

A Comparative Study of Voltage/Frequency Scaling in NoC

Saeeda Usman*, Samee U. Khan*, Sikandar Khan†

**Department of Electrical and Computer Engineering
North Dakota State University
Fargo, ND 58108-6050, USA*

Email: {*firstname.lastname*}@my.ndsu.edu

*†Department of Computer Sciences
COMSATS Institute of Information Technology
Sahiwal 57000, Pakistan.
sikandar@ciitsahiwal.edu.pk*

Abstract— Voltage and frequency is dynamically scaled to produce energy efficient multi-core network-on-chip (NoC). A detailed analysis of the techniques employed for this purpose are studied and based on the optimized performance the most effective ones are reported. We also highlight the most promising high performance and energy minimizing techniques.

Index Terms— Network-on-chip, dynamic voltage and frequency scaling, leakage power

I. INTRODUCTION

Power reduction and energy efficiency are one of the key metrics in the design of upcoming multi-core network-on-chip (NoC) architectures. The NoC is the high performance and scalable alternative to the old-fashioned bus-based architecture. However the NoC consumes a lot of power [3]. The dynamic voltage and frequency scaling (DVFS) is a hardware feature that can in real-time adjust the clock frequency based on the operational voltage of a processor [19, 20]. It is an effective technique for increasing the energy efficiency of a network by reducing the energy dissipation. The key idea is to provide the circuit just enough of speed and voltage that it requires to process the assigned workload. This technique is implemented by a transition of the processor to a low-power state from a high-power state, when the workload reduces [5, 12, 13, 16].

This work is inspired by the emerging technological trends. Modern chips have several cores embedded on them and are designed to achieve optimal power-performance [25]. Each application running on these cores vary in their power-performance requirements to complete an assigned task. A designer needs to keep the variation in mind, and adjust the voltage and frequency judiciously in proportion with the workload. The DVFS is used to improve energy efficiency during low utilization

phase of applications. Lowering the applied voltage and frequency in low-utilization phase prevents thermal violations during sustained high utilization [18].

The remainder of the paper is organized as follows. Section II focuses on the basics of the DVFS technique. Section III dwells on the relationship between the voltage and frequency. Choice of DVFS as design technique in voltage-frequency scaling is discussed in Section IV. Different methodologies of using DVFS are being detailed in Section V. Section 6 concludes the survey.

II. BASICS OF DVFS

Power dissipation in digital Complementary Metal Oxide Semi-conductor (CMOS) circuits takes place due to two major sources. The one arising from the bias and leakage current is known as the static power [26]. Static power is dependent on the process and design technologies [22]. The other form is called as the dynamic power and it commences from the charging and discharging of the voltage saved in the node capacitances of the circuit [11]. The aim of the DVFS is to reduce the dynamic power that can be modeled as [11]:

$$P = V_{DD}^2 f_{CLK} C_{EFF}, \quad (1)$$

where, f_{CLK} is the clock frequency, V_{DD} is the supply voltage, and C_{EFF} is the effective switched capacitance.

III. V/F SCALING LINEARLY

For an efficient control and smooth application of DVFS, a proportionately linear relationship is expected between the frequency and applied voltage [8]. Therefore, due to the linear dependence of frequency on voltage, lowering voltage has the same

proportionate effect as lowering the frequency. Therefore, the power consumption can be modeled in terms of voltage. However, if the ratio of voltage/frequency scaling is not linear, then a re-programming of the delay line has to be incorporated.

Assuming a linear relationship between V_{DD} and f_{CLK} , the power dissipation is given as:

$$P \propto V_{DD}^3 \text{ or } P \propto f_{CLK}^3. \quad (2)$$

The key design idea of DVFS is governed by the above proportional relationship, such that a reduction in the clock frequency or supplied voltage, results in a cubic decrease in the power consumed.

