

Rochester Institute of Technology

## RIT Digital Institutional Repository

---

Theses

---

7-2015

### Dynamic Voltage and Frequency Scaling for Wireless Network-on-Chip

Pratheep Joe Siluvai Iruthayaraj

Follow this and additional works at: <https://repository.rit.edu/theses>

---

#### Recommended Citation

Iruthayaraj, Pratheep Joe Siluvai, "Dynamic Voltage and Frequency Scaling for Wireless Network-on-Chip" (2015). Thesis. Rochester Institute of Technology. Accessed from

This Thesis is brought to you for free and open access by the RIT Libraries. For more information, please contact [repository@rit.edu](mailto:repository@rit.edu).

**Dynamic Voltage and Frequency Scaling for Wireless  
Network-on-Chip**

by

**Pratheep Joe Siluvai Iruthayaraj**

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of  
Master of Science in Computer Engineering

Supervised by

Dr. Amlan Ganguly  
Department of Computer Engineering  
Kate Gleason College of Engineering  
Rochester Institute of Technology  
Rochester, NY  
July 2015

**Approved By:**

---

Dr. Amlan Ganguly  
*Primary Advisor – R.I.T. Dept. of Computer Engineering*

---

Dr. Dhireesha Kudithipudi  
*Secondary Advisor – R.I.T. Dept. of Computer Engineering*

---

Dr. Reza Azarderakhsh  
*Secondary Advisor – R.I.T. Dept. of Computer Engineering*

# Dedication

I would like to dedicate this thesis to my parents Mr. Irudayaraj and Mrs. Regina Rosi who have supported me from the beginning of this journey. I would also like to dedicate this to my mentor and all my friends who have been a great source of motivation and inspiration.

# Acknowledgements

I take this opportunity to express my profound gratitude and deep regards to my primary advisor Dr. Amlan Ganguly for his exemplary guidance, monitoring and constant encouragement throughout this thesis. Dr. Amlan dedicated his valuable time to review my work constantly and provide valuable suggestions which helped in overcoming many obstacles and keeping the work on the right track. I would like to express my deepest gratitude to Dr. Dhireesha Kudithipudi and Dr. Reza Azarderakhsh for sharing their thoughts and suggesting valuable ideas which have had significant impact on this thesis. I am grateful for their valuable time and cooperation during the course of this thesis. I also take this opportunity to thank my research group members for all the constant support and help provided by them.

## **Abstract**

Previously, research and design of Network-on-Chip paradigms were mainly focused on improving the performance of the interconnection networks. With emerging wide range of low-power applications and energy constrained high-performance applications, it is highly desirable to have NoCs that are highly energy efficient without incurring performance penalty. In the design of high-performance massive multi-core chips, power and heat have become dominant constraints. Increased power consumption can raise chip temperature, which in turn can decrease chip reliability and performance and increase cooling costs.

It was proven that Small-world Wireless Network-on-Chip (SWNoC) architecture which replaces multi-hop wireline path in a NoC by high-bandwidth single hop long range wireless links, reduces the overall energy dissipation when compared to wireline mesh-based NoC architecture. However, the overall energy dissipation of the wireless NoC is still dominated by wireline links and switches (buffers).

Dynamic Voltage Scaling is an efficient technique for significant power savings in microprocessors. It has been proposed and deployed in modern microprocessors by exploiting the variance in processor utilization. On a Network-on-Chip paradigm, it is more likely that the wireline links and buffers are not always fully utilized even for different applications. Hence, by exploiting these characteristics of the links and buffers over different traffic, DVFS technique can be incorporated on these switches and wireline links for huge power savings.

In this thesis, a history based DVFS mechanism is proposed. This mechanism uses the past utilizations of the wireline links & buffers to predict the future traffic and accordingly tune the voltage and frequency for the links and buffers dynamically for each time window. This mechanism dynamically minimizes the power consumption while substantially maintaining a high performance over the system. Performance analysis on these DVFS enabled Wireless NoC shows that, the overall energy dissipation is improved by around 40% when compared Small-world Wireless NoCs.

# Table of Contents

Dedication .....	ii
Acknowledgements.....	iii
Abstract.....	iv
Table of Contents .....	vi
List of Figures .....	viii
List of Tables .....	x
Chapter 1 Introduction .....	1
1.1. Network-on-Chip (NoC).....	2
1.2. Dynamic Voltage Frequency Scaling (DVFS) .....	4
1.3. Motivation .....	7
1.4. Thesis Contribution .....	8
Chapter 2 Related Work.....	11
Chapter 3 Wireless NoC Architecture.....	13
3.1. Small World Topology .....	14
3.2. Flow Control and Routing .....	16
3.3. Wireless Interface .....	18
3.4. Antenna and Transceiver .....	19
3.5. Performance Metrics.....	20
3.6. Performance Evaluation .....	22
Chapter 4 Dynamic Voltage Frequency Scaling for WiNoC.....	26
4.1. DVFS Architecture and Modeling.....	26

4.2.	History-based DVFS.....	29
4.2.1	Network Traffic Characteristics .....	29
4.2.2	History-based DVFS Policy .....	33
4.2.3	DVFS Controller and Hardware Implementation .....	35
4.3.	Performance Evaluation of DVFS-enabled WiNoCs .....	37
4.3.1	Energy Dissipation Characteristics .....	37
4.3.2	Bandwidth and Latency Characteristics .....	39
Chapter 5	Conclusion and Future Work.....	43
Bibliography	.....	45



## List of Figures

Fig. 1: Network-on-Chip Architectures [19]. (a) SPIN, (b) CLICHÉ, (c) Torus, (d) Folded torus, (e) Octagon, (f) BFT .....	4
Fig. 2: VF variations in a DVFS System [35].....	5
Fig. 3: Architecture for small-world WiNoC.....	15
Fig. 4: Network Switch with virtual channels .....	17
Fig. 5: (a) On-chip metal zig-zag antenna (reproduced from [3]) (b) On-chip antenna placement on the die (reproduced from [23]) .....	20
Fig. 6: Bandwidth characteristics of mesh and SWNoC.....	23
Fig. 7: Latency characteristics of mesh and SWNoC .....	24
Fig. 8: Energy characteristics of mesh and SWNoC.....	25
Fig. 9: DVFS Mechanism on Switches and Links .....	27
Fig. 10: Power gate transistor width versus processor performance [56].....	28
Fig. 11: Utilization profile for uniform traffic (a) Link Utilization profile (b) Buffer Utilization profile.....	32
Fig. 12: Voltage and Delay Curve .....	34
Fig. 13: Finite State Machine for DVFS Controller .....	35
Fig. 14: Total Network Energy for different NoC Architectures.....	38
Fig. 15: Total Network Energy for Map-reduced traffic .....	39
Fig. 16: Utilization profile for Map-reduced traffic (a) Link Utilization profile (b) Buffer Utilization profile.....	40
Fig. 17: Bandwidth characteristics for different architectures.....	41

Fig. 18: Bandwidth characteristics for Map-reduced traffic.....	41
Fig. 19: Latency characteristics for different architectures .....	42
Fig. 20: Latency characteristics for Map-reduced traffic.....	42

## List of Tables

Table 1: Voltage/Frequency/Threshold Combinations.....	34
Table 2: Characteristics of DVFS Controller.....	35

## Chapter 1 Introduction

Network-on-Chip (NoC) architectures have become a primary focus for researchers for designing high performance and energy efficient multicore processors and System-on-Chip (SoC) architectures that can integrate hundreds of cores in a single chip [1].

Traditional network fabrics suffer from an important performance and power consumption limitation in designing massive multicore chips, where a data transferred between two distant cores because high power consumption and latency issues [2].

However, NoCs have been identified to show increased performance by inserting long-range wired links using the principle of small-world graphs [3]. As the system size scales up, the small-world topology takes advantage of the inherent multi-hop nature of the largely separated communicating cores and reduces the average hop count by introducing a relatively long-distance direct shortcuts. These network fabrics can further be improved by replacing the single hop long-range wired links with energy efficient, wireless links. The wireless shortcuts have been shown to carry a substantial amount of traffic, thus enabling significant energy savings through these low power wireless links.

However, the overall energy savings in the system can further be improved by optimizing the characteristics of the wireline links and associated switches based on the data traffic it takes. Dynamic voltage and frequency scaling (DVFS) is a popular technique that enables power optimization of electronic systems without significantly compromising the system performance. The required voltage and frequency levels in the DVFS is optimized based on the traffic pattern generated by the benchmark applications. Mainly DVFS have been used in power management techniques that addresses issues for energy


savings in processing cores. This paper extends the technique by performing DVFS on NoC platform by characterizing the voltage and frequency levels based on the utilizations of each individual links and switches. These DVFS enabled switches and links helps reducing the energy dissipation of a multi-core chip and consequently provides high energy savings for a network-on-chip platform.

### **1.1. Network-on-Chip (NoC)**

One of the major problems in future SOC designs arises from non-scalable delays in global wires [19]. Global wires carry signals across a chip and typically do not scale in length with technology scaling [20]. Though gate delays scale down with technology, global wire delays typically increase exponentially or linearly by inserting repeaters. Even after repeater insertion [21], the delay may exceed the limit of one or multiple clock cycles. In ultra-deep submicron processes, it is claimed that 80 percent or more of the delay of critical paths will be due to interconnects [22]. As a result, many large designs use as hoc FIFO buffers to synchronously propagate data over large distances to overcome this problem. According to the ITRS report, “Global synchronization becomes prohibitively costly due to process variability and power dissipation, and cross-chip signaling can no longer be achieved in a single clock cycle.” [23]. Thus, system design must work on networking and distributed computation paradigms with functional blocks integrated into the communication backbone. The most frequently used on-chip interconnect architecture is an arbitrated bus, where all communication devices share the same transmission medium. Advantages of such shared-bus architectures are simple topology, low area and extensibility. However, for a long bus line, the intrinsic parasitic resistance and capacitance

can be quite high. Moreover, every additional IP block added to the bus adds to this parasitic capacitance, in turn increasing propagation delay. As the bus length and/or the number of IP blocks increases, the associated bit transfer delay over the bus becomes large and will eventually exceed the targeted clock period. This places a limit on the number of IP blocks that can be connected to a bus and thereby limits the system scalability [24]. One solution to deal with this problem is to split the bus into multiple segments and employ a hierarchical architecture [25]; however, this is ad hoc in nature and has the inherent limitations of the bus-based architecture. In SoCs consisting of several IP blocks, bus-based interconnects will face serious bandwidth problems as all attached devices must share the same medium [24]. To overcome the above-mentioned problems, use of a communication-centric approach to integrate IPs in complex SoCs is advocated. This new model separates the resource elements (i.e., the IPs) from the communication infrastructure (i.e., the network). The need for global synchronization thus disappears. This new approach is explicitly parallel, exhibits modularity to minimize global wires and utilizes locality in power minimization [26]. In a network-centric approach, communication between IPs happens in the form of packets. A common characteristic of such architectures is that the IP blocks communicate with each other using intelligent switches or routers. As such, these switches dubbed infrastructure IPs (I2Ps) [26] provide a robust data transfer medium for the functional IP modules. There is another manner of explaining the relevance of Network on Chips [27]. Reliable communication between circuit components requires a protocol definition that provides some rules describing how the interaction shall take place. These rules ensure that the overall system performance requirements are met, while physical resources like area or energy are minimized. Traditional on-chip communication designs

use ad-hoc 4 approaches that often fail to meet some strict scalability requirement of next-generation SOC designs. Bottlenecks can arise in performance, throughput, power, energy, reliability, synchronization, predictability and concurrency. Designers traditionally stuck to point-to-point connections and bus-based techniques. This approach is acceptable for a small number of blocks when the performance/ latency trade-off is relatively simple.




