Energy-Efficient Multicore Chip Design Through Cross-Layer Approach

Paul Wettin, Jacob Murray, Partha Pande, Behrooz Shirazi
School of Electrical Engineering and Computer Science
Washington State University
Pullman, USA
{pwettin, jmurray, pande, shirazi}@eecs.wsu.edu

Amlan Ganguly
Department of Computer Engineering
Rochester Institute of Technology
Rochester, USA
amlan.ganguly@rit.edu

***Abstract*—Traditional multi-core designs, based on the Network-on-Chip (NoC) paradigm, suffer from high latency and power dissipation as the system size scales up due to the inherent multi-hop nature of communication. Introducing long-range, low power, and high-bandwidth, single-hop links between far apart cores can significantly enhance the performance of NoC fabrics. In this paper, we propose design of a small-world network based NoC architecture with on-chip millimeter (mm)-wave wireless links. The millimeter wave small-world NoC (mSWNoC) is capable of improving the overall latency and energy dissipation characteristics compared to the conventional mesh-based counterpart. The mSWNoC helps in improving the energy dissipation, and hence the thermal profile, even further in presence of network-level dynamic voltage and frequency scaling (DVFS) without incurring any additional latency penalty.**

Keywords—NoC, wireless, mm-wave, small world, DVFS

I. INTRODUCTION

Continuing progress in integration levels in silicon technologies make possible complete end-user systems on a single chip. This massive level of integration makes modern multi-core chips widely adoptable in multiple domains. In the design of high-performance massively multi-core chips, power and heat are the dominant constraints. The increasing power consumption is of growing concern due to several reasons, e.g., cost, performance, reliability, scalability, and environmental impact. Increased power consumption can raise chip temperature, which in turn can decrease chip reliability and performance and increase cooling costs. Performance of a multicore chip is governed by its overall communication infrastructure, which is predominantly a NoC. The existing methods of implementing a NoC with planar metal interconnects are deficient due to high latency, significant power consumption, and temperature hotspots arising out of long, multi-hop wireline links used in data exchange. It is possible to design high-performance, robust, and energy-efficient multi-core chips by adopting novel architectures inspired by complex network theory in conjunction with on-chip wireless links. Networks with the small-world property have a very short average path length, making them particularly interesting for efficient communication with minimal resources. Using the small-world approach [1], we can build a highly efficient NoC with both wired and wireless links. Neighboring cores should be connected through normal metal wires while widely separated cores will communicate through long-range, single-hop, wireless links. Wireless Network-on-Chip (WiNoC) is an enabling technology to integrate very high numbers of cores in a single chip [2]. By reducing the hop count between largely separated communicating cores, wireless shortcuts have been shown to carry a significant amount of the overall traffic within the network. The amount of traffic detoured in this way is substantial and the low power wireless links enable energy savings. The overall energy dissipation of the WiNoC can be improved even further if the characteristics of the wireline links are optimized according to the traffic patterns. 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 [3]. In this work our aim is to demonstrate how the power and thermal efficiencies of multi-core chips can be improved by adopting a cross-layer design methodology by deploying long-range wireless interconnects as well as incorporating suitable DVFS schemes on the wireline links.

II. 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 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 Island (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 versus off-chip DC-DC converters [4], the number of allowed discrete voltage levels, and centralized versus distributed control techniques [3]. 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 addresses power and thermal management strategies for the processing cores only. Networks consume a significant part of the chip's power budget; greatly affecting overall temperature. However, there is

This work was supported in part by the US National Science Foundation (NSF) grants CCF-0845504, CNS-1059289, and CCF-1162202, and Army Research Office grant W911NF-12-1-0373.

978-3-9815370-0-0/DATE13/©2013 EDAA

little research on how they contribute to 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 simultaneous dynamic voltage scaling of processors and communication links for 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 networks 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 [2]. It is already shown that small-world network architectures with long-range wireless shortcuts can significantly improve the energy consumption and achievable data rate of massive multi-core computing platforms [2]. 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.