IV. DESIGN OBJECTIVE

The design objective in applying DVFS is to propose a design that has the lowest power leakage value of NoC nodes, during the inactive phase. The low power leakage can be obtained by assigning sleep mode to the inactive cores in a NoC. The aforementioned can be achieved by reducing the applied voltage to the node. However, the node in the sleep mode is not completely turned off and remains in its functionality state at a minimum voltage cost [6]. Performance worsens with the application of aggressively scales the supply voltage to improve the energy efficiency [27].

Energy-saving opportunities are dependent on the workload characteristics and vary accordingly. This is due to the fact that the node count increases in the allotted problem size the energy efficiency decrease. Energy-saving opportunities increases as the workload of a node falls below its utilizable threshold [15]. Using DVFS, the energy efficiency is inclined to a higher value as compared to other voltage and frequency (V/F) scaling techniques, such as Dynamic Power Management (DPM), Dynamic Voltage Scaling (DVS), and Dynamic Frequency Scaling (DFS). Due to the exponential increase in the density of the transistors per core and the level of compaction, a mechanism is required to yield high performance while maintaining good energy budget. The aforementioned technique of DVFS proves to be a need of the day and has a positive impact on the growth of the multi-core architecture.

V. EVOLUTION OF DVFS

A lot of research has been done in the domain of the improvement of energy efficiency in today's multi-core architecture [1]. A brief overview of few recent evolutionary inspired approaches to power optimization using DVFS in multi-core NoC is presented here. Lee *et al.* [14] emphasized on the need of a Power Management Unit (PMU) that controls the generation of the supply voltage and

clock to fine-tune the speed of the target domain. The occupational level and performance requirement of each domain is generated to the PMU that decides the upgrade of frequency and voltage for each of the power domain. V/F line usage by the PMU causes the system to lock the V/F line [14]. The lock essentially discourages multi-threading.

Malkowski *et al.* [17] exploits the fact that the memory bound code can be power optimized for achieving energy saving. As soon as the application enters a memory bound phase, prefetchers are activated. In the memory bound phase, the execution time is dominated by the latency caused in the access to the main memory. Applications in the memory bound phase tend to be idle, waiting for the required set of data. Moreover, in the memory bound phase, a reduction in the CPU voltage and frequency does not lower the performance. Therefore, power optimization can be achieved.

Beigne *et al.* [2] proposed a fully power-aware locally-synchronous and globally-asynchronous NoC circuit for the implementation of DVFS mechanism. Adaptive design techniques are used for each of the synchronous NoC unit. The main CPU is required to directly disable/enable the units for entering the power off mode and executing the reset phase [2]. A multiple supply line for V/F decreases the voltage-frequency transition delay. Therefore, the system becomes more time efficient.

Howard *et al.* [10] described an evolutionary approach for a fine-grained power management of multi-core NoC using Voltage and Frequency Islands (VFI). The power management protocols in the software exploit the benefits of the VFI through the DVFS. With an active DVFS, a reduction of 80% in the measured power can be achieved, as compared to the phase when DVFS is inactive.

Performance of a many-core processor can be increased by using more cores and parallelizing the workloads [7]. Many core processors with DVFS can optimize power-performance of parallel workloads by varying V/F level of VFIs.

Per-core DVFS if applied to chips achieves better application performance and control accuracy as compared to per-chip DVFS. It is known that DVFS can allow a cubic reduction in power density for each of the core in a Chip Multi-processor (CMP) employing DVFS [23]. Performance of the CMP can further be enhanced by minimizing the transition time of V/F scaling. Inter band Tunnel Field Effect Transistors (TFETs) cores are more energy-efficient at low frequencies as compared to the CMOS cores. Therefore, combining the two types of cores in a multi-core architecture and applying DVFS for V/F scaling, a thread migration scheme is proposed. Low frequency applications are migrated to the TFET

cores, while high frequency applications are more efficiently handled by the CMOS core. The scheme achieves average dynamic energy and leakage energy savings of 17% and 30%, respectively. However, the scheme has an overhead of performance degradation of 1% [21].