**Fig. 1: Network-on-Chip Architectures [19]. (a) SPIN, (b) CLICHÉ, (c) Torus, (d) Folded torus, (e) Octagon, (f) BFT**

## 1.2. Dynamic Voltage Frequency Scaling (DVFS)

One active area of work on NoC has focused on dynamically varying operating voltage and frequency levels to achieve a balance between power and performance [29]. This technique, referred to as DVFS, is used quite often in SoC designs [30]. Dynamic voltage and frequency scaling (DVFS) was introduced in the 90's [31] to dramatically reduce power consumption in large digital systems by varying both voltage and frequency

of the system with respect to changing workloads [32, 33, 34, 35]. Fig. 2 shows the time varying pattern of voltage and frequency in a system exhibiting DVFS [35]. Alternative techniques using voltage/frequency islands (VFIs) for IP blocks are used in achieving fine grain system-level power management [36]. Use of VFIs in the NoC context can provide better power-performance tradeoffs than single voltage, single clock frequency case as it benefits from the natural partitioning and mapping of applications onto the NoC platform. Despite the huge potential for energy savings with 6 VFIs, the NoC design methodologies considered so far are limited to a single voltage clock domain [37, 38, 39]. Studies that do consider multiple VFIs assume that each module/core in the design belongs to a different island and different islands are connected by point-to-point (P2P) links [40, 41].



**Fig. 2: VF variations in a DVFS System [35]**

Power-gating is a standby-leakage reduction method developed in [42, 43, 44, 45, 46]. In a power gating design, sleep transistors are used as switches to shut off power



supplies to parts of a design in standby mode [47]. Clock gating was also proposed as a power saving technique [48, 49, 40]. Some studies indicate that the clock signals in digital computers consume a large (15–45%) percentage of the system power [51]. Thus, the circuit power can be greatly reduced by reducing the clock power dissipation. Many clock power reduction techniques have focused on reduced voltage swings, buffer insertion, and clock routing [52]. In many cases, switching of the clock causes a huge 7 gate activity. In circuits with controllable clocks, master clock is used to derive all other clocks which, based on certain conditions, can be slowed down or stopped completely with respect to the master clock [53].

In this thesis mainly the DVFS circuits and techniques are applied to each switch and wireline links in the NoC such that each switch and link is contained in its own independent clock and voltage domain. Lowering the supply voltage leads to a square reduction in dynamic power based on the dynamic power-voltage relationship given as,

$$P_{dyn} = \alpha CV^2 f$$

Where,  $\alpha$  is the switching probability or activity,  $C$  is the total load capacitance,  $V_{dd}$  is the supply voltage, and  $f$  is the clock frequency. Without altering the supply voltage, power can be reduced with frequency reduction, but the energy consumption per operation remains the same. Supply voltage reduction on the other hand, contributes directly to energy reduction, where the dynamic energy consumption of a gate is a direct function of the supply voltage:  $E = CV_{dd}^2$ . Leakage power is reduced as well with reduced supply voltage under normal circumstances. DVFS becomes increasingly important as leakage power becomes dominant contribution to power consumption in very deep-submicron CMOS technologies [54]. Benefits of DVFS also include counteracting process variation

and thermal effects [55]. Slower parts of the chip can be speeded up with higher voltages, and hotter parts can be cooled with lower voltages.

Reduction in supply voltage results in an increased gate propagation delay ( $t_d$ ),

$$t_d = CV_{dd}/(V_{dd} - V_t)^\alpha$$

Where,  $V_t$  is the threshold voltage, and  $\alpha$  is the velocity saturation index. To guarantee correct operation of a synchronous system, the frequency must normally be scaled along with the voltage. The performance overhead of frequency and voltage scaling can be mitigated in a multi-core network-on-chip architecture by taking advantage of the variation in workloads across the buffers and wireline links in the network. Switches (buffers) can operate at a higher voltage during periods of high utilization, and at lower voltages during periods of low utilization to minimize energy dissipation.

### **1.3. Motivation**

The limitations and design challenges associated with existing NoC architectures has led to the emergence of Wireless Network-on-Chip, that enables technology to design high bandwidth and low power multicore architectures[13],[11]. It is shown in [13] that the network consumes a significant part of the chip's power budget and it can be almost 50% depending on the application. Most of the existing works related to the design of wireless NoC demonstrate its advantage in terms of latency and energy dissipation provided by the wireless channels only. The main emphasis always has been on the characteristics of the wireless links. However, the overall energy dissipation of the wireless NoC can be improved even further if the characteristics of the wireline links and buffers are optimized depending on their utilization requirements for different traffic patterns.

Dynamic Voltage Frequency Scaling (DVFS) is known to be an efficient power management technique. Taking advantage of the characteristics of the wireline links and buffers in the interconnection network, DVFS can be incorporated for significant energy savings in Wireless NoC system. This thesis is aimed at implementing DVFS scheme in the interconnection network of small world wireless NoC system with efficient history based DVFS controller for wireline links and buffers in the network, for significant energy savings.

#### **1.4. Thesis Contribution**

In this thesis, it will be demonstrated that by implementing Dynamic Voltage and Frequency Scaling on interconnection networks for Wireless NoCs with efficient DVFS controller, significant power savings can be achieved when compared to existing power management techniques in Wireless NoCs. The history based DVFS technique will be implemented on wireline links and buffers by taking advantage of the utilization characteristics over a time window to predict the future utilization and vary the voltage and frequency accordingly for significant energy savings. This proposed system will prove to be more energy efficient than previous NoC paradigms. The proposed system is implemented with other existing topologies and will be evaluated for performance characteristics and energy savings. Furthermore, the trade-off between performance and energy will be established for different traffic conditions. The following points will summarize the contributions made during this work.

- **Proposed Power Management System**

- Design and implementation of Dynamic Voltage Frequency Scaling for Wireless NoCs
- Design and development of history-based DVFS controller.
- **Evaluation of Wireless NoC Schemes**
  - Evaluation of energy savings over Wireless NoCs and DVFS-enabled Wireless NoCs
  - Evaluation of performance trade-off over DVFS-enabled Wireless NoCs
  - Evaluation of DVFS-enabled Wireless NoCs for different traffic patterns
- **Development of simulation framework**
  - Develop a cycle accurate simulator to implement the wireless NoC architectures with Dynamic Voltage and Frequency Scaling for wireline links and buffers.
  - Develop an efficient algorithm to implement the DVFS controller for history based prediction mechanism.
  - Obtain experimental results of DVFS-enabled Wireless NoC architecture with other wired and wireless architectures with respect to the following parameters using the cycle accurate simulator.
    - Peak achievable bandwidth
    - Latency Overheads
    - Packet energy dissipation
    - Non-uniform traffic patterns

- **Publication**

- Nassef Mansoor, **Pratheep Joe Siluvai Iruthayaraj**, Amlan Ganguly, “Design Methodology for a Robust and Energy-Efficient Millimeter-wave Wireless Network-on-Chip”, IEEE Trans. on Multi-Scale Computing System, June 2015.

## Chapter 2 Related Work

Various research groups have investigated power and thermal management of multicore-based computing platforms. Dynamic voltage and frequency scaling (DVFS) is a popular methodology to optimize the power usage/heat dissipation of electronic systems without significantly compromising overall system performance. Hence DVFS can be applied to multi-core processors; to all cores or to individual cores independently [4]. Multi-core chips implemented with multiple Voltage Frequency Islands (VFI) design styles are other promising alternatives. VFI is shown to be effective in reducing on-chip power dissipation [5] [6]. Various research groups have addressed designs of appropriate DVFS control algorithms for VFI systems [7]. Some researchers have also recently discussed the practical aspects of implementing DVFS control on a chip, such as tradeoffs between on-chip verses off-chip DC-DC convertors [4], the number of allowed discrete voltage levels, and centralized verses distributed control techniques [11]. Thermal-aware techniques are principally related to power-aware design methodologies using DVFS [8]. It is shown that distributed DVFS provides considerable performance improvement under thermal duress [8].

Most of the existing works principally address power and thermal management strategies for the processing cores only. Network consume a significant part of the chip's power budget; generally affecting overall temperature. However, there is little research on how they contribute to the thermal issues [9]. Thermal Herd, proposed in [9], provides a distributed runtime scheme for thermal management that allows routers to collaboratively regulate the network temperature profile and work to avert thermal emergencies while

minimizing performance impact. For the first time, [10] addressed the problem of simultaneously dynamic voltage scaling of processors and communication links for the real-time distributed systems. Intel's recent multi-core-based single chip cloud computers (SCC) incorporate DVFS both in the core and the network levels. However, all of the above-mentioned works principally consider standard multi-hop interconnection network-levels for the multi-core chips; the limitations of which are well known.

A comprehensive survey regarding various WiNoC architectures and their design principles are presented in [12]. It is already shown that the small-world network architecture with long-range wireless shortcuts can significantly improve the energy consumption and achievable data rate of massive multicore computing platforms [12]. Here, we complement that effort by simultaneously addressing the power and thermal management of WiNoC-based multi-core processing platforms by incorporating network-level DVFS.

## Chapter 3 Wireless NoC Architecture

The earlier interconnect technologies have been used in existing NoC platforms without significant architectural innovations, which undermines the performance gains. However, the emerging technologies make direct connections between physically distant cores on the chip viable due to their high communication bandwidth and low power dissipation characteristics. This allows innovation in the design of the NoC architecture to maximize the utilization of the performance benefits of these emerging interconnects, specifically the wireless communication channels. Many naturally occurring networks are known to have the so-called small-world property. Networks with the small-world property have a very short average path length, which is commonly measured as the number of hops between any pair of nodes. The average shortest path length of small-world graphs is bounded by a polynomial in  $\log(N)$ , where  $N$  is the number of nodes, which makes them particularly interesting for efficient communication with minimal resources [16, 17]. This feature of small-world graphs makes them particularly attractive for constructing scalable WiNoCs. Most complex networks, such as social networks, the Internet, as well as certain parts of the brain exhibit the small-world property. This makes them scalable with increase in system size. Thus such connection topologies are suitable for modern multi-core systems, which have hundreds of cores on a single die. The adopted small-world topology essentially inserts long-range links in the NoC. However, long wireline interconnects incur high energy dissipation and latency in data transfer. So as many long-range links as possible are replaced with wireless interconnects based on the scalable small-world wireless NoC architecture.



