

An Efficient Hardware Implementation of DVFS in Multi-Core System with Wireless Network-on-Chip

Hemanta Kumar Mondal
IIT-Delhi, ECE
New Delhi, India
hemanta@iitd.ac.in

Gade Narayana Sri Harsha
IIT-Delhi, ECE
New Delhi, India
Narayana1294@iitd.ac.in

Sujay Deb
IIT-Delhi, ECE
New Delhi, India
sdeb@iitd.ac.in

Abstract—Networks-on-Chip (NoC) have emerged as communication backbones for enabling high degree of integration in future many-core chips. Despite their advantages, the communication is multi-hop and causes high latency and power dissipation, especially in larger systems. Wireless Network-on-Chip (WNoC) significantly improves the latency over traditional wired NoCs for multi-core systems. But on-chip wireless interfaces (WIs) have their own power and area overhead. In this paper we design and implement a Dynamic Voltage Frequency Scaling (DVFS) technique and extend it to provide power gating to the WIs. This approach effectively reduces the energy consumption in multi core systems. A centralized controller with dual-band wireless transceiver implements per-core DVFS. The scheme ensures balanced workload and energy consumption of the chip and efficient power gating for the WIs. It helps to alleviate the power consumption up to 33.085 % for on-chip communications infrastructure with little overheads.

Keywords—wireless network-on-chip; dual-band transceiver and antenna; dynamic voltage and frequency; low power

I. INTRODUCTION

With progress in CMOS technologies and Chip Multiprocessors (CMPs), multiple processing cores are embedded on a single chip to achieve higher performances. But this has resulted in increased power density and energy consumption across the chip, increased complexity of the communication architecture and large interconnect delays. Network-on-Chip with planer metal interconnects can significantly improve performance in comparison with the traditional bus architectures. But latency and power consumption of NoC based systems is increasing due to multi-hop long distance communication in many core systems. One possible solution to reduce the latency of the NoC systems is the use of high bandwidth single hop wireless links. Different wireless NoC architectures have been proposed ([1],[2],[3],[4]), but these wireless links have its own power and area overhead.

Millimeter-wave based wireless NoC architecture proposed in [5] uses few WIs spread across many-core systems. As all the WIs work in the same frequency, the power consumption related to WIs can be further improved by switching off the WIs that are not currently using the wireless channel. Power-gating techniques applied for multi core systems ([6],[7],[8]) can be adopted for making the WIs more power efficient.

Dynamic Voltage/Frequency Scaling (DVFS) techniques have been proposed to reduce the power consumption of the chip with little affect on the performance of the system. By exploiting the processing insensitive idle phases of an application/task, DVFS methods reduce the supply voltage/frequency to achieve large reductions in power with little performance loss. But most of the proposed techniques are specifically designed to tackle the power consumption in cores. Few DVFS methods proposed for NoC ([9],[10]) employ distributed controllers which have large area and power overheads. To overcome this, cluster based DVFS is proposed for many cores systems [11]. Each cluster consists of DVFS controller, voltage regulators and frequency generators. This approach can still have a large hardware overhead for multiple controllers.

To address these issues, we propose a robust centralized controller unit (CCU) for improving the power efficiency of multi core systems with hybrid hierarchical NoC architecture as shown in Fig. 1. The term hybrid indicates that both wired and wireless interconnects are used in the NoC. The system is divided into multiple clusters or subnets. DVFS regulators are placed at the center of each cluster. CCU is at the top of the architecture and observes and regulates all the DVFS regulators. The CCU implements algorithms to control the core voltage/frequency along with power gating control signals for the WIs. The time critical control signals should reach the CCU from different parts of the chip and the response from CCU should also travel to different parts of the chip in a timely and energy efficient way. In this paper we proposed to use wireless links for communicating these control signals to and from the CCU. The existing WIs [5] use single band transceivers for data transmission. We propose to use a dual-band transceiver and antenna in each WIs for data and control signal communication separately. We demonstrate the advantages of our proposed design and discuss the overheads and trade-offs.