Herbert *et al.* [8] proposed a process variation parameter to shift work to efficient dies or cores. The work shifting can be divided into two levels: **(a)** among dies belonging to a given speed bin and **(b)** between VFI on a given die. The traditional DVFS schemes have neglected the variations in both the spatially-correlated within-die process and die-to-die, which if addressed can improve energy efficiency [8]. A balance between performance and power in CMP can efficiently be achieved using fine-grained DVFS [9]. Therefore, from full-chip V/F control, the CMP is moving towards finer-grained methods. By shifting work from the inefficient, leaky processing units to less leaky and more efficient ones, more power conservation can be achieved [8].

Thread packing along with DVFS is used to maximize efficiency [4]. Reda *et al.*[18] proposed a novel scheme specifies the number of threads running on a core and uses the power cap mechanism to maximize the performance for multi-threaded workloads onto a variable number of cores. The thread reduction methodology along with the power cap enabled, limits the maximum power that a

particular core can use at any instant of time. Power conservation is achieved by assigning the sleep mode state to the idle cores [18]. The aforementioned technique was tested on a quad-core (Intel's Core i7) based server and reduced the workload energy consumption by an average of 51.6% compared to existing control techniques, that implements DVFS alone.

Energy consumption can be reduced further in comparison to all prior discussed DVFS techniques if the processor is put into the deep sleep state [24]. In the deep sleep mode, the processor is paused, and the consumption of power in is significantly low [24]. As the intensity of the workload increases from light to medium, the DVFS is used to guarantee high performance by scaling the processor frequency. Table 1 summarizes the research regarding energy efficiency using DVFS in multi-core NOC. The comparison in the table provides a reference to the parameters required to be considered for choosing the best design for implementation technique using DVFS. A check mark (✓) in the table refers to the fulfillment of the parameter, a cross mark (✗) corresponds that the parameter under consideration is unfulfilled, and a negative 1 reveals a performance degradation of -1%. Whereas, a N/A entry in the table refers to the fact that either the parameter was inapplicable or not discussed.

Table 1: Energy Efficiency Fulfilled/Not-fulfilled by Various DVFS Approaches

Works	Power Reduction	Leakage Energy Saving	Variation Aware	Performance Up-gradation	Multiple V/F lines
[7]	65%	✓	✗	✗	✗
[9]	50%	✓	✗	✗	✗
[1]	N/A	✓	✓	✓	✗
[5]	80%	✗	✓	✓	✓
[12]	19%	✓	✓	-1%	✗
[4]	10%	✓	✓	✓	✓
[3]	20-60%	✓	✓	✗	✓
[13]	Info. N/A	✗	✓	✓	✓
[2]	22-86%	✓	✓	✗	✓
[6]	61.69%	✗	✓	✗	✓
[14]	63%	✗	✓	✓	✓

VI. CONCLUSIONS

Techniques using DVFS have been studied in detail to find the most energy efficient technique of implementation of DVFS. Implementation of the traditional DVFS operations can be disintegrated into two parts: **(a)** clock transition and **(b)** voltage transition depending on the workload of the processor. Optimal core mapping also contributes to boost the performance. The DVFS has become an integral part of the NoC for saving power as well as for reducing congestion by increasing the frequency of the

congested router. Therefore, DVFS is essential for increasing the throughput of the system.

ACKNOWLEDGMENTS

The authors are grateful to Muhammad Awais, Osman Khalid, and Assad Abbas for their valuable suggestions.