### 3.1. Small World Topology

In this type of topology, each core is connected to a NoC switch and the switches are interconnected using wireline and wireless links. The topology is a small-world network where the links between switches are established following a power law distribution as shown below.

$$P(i, j) = \frac{l_{ij}^{-\alpha} f_{ij}}{\sum_{\forall i} \sum_{\forall j} l_{ij}^{-\alpha} f_{ij}} \quad (1)$$

Where, the probability of establishing a link, between two switches,  $i$  and  $j$ ,  $P(i, j)$ , separated by an Euclidean distance of  $l_{ij}$  is proportional to the distance raised to a finite power [17]. The distance is obtained by considering a tile-based floorplan of the cores on the die. The frequency of traffic interaction between the cores,  $f_{ij}$ , is also factored into (1) so that more frequently communicating cores have a higher probability of having a direct link. This frequency is expressed as the percentage of traffic generated from  $i$  that is addressed to  $j$ . This frequency distribution is based on the particular application mapped to the overall NoC and is hence known prior to wireless link insertion. Therefore, the apriori knowledge of the traffic pattern is used to establish the topology with a correlation between traffic distribution across the NoC and network configuration as in [18]. This optimizes the network architecture for non-uniform traffic scenarios. The parameter  $\alpha$  govern the nature of connectivity. Higher the value of alpha, lesser the number of longer links which brings down the total wiring cost for the system. Also, it is established in [17] that choosing a value of  $\alpha < D+1$ , where  $D$  is the dimension of the network a small-world network connectivity can be established. In our case the NoC is arranged in a 2D tile and consequently,  $D=2$ . The value of  $\alpha$  was chosen to be 1.8 to establish a small-world

connectivity [17] for which it also noticed that the system has maximum throughput with minimum wiring cost. As the links are established probabilistically following (1) the number of ports of each switch may not be the same. The average number of ports per switch is however constrained to be 5 to have the total number of connections same as that of a mesh. Fig. 1. Shows the small-world WiNoC architecture of 25 core system where each core is associated with a NoC Switch, connected using wireline links and few long distant cores using wireless links.

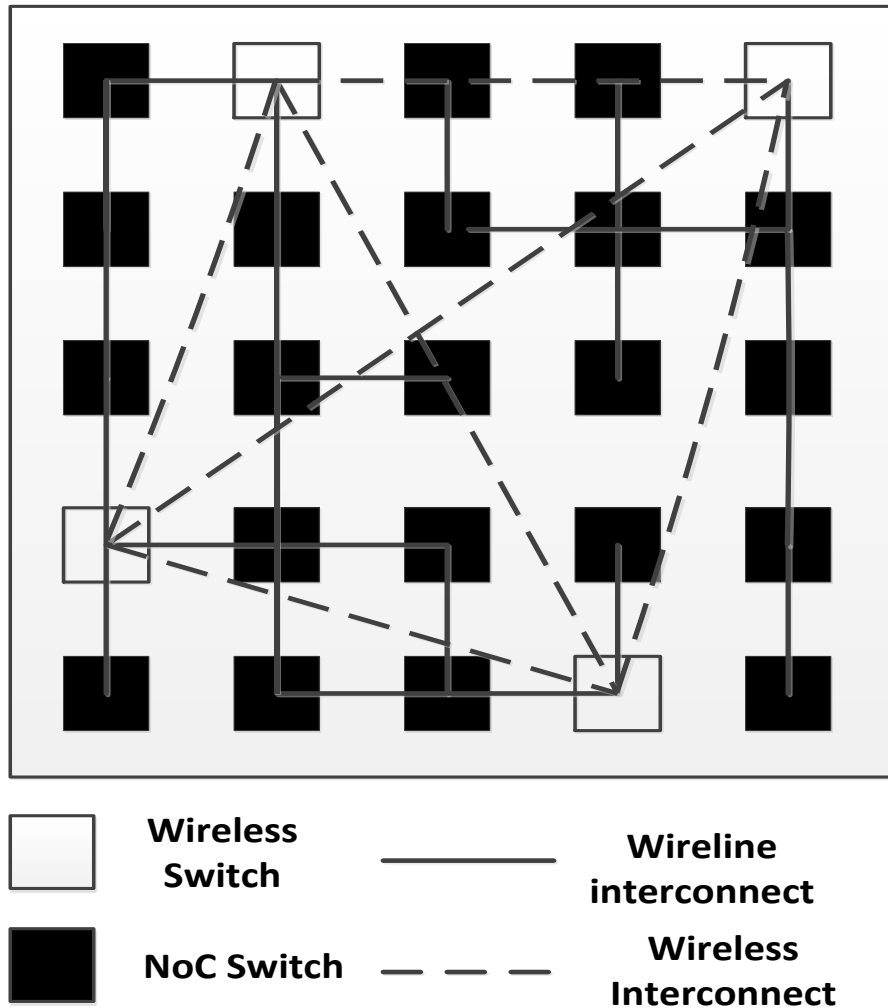


Fig. 3: Architecture for small-world WiNoC

### **3.2. Flow Control and Routing**


For a conventional NoC system, there can be basically three types of switching that can be considered for data routing. Namely, Circuit Switching, Packet Switching and Wormhole Switching.

In case of circuit switched networks, a dedicated path is reserved for the complete duration of the transmission. Even though the network bandwidth is reserved during the transmission it is highly inefficient when there are many nodes waiting for transmission along the same path which eventually degrades the system performance.

In case of packet switching, data is divided into packets and sent over the network to the destination. Even though there is no reservation of path for transmission, the packets needs to be buffered in the switches along the path to the destination. In an SOC, this means more area overhead for the switches which are not acceptable as on-chip silicon real estate is limited.

In this research work, wormhole switching is adopted wherein packets are divided into small units called flow control units or flits. The size of flit is chosen such that a single flit can traverse a single hop in a single clock cycle. These flits are transmitted along the network across switches .Hence the large buffer requirement for the switches are avoided. The first flit or the header flit of a packet contains the routing information. This information enables the switches to setup the path and the rest of the flits follow this path in a pipelined fashion [2]. But a problem associated with such a switching technique is that distinct messages cannot be sent over a switch at the same time, as the path would be reserved for a particular packet till it is completely transmitted. Hence to solve this problem a concept called virtual channels was introduced.

Basically a virtual path is reserved for each distinct message. This is accomplished by reserving separate buffers for each message in all the switches along the path, forming a distinct virtual path for each message. Fig. 4 shows a block diagram of how this is accomplished. Here node A and node B are allocated separate buffers along the path which enables the switch to receive and send messages from both the nodes, simultaneously using a multiplexer.



**Fig. 4: Network Switch with virtual channels**

WiNoC has adopted wormhole routing in which data is transferred via flits using virtual channels (VCs) [14]. WiNoC is essentially an irregular architecture and in irregular architectures it is important to achieve distributed and deadlock-free routing of data flits. This is achieved through a layered shortest path routing policy (LASH) [15]. In LASH, shortest paths between different source-destination pairs are separated into multiple virtual layers with specific VCs dedicated for each layer. This avoids cyclic dependencies between paths in a particular layer. Computation of the path for each packet would result in a large overhead hence, the shortest path between any source and destination is pre-computed offline. Each switch has a routing table, which contains only the identity of the next switch

corresponding to all possible final destinations. As a result, the memory required to store the routing table is linearly proportional to system size. When a header flit arrives at a particular switch the next switch is determined from the routing table based on the final destination of the packet. The header flit is then routed to the appropriate port along the particular VC reserved for its source/destination pair. Only the next switch is determined at each intermediate switch making the routing decision fast and efficient. Since the routing paths are the shortest paths, high data rates can be achieved with moderate number of VCs to avoid deadlock [15]. In order to grant access to the wireless channel to multiple WIs in a distributed manner, token flow control is adopted. Only after all the flits belonging to a particular packet are transmitted, the token is forwarded to the next WI. Since WIs provide shorter paths to route packets, many messages would try to access them leading to congestion. To avoid congestion at the WIs, if no buffer space is available at the wireless port of a switch then the packet is routed through the shortest available wired path.

### **3.3. Wireless Interface**


The two important WI components are the antenna and the transceiver. The on-chip antenna for the mSWiNoC has to provide the best power gain for the smallest area overhead. A metal zigzag antenna has been demonstrated to possess these characteristics and hence is used for this work [11]. To ensure high throughput and energy efficiency, the WI transceiver circuitry has to provide a very wide bandwidth as well as low power consumption.

### **3.4. Antenna and Transceiver**

Suitable on-chip antennas are necessary to establish wireless links for WiNoCs. In [13] the authors demonstrated the performance of silicon integrated on-chip antennas for intra- and inter-chip communication. They have primarily used metal zig-zag antennas operating in the range of tens of GHz. Design of an ultra-wideband (UWB) antenna for inter- and intra-chip communication is elaborated in [19]. This particular antenna was used in the design of a wireless NoC [9] mentioned earlier in chapter 1. The above mentioned antennas principally operate in the millimeter wave (tens of GHz) range and consequently their sizes are on the order of a few millimeters. If the transmission frequencies can be increased to THz/optical range then the corresponding antenna sizes decrease, occupying much less chip real estate. Characteristics of metal antennas operating in the optical and near-infrared region of the spectrum of up to 750 THz have been studied [20]. Antenna characteristics of carbon nanotubes (CNTs) in the THz/optical frequency range have also been investigated both theoretically and experimentally [21-22]. Although CNT antennas will support higher data bandwidth but significant manufacturing challenges need to be overcome to make them feasible for adoption in mainstream chip fabrication processes. That is why a metal based CMOS process compatible antenna structure is used in this work which can be adopted in the near future.

The on-chip antenna for the proposed wireless NoC has to provide the best power gain for the smallest area overhead. A metal zig-zag antenna [23] has been demonstrated to possess these characteristics. This antenna also has negligible effect of rotation (relative angle between transmitting and receiving antennas) on received signal strength, making it most suitable for on-chip wireless interconnects. This thesis work uses the zig-zag antenna

used in [3] designed with  $10\mu\text{m}$  trace width,  $60\mu\text{m}$  arm length and  $30^\circ$  bend angle. The axial length depends on the operating frequency of the antenna. The characteristics of the antennas are simulated using the ADS momentum tool. High resistivity silicon substrate ( $\rho=5\text{k}\Omega\text{-cm}$ ) is used for the simulation. The details of the antenna simulation setup and antenna structure are shown in Fig. 5(a) [24]. To represent a typical inter-subnet communication range the transmitter and receiver were separated by  $20\text{ mm}$ . The forward transmission gain ( $S_{21}$ ) of the antenna obtained from the simulation is shown in Fig. 5(b). As shown in Fig. 5(b), we are able to obtain a  $3\text{ dB}$  bandwidth of  $16\text{ GHz}$  with a center frequency of  $57.5\text{ GHz}$ . For optimum power efficiency, the quarter wave antenna needs an axial length of  $0.38\text{ mm}$  in the silicon substrate.



