.0v).

C. Traffic Model of the 3x3 Mesh Architecture

For evaluating energy dissipation and throughput of the fully asynchronous NoC, we have focused on a network containing a homogeneous 3x3 set of clusters. The IPs are connected to the local input ports via the SA_FIFO and AS_FIFO used for traffic generation. Each IP consists of two parts, traffic generator (TG) and network analyzer (NA). The TG is connected to the router's local input port and used to inject packets to the network. The NA is connected to the router's local output ports and consumes the generated traffic and check the delivery of packets. HERMES [22] is the reference NoC used to evaluate traffic modeling (a packet-switched mesh-based NoC). The first flits of a packet contain header information and the address of the target nodes. The remaining portion of the packet is the data payload.

The next step is to construct a traffic pattern (i.e [23]) to measure NoC throughput and energy dissipation. In this traffic pattern, all IPs generate packets continuously in time. Streaming applications for video-generator IPs as synchronous subsystems are modeled with 82,000 bytes frames size. As illustrated in Figure 7, IP0 and IP2 generate two video traffic flows for IP7 and IP4, respectively. Each voice-generator IP randomly transmits 14,000 byte packets to other IPs for characterizing noise traffic. The goal of noise traffic is to disturb the video flows. Figure 7 illustrates the spatial distribution of IPs.



Figure 7. Spatial distribution of packet used in 3x3 mesh.

IV. EXPERIMENTAL RESULTS

DVFS techniques are widely used for optimizing the energy in synchronous domains. To use these techniques for the introduced fully asynchronous NoC, as the global clock signal is removed; the concept of clock period should be redefined. We introduce the throughput as the number of bits per unit time that are transferred through a specific link of the switch. Therefore, the throughput is evaluated according to the flit latency and the packet generation interval. By flit latency we mean the difference between the departure time and arrival time of 32-bit flits in the NoC switch which depends on the switch architecture and also the operating voltage. Therefore, for switch architectures, throughput scaling is performed by scaling the packet generation interval. Using this concept, we can exploit traditional DFS techniques in synchronous circuits by dynamic scaling of packet generation intervals in our asynchronous switch.

The energy dissipation of a circuit is proportional to $\alpha C_l v^2$. According to this formula, energy dissipation depends on the operating voltage, total capacitance, and number of transitions. Therefore, a DFS technique cannot reduce the energy dissipation of a synchronous circuit. Figure 8 presents energy saving and throughput degradation of a 3x3 mesh which used the provided traffic model. Energy and throughput are evaluated in the nominal voltage (1.0v) based on different packet generation intervals. According to the results, for interval values below 4ns energy savings and throughput degradations are almost equal (below 1%). However, when the interval is above 4ns, energy saving is saturated in 0.6% while throughput is significantly degraded for the asynchronous 3x3 mesh. Although DFS techniques can improve power consumption in synchronous circuits, interval scaling and consequently throughput scaling in not recommended for energy saving in the presented asynchronous NoC.



Figure 8. Energy savings and throughput degradations in the nominal voltage, based on different packet generation intervals.

Although dynamic throughput scaling fails to achieve considerable amount of energy saving for the fully asynchronous NoC, voltage scaling is able to save significant amount of energy for such structure. This voltage scaling technique is manually applied by scaling the voltage of NoC including switches and links, from 1.0v to 0.4v. Table 1 presents energy saving and throughput degradation of the 3x3 mesh architecture based on different voltage values. According to the quadratic dependency of energy to voltage we suggest a dynamic voltage scaling (DVS) technique for reducing energy at the expense of acceptable performance degradation.

Voltage (V)	Throughput	Energy
	degradation (%)	saving (%)
1.0	0	0
0.9	2.6	20.7
0.8	12.8	37.9
0.7	30.2	51.2
0.6	47.2	64.8
0.5	63.4	77.4
0.4	79.8	91.9

TABLE 1. ENERGY CONSUMPTION AND MAXIMUM THROUGHPUT IN

As shown in Table 1, lower voltages yields lower throughput and energy. Figure 9 shows the effects of voltage and packet generation interval on ED. As indicated in this figure, higher throughputs yields better ED values. This observation limits the range of throughput scaling and proposes packet generation intervals below 4ns and voltage scaling between 1.0v to 0.7v as the optimum ranges for a DVS scheme. To the best of our knowledge, it is the first voltage scaling range which is suitable for dynamic voltage scaling scheme for the fully asynchronous NoC for GALS-based MPSoC architectures.



Figure 9. Effect of voltage and packet generation interval on ED.