The main contributions of this paper include a novel centralized DVFS controller to balance workload and power consumption of the chip that also applies power gating to WIs to ensure power efficient utilization of the wireless channel. The paper is organized into following sections. Section II briefly describes the related work. The hierarchical NoC architecture and centralized DVFS controller are explained in section III. The hardware implementation of controller, different emerging interconnects and dual-band transceiver are

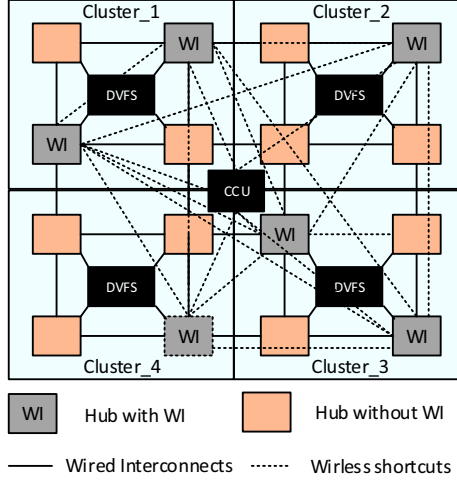


Fig. 1. Proposed cluster-based NoC architecture with DVFS

explained in section IV. Section V describes experimental results. Section VI concludes this work.

II. RELATED WORK

To reduce the power consumption in multi core systems, different DVFS approaches have been proposed. An optimization framework based on thermal constraints to devise frequency plans for optimal performance is discussed in [12]. Authors of [13] propose a method to optimize the voltage/frequency islands in multi core system with NoC to reduce the power consumption. The performance and overheads of a distributed DVFS in Globally Ratiochronous Locally Synchronous (GRLS) systems is discussed in [14]. The implemented DVFS approaches include per core, clustered and centralized methods. Centralized controllers, even though requires complex hardware reduces the overall area and power consumption compared to distributed approaches.

A wireless NoC architecture using millimeter wave interconnects is discussed in [1]. A Token-based Adaptive Power gating (TAP) approach to actively power gate the core during memory access in many core systems is discussed in [15]. TAP tracks the memory request and its estimated arrival time for power gating the core. A dual-band (900-MHz/1.8-GHz) transmitter is proposed in [16]. For Bluetooth and 802.11b applications, dual-mode transceiver designs are discussed in [17] and [18].

Our aim is to design an efficient wireless centralized controller to enhance the performance of wireless NoC and reduce the power consumption of the system. We also designed dual-band transceiver and antenna for WIs to ensure power efficient and low latency on-chip data and control signals communications.

III. PROPOSED ARCHITECTURE

A. Hierarchical NoC Architecture

Hierarchical NoC architecture using WIs is discussed in [5], where the entire system is divided into multiple smaller modules called subnets. Each subnet, shown in Fig. 2, is a cluster of neighboring cores and has NoC switches and links.

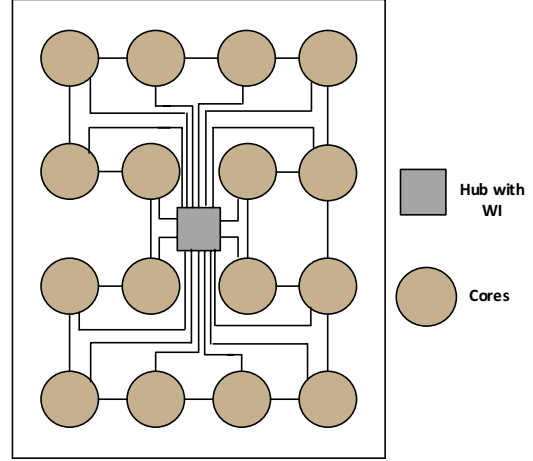


Fig. 2. Subnet structure of hierarchical wireless NoC architecture

All the cores in a subnet are connected to a central hub through wired interconnects and hubs from multiple subnets are connected to each other forming a hierarchical network. Few such hubs have WIs and only the hubs separated by long distances use WIs, thereby limiting the number of wireless interfaces and their associated overheads. In this paper we use StarRing topology i.e. ring with a central hub for the subnet. Flit-based wormhole routing technique is used for data transfer. For intra subnet communication, if the hop count is more than two, data is transferred through the central hub otherwise data is transferred along the ring. All the WIs operates at the same frequency for on-chip data communication. In the proposed implementation, we used dual-band transceivers for WIs which uses different frequencies for data and control signals. This improves the performance of the WIs as discussed in the later section.