III. WIRELESS NOC ARCHITECTURE

Many naturally occurring complex networks, such as social networks, the Internet, the brain, as well as microbial colonies exhibit the small-world property [11]. Small-world graphs are characterized by many short-distance links between neighboring nodes as well as a few relatively long-distance direct shortcuts. Networks with the small-world property have a very short average path length. This makes them interesting for efficient communication with minimal resources. These features of small-world graphs make them particularly attractive for constructing scalable WiNoCs because the long-distance shortcuts can be realized using high-bandwidth, low-energy, direct wireless interconnects while the local links can be realized with regular wireline links. In this work we consider a small-world NoC (SWNoC) architecture where the shortcuts are implemented through wireless millimeter-wave (mm-wave) links. In the following subsections we discuss the characteristics of the proposed mm-wave SWNoC (mSWNoC) architecture whose performance and temperature profiles with and without network-level DVFS are analyzed in this paper.

A. mSWNoC Topology

In the proposed mSWNoC topology, each core is connected to a NoC switch and switches are interconnected using both wireline and wireless links. The topology of the mSWNoC is a small-world network where the wireline links between switches are established following a power-law distribution as shown in (1),

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

where, the probability, $P(i,j)$, of establishing a link, between two switches, i and j separated by a Euclidean distance of l_{ij} is proportional to the distance raised to a finite power, α [11]. The distance is obtained by considering a tile-based floorplan of the cores on the die. The frequency of traffic interaction between the switches, f_{ij}, is also factored into (1) so that more frequently communicating switches 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. f_{ij} is based on the particular application mapped to the overall mSWNoC and is hence, set prior to link insertion. This can optimize the network architecture for non-uniform traffic scenarios. The parameter α governs the nature of connectivity. It has been shown that α being less than $D + 1$, D being the dimension of the network, ensures the small-world characteristic [11]. Most naturally occurring small-world networks are created by self-assembly and do not guarantee a fully connected topology, as it is fundamentally stochastic in nature. So, we adopt an iterative algorithm to ensure that the mSWNoC is connected without isolated clusters. To establish the network connectivity each pair of switches in the NoC is selected and a wireline link is established between them with the probability given in (1). The network setup is repeated until a fully connected system is formed. We have assumed an average number of connections from each switch to other switches, $<k>$, which was chosen to be equal to that of a conventional mesh as to not introduce any additional switch area overhead. Also, an upper bound, k_{max}, is imposed on the number of wireline links attached to a particular switch so that no particular switch becomes unrealistically large in the mSWNoC. This also reduces the skew in the distribution of the links among the switches. Both $<k>$ and k_{max} do not include the local NoC switch port connected to the core.

The wireline small-world architecture is then augmented with mm-wave wireless links between switches that are separated by long distances. It is shown in [2] that mm-wave wireless shortcuts are always more energy efficient over their wireline counterparts whenever the link length is 7 mm or more in the 65 nm technology node. Hence, the wireless links are inserted between switches separated by at least 7 mm. It has already been demonstrated that using on-chip mm-wave wireless links, it is possible to create three non-overlapping channels [2]. Using these three non-overlapping channels we overlay the wireline small-world connectivity with the wireless channels such that a few switches get a wireless port. Each of these wireless ports will have wireless interfaces (WIs). The WIs are then assigned one of the three channels; repeating this for all the WIs in the network. One WI is replaced by a gateway WI, which has transceivers tuned to all three channels. The gateway WI facilitates data exchange between the different non-overlapping wireless channels. The WIs are then placed onto the small-world architecture using a simulated annealing based optimization technique incorporating the above-mentioned length constraint to minimize the average hop count weighted by the probability of traffic interactions between the cores as proposed in [12].

B. Communication and Channelization

This section describes the overall communication mechanism, which includes routing and flow control, and the WI components for mSWNoC.

