
An overview of achieving energy efficiency in on-chip networks

Masud Al Aziz and Samee Ullah Khan*

NDSU-CIIT Green Computing and Communications Laboratory,
Department of Electrical and Computer Engineering,
North Dakota State University,
Fargo, ND 58108-6050, USA
E-mail: masud.aziz@ndsu.edu E-mail: samee.khan@ndsu.edu
*Corresponding author

Thanasis Loukopoulos

Department of Informatics and Computer Technology,
Technological Educational Institute (TEI) of Lamia,
35100 Lamia, Greece
E-mail: luke@teilam.gr

Pascal Bouvry

Faculty of Science, Technology and Communication,
University of Luxembourg,
6 rue Coudenhove Kalergi,
L-1359 Luxembourg, Luxembourg
E-mail: pascal.bouvry@uni.lu

Hongxiang Li

Department of Electrical and Computer Engineering,
North Dakota State University,
Fargo, ND 58108-6050, USA
E-mail: hongxiang.li@ndsu.edu

Juan Li

NDSU-CIIT Green Computing and Communications Laboratory,
Department of Computer Science,
North Dakota State University,
Fargo, ND 58108-6050, USA
E-mail: j.li@ndsu.edu

Abstract: Due to the increasing bandwidth demand for the network-on-chip (NoC), interconnection networks become a dominant source of energy consumption in systems-on-chip (SoCs) and chip multi processors (CMPs). Therefore, energy efficient NoC is key to a successful SoC development. This paper presents an overview of different techniques to achieve energy efficiency at the different levels of NoC design including:

- a component level where dynamic voltage scaling (DVS) and dynamic link shutdown (DLS) techniques are reviewed
- b circuit level, e.g., voltage swinging of signals
- c architectural level, where specialised tools, such as Wattch and Orion are discussed.

We also summarise research on thermal optimisation issues. To the best of our knowledge, this is the first survey of recent research results on the area.

Keywords: network on chip; NoC; energy-efficiency; power-aware communications; interconnection networks.

Reference to this paper should be made as follows: Aziz, M.A., Khan, S.U., Loukopoulos, T., Bouvry, P., Li, H. and Li, J. (2010) 'An overview of achieving energy efficiency in on-chip networks', *Int. J. Communication Networks and Distributed Systems*, Vol. 5, No. 4, pp.444-458.

Biographical notes: M.A. Aziz received his Bachelor degree from Bangladesh. He is currently a graduate student at the North Dakota State University.

S.U. Khan received his BS from Ghulam Ishaq Khan Institute of Engineering Sciences and Technology, Topi, Pakistan and his PhD from the University of Texas, Arlington, TX, USA in 2007. He is currently an Assistant Professor of Electrical and Computer Engineering at the North Dakota State University, Fargo, ND, USA. He has extensively worked on the general topic of resource allocation in autonomous heterogeneous distributed computing systems. He is the Associate Editor of the International Journal of Communication Systems and the Security and Communication Networks Journal.

T. Loukopoulos received his Diploma in 1997 from the Computer Eng. Dept. of Univ. of Patras, Greece and PhD degree in 2002 from the Computer Science Dept. of the Hong Kong Univ. of Science and Technology (HKUST). He is currently an Assistant Professor at the Dept. of Informatics and Computer Technology of the Technological Educational Institute (TEI) of Lamia, Greece and a Cooperating Researcher for the Distributed and Ubiquitous Computing Systems Laboratory of the Center for Research and Technology-Thessaly (CERETETH). He has published 20 papers in the broad area of resource allocation and scheduling in parallel and distributed systems and received the best paper award in ICPP 2001. He is an IEEE member.

P. Bouvry obtained his undergraduate degree in Economical and Social Sciences and his Masters in Computer Science with distinction in 1991 from the University of Namur, Belgium. He obtained his PhD in Computer Science with great distinction at the University of Grenoble (INPG), France in 1994. He is currently heading the Computer Science and Communications (CSC) research unit of the Faculty of Sciences, Technology and Communications of Luxembourg University and serving as a Professor. He is also the Treasurer of the Administration Board of CRP-Tudor and a member of various scientific committees (ERCIM WG, COST TIST, LIASIT, etc.).