V. CONCLUSION

We presented a fully asynchronous NoC architecture for GALS-based MPSoC architectures used in low power and high performance applications. Energy/throughput trade-off was analyzed for single switch as well as 3x3 mesh. The proposed voltage and throughput scaling guidelines can be deployed in dynamic voltage scaling schemes for further throughput-aware energy saving. According to the results, although DFS techniques can improve power consumption in synchronous circuits, interval scaling and consequently throughput scaling is not recommended for energy saving in the presented asynchronous NoC, and the best energy-delay is achieved in high throughput regions. In contrast, with 13% throughput degradation a DVS technique is able to save energy up to 40% while a throughput scaling technique only achieves 0.6% energy saving with the same amount of throughput degradation.

VI. REFERENCES

- [1] Keutzer, K. et al. "System-level design: orthogonalization of concerns and platform-based design". *IEEE Transactions on Computer-Aided Design*, v.19(12), 2000, pp. 1523-1543.
- [2] Benini, L. De Micheli,G. "Networks on chips: a new SoC paradigm". *IEEE Comp.*, v.35(1), 2002, pp. 70-78.
- [3] W. J. Dally and B. Towles, "Route packets, not wires: Onchip interconnection networks," *in Proc. DAC*, Jun. 2001, pp. 684–689.
- [4] D. M. Chapiro, "Globally asynchronous locally synchronous systems," *PhD thesis, Stanford University*, 1984.
- [5] C. Xian, Y. H. Lu, and Z. Li, "Dynamic voltage scaling for multitasking real-time systems with uncertain execution time," *IEEE Transactions on computer-aided design of integrated circuits and systems*, vol. 27, no. 8, august 2008, pp. 1467-1488.
- [6] G. Magklis, P. Chaparro, J. Gonzalez, A. Gonzalez, "Independent front-end and back-end dynamic voltage scaling for a GALS microarchitecture," *Proc. of ISLPED* '06, October 2006, pp. 49-54.
- [7] Q. Wu, P. Juang, M. Martonosi, and D. W. Clark "Formal online methods for voltage/frequency control in multiple clock domain microprocessors," *Proc. of ASPLOS'04*, October, 2004, pp. 248-259.
- [8] P. Guerrier, A. Greiner. "A generic architecture for on chip packet-switched interconnections," *Proc. of DATE 2000*, pp. 250-256.
- [9] A. Adriahantenaina, A. Greiner, "Micro-network for SoC: implementation of a 32-port SPIN network," *Proc. of DATE 2003*, pp. 11128.
- [10] J. Dielissen, A. Rădulescu, K. Goossens, E. Rijpkema, "Concepts and implementation of the philips network-on-chip", *IP-SOC 2003.*
- [11] M. Dall'Osso, G. Biccari, L. Giovannini, D. Bertozzi, L. Benini, "Xpipes: a latency insensitive parameterized

network-on-chip architecture for multi-Processor SoCs," *Proc. of the 21st ICCD*, 2003, pp. 536.

- [12] M. Millberg, E. Nilsson, R. Thid, A. Jantsch, "Guaranteed bandwidth using looped containers in temporally disjoint networks within the nostrum network on chip," *Proc. of DATE 2004.*
- [13] T. Bjerregaard, J. Sparsø, "A router architecture for connection-oriented service guarantees in the MANGOclockless Network-on-Chip," *Proc. of DATE* 2005, pp. 1226 – 1231.
- [14] D. (R.) Rostislav, V. Vishnyakov, E. Friedman, R. Ginosar, "An asynchronous router for multiple service levels networks on chip," *Proc. of the ASYNC 2005.*
- [15] E. Beigne, F. Clermidy, P. Vivet, A. Clouard, M. Renaudin, "An asynchronous NoC architecture providing low latency service and its multi-Level design framework," *Proc. of the 11th ASYNC*, 2005, pp. 44 – 53.
- [16] A. Sheibanyrad, A. Greiner, I. Miro-Panades, "Multisynchronous and fully asynchronous NoCs for GALS architectures," *IEEE Design & Test*, vol. 25, Issue 6, November 2008, pp. 572-580.
- [17] 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. DATE 2007*, pp. 1090-1095.
 [18] L. Shang, L.-S. Peh, N.K. Jha, "Dynamic voltage scaling
- [18] L. Shang, L.-S. Peh, 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.
- [19] E. Beigné, F. Clermidy, S. Miermont, P. Vivet, "Dynamic voltage and frequency scaling architecture for Units integration within a GALS NoC," *Proc. of the Second* ACM/IEEE International Symposium on Networks-on-Chip, 2008, pp. 129-138.
- [20] 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.
- [21] M. Pedram, J. M. Rabaey, "Power aware design methodologies," Kluwer: Academic, 2002.
- [22] Moraes, F. et al. "Hermes: an infrastructure for low area overhead packet-switching networks on chip". *Integration, the VLSI Journal*, vol. 38(1), 2004, pp. 69-93.
- [23] L. Tedesco, A. Mello, L. Giacomet, N. Calazans, F. Moraes, "Application Driven Traffic Modeling for NoCs", *Proc. of the 19th annual symposium on Integrated circuits* and systems design, 2006, pp. 62–67.