1) Routing and Flow Control

In the mSWNoC, data is transferred via flit-based, wormhole routing. Between a source and destination pair, the wireless links, through the WIs, are only chosen if the wireless

path reduces the total path length compared to the wireline path. This can potentially give rise to hotspot situations in the WIs. Many messages will try to access the WI shortcuts simultaneously, thus overloading the WIs, which would result in higher latency and energy dissipation. Token flow control [13] and distributed routing are used to alleviate this problem. An arbitration mechanism is designed to grant access to the wireless medium to a particular WI, including the gateway WI, at a given instant to avoid interference and contention between the WIs that have the same frequency. To avoid centralized control and synchronization, the arbitration policy adopted is a wireless token passing protocol [14]. In this scheme, a single flit circulates as a token in each frequency channel. The particular WIs possessing the wireless tokens can broadcast flits into the wireless medium in their respective frequencies. The wireless token is forwarded to the next WI operating in the same frequency channel after all flits belonging to a message at the WI are transmitted. Packets are rerouted, through wireline links, if the WI buffers are full and it does not have the token. To ensure deadlock-free routing in this small-world topology we adopted the Tree-based Routing Architecture for Irregular Networks (TRAIN) algorithm [15].

2) Wireless Interface

The two principal WI components are the antenna and the transceiver, whose characteristics are outlined below.

The on-chip antenna for the mSWNoC has to provide the best power gain for the smallest area overhead. A metal zigzag antenna has been demonstrated to possess these characteristics [16]. This antenna also has negligible effect of rotation (relative angle between transmitting and receiving antennas) on received signal strength, making it most suitable for mm-wave NoC applications. Zigzag antenna characteristics depend on physical parameters like axial length, trace width, arm length, bend angle, etc. By varying these parameters, the antennas are designed to operate on different frequency channels.

To ensure high throughput and energy efficiency, the WI transceiver circuitry has to provide a very wide bandwidth, as well as low power consumption. Non-coherent on-off keying modulation is chosen, as it allows relatively simple and low-power circuit implementation. The transmitter consists of an up-conversion mixer and a power amplifier. In the receiver, a direct-conversion topology is used, consisting of a low noise amplifier, a down-conversion mixer and a baseband amplifier. An injection-lock voltage-controlled oscillator is reused for both the TX and RX.

IV. DVFS METHODOLOGY

The wireless links in the mSWNoC establish one-hop shortcuts between the far apart switches and facilitate energy savings in data exchange. However, the overall energy dissipation within the network is still dominated by the flits traversing the wireline links. Hence, we propose to incorporate DVFS on these wireline links to save more energy. A method for history based network-level DVFS was proposed in [17]. In this scheme, every NoC switch predicts future traffic patterns based on what was seen in the past. The relevant metric to determine whether DVFS should be performed is the link utilization. The short term link utilization is characterized by (2),

$$U_{short} = \frac{1}{H}\sum_{i=1}^{H} f_i \tag{2}$$

where, H is the history window, and f_i is 1 if a flit traversed the link on the i^{th} cycle of the history window and 0 otherwise. The predicted future link utilization, $U_{predicted}$, is an exponential weighted average determined for each link according to (3),

$$U_{predicted} = \frac{W \cdot U_{short} + U_{predicted}}{W+1} \tag{3}$$

where, W is the weight given to the short term utilization over the long term utilization. After T cycles have elapsed, where $1/T$ is the maximum allowable switching rate, the NoC switch determines whether a given link's predicted utilization meets a specific threshold. By allowing thresholds at several different levels of $U_{predicted}$, a finer-grain balance between energy savings, due to lowering the voltage and frequency, and latency penalty, due to mispredictions, can be obtained.

Voltage regulators are required to step up or step down voltage in order to dynamically adjust voltage, and hence frequency. By using on-chip voltage regulators with fast transitions, latency penalties and energy overheads due to voltage transitions can be kept low. We estimate the energy overhead introduced by the regulators due to voltage transition as:

$$E_{regulator} = (1-\eta)\cdot C_{filter} \cdot \left|V_2^2 - V_1^2\right| \tag{4}$$

where, $E_{regulator}$ is the energy dissipated by the voltage regulator due to a voltage transition, η is the power efficiency of the regulator, C_{filter} is the regulator filter capacitance, and V_2 and V_1 are the two voltage levels.