H. Li received his BS from Xi'an Jiaotong University, China in 2000, his MS from Ohio University in 2004 and his PhD from University of Washington, Seattle in 2008, all in Electrical Engineering. He is currently an Assistant Professor at the Department of Electrical and Computer Engineering, North Dakota State University, Fargo. His research interests include broadband mobile communications, cognitive radio and intelligent hybrid network. He has published over 20 articles. He received the Chinese Government Award for Outstanding Self-Financed Students Abroad and UWEE Outstanding Research Assistant Award in 2008.

Juan Li received her PhD in Computer Science in 2008 from the University of British Columbia, Vancouver, Canada. She is currently an Assistant Professor at North Dakota State University, USA. Her research interests include distributed system, mobile wireless network, and semantic web.

1 Introduction

With an ever growing number of transistors over a single chip (millions), the chances of integration offered today are unprecedented (Brooks and Martonoshi, 2001a; Jayaseelan and Mitra, 2009). The following two are the most popular directions that take advantage of such possibilities. The first approach, chip multi-processors (CMP), integrates multiple cores over a single chip, while the second approach, systems-on-chip (SoC), integrates all of the necessary components of a computing system, e.g., memory (Adler and Friedman, 1997). Usually, CMPs are of general purpose, while SoC design is related to specific applications, such as consumer electronics (Benini and Micheli, 2001) and multimedia (Chen et al., 2006; Koziri et al., 2007).

Both SoC and CMP involve large volumes of interprocessor communication. The network-on-chip (NoC) paradigm has been widely proposed as the interconnect fabric for SoC (Benini and Micheli, 2001; Dally and Peh, 2000; Dally and Towles, 2001; Sgroi et al., 2001) and CMP (Sankaralingam et al., 2003; Taylor et al., 2002b). Under the NoC paradigm, network components consisting of multiple point-to-point links and switches are integrated in the chip to handle communication. The NoC paradigm has multiple advantages over the traditional global-wiring bus approach. First, NoC handles interconnection in a systematic way according to networking theory. Second, NoC allows multiple asynchronous clocking that is particularly useful in SoCs. Third, NoC scales better compared to the bus approach, and is closer to achieve the 'linear effort property' principle (Tenhunen and Jantsch, 2003), under which, adding a new component should incur an integration cost that depends only on the number of existing components and not on the complexity of the added component.

NoC must be able to provide tight-delay guarantees (Taylor et al., 2002a). For this reason, prior micro-architectures were performance-aware rather than energy-aware (Patel et al., 1997). However, as the number of on-chip components increases, the tight-delay guarantees lead to more bandwidth requirements and as a consequence the NoC design complexity increase. With higher complexity, the energy consumed by NoC also becomes significant (Ghosh et al., 2009; Shang et al., 2002). For instance, MIT RAW NoC consumes 36% of the total system power (Shang et al., 2004). Therefore, power-aware NoC design is of paramount importance to energy-efficient SoC and CMP

development. In this paper, we present an overview of the key techniques used in achieving the aforementioned objectives.

For an older survey on the field, the interested reader is referred to Raghunathan et al. (2003). Compared to Tenhunen and Jantsch (2003), our contributions are twofold:

- a we include more recent research results
- b we cover thermal issues.

Concerning the latter, many researchers, e.g., Shang et al. (2004) have suggested including thermal issues in power-aware SoC design. This becomes more important when considering that temperature rises exponentially with the increase in power consumption. Therefore, a small increase in power consumption results in a rapid temperature rise (Brooks and Martonoshi, 2001b) that may not be efficiently tackled by the current state-of-the-art cooling systems (Shang et al., 2004).

The rest of the paper is organised as follows. Section 2, discusses on-chip architecture and NoC building blocks. In Section 3 optimisation methods at the component level are summarised, with components being routers and links. Circuit-level power reduction techniques are presented in Section 4, while Section 5 surveys architectural-level tools. Thermal efficiency issues are included in Section 6. Finally, Section 7 concludes the paper. To increase readability, Table 1 lists all the acronyms used in the paper.