B. Centralized DVFS Controller

In the proposed DVFS controller a chip level unit that observes various system parameters to determine the operating voltage and frequency of all cores in the system and controls

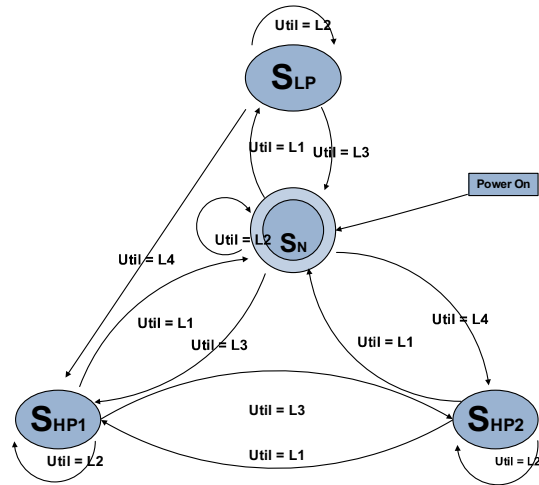


Fig. 3. Core state machine diagram

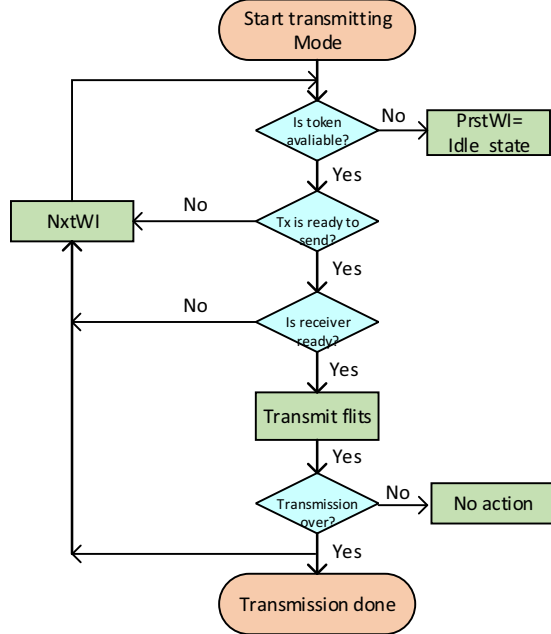


Fig. 4. Flow diagram showing power-gating control logic for WIs

the sleep mode operation of the WIs according to the availability of data. It generates the necessary control signals for DVFS regulators and to apply power gating to the WIs.

The core control unit of the DVFS controller uses a time slice based approach, where the system parameters and core state in current time slice are used to set the voltage and frequency for the next time slice. The duration of the time slice can be chosen based on general purpose of the system. Choosing too small duration may result in increased operation overhead of the controller whereas choosing very long duration may result in skipping significant changes in the system state. At each time slice, the core control unit observes the current utilization of a core, core busy/idle pattern and the overall temperature of the cluster. Using the core busy/idle state changes, the average duration for which a core is busy performing any task. Using the past utilizations of the core during each time slice, a probabilistic state change model is developed which assigns a probability value between all possible state changes based number of core states available. These values are updated dynamically at each time slice based on current observed core utilization values. Based on the values from state change model and the average busy time of the core, the utilization and state of the core for next time slice is predicted. If the observed temperature of the cluster is within the acceptable levels for a particular state, this predicted state is applied to the core for next time slice. In case the temperature is moderately high and the predicted state is a higher power state, we prevent the transition and continue the core in current state. If the temperature is very high, we assign a low power state to the core irrespective of the predicted state. This is done to prevent excessive use of a particular core and to maintain the thermal density of a cluster within reliable levels. A manual user level input is also incorporated which forces a high performance state or a high power savings state irrespective of the predicted state of the core. In either case, the normal

controller operation is disabled. For the current implementation, we considered four states, a low power state (S_{LP}), normal state (S_N) and two high power states (S_{HP1} and S_{HP2}). At any time slice, the processor can be in one of the four states. The processor utilization for a time slice is also divided into four levels as 0-30 %, 30-50 %, 50-80 % and above 80 %. The state machine diagram representing the possible state change scenarios and corresponding conditions is shown in the Fig. 3. S_{LP} , S_N , S_{HP1} , S_{HP2} denote the different core states corresponding to low power, normal and high performance. Once the system is powered, all the cores start in normal state. The utilization L1 corresponds to the lowest predicted core utilization for the next time slice and the immediate lower state is assigned to the core. For example, if a core is in S_N state and the predicted utilization is L1, the assigned state for the next time slice will be S_{LP} . Similarly L2, L3 and L4 correspond to the normal and high utilizations for the next time slice and the corresponding state changes are as shown.