Similar to [17], a DVFS algorithm was developed using (2) and (3). After T cycles, the algorithm determines if DVFS should be performed on the link based on the predicted bandwidth requirements of future traffic. Depending on which threshold was crossed, if any, the switch then determines whether or not to tune the voltage and frequency of the link. In order to prevent a direct multi-threshold jump, which would cause high delay and energy overhead, the voltage and frequency can step up once, step down once, or remain unchanged during one voltage/frequency transition. After each adjustment of the voltage/frequency pair on a given link, energy savings and latency penalty was determined.

A misprediction penalty is caused when the adjusted voltage/frequency pair did not meet the bandwidth requirements of the traffic over the given switching interval. The bandwidth requirement of the link was obtained by viewing the current link utilization over a smaller window whose size was determined as the average latency of a flit in the non-DVFS network.

V. EXPERIMENTAL RESULTS

In this section, we evaluate the performance and temperature profile of the mSWNoC and compare those with the conventional mesh-based NoC by incorporating the network-level DVFS elaborated in section IV. We use GEM5 [18], a full system simulator, to obtain detailed network-level information. We consider a system of 64 alpha cores running Linux within the GEM5 platform for all experiments. Three SPLASH-2 benchmarks, FFT, RADIX, LU, [19] and the

PARSEC benchmark CANNEAL [20] are considered as they vary in characteristics from computation intensive to communication intensive in nature.

The initial frequency of traffic interaction between the switches, f_{ij}, is obtained from GEM5. The f_{ij} matrices are used to generate the traffic patterns for each benchmark in a NoC simulator to obtain the NoC performance in terms of packet latency, packet energy, and link and switch power. The width of all wired links is considered to be same as the flit width, which is 32 bits in this paper. Each packet consists of 64 flits. Similar to the wired links, we have adopted wormhole routing in the wireless links too. The NoC simulator uses switches synthesized from an RTL level design using TSMC 65-nm CMOS process, using Synopsys™ Design Vision. All ports except those associated with the WIs have a buffer depth of two flits. Each switch port has four virtual channels. The ports associated with the WIs have an increased buffer depth of eight flits to avoid excessive latency penalties while waiting for the token. Increasing the buffer depth beyond this limit does not produce any further performance improvement for this particular packet size, but will give rise to additional area overhead [14]. Energy dissipation of the network switches were obtained from the synthesized netlist by running Synopsys™ Prime Power, while the energy dissipated by wireline links was obtained through HSPICE simulations taking into consideration length of the wireline links. Each wireless link can sustain a data rate of 16Gbps and has an energy dissipation of 2.3pJ/bit [2].

After obtaining network power values, the network switches and links are arranged on a 20mm x 20mm die. The floor plans, along with the power values, are used in HotSpot [21] to obtain steady state thermal profiles. The architecture-dependent network power values with and without DVFS in presence of the specific benchmarks, are fed to the HotSpot simulator to obtain the temperature profiles of each scenario.

First, we discuss the features of the mSWNoC architecture and then we present its latency, energy dissipation and thermal characteristics by incorporating DVFS.

A. Wireless Channel Characteristics

The metal zig-zag antennas described in section III are used to establish the on-chip wireless links. We are able to obtain three different channels with 3dB bandwidths of 16 GHz and center frequencies of 31, 57.5, and 120 GHz respectively with a communication range of 20 mm. For optimum power efficiency, the quarter wave antennas use axial length of 0.73, 0.38, and 0.18 mm respectively. The antenna design ensures that signals outside the communication bandwidth, for each channel, are sufficiently attenuated to avoid inter-channel interference. The total area overhead per wireless transceiver turns out to be 0.3 mm^2.