Table 1 List of acronyms used in this paper

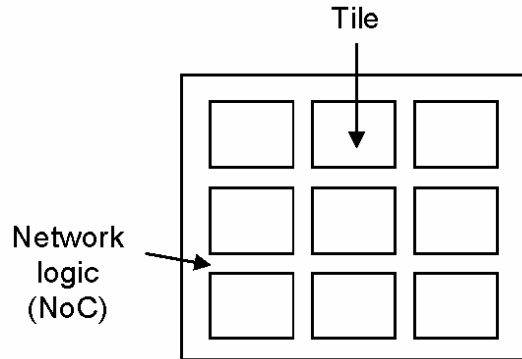
<i>Acronyms</i>	<i>Meaning</i>
CMP	Chip multi-processor
DDCD	Dynamic diode connected driver
DLL	Delay locked loop
DLS	Dynamic link shutdown
DPM	Dynamic power management
DVS	Dynamic voltage scaling
DTM	Dynamic thermal management
MP3	Moving picture experts group layer-3 audio
MPEG4	Moving picture experts group 4
NoC	Network-on-Chip
OFDM	Orthogonal frequency division multiplexing
SNR	Signal-to-noise ratio
SoC	System-on-Chip

2 Preliminaries

Under the NoC framework, all inter-module SoC or CMP communications, such as between processors, memories, and peripherals is done using packet transmissions. Therefore, the modular design of SoC and CMP processors is facilitated. Moreover, the structured networks in a NoC, experience well-controlled electrical parameters, such as low and predictable cross-talk (Dally and Poulton, 1998). These controlled electrical parameters eliminate timing iterations, increase bandwidth, and enable the use of

high-performance circuits. Furthermore, these signalling circuits also can reduce power consumption by a factor of ten and increase propagation velocity by three times (Dally and Towles, 2001). To further illustrate the inner workings of a NoC, an example NoC architecture specified in Tiwari et al. (1998) is discussed in the subsequent text.


Figure 1 A chip example using NoC



A chip is assumed to be split into tiles that are pieces of silicon die. Each tile is used to accommodate one or more client modules, such as processors, DSPs, peripheral controllers and memory subsystems. Figure 1 shows an example division of the chip into nine (3 mm × 3 mm) tiles. The client modules communicate with each other through the network. There are no top-level connections other than the network wires. The network logic occupies a small amount of area between the tiles and consumes a portion of the top two metal layers of the chip, that is, the layers of a chip that are used for accommodating interconnection wires.

The NoC components are routers and links. An example NoC architecture is shown in Figure 2 where n client modules ($A_1..A_n$) are directly connected in a single router. The router is responsible for receiving and directing packets towards their final destination.

Figure 2 An example NoC



To understand the energy efficient NoC models, we first define ‘flit’. A flit is the smallest unit of information that comprises of the header, encoding, and packet size information (Chen and Peh, 2003). As per Wang et al. (2003) the energy consumed by a flit during transmission comprises of:

$$\begin{aligned} E_{flit} &= (E_{wrt} + E_{rd} + E_{arb} + E_{xb} + E_{link}) \times H, \\ &= (E_{buf} + E_{arb} + E_{xb} + E_{link}) \times H, \end{aligned} \quad (1)$$

where E_{wrt} is the average energy dissipated when a flit is written on buffer, E_{rd} is the average energy dissipated when reading a flit from the input buffer, $E_{buf} = E_{wrt} + E_{rd}$ is the average buffer energy, E_{arb} is the average arbitration energy, E_{xb} is the average crossbar traversal energy, E_{link} is the average link traversal energy, and H is the number of hops traversed by this flit. Because buffering, arbitration, and crossbar traversal are all part of the routing mechanism, we can say $E_R = E_{buf} + E_{arb} + E_{xb}$. Equation (1) can be reformulated as:

$$E_{flit} = (E_R \times H) + (E_{link} \times D), \quad (2)$$

where E_R is the average router energy, E_{wire} is the average link wire transmission energy per unit length assuming optimally-placed repeaters, and D is the Manhattan distance between source and destination. The Manhattan distance is the shortest distance between two routers along the perimeter of the chip. We can observe from Equation (2) that minimising E_R will result in minimising E_{flit} , when all other variables are held constant.