The flow diagram for power-gating control logic in the CCU for WIs is shown in Fig. 4. The WIs are operated in a round robin fashion. Initially all the transceivers are kept in sleep mode, denoting them as WI_1 to WI_N for N WIs. This block is event triggered, which passes a token from WI_1 to WI_N to check for availability of data to be transmitted. If there is no data to be transmitted at present WI ($PrstWI$) the controller puts it in idle state and the token is passed to next WI ($NxtWI$). If data is available, the transmitter sends the receiver address to the controller. If the receiver is ready, the controller then generates the appropriate signals for the DVFS regulators to turn on the transmitter and receiver pair and data transmission starts. After the transmission is over, the controller passes the token to next WI ($NxtWI$). This process continues for all the WIs and at any point of time, only two WIs are active at the most. In this way, significant power can be saved with a little overhead.

IV. HARDWARE IMPLEMENTATION

A. Controller

The core control unit consists of four major blocks; the Current State block, State Change Model, Busy/Idle pattern and the Temperature Control block as shown in Fig. 5. The

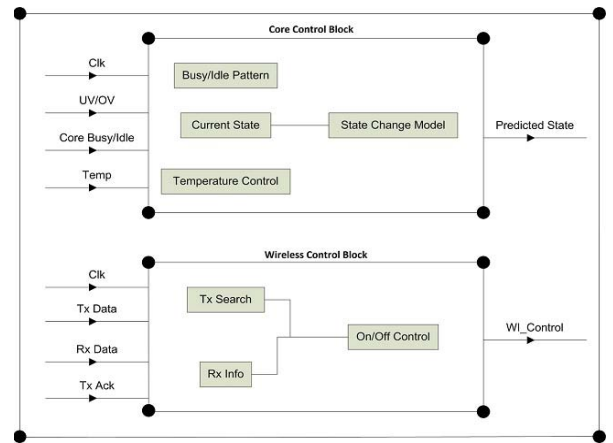


Fig. 5. Centralized DVFS controller

Current State block counts the number of clock cycles for which the core is busy in a given time slice to calculate the utilization of the core in that slice. The output is two bit value representing one of the four utilization levels mentioned in the previous section. All the counters and data is reset at the end of each time slice. The Busy/Idle pattern block also reads in the core busy/idle state at each clock cycle and counts the number of clock cycles for which the core is continuously busy/idle and uses this data to update the average duration for which the core remains busy. This block operates irrespective of the time slice and using the average values, estimates the duration for which the core can be busy in the next time slice. The State Change Model block reads in the current state of the core from Current State block and using the state change model estimates the next state of the core. By taking into considering all previous transitions between the four states, probabilities for transition to the four states from any state are developed. The probabilities are represented as number of transitions for every 100 times the core in a particular state to avoid using floating point numbers and division operations. The values are updated at each time slice if the current state is not same as the estimated at previous time slice. Using the data from this block and busy/idle block, we predict the state for next time slice. The Temperature Control block reads the temperature of the cluster at each time slice and compares with the predefined allowable temperature ranges for each state. The temperature is divided into multiple levels according to the thresholds defined for each state. Depending on the level of the current temperature, this block sends a two bit signal to the controller according to the method described in previous section to determine whether to apply the predicted state or not.

In case of wireless controller, a single instance controls all the available transceivers. But the core control block can have multiple instances with higher number of instances creating area overhead and lower number of instances creating delay overhead. For testing purposes, we created for instances of core control unit for each of the clusters available in the system. Each control unit handles the data from four cores of the corresponding cluster. Since, the centralized unit needs to communicate to different parts of the chip in the fastest possible way to obtain best performance, we explored the possibility of using on chip wireless links to communicate the control signals.

B. Emerging Interconnects

The multi-hop delay of the wired interconnects adds a significant delay in transmitting the control signals from the CCU to various clusters in the system. To tackle this issue, various emerging interconnect options are considered and compared for optimum performance.