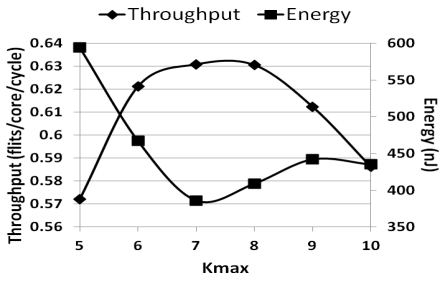

Figure 1. Variation of throughput and energy dissipation for a 64-core

B. Determination of mSWNoC Topology

We determine the exact topology of the mSWNoC based on the principles of small-world graph as discussed in section III. It is already shown that in a small-world network, the average hop count decreases with increases in the power law exponent α for a fixed wiring cost [11]. With α=1.8, the average hop count is minimum with a fixed wiring cost. Hence while designing the topology we considered this value of α. The other parameter that we need to determine is the maximum number of wireline links attached to a particular switch, k_{max}. We perform system level simulations using uniform random traffic before placing any WIs to obtain the optimal value of k_{max}. Fig. 1 shows the variation of the throughput and packet energy with respect to k_{max} for the 64-core SWNoC (mSWNoC without WIs.)

We chose the value of k_{max}, which optimizes both throughput and energy, and from Fig. 1 it can be concluded that k_{max} of 7 maximizes the throughput with minimum energy dissipation. After the topology was setup using (1), the WIs were then placed using simulated annealing, as mentioned in section III, by varying the number and location of the WIs.

Increasing the number of WIs improves the connectivity of the network as they establish one-hop shortcuts. However, the wireless medium is shared among all the WIs and hence, as the number of WIs increases beyond a certain limit, performance starts to degrade due to the large token returning period. Moreover, as the number of WIs increases, the overall energy dissipation from the WIs becomes higher, and it causes packet energy to increase as well. Fig. 2 shows the variation of bandwidth and energy dissipation with varying number of WIs for a 64-core system size for uniform random traffic. It is clear that if the number of WIs increases beyond 12, then the achievable bandwidth actually degrades and the energy dissipation starts to increase. Hence, the optimum number of WIs is 12, and as such, these WIs are placed in the proposed architecture. As mentioned in section III, three frequency channels are distributed among these WIs where more frequently communicating WIs use the same channel.

C. DVFS Setup

To determine the appropriate prediction and switching windows, we consider the energy latency product by varying these two parameters. The optimum point is found for all the benchmarks considered in this work. A small switching window may catch data bursts, which do not represent a long-term trend of the benchmark's traffic. Consequently, widely

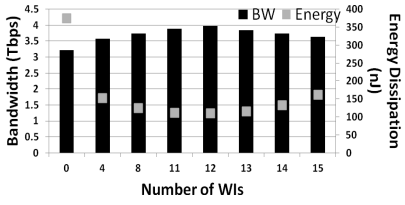

Figure 2. Bandwidth/energy dissipation tradeoff for 64-core mSWNoC by varying the number of WIs.

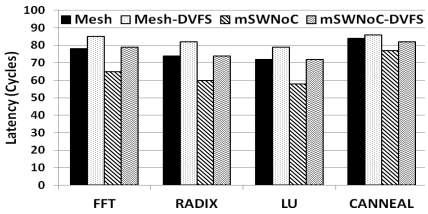

Figure 3. Average packet latency of mesh and mSWNoC with/without DVFS.

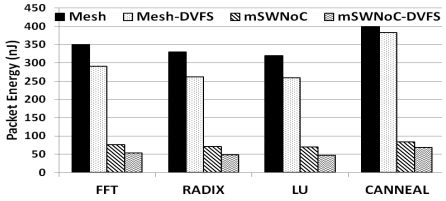

Figure 4. Average packet energy of mesh and mSWNoC with/without DVFS.

varying short-term traffic utilizations will cause the voltage/frequency to change often. In this case, the regulator energy overhead may outweigh the benefits of a lower misprediction penalty. As the switching window widens, the regulator energy overhead impact is decreased, while the latency penalty increases. A switching window of T = 450 cycles, and a history window of H = 450 cycles were selected for all the SPLASH-2 benchmarks, as this optimizes the energy-latency tradeoff for them. For CANNEAL however, $T = H = 900$ cycles.