**Fig. 5: (a) On-chip metal zig-zag antenna (reproduced from [3]) (b) On-chip antenna placement on the die (reproduced from [23])**

### 3.5. Performance Metrics

The experiments are carried out using a cycle accurate simulator implementing the NoC architectures with 3-stage switches namely, input, output arbitrations and routing [2].

The number of VCs in the Small-world WiNoC switches depends on the system size and

the number of interconnects. As shown in [30] irregular networks of size 64, 128 and 256 cores require 4, 6 and 9 layers for deadlock-free routing. Each layer is considered to have a single VC reserved. The mesh architecture is considered to have 4 VCs in each input and output port. Each VC has a buffer depth of 2 flits. A uniform random spatial distribution of traffic is used for the all experiments. All the NoC components are driven with a 2.5GHz clock. All simulations are performed for ten thousand cycles allowing for transients to settle in the first few thousand cycles. If the wireline links are long enough to take more than 1 clock cycle for transmission of a flit they are pipelined by insertion of FIFO buffers such that between any two stages it is possible to transfer an entire flit in 1 clock cycle. The on-chip zig-zag antennas are able to provide a bandwidth of 16GHz around a center frequency of 60GHz [3] while the transceivers [23] are able to sustain a maximum data rate of 6Gbps. All the wireless switches are equipped with the same transceivers. We have considered a flit size of 32 bits and a packet size of 64 flits.

The metrics for performance evaluation are maximum achievable bandwidth and packet energy dissipation. Maximum achievable bandwidth is the peak sustainable data rate in number of bits successfully routed per second. Bandwidth,  $B$  can be determined as,

$$B = t\beta Nf \quad (2)$$

Where,  $t$  is the maximum throughput in number of flits received per core per clock cycle at network saturation,  $\beta$  is the number of bits in a flit,  $N$  is the number of cores in the NoC and  $f$  is the clock frequency. The throughput is directly obtained from system level simulations performed by the NoC simulator.

The packet energy dissipation,  $E_{pkt}$  is the average energy dissipated in transmission of a packet from source to destination over the NoC. It can be measured as,



$$E_{pkt} = \frac{(\sum_{i=1}^{N_{pkt}} (L_i - h_i \lambda) E_{buf} + h_i E_{wire} \lambda) + N_{sim} E_{wireless}}{N_{pkt}} \quad (3)$$

Where,  $N_{pkt}$  is the number of packets routed in the NoC,  $L_i$  is the latency of the  $i^{\text{th}}$  packet,  $h_i$  is the number of hops in the path of the packet and  $E_{buf}$  is the energy dissipation of a flit in the NoC switch buffers. The energy dissipation of a wireline hop is  $E_{wire}$  and  $\lambda$  is the packet length in number of flits.  $N_{sim}$  is the duration of the simulation and  $E_{wireless}$  is the energy dissipated by all the wireless transceivers in the WiNoC in one cycle.

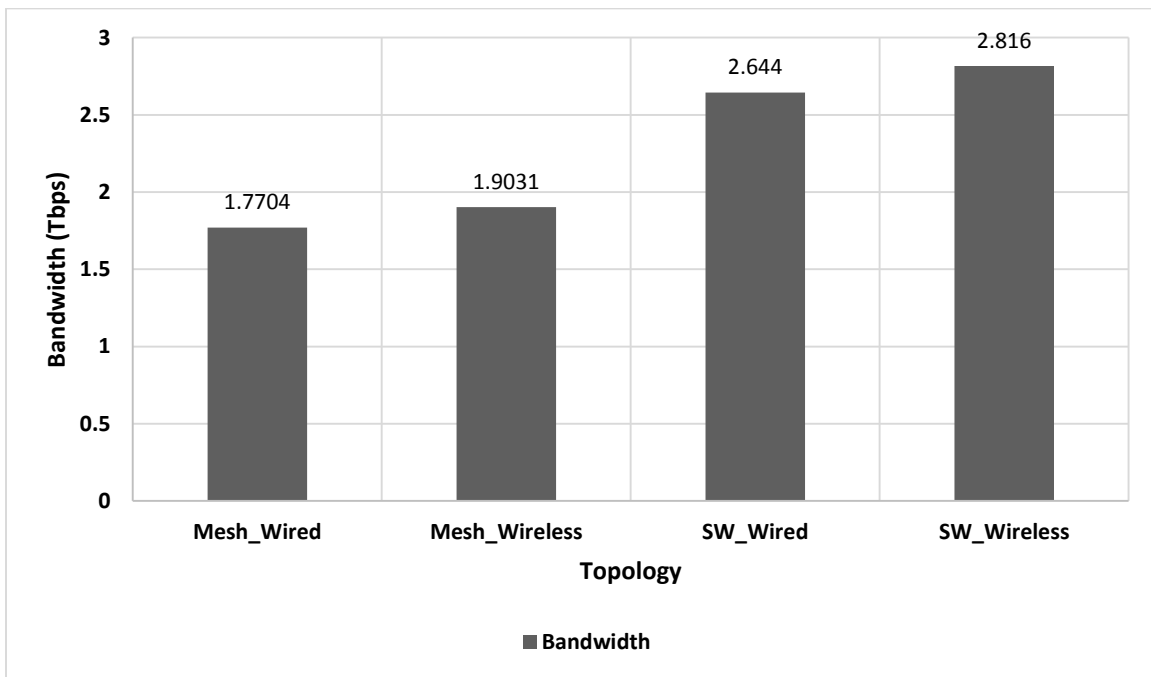
### 3.6. Performance Evaluation

#### Experimental Setup:

In this section, a complete evaluation on the basis of bandwidth, latency and energy dissipation is carried out comparing the wired and wireless network architectures of SWNoC and mesh-based NoC. GEM5, a full system simulator is used to obtain detailed network-level information. An 8x8 core network with 64 routers, each with 8 virtual channels and 16 flit buffers per input port is assumed. Fixed length packets of 64 flits where, the head flit leading 63 body flits, and each flit being 32-bits wide are assumed. Similar to the wired links, wireless links are also incorporated with warm-hole routing. The NoC simulator uses switches synthesized from an RTL level design using TSMC 65nm CMOS process, using Synopsys Design Vision. Energy dissipation of the network switches were obtained from the synthesized net-list by running Synopsis Prime Power, while energy dissipated by wireline links were obtained through HSPICE simulations taking into consideration length of the wireline links. Each wireless link can sustain data rate of 16Gbps and has an energy dissipation of 2.3pJ/bit [2].

### Performance Characteristics:

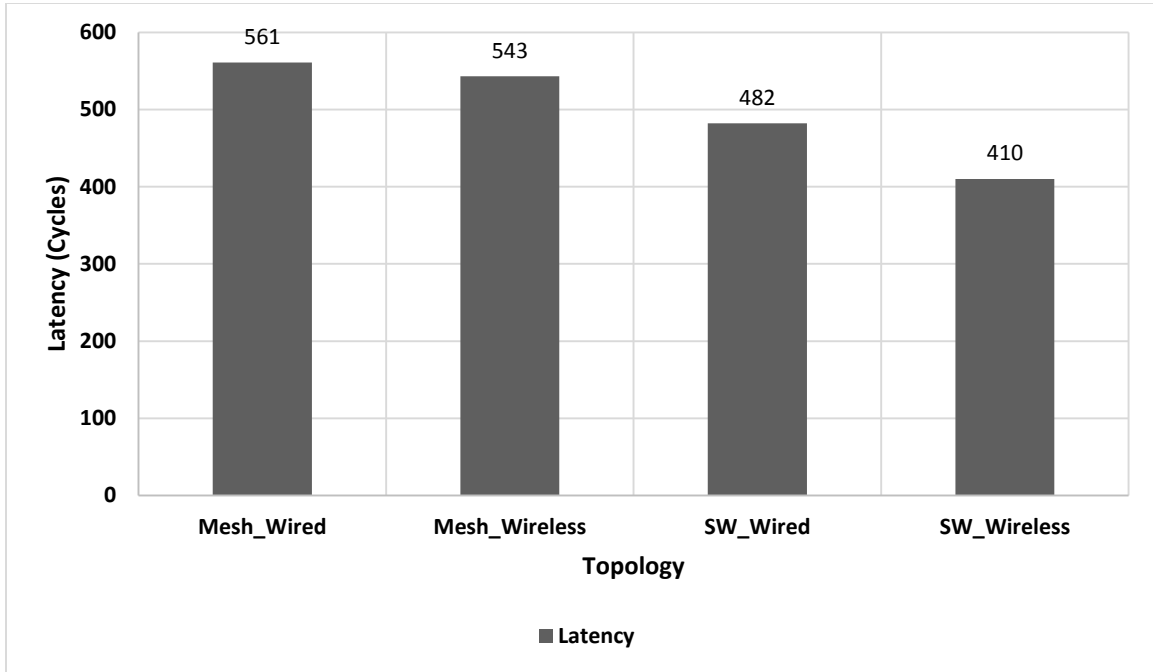
Here, presented is the bandwidth, latency and packet energy profiles of the wired and wireless implementation of small-world and mesh based NoCs. Fig. 6 shows the bandwidth characteristics, Fig. 7 shows the latency characteristics and Fig. 8 shows the energy dissipation profile for different NoC topologies. It can be observed from Fig. 6 that when comparing traditional mesh based NoCs, small-world topology has better performance in terms of bandwidth and particularly, small-world with 10 wireless nodes have the superior performance over other implementations.