To communicate with various clusters, three novel interconnect options; wire based, G-lines based [19] and mm-wave wireless are considered. For a large multi core chip sizes of 20 mm x 20 mm, there can be a maximum wire length of 15 mm between controller and clusters in the proposed topology. And the delay due to a wired interconnect significantly affects the performance of the controller. This delay can be reduced by using G-lines for all communications between the controller and clusters. To further reduce this delay, we can use the WIs of the existing NoC architecture for transmitting the control

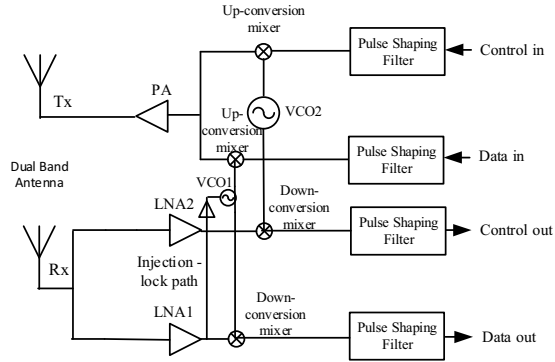


Fig. 6. Dual-band transceiver block

signals also. Towards this goal, we propose a dual-band wireless transceiver which can be used for both data and control signal transmission.

C. Dual-Band Wireless Transceiver Circuit with Antenna

To ensure low latency and energy efficiency of WNoC, the transceiver circuit has to provide a wide bandwidth as well as low power consumption. In designing the on-chip dual-band transceiver, design considerations are taken into account at architecture level with a design adopted from ([15],[16],[17]). Non-coherent on-off keying (OOK) modulation scheme is used here. The transmitter consists of pulse shaping filter, up-conversion mixer, and a power amplifier as shown in Fig. 6. At the transmitter (Tx) side, the data and control signals are fed into a pulse shape filter and the filtered signal is then amplified by a power amplifier (PA). At the receiver (Rx) end, the received RF signals are fed into the low noise amplifier (LNA) and mixer through pulse shaping filter. Two LNA; LNA₁ and LNA₂ are required for two different frequencies of operation. LNA₁ is used for data and LNA₂ is for control signals. Power hungry PLL can be replaced by injection-lock voltage controlled oscillator (VCO) and direct conversion topology. The VCO is used to generate the carrier frequencies; VCO₁ is used for data and VCO₂ is used for control signals.

For the on-chip antenna associated with each WI, we used the planar log-periodic antenna design [20] at millimeter frequency range to improve the performance of wireless interconnects. This antenna is operated as a dual-band antenna with 44 GHz for control signals and 60 GHz for data transmissions. The implementation details are discussed in the results section.

V. EXPERIMENTAL RESULTS

In this section we discuss the experimental result that demonstrates the performance and overheads of the proposed centralized controller in WNoC architecture. We discuss the simulation setup, performance of the DVFS controller and WIs and the overheads of each component.

A. Simulation Setup

In this paper, we consider two different system sizes of 256 and 512 cores each, and die area is kept fixed at 20 mm x 20 mm. The centralized controller is synthesized using the

TABLE 1. VOLTAGE AND FREQUENCY COMBINATION

Operation	Voltage (V)	Frequency (GHz)	Threshold
Normal	1	2.5	≥ 0.9
Optimal	0.866	2.0	≥ 0.7
Optimal	0.733	1.75	≥ 0.6
Optimal	0.600	1.5	≥ 0.5
Sleep mode	0	0	0

Design Compiler tool from Synopsys with 65-nm standard cell library from TSMC at a clock frequency of 2.5 GHz. The millimeter-wave dual-band wireless transceiver is designed and simulated using Cadence tool with TSMC 65-nm CMOS process. The gain and bandwidth performance of the antennas is obtained using the HFSS tool.

B. Centralized Controller

The voltage and frequencies used for testing are presented in Table 1. The optimum number of WIs for these two system sizes are 6 and 13 respectively [1]. The transceivers operate at 1 V when active. When all the transceivers are always active and operated without using DVFS, they consume a total power of 440.4 mW for a 256 core system. By applying the proposed method, at any point of time, a maximum power of 147.35 mW is consumed for a single channel communication. So the proposed technique can save up to 66.54 % of the total power consumed. The hierarchical wireless NoC architecture proposed in [5] consumes 220.2 mW of power for all the transceivers. Hence the proposed controller can improve the power consumption up to 33.085 % compared to the existing architecture.

As the number of cores in the system increase, the number of WIs needed for optimal performance also increases. We have used 6 WIs for 256 cores and 13 WIs in case of 512 cores. The power saving for different system sizes are different and vary from 33.085 % for 6 WIs to 69.12 % for 13 WIs.