The selected voltage/frequency pairs corresponding to each frequency are shown in Table 1. Based on these pairs, required bandwidth thresholds as a fraction of the maximum allowable value were determined to be proportional to the scaled voltage/frequency pairs.

Table 1. DVFS Voltage/Frequency/Threshold combinations.

Threshold	≥0.9	≥0.8	≥0.7	≥0.6	≥0.5	<0.5
Voltage (V)	1	0.9	0.8	0.7	0.6	0.5
Frequency (GHz)	2.5	2.25	2.0	1.75	1.5	1.25

D. *Performance Characteristics*

In this section we first present the latency, packet energy, and thermal profiles of the mSWNoC and mesh-based wireline NoC architectures. Figs. 3 and 4 show the latency and energy characteristics for the mSWNoC and mesh respectively considering the four different benchmarks, with and without DVFS. It can be observed from Fig. 3 that for all the benchmarks there is an associated latency penalty when implementing DVFS. As mentioned in section IV, the latency penalty arises when the predicted utilization did not meet the bandwidth requirements of the flits traversing the link. The latency of mSWNoC is lower than that of the mesh architecture. This is due to the small-world architecture of mSWNoC with direct long-range wireless links that enables a smaller average hop-count than that of mesh [22]. The important point to note here is that by utilizing the benefits of the mSWNoC architecture, we can implement DVFS without incurring an overall latency penalty with respect to the conventional wireline mesh.

Fig. 4 shows the packet energy dissipation for both the mSWNoC and mesh architectures. The packet energy dissipation is defined as the average energy dissipated in transmitting one packet between source and destination switches, thereby reflecting the network characteristics. While determining the packet energy in presence of DVFS, the energy overhead due to the voltage regulators are incorporated by considering η to be 0.77 and C_{filter} to be 5nF according to [23]. From Fig. 4, we can observe that the mSWNoC architecture provides a significant reduction in energy. The two contributors of the energy dissipation are from the switches and the interconnect infrastructure. In the mSWNoC, the overall switch energy decreases significantly compared to a mesh as a result of the better connectivity of the architecture. In this case, the hop-count decreases significantly and hence the packet has to traverse through less number of switches and links. In addition, a significant amount of traffic traverses through the energy efficient wireless channels; consequently allowing the interconnect energy dissipation to decrease. By implementing DVFS, for both the architectures, energy savings are obtained. The energy savings between mesh and DVFS-enabled mesh vary between 4.2% for CANNEAL to 20.6% for RADIX. Between the mSWNoC and DVFS-enabled mSWNoC the savings are between 17.9% for CANNEAL and 32.8% for LU. As the mSWNoC architecture lowers the amount of traffic traversing wired links, there is more opportunity to perform DVFS on the mSWNoC architecture. A histogram of the link level traffic utilizations is shown in Fig. 5. From this, it is clear that in the mesh architecture, a significant amount of links have more than 90% utilization for the considered benchmarks, and hence, the savings in energy are lower. In this case, there is not significant room for improvement with DVFS as the voltage and frequency cannot be tuned often on the links with high utilization. On the other hand, the mSWNoC reduces traffic on wireline links, which can also be seen clearly in Fig. 5. As the majority of links fall under 50% utilization in the mSWNoC architecture, there is significant opportunity for implementing DVFS. Because of this, there is room for more energy savings in mSWNoC in presence of DVFS compared to the mesh.

E. *Network Temperature Profile*

In this subsection we evaluate the thermal profile of the mSWNoC incorporating DVFS. To further understand architectural-based improvements in the thermal profile, we have obtained the communication density within the region that is hottest in the network. The communication density is defined as the average number of flits per cycle in a given switch or link. We first obtain the communication densities of the mSWNoC and traditional mesh architectures in presence of the different traffics. The average communication density was determined over the hotspot area of the network for each benchmark to analyze the thermal profile. Here, the hotspot area is defined as the original mesh location, which is forming

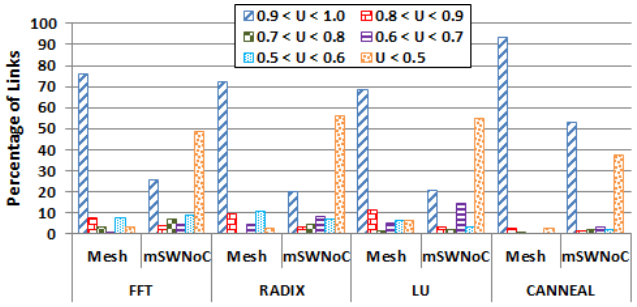

Figure 5. Wireline link utilization histogram.

a localized region of high communication density, and hence, high temperatures.