**Fig. 6: Bandwidth characteristics of mesh and SWNoC**

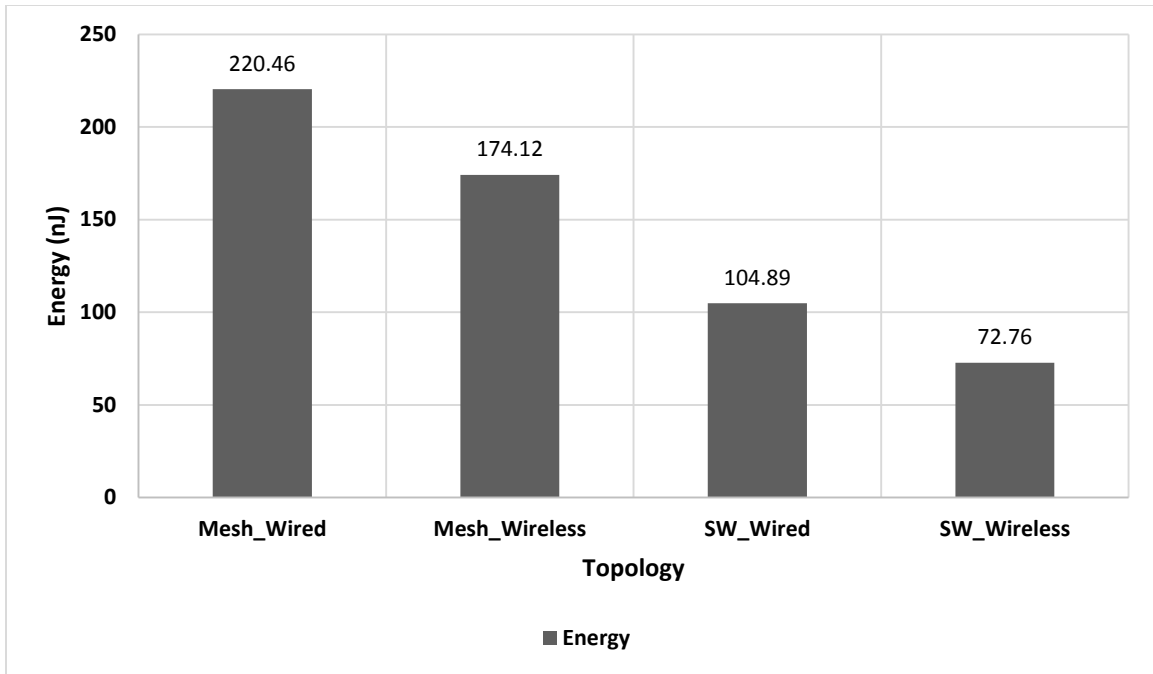
It can be observed from Fig. 7 that for all of the network topologies considered, the latency of SWNoC is lower than that of the mesh architectures. This is due to the small-

world architecture of SWNoC with direct long-range, one-hop wireless links that enables a smaller average hop-count than that of mesh. With this significant decrease to the overall latency in the SWNoC architecture, an opportunity is formed to further increase energy savings to match the performance of SWNoC with that of the baseline mesh architectures.



**Fig. 7: Latency characteristics of mesh and SWNoC**

Now, evaluating the energy dissipation characteristics of the mesh and small-world with wired and wireless counterparts, it is evident that introduction of wireless links in the topology could avoid energy dissipation of few long range wireline links. In addition, the small-world topology reduces the multi-hop communication and thus ensures less energy dissipation due to long range wireline links. Further, by replacing long range wireline links by wireless links, the energy dissipation is significantly reduced as shown in Fig. 8



**Fig. 8: Energy characteristics of mesh and SWNoC**

Hence, it was clearly seen from these performance analysis that small-world wireless network-on-chip architectures significantly improves overall energy dissipation and performance when compared to mesh-based architectures.

## **Chapter 4    Dynamic Voltage Frequency Scaling for WiNoC**

It is established that Small-world Wireless Network-on-Chip is an enabling architecture to improve power efficiency and performance characteristics of multi-core architectures. The inherent SWNoC architecture modifies the distribution of network traffic patterns among network elements significantly. The execution flow of a program on a multicore NoC generally contains periods of heavy computation followed by periods of inter-core data exchange. During periods of high computation, network usage may be at a minimum, allowing the voltage and frequency of links and switches to be tuned down in order to save energy. Hence, it is possible to vary the voltage and frequency of the SWNoC switches and links depending on the traffic-dependent bandwidth requirements.

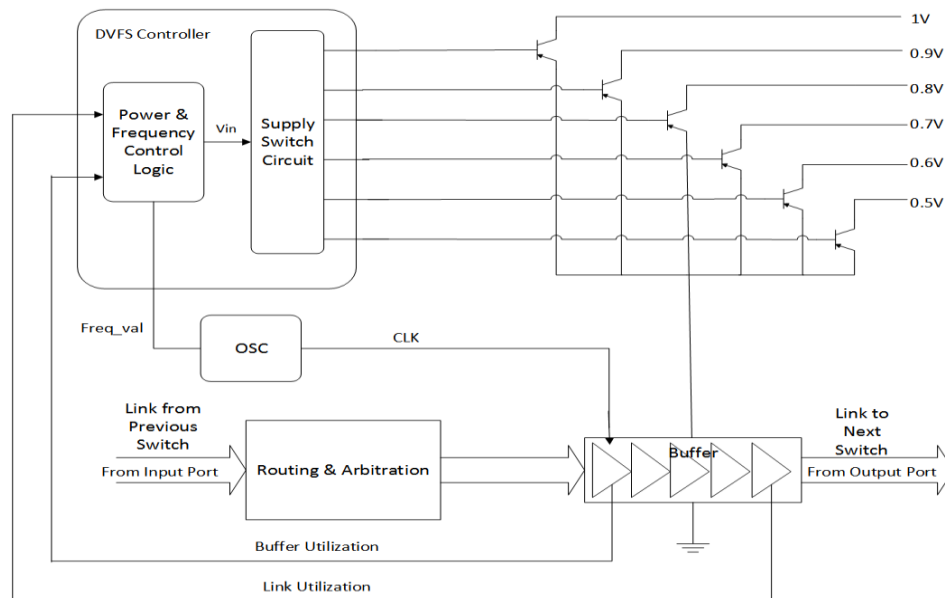
Here in this thesis, a fully distributed fine-grain DVFS is employed on switches and links, where the ports and links are tuned according to their utilizations following a history based algorithm that predicts the future traffic characteristics on the network based on what was seen in the past. The utilization characteristic is chosen to be a relevant metric to determine whether DVFS should be performed.

### **4.1.    DVFS Architecture and Modeling**


Most DVFS architectures apply only a single DVFS controller to an entire chip using an on-chip or off-chip DC-DC converter. A fine grain DVFS implementation can increase the effectiveness of DVFS by tuning the supply voltage to individual parts of the chip. One way to achieve this is to supply discrete voltages to the chip, and have the individual switches and links to switch between these voltages.

Fig. 9 shows a concept diagram of DVFS with five voltage supplies using PMOS power gates [8]. Current flow through power gate transistors result in a voltage drop that negatively impacts performance. The amount of voltage drop,  $V_{PG}$ , is related to the dimensions of the power gates:  $V_{PG} = I_{PG}R_{PG}$ , where  $I_{PG}$  is the current through the power gates,  $R_{PG}$  is directly related to  $L/W$  where  $L$  and  $W$  are the length and width of the power gate transistors respectively. The voltage drop causes an increase in the power gate's delay. Voltage drop can be reduced by making  $W/L$  as large as possible, which can be accomplished by adding power gates in parallel.

To accurately measure the performance loss associated with the power gates, a precise current profile from the processor core is first obtained with SPICE simulations. This current waveform is then used to create the voltage drop across the power gates, and the resulting increase in delay can be measured. In 65nm technology, the relationship between power gate width and performance is shown in Fig.10.



**Fig. 9: DVFS Mechanism on Switches and Links**



**Fig. 10: Power gate transistor width versus processor performance [56]**

The DVFS controller in Fig. 9 contains the logic to estimate the voltage and frequency that is required by the buffers and the links. This estimate is obtained based on the Buffer and Link Utilization values that is obtained from the buffers in the switch. The decision of voltage and frequency is made using an algorithm that is explained later in this thesis. Frequency scaling is performed by incrementing or decrementing the clock frequency based on the utilization information. A range of allowable frequencies is assigned for each voltage setting which is shown in Table 1, where the setting of  $V_{in}$  are mapped to settings of Freq\_val. Therefore frequency scaling is performed automatically depending on the voltage setting of the buffers and links.

## 4.2. History-based DVFS

The mechanism controlling the DVFS has to carefully trade off power and performance, minimizing the power consumption of the network while maintaining high performance. Hence, in this thesis a distributed history-based DVFS policy is proposed. In this policy, the router port predicts future communication workload based on the analysis of prior traffic, then dynamically adjusts the voltage and the corresponding frequencies of its buffers and links to accommodate the network load.

### 4.2.1 Network Traffic Characteristics

Network communication traffic characteristics can be captured with various network traffic measures. In order to predict the network load based on what was seen in the past, a suitable indicator has to be explored over a fixed time window. The metric to determine whether DVFS should be performed is utilization. Tuning a given link and buffer's voltage and frequency is determined by the link utilization and buffer utilization respectively.

#### Link Utilization

$$LU = \frac{\sum_{t=1}^N F(t)}{N} \quad 0 < LU < 1$$

Where,

$LU$  is the Link Utilization.

$N$  is the number of clock cycles, which is sampled within a history window size  $H$ .

$F(t)=1$  if traffic passes the link  $i$  in cycle  $t$ , else  $F(t)=0$  if no traffic passes link  $i$  in cycle  $t$ .



Link Utilization is a direct measure of the traffic workload in the links of the network. First, the number of flits that gets transferred using a particular link is captured over a time window, then the Utilization is measure as a ratio of total flits transferred to the number of clock cycles that is sampled within a history time window. Assuming that, a flit takes one time unit to get transferred over the link. Link utilization can take any value between 0 and 1. A higher link utilization ( $\geq 0.5$ ) reflects that more data are sent to the next router. As the history is predictive, this indicates a higher link voltage and frequency is needed to meet the performance requirement. Conversely, lower link utilization ( $< 0.5$ ) implies the existence of more idle cycles. Hence, decreasing the link frequency can lead to power savings without significantly affecting performance.

To investigate how the predicted link utilization vary the frequency based on the predictions, the utilization values and the clock speed is tracked with a two dimensional 8x8 mesh network. Network traffic is generated based on the bench mark traffic patterns, for this study a uniform traffic pattern was considered. Fig. 10(a) shows the utilization of a single link as sampled every 100 cycles ( $H = 100$ ), across the entire timing simulation of uniform load. It can be seen that at low traffic workloads ( $LU < 0.5$ ), contention of the link is less and hence the utilization is lower, correspondingly the operating frequency of the links is also lowered. When contention starts to build in the link and when the utilization crosses the threshold 0.5, the link operating frequency is increased and reaches its maximum operating frequency as work load increases. The graph clearly shows the two extreme network scenarios of lightly congested and heavily loaded along with its frequency requirements for DVFS operation. At low network loads, since the flit will not be stalled in the succeeding router, any increase in link delay directly contributes to overall packet

latency. At high network loads, flits will be stalled in the next router for a long time, hence getting there faster will not be significant. In this case, link frequency can be decreased more aggressively with minimum delay constraints. Hence, link utilization alone will not be sufficient for guiding the history-based DVFS policy. One more measure, input buffer utilization is investigated.

### **Input Buffer Utilization**

$$BU = \frac{\sum_{t=1}^N (E(t) / B)}{H} \quad 0 < BU < 1$$


Where,

BU is the Buffer Utilization


$E(t)$  is the number of VC's in input buffer that are occupied at time  $t$ .