C. Emerging Interconnects

As discussed in the previous section, there can be a maximum wire length of 15 mm between controller and clusters. Using a wired interconnect, the maximum latency is

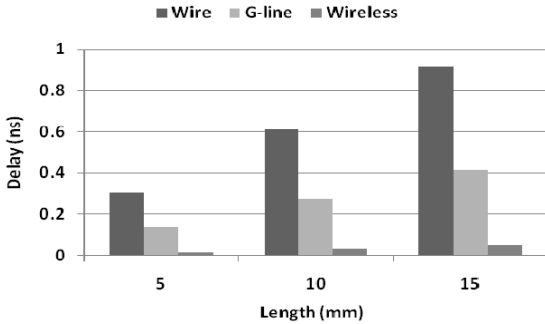


Fig. 7. Delay for different lengths of emerging interconnects




Fig. 8. Energy per bit for different lengths of emerging interconnects

917 ps which requires multiple clock cycles. Latency is reduced to 414 ps by using the G-lines. But this still requires more than one clock cycle to transmit the control signals from the controller to cluster. By using the WI for control signals as well, the delay is reduced to 50 ps and hence the signals can be transmitted within a single clock cycle. The delay using different interconnects for different lengths is presented in Fig. 7 and it is evident that using the existing WNoC architecture for control signals the latency can be significantly reduced.

The energy per bit for control signals through wireless channel is also significantly less compared to the other interconnects. The per bit energy for wire and G-line are 5.025 pJ/bit and 7.038 pJ/bit respectively. We have used wireless interconnects which dissipates 0.459 pJ/bit. The energy per bit for different lengths of emerging interconnects is presented in Fig. 8.

D. Dual-Band Transceiver with Antenna

The dual-band wireless transceiver (inclusive of OOK modulator/demodulator, LNA₁ and LNA₂, PA, VCO₁ and VCO₂, Mixer) occupies a total area of 0.6 mm². The wireless link dissipates 0.459 pJ/bit. It can sustain a data rate of 16 Gbps with power consumption 73.4 mW.

Two on-chip log-periodic planar antennas are simulated and integrated on the same substrate separated by a distance 20 mm. The antenna is 1.1825 mm long. The return loss of planner log periodic antenna (S₁₁ parameter) is shown in Fig. 9. Dual-band antenna is operated at 44 GHz and 60 GHz with a 10 % fractional bandwidth at 60 GHz. The gain of this antenna is 38.65 dB at 60 GHz.

E. Overheads

The proposed centralized controller with DVFS adds an area overhead of 12.5 % of the total chip area for 256 cores.




Fig. 9. Planer log-periodic antenna characteristics




Fig. 10. Wireless power savings and area trade-offs for WI

The area overhead increases linearly with the increase in number of WIs. Area overhead for transceiver (inclusive of OOK modulator/demodulator, LNA₁ and LNA₂, PA, VCO₁, and VCO₂) is 0.6 mm². A 256-core system with 6 wireless transceivers consumes 9.6 % of total Si area. Power consumption by centralized controller with DVFS is 0.5469 mW. Dual-band transceiver consumes 73.4 mW. For single channel communication, the overall power consumption for WI with control unit is determined to be 147.35 mW. The tradeoff between percentage wireless power savings and area overhead for different number of WIs is shown in the Fig. 10.

VI. CONCLUSION AND FUTURE WORK

In this paper we implement a centralized controller with DVFS module for cores and power gating controller for wireless interfaces to achieve optimal system wide energy savings and performance. A dual-band transceiver with appropriate antenna is designed for power and latency efficient communication between controller and various clusters of the chip. The controller is shown to achieve up to 33.085 % of energy savings over the wireless channel.

In ongoing and future investigations, we intend to establish a detailed full system analysis along with thermal modeling of our proposed architecture.

ACKNOWLEDGMENT

This work is supported by the DST Inspire faculty fellowship granted by the Department of Science and Technology (DST), Govt. of India.