REFERENCES

- [1] M. A. Aziz, S. U. Khan, T. Loukopoulos, P. Bouvry, H. Li, and J. Li, "An overview of achieving energy efficiency in on-chip networks," *International Journal*

- of Communication Networks and Distributed Systems*, Vol. 5, No. 4, July 2010, pp. 444-458.
- [2] E. Beigne, F. Clermidey, H. Lhermet, S. Miermont, Y. Thonnart, X. Tran, A. Valentin, D. Varreau, P. Vivet, X. Popon and H. Lebreton, "An asynchronous power aware and adaptive NOC based circuit," *IEEE Journal of solid-state Circuits*, Vol. 44, NO. 4, April 2009, pp. 1167 – 1177.
- [3] K. Chakroborty and S. Roy, "Architecturally homogeneous power-performances heterogeneous multicore systems," *IEEE Transactions on VLSI Systems*, Vol. 21, No. 4, June 2012, pp. 670-679.
- [4] R. Cochran, C. Hankendi, A. K. Coskun, and S. Reda, "Pack & Cap: Adaptive DVFS and thread packing under power caps," *IEEE/ACM International symposium on Microarchitecture*, Dec. 2011, pp. 175-185.
- [5] C. O. Diaz, M. Guzek, J. E. Pecero, P. Bouvry, and S. U. Khan, "Scalable and energy-efficient scheduling techniques for large-scale systems," *International Conference on Computer and Information Technology (CIT '11)*, Sept. 2011, pp. 641-647
- [6] C. O. Diaz, M. Guzek, J. E. Pecero, G. Danoy, P. Bouvry, and S. U. Khan, "Energy-aware fast scheduling heuristics in heterogeneous computing systems," *International Conference on High Performance Computing and Simulation (HPCS 2011)*, July 2011, pp. 478-484.
- [7] S. Dighe, S. R. Vangal, P. Aseron, S. Kumar, T. Jacob, K.A. Bowman, J. Howard, J. Tschanz, V. Erraguntla, N. Borkar, V.K. De, and S. Borkor, "Within-die variation-aware dynamic-voltage-frequency-scaling with optimal core allocation and thread hopping for the 80-core TeraFLOPS processor," *IEEE Journal Of Solid-State Circuits*, Vol. 46, No. 1, Jan. 2011, pp. 184-193.
- [8] S. Herbert, S. Garg, and D. Marculescu, "Exploiting process variability in voltage/frequency control," *IEEE Transactions on VLSI systems*, Vol. 20, No. 8, Aug 2012 pp. 1392-1404.
- [9] S. Herbert and D. Marculescu, "Variation-aware dynamic voltage/frequency scaling," *International Symposium on High-Performance Computer Architecture (HPCA 2009)*, Feb. 2009, pp. 301-312.
- [10] J. Howard, S. Dighe, S. R. Vangal, P. Aseron, S. Kumar, T. Jacob, K.A. Bowman, J. Howard, J. Tschanz, V. Erraguntla, N. Borkar, V.K. De, and S. Borkor, "A 48-core ia-32 processor in 45 nm CMOS using on-die message-passing and DVFS for performance and power scaling," *IEEE Journal of Solid-State Circuits*, Vol. 46, No. 1, Jan. 2011, pp. 173-183.
- [11] X. Jin and S. Goto, "Hilbert transform-based workload prediction and dynamic frequency scaling for power efficient video encoding," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 31, No.5, May 2012, pp. 649-661.
- [12] S. U. Khan and I. Ahmad, "A cooperative game theoretical technique for joint optimization of energy consumption and response time in computational grids," *IEEE Transactions on Parallel and Distributed Systems*, Vol. 20, No. 3, Mar. 2009, pp. 346-360.
- [13] S. U. Khan and N. Min-Allah, "A goal programming based energy efficient resource allocation in data centers," *Journal of Supercomputing*, Vol. 61, No. 3, Sept. 2012, pp. 502-519.