B is the total input buffer size.

Buffer Utilization tracks how many VC's in each buffer switch are occupied over a time window. As traffic increases, more flits are stored in the buffers which may lead to contention and increase in the utilization measure. Buffer utilization is calculated over a time window as a fraction of number of flits occupied in a buffer to the total capacity of the buffer. Buffer utilization can take any value between 0 and 1. The input buffer utilization is tracked downstream from the same link shown in Fig. 10(a) to investigate how frequency scaling behaves with network traffic and contention. Fig. 10(b) shows the graph that plots the buffer utilization and the frequency of operation that scales based on the utilization values.



(a)



(b)

**Fig. 11: Utilization profile for uniform traffic (a) Link Utilization profile (b) Buffer Utilization profile**

## 4.2.2 History-based DVFS Policy

Network traffic exhibits two dynamic trends: transient fluctuations and long-term transitions. This history-based DVFS policy filters out short-term traffic fluctuations and adapts link frequencies and voltages carefully to traffic transitions. This is carried out by sampling link and input buffer utilization within a fixed history window and using the exponential weighted average utilization to combine current and past utilizations history.

### Algorithm 1:

Assume that initially,  $LU_{past} = 1$  &  $W =$  any finite number.

*for each(time window) begin*

$$LU_{predicted} = (W * LU_{current} + LU_{past}) / (W + 1)$$

$$LU_{past} = LU_{predicted}$$

$$BU_{predicted} = (W * BU_{current} + BU_{past}) / (W + 1)$$

$$BU_{past} = BU_{predicted}$$

$$LinkVolt = 1v$$

$$BufferVolt = 1v$$

$$ThresholdVolt = 0.5v$$

*if*( $LU_{predicted} < ThresholdVolt$  &&  $LinkVolt < 1$ )

$$LinkVolt = LinkVolt + 0.1$$

*elseif*( $LU_{predicted} \geq ThresholdVolt$  &&  $LinkVolt > 0.5$ )

$$LinkVolt = LinkVolt - 0.1$$

*end*

*if*( $BU_{predicted} < ThresholdVolt$  &&  $BufferVolt < 1$ )

$$BufferVolt = BufferVolt + 0.1$$


*elseif*( $BU_{predicted} \geq ThresholdVolt$  &&  $Buffer Volt > 0.5$ )

$$BufferVolt = BufferVolt - 0.1$$

*end*

*endfor*

Given the predicted communication link utilization  $LU_{predicted}$  and input buffer utilization  $BU_{predicted}$ , the DVFS policy dynamically adapts its voltage scaling and frequency scaling to achieve power savings with minimal impact on performance. It determines whether to increase link voltage and frequency to the next higher level, decrease link voltage and frequency to the next lower level, or do nothing. So, when a link is going to be highly utilized, voltage frequency scaling is carried out to handle load. Similarly, if a link is mostly idle, DVFS is carried out to save power. Otherwise, voltage frequency scaling is conservatively carried out to minimize impact on performance. The states depends on five voltage and frequency levels shown in Table 1. The pseudo-code of the proposed DVFS policy is shown in Algorithm 1. The threshold voltage 0.5v is chosen based on the plot shown in Fig.12. Since beyond 0.5v there much increase in delay which can considerably affect the system performance.



**Fig. 12: Voltage and Delay Curve**

### 4.2.3 DVFS Controller and Hardware Implementation

The proposed DVFS policy relies only on local link and buffer utilizations information. This avoids communication overheads in relying global information, and permits a simple hardware implementation. To measure link utilization, a counter at each output port gathers the total number of cycles that are used to relay flits in each history interval. Another counter captures the ratio between the router and link clocks. A simple multiplier combines these two counters to calculate the link utilization. Most interconnection network routers use credit-based flow control. Current buffer utilization is already available. Two registers store  $LU_{past}$  and  $BU_{past}$ , which feed the circuit module calculating the exponential weighted average. Finally some combinational logic performs the threshold comparisons and outputs signal that control the DVFS links and buffers.

Fig. 11 shows the finite state machine that was built inside the controller to switch between different voltage and frequency levels based on the utilization values of the link and buffers. Where, LU specifies Low Utilization (Utilization  $< 0.5$ ) and HU specifies High Utilization (Utilization  $\geq 0.5$ ).

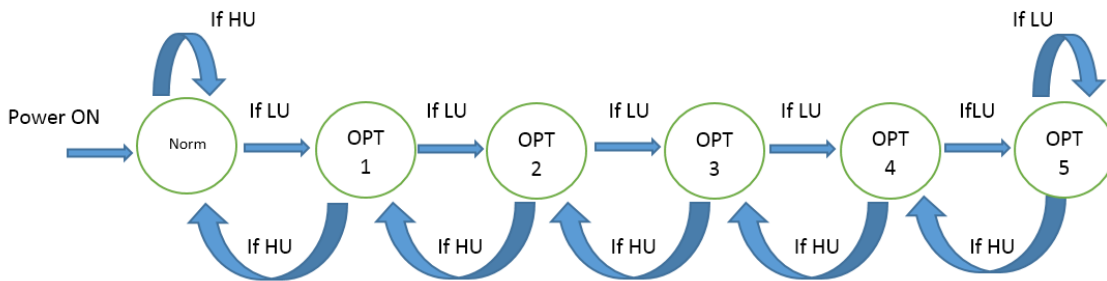


Fig. 13: Finite State Machine for DVFS Controller

**Table 1: Voltage/Frequency/Threshold Combinations**

<b>States</b>	<b>Voltage (Volts)</b>	<b>Frequency (GHz)</b>
<b>Normal</b>	1	2.5
<b>OPT 1</b>	0.9	2.25
<b>OPT 2</b>	0.8	2
<b>OPT 3</b>	0.7	1.75
<b>OPT 4</b>	0.6	1.5
<b>OPT 5</b>	0.5	1.25

The DVFS controller were synthesized from RTL level design using 65nm standard cell libraries from CMP [23], using Synopsys. The delay, area and power numbers are shown in Table 2. The controller results with an area of about 300 equivalent logic gates per router port. As the circuit does not lie on the critical path of the router, its delay can be ignored. With power estimation for the circuit, it was found that the power overhead is negligible which approximately 100nW per router port.

**Table 2: DVFS Controller in Overall System metrics**

	<b>DVFS Controller</b>
<b>Power</b>	15.332 $\mu$ W (~2% of Overall Chip)
<b>Area</b>	44850 $\mu$ m <sup>2</sup> (~3% of Overall Chip)
<b>Delay</b>	0.11ns (within max clock period 0.4ns)

### **4.3. Performance Evaluation of DVFS-enabled WiNoCs**

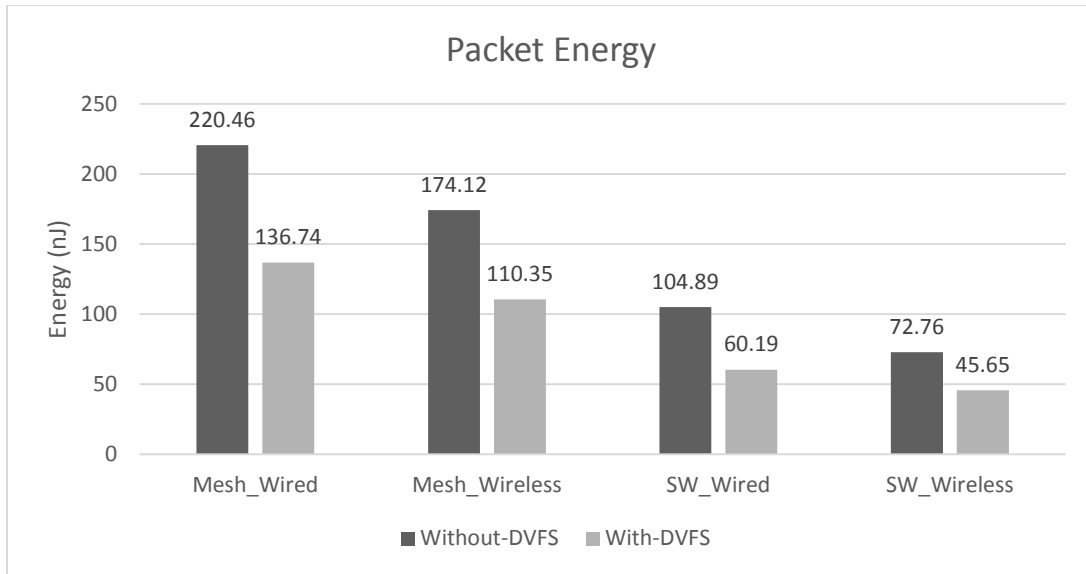
This section evaluates the performance of the history-based DVFS policy with conservative DVFS enabled wireline links and input buffers on a NoC environment. With the same experimental setup as explained in chapter 3, considering a buffer depth of 16 and 8 virtual channels, the DVFS enabled NoC architectures are evaluated by studying the trade-off between network latency/throughput degradation and dynamic power savings. One of the key goals of this thesis is to uncover the effect of DVFS enabled links and input buffers on network power and performance, providing insights that will guide future design of DVFS enabled Wireless NoCs.

#### **4.3.1 Energy Dissipation Characteristics**

This subsection presents the network-level energy dissipation of the SWNoC by incorporating the DVFS technique described earlier. For completeness, the characteristics of the conventional wireline and wireless mesh architecture incorporating DVFS technique is shown.

Fig. 12 shows the network energy dissipation for the various architectures for 16 buffer depth and 8 virtual channels. It can be observed from Fig. 12 that among different architectures, the network energy is much lower for the SWNoC. The two main contributors of the energy dissipation are the switches and the interconnect infrastructure, In the SWNoC, the overall switch energy decreases significantly compared to mesh as a result of the better connectivity of the architecture.






**Fig. 14: Total Network Energy for different NoC Architectures**

In this case, the hop count decreases significantly, and hence, on an average, packets have to traverse through less number of switches and links. In addition, a significant amount of traffic traverses through the energy efficient wireless channels in SWNoC; consequently decreasing the interconnect energy dissipation. With the addition of DVFS, total network energy can be further reduced. As the traffic traversing through the wireline links is heavily reduced in the SWNoC, the opportunity for implementing DVFS is significant. From Fig. 12 it is clear that DVFS-enabled SWNoC saves 40% of energy with respect to the Non\_DVFS SWNoC implementation.


Fig. 13 shows the network energy savings on DVFS-enabled SWNoC architecture for different traffic benchmarks. It is clear that for different traffic patterns, the DVFS-enabled SWNoC architecture can still provide approximately 40% savings in energy dissipation for all traffic patterns.