The average wireline link communication densities in mSWNoC within the hotspot region of interest were reduced by 48.6%, 44.2%, 44.5%, and 50.5% for the FFT, RADIX, LU, and CANNEAL traffics, respectively, compared to the traditional mesh topology. The long-range wireless links significantly reduce the communication density of the hotspot region by efficiently routing a major portion of the data, which would otherwise travel through the wireline links of the hotspot region. Similarly, savings of average router communication densities were 47.8%, 45.1%, 41.6%, and 55.3% for the FFT, RADIX, LU, and CANNEAL traffics, respectively. This directly relates to the thermal profile of the chip, as reduction in communication density will produce less heat.

To demonstrate the thermal profile, we consider the CANNEAL benchmark, as its communication density is the largest among all the benchmarks. As explained at the beginning of section V, we use the HotSpot tool to determine the thermal profile. To quantify the thermal improvements within the hotspot region, we define the average reduction in switch or wireline link temperature between mesh and DVFS-enabled mSWNoC as ΔT_{avg}, and the difference in maximum switch or link temperature between mesh and DVFS-enabled mSWNoC as $\Delta T_{hotspot}$,. The average reduction in switch and link temperatures between mesh and DVFS-enabled mSWNoC were 14.34°C and 9.78°C, respectively and $\Delta T_{hotspot}$ for the switches and links were 24.50°C and 10.98°C, respectively. As mentioned above, the improvements to the hotspot area are because of two reasons. First, due to the mSWNoC, the communication density within the centralized hotspot region is reduced drastically. Secondly, the DVFS additionally has the ability to reduce energy produced by network elements within the hotspot location, directly impacting the thermal profile as well. The values of ΔT_{avg} and $\Delta T_{hotspot}$ for all of the other benchmarks are shown in Table 2. The SPLASH-2 benchmarks are highly computational, and hence the communication among the cores is less than that in CANNEAL. Consequently, the room for improving the network-level temperatures for these benchmarks is much less than CANNEAL.

Table 2. Change in temperatures between mesh and DVFS-enabled mSWNoC

		FFT	RADIX	LU	CANNEAL
Switch	ΔT_{avg}	4.05°C	1.92°C	2.34°C	14.34°C
	$\Delta T_{hotspot}$	9.76°C	3.69°C	5.81°C	24.50°C
Link	ΔT_{avg}	2.75°C	1.25°C	1.59°C	9.78°C
	$\Delta T_{hotspot}$	3.50°C	1.41°C	2.06°C	10.98°C

VI. CONCLUSION

Millimeter-wave wireless small-world NoC (mSWNoC) is an enabling technology to address the inherent performance limitations of conventional multi-hop wired counterparts. The mSWNoC is capable of achieving better latency and energy dissipation compared to its conventional wireline mesh-based counterpart. Implementing DVFS on the wireline links further improves the energy dissipation profile of mSWNoC without paying any latency penalty with respect to the mesh. We also demonstrate that the switch and link temperatures in the hotspot region of the mesh are significantly improved in a DVFS-enabled mSWNoC. The temperature improvement is more pronounced for a communication-intensive benchmark than the computationally extensive ones.

REFERENCES

[1] U. Y. Ogras and R. Marculescu, "It's a small world after all: NoC performance optimization via long-range link insertion," IEEE Trans. Very Large Scale Integr. Syst., vol. 14, no. 7, 2006, pp. 693-706.

[2] S. Deb, et. al, "Wireless NoC as interconnection backbone for multicore chips: promises and challenges," IEEE J. Emerg. Sel. Topic Circuits Syst., vol. 2, no. 2, 2012, pp. 228-239.