REFERENCES

- [1] Deb, S. et al., "Design of an Energy Efficient CMOS Compatible NoC architecture with Millimeter-Wave Wireless Interconnects," *IEEE Transactions on Computers*, vol. 62, pp. 2382-2396, issue: 99, September 2012.
- [2] D. DiTomaso, A. Kodi, S. Kaya, D. Matolak, "iWISE: Inter-router Wireless Scalable Express Channels for Network-on-Chip (NoCs) Architectures," *19th Annual IEEE Symposium on High Performance Interconnects*, pp. 11-18, Aug, 2011.
- [3] S. B. Lee, S. W. Tam, I. Pefkianakis, S. Lu, M. F. Chang, C. Guo, G. Reinman, C. Peng, M. Naik, L. Zhang, and J. Cong, "A scalable micro

- wireless interconnect structure for cmps," *Mobicom '09*, pp. 217- 228, September, 2009.
- [4] C. Wang, W. Hu, N. Bagherzadeh, "A Wireless Network-on-Chip Design for Multicore Platforms," *19th International Euromicro Conference on Parallel, Distributed and Network-Based Processing*, pp. 409-416, Feb. 2011.
- [5] S. Deb, et al., "Design of an Efficient NoC Architecture using Millimeter-Wave Wireless Links," *13th Int'l Symposium on Quality Electronic Design (ISQED)*, pp. 165-172, March 2012.
- [6] Annavaram, M., et al. "A case for guarded power gating for multi-core processors," *IEEE 17th International Symposium on High Performance Computer Architecture (HPCA)*, pp. 291-300, Feb. 2011.
- [7] Roy, Soumyaroop, Nagarajan Ranganathan, and Srinivas Katkoori. "State-retentive power gating of register files in multicore processors featuring multithreaded in-order cores," *IEEE Transactions on Computers*, vol.60 pp. 1547-1560, November 2011.
- [8] Mondal, H.K., Deb, S., "Energy Efficient On-chip Wireless Interconnects with Sleepy Transceivers," *IEEE International Design and Test Symposium (IDT)*, pp.1-6, December 2013.
- [9] Murray, Jacob, et al. "Sustainable dual-level DVFS-enabled NoC with on-chip wireless links," *14th Int'l Symposium on Quality Electronic Design (ISQED)* ISQED, pp. 135-142, March 2013.
- [10] Murray, Jacob, et al. "Evaluating effects of thermal management in wireless NoC-enabled multicore architectures," *IEEE International Green Computing Conference (IGCC)*, pp. 1-8, June 2013.
- [11] Kolpe, Tejaswini, Antonia Zhai, and Sachin S. Sapatnekar. "Enabling improved power management in multicore processors through clustered dvfs," *IEEE Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pp. 1-6, March 2011.
- [12] Kadin, Michael, and Sherief Reda. "Frequency planning for multi-core processors under thermal constraints," *Low Power Electronics and Design (ISLPED)*, 2008 ACM/IEEE International Symposium on. IEEE, pp. 463-470, Oct. 2008.
- [13] Jang, Wooyoung, Duo Ding, and David Z. Pan. "Voltage and frequency island optimizations for many-core/networks-on-chip designs," *Green Circuits and Systems (ICGCS)*, 2010 International Conference on. IEEE, pp. 217-220, June 2010.
- [14] Chabloz, Jean-Michel, and Ahmed Hemani. "Distributed DVFS using rationally-related frequencies and discrete voltage levels," *Low-Power Electronics and Design (ISLPED)*, 2010 ACM/IEEE International Symposium on. IEEE, pp. 247-252, Aug. 2010.
- [15] Kahng, Andrew B., et al. "Many-Core Token-Based Adaptive Power Gating," *IEEE Transactions On Computer-Aided Design of Integrated Circuits And Systems*, 2012.
- [16] Behzad Razavi, "A 900-MHz/1.8-GHz CMOS Transmitter for Dual-Band Applications," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 5, May 1999.
- [17] TB Cho, "A 2.4-GHz Dual-Mode 0.18-um CMOS Transceiver for Bluetooth and 802.11b," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 11, November 2004.
- [18] YJ Jung et al., "A 2.4-GHz 0.25-um CMOS Dual-Mode Direct-Conversion Transceiver for Bluetooth and 802.11b," *IEEE Journal of Solid-State Circuits*, vol. 39, no. 7, July 2004.
- [19] T. Krishna et al., "NoC with Near-Ideal Express Virtual Channels Using Global-Line Communication," *Proc. IEEE Symp. High Performance Interconnects (HOTI)*, pp. 11-20, Aug, 2008.
- [20] Samaiyar, A., et al. "Millimeter-Wave Planar Log Periodic Antenna for On-Chip Wireless Interconnects", *EUCAP*, 2014.