**Fig. 15: Total Network Energy for Map-reduced traffic**

#### 4.3.2 Bandwidth and Latency Characteristics

This section accounts the bandwidth and latency penalty for implementing DVFS mechanism on NoC architectures. Fig. 14 shows the peak achievable bandwidth for NoC architectures with and without DVFS. It is clear that there is bandwidth penalty with the dynamic voltage scaling on wired and wireless system regardless. However, the penalty that is to be accounted for is only 3-4% of the peak bandwidth that was achieved without implementing DVFS. The difference in bandwidth is mainly due to the fact that the links and input buffers are switched between different voltages and frequency values. The bandwidth reduction is the penalty for frequent switching characteristics of the DVFS policy. Fig.16 clearly shows the voltage fluctuations that does not follow the predictions because of the smaller steps. Hence there is bandwidth and latency limitations for different traffic patterns.

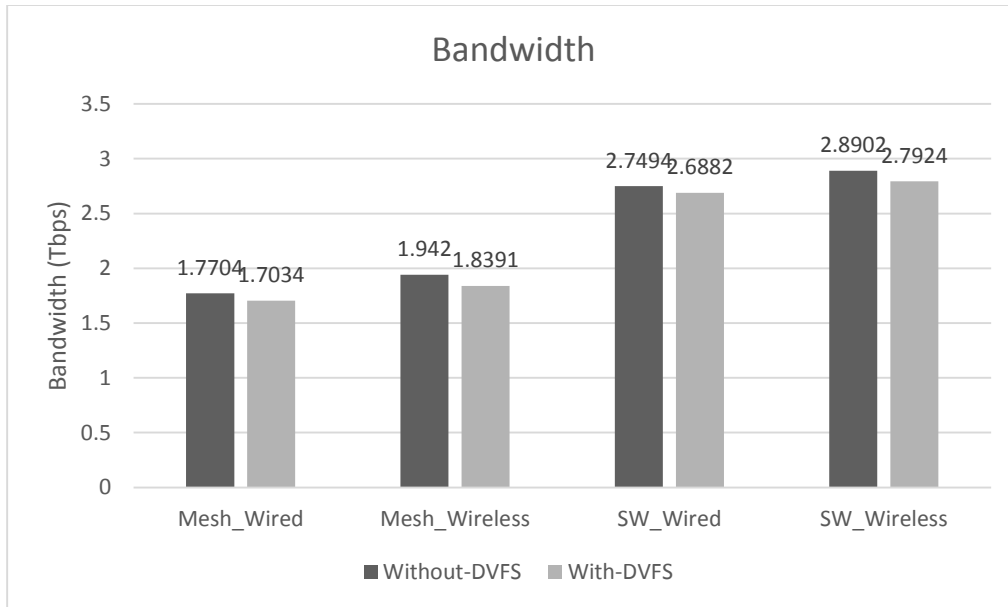


(a)



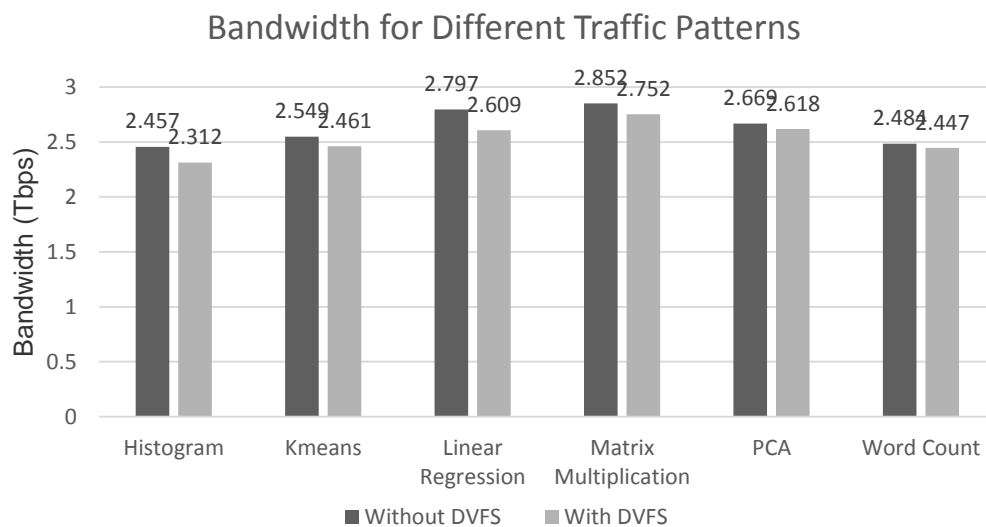
(b)

**Fig. 16: Utilization profile for Map-reduced traffic (a) Link Utilization profile (b) Buffer Utilization profile**




**Fig. 17: Bandwidth characteristics for different architectures**

This degradation will not be a dominant concern when there is 40% energy reduction in the system. Still the DVFS-enabled small-world wireless performance is better than the small-world wired implementation. In low-power applications, 3-4% performance compromise is acceptable with 40% energy savings




**Fig. 18: Bandwidth characteristics for Map-reduced traffic**

Latency penalty is not very significant as shown in Fig. 15. However, the latency is high in few architectures because of slowing down of the links and buffers by frequency scaling, where one flit takes 2 cycles to be transfers from one link to another instead of 1 flit per cycle. Hence, there is a negligible amount of increase in cycles in few NoC architectures. However, the DVFS-enabled SWNoC shows very less change in the latency numbers which still proves to be a better architecture for maintain the desired frequency with significant amount of energy savings.



**Fig. 19: Latency characteristics for different architectures**



**Fig. 20: Latency characteristics for Map-reduced traffic**

## Chapter 5 Conclusion and Future Work

In this thesis, it is demonstrated that how a small-world DVFS-enabled wireless network-on-chip improves the energy dissipation of a multi-core chip. By adopting a small-world interconnection infrastructure, where long distance communications will be predominantly achieved through high performance specialized single hop wireless links, communications can be made significantly more energy efficient. To further extend the energy savings, implementing network level DVFS on wireline links and input buffers, a significant energy savings can be achieved. Just as DVFS in microprocessors exploits the variance in processor utilization to tune processor voltage and frequencies for power savings, network level DVFS allows the links and buffers to be tuned for power efficiency as network utilization varies. To incorporate a promising DVFS policy that maximizes dynamic power saving while minimizing latency and throughput degradation, a history-based DVFS policy was proposed. This policy uses the past network history to predict future network needs and carefully control the frequency and voltage of the links and buffers. The performance evaluation of this history-based DVFS policy on SWNoC architecture showed that it could achieve 40% energy savings over the traditional SWNoC without network level DVFS. As power becomes increasingly as, if not more, important than performance in interconnection networks, there is a clear need for power management mechanisms that target network power efficiency while maintaining good latency/throughput performance. This research demonstrates the effectiveness of network level DVFS as a power optimization mechanism for interconnection networks.

This work can be extended by complementing the energy savings at the network level with suitable methodologies to improve the energy dissipation of the computational cores within the DVFS-enabled SWNoC framework. In addition, a suitable on-chip wireless transceiver design with power gating can also contribute a significant amount of energy savings. With DVFS implemented both at network level and processor level along with power gated wireless transceivers, larger energy savings can be realized over WiNoCs.

## Bibliography

- [1] Mishra, A.K.; Das, R.; Eachempati, S.; Iyer, R.; Vijaykrishnan, N.; Das, C.R., "A case for dynamic frequency tuning in on-chip networks," *Microarchitecture*, 2009. MICRO-42. 42nd Annual IEEE/ACM International Symposium on , vol., no., pp.292,303, 12-16 Dec. 2009
- [2] Murray, J.; Pande, P.P.; Shirazi, B., "DVFS-enabled sustainable wireless NoC architecture," *SOC Conference (SOCC)*, 2012 IEEE International , vol., no., pp.301,306, 12-14 Sept. 2012
- [3] Ogras, U.Y.; Marculescu, R., "'It's a small world after all': NoC performance optimization via long-range link insertion," *Very Large Scale Integration (VLSI) Systems*, *IEEE Transactions on* , vol.14, no.7, pp.693,706, July 2006
- [4] Wonyoung Kim; Gupta, M.S.; Gu-Yeon Wei; Brooks, D., "System level analysis of fast, per-core DVFS using on-chip switching regulators," *High Performance Computer Architecture*, 2008. HPCA 2008. IEEE 14th International Symposium on , vol., no., pp.123,134, 16-20 Feb. 2008
- [5] Wooyoung Jang; Pan, D.Z., "A Voltage-Frequency Island Aware Energy Optimization Framework for Networks-on-Chip," *Emerging and Selected Topics in Circuits and Systems*, *IEEE Journal on*, vol.1, no.3, pp.420,432, Sept. 2011
- [6] Ogras, U.Y.; Marculescu, R.; Marculescu, D., "Variation-adaptive feedback control for networks-on-chip with multiple clock domains," *Design Automation Conference*, 2008. DAC 2008. 45th ACM/IEEE , vol., no., pp.614,619, 8-13 June 2008
- [7] Niyogi, K.; Marculescu, D., "Speed and voltage selection for GALS systems based on voltage/frequency islands," *Design Automation Conference*, 2005. Proceedings of the ASP-DAC 2005. Asia and South Pacific , vol.1, no., pp.292,297 Vol. 1, 18-21 Jan. 2005
- [8] Donald, J.; Martonosi, M., "Techniques for Multicore Thermal Management: Classification and New Exploration," *Computer Architecture*, 2006. ISCA '06. 33rd International Symposium on , vol., no., pp.78,88, 0-0 0



- [9] Shang, L.; Li-Shiuan Peh; Kumar, A; Jha, N.K., "Temperature-Aware On-Chip Networks," *Micro, IEEE* , vol.26, no.1, pp.130,139, Jan.-Feb. 2006  
doi: 10.1109/MM.2006.23
- [10] Jiong Luo; Li-Shiuan Peh; Niraj Jha, "Simultaneous dynamic voltage scaling of processors and communication links in real-time distributed embedded systems," *Design, Automation and Test in Europe Conference and Exhibition, 2003* , vol., no., pp.1150,1151, 2003
- [11] Deb, S.; Ganguly, A.; Pande, P.P.; Belzer, B.; Heo, D., "Wireless NoC as Interconnection Backbone for Multicore Chips: Promises and Challenges," *Emerging and Selected Topics in Circuits and Systems, IEEE Journal on* , vol.2, no.2, pp.228,239, June 2012
- [12] Garg, Siddharth; Marculescu, D.; Marculescu, R.; Ogras, U., "Technology-driven limits on DVFS controllability of multiple voltage-frequency island designs: A system-level perspective," *Design Automation Conference, 2009. DAC '09. 46th ACM/IEEE* , vol., no., pp.818,821, 26-31 July 2009
- [13] Wettin, Paul; Murray, Jacob; Pande, Partha; Shirazi, Behrooz; Ganguly, Amlan, "Energy-efficient multicore chip design through cross-layer approach," *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2013* , vol., no., pp.725,730, 18-22 March 2013
- [14] J. Duato, S. Yalamanchili, and L. Ni, *Interconnection Networks: An Engineering Approach*. Morgan Kaufmann, 2003, p. 600.
- [15] O. Lysne, T. Skeie, S. -a. Reinemo, and I. Theiss, "Layered routing in irregular networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 1, pp. 51–65, Jan. 2006.
- [16] J. Lin, H. Wu, and Y. Su, "Communication using antennas fabricated in silicon integrated circuits," *Solid-State Circuits*, vol. 42, no. 8, pp. 1678–1687, 2007.
- [17] A. Ganguly and Vineeth vijayakumaran, Manoj Prashanth Yuvaraj, Naseef Mansoor, "CDMA Enabled Wireless Network-on-Chip,".
- [18] R. Jotwani and S. Sundaram, "An x86-64 core implemented in 32nm SOI CMOS," *Proc. IEEE Int. Solid-State Circuits Conf.*, pp. 106–107, 2010.