[3] S. Garg, D. Marculescu, R. Marculescu, "Technology-driven limits on run-time power management algorithms for multi-processor systems on chip," ACM J. Emerg. Technol. Compt. Syst., vol. 8, no. 4, 2012.

[4] W. Kim, et al., "System level analysis of fast, per-core DVFS using on-chip switching regulators," Proc. of the International Symposium on High Performance Computer Architecture, 2008, pp. 123-134.

[5] W. Jang, et al., "A voltage-frequency island aware energy optimization framework for networks-on-chip," IEEE J. Emerg. Sel. Topic Circuits Syst., vol. 1, no. 3, 2011, pp. 420-432.

[6] U. Ogras, R. Marculescu, D. Marculescu, "Variation-adaptive feedback control for networks-on-chip with multiple clock domains," Proc. Design Automation Conference, 2008, pp. 614-619.

[7] K. Niyogi, D. Marculescu, "Speed and voltage selection for GALS systems based on voltage/frequency islands," Proc. Asia and South pacific Design Automation Conference, 2005, vol. 1, pp. 292-297.

[8] J. Donald, M. Martonosi, "Techniques for multicore thermal management: classification and new exploration," Proc. of Symposium on Computer Architecture, 2006.

[9] L. Shang, et al., "Temperature-aware on-chip networks," IEEE Micro: Micro's Top Picks from Computer Architecture Conferences, 2006.

[10] J. Luo, N. K. Jha, L.-S.Peh, "Simultaneous dynamic voltage scaling of processors and communication links in real-time distributed embedded systems," IEEE Trans. on Very Large Scale Integration Systems, vol. 15, no. 4, 2007.

[11] T. Petermann and P. De Los Rios, "Spatial small-world networks: a wiring cost perspective," arXiv: cond-mat/0501420v2.

[12] A. Ganguly, et al., "Scalable hybrid wireless network-on-chip architectures for multi-core systems," IEEE Trans. on Compt., vol. 60, no. 10, pp. 1485-1502.

[13] A. Kumar, et al., "Token flow control," Proc. of the 41st IEEE/ACM International Symposium on Microarchitecture, 2008, pp. 342-353.

[14] K. Chang, et al., "Performance evaluation and design trade-offs for wireless network-on-chip architectures," ACM J. Emerg. Technol. Comput. Syst., vol. 8, no. 3, 2012.

[15] H. Chi and C. Tang, "A deadlock-free routing scheme for interconnection networks with irregular topology," Proc. of ICPADS, pp. 88-95.

[16] A. B. Floyd, et al., "Intra-chip wireless interconnect for clock distribution implemented with integrated antennas, receivers, and transmitters," IEEE J. Solid-State Circuits, vol. 37, no. 5, pp. 543-552.

[17] L. Shang et al., "Dynamic Voltage scaling with Links for Power Optimization of Interconnection Networks", Proc. of HPCA, 2003.

[18] N. Binkert, et al., "The GEM5 simulator," ACM SIGARCH Computer Architecture News, 39(2), 2011, pp. 1-7.

[19] S.C. Woo, et al., "The SPLASH-2 programs: characterization and methodological considerations," Proc. of ISCA, 1995, pp. 24-36.

[20] C. Bienia, "Benchmarking modern multiprocessors," Ph.D. Dissertation, Princeton Univ., Princeton NJ, Jan. 2011.

[21] K. Skadron, et al., "Temperature-aware microarchitecture," Proc. of the International Symposium on Computer Architecture, 2003, pp. 2-13.

[22] A. Ganguly, P. Wettin, K. Chang, and P. P. Pande, "Complex network inspired fault-tolerant NoC architectures with wireless links", Proc. of ACM/IEEE NOC Symposium, 2011.

[23] S. S. Kudva and R. Harjani, "Fully-integrated on-chip DC-DC converter with a 450X output range," IEEE J. Solid-State Circuits, vol. 46, no. 8, 2011, pp. 1940-1951.