- [14] J. Lee, B. Nam and H. Yoo, "Dynamic voltage and frequency scaling (DVFS) scheme for multi-domains power management," *IEEE Asian Solid-State Circuits Conference (ASSCC '07)*, Nov. 2007, pp. 360 – 363.
- [15] D. Li, B. R. D. Supinski, M. Schulz, D. S. Nikolopoulos, and K.W. Cameron, "Strategies for energy efficient resource management of hybrid programming models," *IEEE Transactions on parallel and distributed Systems*, Vol. 24, Jan. 2013, pp. 144-157.
- [16] P. Lindberg, J. Leingang, D. Lysaker, K. Bilal, S. U. Khan, P. Bouvry, N. Ghani, N. Min-Allah, and J. Li, "Comparison and analysis of greedy energy-efficient scheduling algorithms for computational grids," in *Energy Aware Distributed Computing Systems*, A. Y. Zomaya and Y.-C. Lee, Eds., John Wiley & Sons, Hoboken, NJ, USA, 2012, ISBN 978-0-470-90875-4, Chapter 7.
- [17] K. Malkowski, P. Raghavan, M. Kandemir, and M. J. Irwin, "Phase-aware adaptive hardware selection for power-efficient scientific computation," *International Symposium on Low Power Electronics and Design (ISPLED'07)*, Aug 2007, pp. 403-406.
- [18] S. Reda, R. Cochran, and A. K. Coskun, "Adaptive power capping for servers with multi-threaded workloads," *International Symposium on Microarchitecture (MICRO '12)*, Vol. 32, No. 5, Aug. 2012, pp. 64-75.
- [19] Y. Seo, J. Kim, and E. Seo, "Effective analysis of DVFS and DPM in mobile devices," *Journal of Computer Science and Technology*, Vol. 27, No. 4, July 2012, pp. 781-790.
- [20] J. Shuja, S. A. Madani, K. Bilal, K. Hayat, S. U. Khan, and S. Sarwar, "Energy-efficient data centers," *Computing*, Vol. 94, No. 12, , 2012, pp. 973-994.
- [21] K. Swaminathan, E. Kultursay, V. Saripalli, V. Narayanan, M. Kandemir, and S. Datta, "Improving energy efficiency of multi-threaded applications using heterogeneous CMOS-TFET multicores," *International Symposium on Low Power Electronics and Design (ISPLED '2011)*, Aug 2011, pp. 247-252.
- [22] G. L. Valentini, W. Lassonde, S. U. Khan, N. Min-Allah, S. A. Madani, J. Li, L. Zhang, L. Wang, N. Ghani, J. Kolodziej, H. Li, A. Y. Zomaya, C.-Z. Xu, P. Balaji, A. Vishnu, F. Pinel, J. E. Pecero, D. Kliazovich, and P. Bouvry, "An overview of energy efficiency techniques in cluster computing systems," *Cluster Computing*, Vol. 16, No. 1, Mar. 2013, pp. 3-15.
- [23] X. Wang, K. Ma, and Y. Wand, "Adaptive power control with online model estimation for chip multiprocessors," *IEEE Transactions on parallel and*

- Distributed Systems*, Vol. 22, No. 10, Oct. 2011, pp. 1681-1696.
- [24] Y. Wang and X. Wang, "Virtual Batching: request batching for server energy conservation in virtualized data centers," *IEEE Transactions on Parallel and Distributed Systems*, Aug. 2012, pp. 1.
- [25] K. Yadav, M. Casu, M. R. Zamboni, and Maurizio, "A simple DVFS controller for a NOC switch," *Ph. D. Research in Microelectronics and Electronics (PRIME '12)*, June 2012, pp. 1-4.
- [26] S. Zeadally, S. U. Khan, and N. Chilamkurti, "Energy-efficient networking: past, present, and future," *Journal of Supercomputing*, Vol. 62, No. 3, Dec. 2012, pp. 1093-1118.
- [27] B. Zhai, D. Blaauw, D. Sylvester, K. Flautner, "Theoretical and practical limits of dynamic voltage scaling," *Design Automation Conference (DAC '04)*, July 2004, pp. 868-873.