- [19] P.P. Pande, C. Grecu, M. Jones, A. Ivanov, and R. Saleh, "Performance Evaluation and Design Trade-Offs for Network-on-Chip Interconnect Architectures," *IEEE Tran. Computers*, vol. 54, no. 8, pp. 1025-1040, Aug. 2005.
- [20] R. Ho, K.W. Mai, and M.A. Horowitz, "The Future of Wires," *Proc. IEEE*, vol. 89, no. 4, pp. 490-504, Apr. 2001.
- [21] P. Kapur, J.P. Mc Vittie, and K.C. Saraswat, "Technology and Reliability Constrained Future Copper Interconnects—Part II: Performance Implications," *IEEE Tran. Electron Devices*, vol. 49, no. 4, pp. 598-604, Apr. 2002.
- [22] D. Sylvester and K. Keutzer, "Impact of Small Process Geometries on Microarchitectures in Systems on a Chip," *Proc. IEEE*, vol. 89, no. 4, pp. 467-489, Apr. 2001.
- [23] Semiconductor Industry Association (SIA). (2003). International Roadmap for Semiconductors, 2003 edition, Austin, TX. International SEMATECH, 2003. [Online]. Available: <http://www.itrs.net/links/2003itrs/home2003.htm>
- [24] C. Grecu, P.P. Pande, A. Ivanov, and R. Saleh, "Structured Interconnect Architecture: A Solution for the Non-Scalability of Bus-Based SoCs," *Proc. Great Lakes Symp. VLSI*, pp. 192-195, Apr. 2004.
- [25] C. Hsieh and M. Pedram, "Architectural Energy Optimization by Bus Splitting," *IEEE Tran. Computer-Aided Design*, vol. 21, no. 4, pp. 408-414, Apr. 2002. 41
- [26] M. Horowitz and B. Dally, "How Scaling Will Change Processor Architecture," *Proc. Int. Solid-State Circuits Conf.*, pp. 132-133, Feb. 2004.
- [27] J. Nurmi, *Interconnect-Centric Design for Advanced SoC and NoC*. Springer Science + Business Media Inc., Germany, 2005.
- [28] M. Sgroi, M. Sheets, A. Mihal, K. Keutzer, S. Malik, J. Rabaey, and A. Sangiovanni-Vincentelli, "Addressing the System-on-a-Chip Interconnect Woes through Communication based Design," *Proc. 38th Design Automation Conf.*, Las Vegas, pp. 667-72, Jun. 2001.
- [29] P. Macken, M. Degrauwe, M. V. Paemel, and H. Oguey, "A Voltage Reduction Technique for Digital Systems," *IEEE Int. Solid-State Circuits Conf.*, pp. 238-239, Feb. 1990.

- [30] C. Lai, J. H. Lin, and Y. F. Wang, "DVFS SoC Architecture and Implementation," *SoC Technology Journal*, vol. 3, pp. 84–91, Nov. 2005.
- [31] C. Isci, A. Buyuktosunoglu, C.-Y. Cher, P. Bose, and M. Martonosi, "An Analysis of Efficient Multi-core Global Power Management Policies: Maximizing Performance for a Given Power Budget," *Proc. 39th Annu. IEEE/ACM Int. Symp. Microarchitecture*, vol. 26, no. 1, pp. 119-129, Feb. 2006.
- [32] G. Semeraro, G. Magklis, R. Balasubramonian, D. H. Albonesi, S. Dwarkadas, and M. L. Scott, "Energy-efficient Processor Design Using Multiple Clock Domains with Dynamic Voltage and Frequency Scaling," *Int. Symp. High-Performance Computer Architecture*, pp. 29-40, Feb. 2002. 42
- [33] T. Simunic, L. Benini, A. Acquaviva, P. Glynn, and G. D. Micheli, "Dynamic Voltage Scaling and Power Management for Portable Systems," *Design Automation Conf.*, pp. 524-529, Jun. 2001.
- [34] Q. Wu, P. Juang, M. Martonosi, and D. W. Clark, "Voltage and Frequency Control with Adaptive Reaction Time in Multiple-Clock-Domain Processors," *11th Int. Symp. High-Performance Computer Architecture*, pp. 178-189, Feb. 2005.
- [35] W. Kim, M. Gupta, G. Y. Wei, and D. Brooks, "System Level Analysis of Fast, Per-core DVFS Using On-chip Switching Regulators," *Int. Symp. High Performance Computer Architecture*, pp. 123–134, Feb. 2008.
- [36] U.Y. Ogras, R. Marculescu, P. Choudhary, and D. Marculescu, "VoltageFrequency Island Oartitioning for GALS-Based Networks-on-Chip," *Proc. 44th Annu. Design Automation Conf.*, pp. 110-115, Jun. 2007.
- [37] D. Bertozzi, "NoC Synthesis Flow for Customized Domain Specific Multiprocessor Systems-on-Chip," *IEEE Tran. Parallel and Distributed Systems*, vol. 16, no. 2, pp. 113-129, Feb. 2005.
- [38] J. Dielissen, A. Radulescu, K. Goossens, and E. Rijpkema, "Concepts and Implementation of the Philips Network-on-Chip," *IP-based SoC Design*, Nov. 2003.
- [39] 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. Design Automation and Test in Europe (DATE), pp. 890-895, Feb. 2004.  
43

[40] Y. S. Dhillon, A. U. Diril, A. Chatterjee, and H. S. Lee, “Algorithm for Achieving Minimum Energy Consumption in CMOS Circuits Using Multiple Supply and Threshold Voltages at the Module Level,” Proceedings of ICCAD, pp. 693-700, Nov. 2003.

[41] K. Niyogi and D. Marculescu, “Speed and Voltage Selection for GALS Systems Based on Voltage/Frequency Islands,” Proceedings of ASP-DAC, pp. 292-297, Jan. 2005.

[42] M. Powell, S.-H Yang, B. Falsafi, K. Roy, and T.N. Vijaykumar, “Reducing Leakage in a High-Performance Deep-Submicron Instruction Cache,” IEEE Tran. VLSI Systems, vol. 9, no. 1, pp. 77-89, Feb. 2001.

[43] S. Shigematsu, S. Mutoh, Y. Matsuya, and J. Yamada, “A 1-V High-Speed MTCMOS Circuit Scheme for Power-Down Application Circuits,” IEEE Journal on Solid-State Circuits, vol. 32, no. 6, pp. 861-869, Jun. 1997.

[44] B.H. Calhoun, F.A. Honore, and A.P Chandrakasan, “A Leakage Reduction Methodology for Distributed MTCMOS,” IEEE Journal on Solid-State Circuits, vol. 39, no. 5, pp. 818-826, May 2004.

[45] C. Long and L. He, “Distributed Sleep Transistor Network for Power Reduction,” Proc. IEEE/ACM Design Automation Conf., pp. 181-186, Jun. 2003.

[46] A. Ramalingam, B. Zhang, A. Davgan, and D. Pan, “Sleep Transistor Sizing Using Timing Criticality and Temporal Currents,” Proc. ASP-DAC, pp. 1094-1097, Jan. 2005. 44

[47] K. Shi and D. Howard, “Challenges in Sleep Transistor Design and Implementation in Low-Power Designs,” Proc. 43rd Annu. Design Automation Conf., pp. 113-116, Jul. 2006.

[48] D.E. Lackey, P.S. Zuchowski, T.R. Bednar, D.W. Stout, S.W. Gould, and J.M. Cohn, “Managing Power and Performance for System-on-Chip Designs Using Voltage Islands,” IEEE/ACM Int. Conf. Computer-Aided Design, pp. 195–202, Nov. 2002.

- [49] J. Tschanz, S. Narendra, Y. Yibin, B. Bloechel, S. Borkar, and V. De, "DynamicSleep Transistor and Body Bias for Active Leakage Power Control of Microprocessors," IEEE Int. Solid-State Circuits Conf., vol. 1, pp. 102–481, Feb. 2003.
- [50] Q. Wu, M. Pedram, and X. Wu, "Clock-gating and Its Application to Low Power Design of Sequential Circuits," IEEE Custom Integrated Circuits Conf., pp. 479–482, May 1997.
- [51] M. Pedram, "Power Minimization in IC Design: Principles and Applications," ACM Tran. Design Automation, vol. 1, no. 1, pp. 3–56, Jan. 1996.
- [52] G. Friedman, "Clock Distribution Design in VLSI Circuits: An Overview," Proc. IEEE ISCAS, San Jose, CA, pp. 1475–1478, May 1994.
- [53] Q. Wu, M. Pedram, and X. Wu, "Clock-Gating and Its Application to Low Power Design of Sequential Circuits," Proc. IEEE Custom Integrated Circuits Conf., vol. 47, pp. 415–420, May 2000.
- [54] "International technology roadmap for semiconductors", in <http://www.itrs.net/reports.html>, 2006.
- [55] Yongpan Liu; Huazhong Yang; Dick, R.P.; Hui Wang; Li Shang, "Thermal vs Energy Optimization for DVFS-Enabled Processors in Embedded Systems," Quality Electronic Design, 2007. ISQED '07. 8th International Symposium on , vol., no., pp.204,209, 26-28 March 2007.
- [56] Cheng, W.H.; Baas, B.M., "Dynamic voltage and frequency scaling circuits with two supply voltages," Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on , vol., no., pp.1236,1239, 18-21 May 2008
- [57] Murray, J.; Tang, N.; Pande, P.P.; Deukhyoun Heo; Shirazi, B.A., "DVFS Pruning for Wireless NoC Architectures," Design & Test, IEEE , vol.32, no.2, pp.29,38, April 2015
- [58] Li Shang; Li-Shiuan Peh; Jha, N.K., "Dynamic voltage scaling with links for power optimization of interconnection networks," High-Performance Computer Architecture, 2003. HPCA-9 2003. Proceedings. The Ninth International Symposium on , vol., no., pp.91,102, 8-12 Feb. 2003