In the subsequent text, we present existing research on energy-efficient NoC component design. Before continuing though, we would like to mention that the NoC terminology has been used in the past to characterise aspects varying from gate level physical implementation, across system layout aspects and applications, to design methodologies and tools. One of the reasons for the widespread adaptation of this terminology lies in the readily available and widely accepted abstraction models for networked communications (Bierrgaard and Mahadevan, 2006). The open system interconnection (OSI) model of layered network communication can be easily adapted for NoC usage as proposed in Benini and Micheli (2002).

3 Achieving energy efficiency in NoC components

The first work recognising the need to consider power consumption constraints in interconnection network design was Benini and Micheli (2002) that proposed a power consumption model for routers and links. Wang et al. (2003) studied power models for different NoCs and proposed several micro-architectures for key router components, such as segmented crossbar, cut-through crossbar, and write-through buffer. They also studied the power saving potential of an existing NoC architecture, termed ‘Express Cube’. Maheshwari and Burleson (2001) proposed splitting monolithic bus architecture into a layered architecture that is more energy efficient due to reduced capacitive load during bus transfers.

A number of papers use dynamic voltage scaling (DVS) to reduce power consumption by NoC. The DVS technique was first proposed and widely used in

microprocessors (Burd, 1998). DVS exploits the variance in processor utilisation to lower the frequency (by means of lower the supply voltage) when the processor is lightly loaded, while increasing it to maximum when the processor is overloaded with instruction sets. Shang et al. (2002) illustrated a power optimisation mechanism for interconnection networks by applying DVS to network links. As with the instruction workload of a microprocessor, link utilisation also varies, depending on applications' communication patterns.

Although DVS might reduce energy consumption, yet there are two pitfalls to be avoided. First, by reducing supply voltage and frequency we also increase leakage current that might in fact lead to higher power consumption (Jejuriker et al., 2004). To minimise leakage current, processors support various shutdown modes. For example, Transmeta Crusoe processor supports various sleep modes, such as normal, auto-halt, quick start, deep sleep, and off for various types of workload (Jayaseelan and Mitra, 2009). The second challenge has to do with one of the primary reasons that made digital circuits more popular compared to analogue ones, i.e., noise immunity. Digital circuits exhibit non-linear voltage transfer characteristics. However, due to smaller supply voltage, noise immunity becomes extremely difficult to maintain. This is the case for instance in the deep submicron (DSM) technology (Hedge and Shanbhag, 2000). Therefore, operating at maximum or minimum operating voltage (and as a consequence, frequency) is seldom the optimal option. In fact, the optimal voltage and frequency values can be defined to be the ones corresponding to a computational speed at which energy consumption per workload is the minimum (Irani et al., 2003). Last, to achieve full benefits of the DVS technique, digital circuits must be designed to accommodate large ripple in supply voltage (Burd and Brodersen, 2000; Burd et al., 2000).

To minimise the communication links' power consumption, Taylor et al. (2002b) proposed the dynamic link shutdown (DLS) technique. DLS is based on the premise that shifting the load that goes through underutilised links towards a subset of highly utilised links will allow the complete shutdown of the first ones. To benefit from DLS, Taylor et al. (2002a) present an adaptive routing strategy that intelligently uses a subset of links for communication, thereby facilitating dynamic link shutdowns for minimising energy consumptions. They also showed that the proposed DLS technique can provide moderate energy savings with minimal degradation in average network latency.

Finally, Chen and Peh (2003) proposed a model for managing leakage power in the interconnection network. They demonstrated that router buffers are the prime candidates for leakage power optimisation and explored power-aware buffer policies that managed to save up to 96.6% of the total buffer leakage power.

Table 2 summarises the primary features of the presented techniques for energy optimisation at the NoC component level. Wang et al. (2003) compared several crossbar (router based) power reduction techniques, namely, segmented crossbar, cut-through crossbar, and write-through crossbar. The aforementioned techniques differ in the methodology utilised to access the crossbar switch. Although achieving the goal of reducing power consumption, DVS (Shang et al., 2002), Segmented crossbar (Wang et al., 2003), and cut-through crossbar (Wang et al., 2003) suffer from poor signal-to-noise ratio (SNR). On the other hand, DLS (Taylor et al., 2002b) and Express cube (Wang et al., 2003) reduce power consumption without sacrificing SNR. However, both Express cube and DLS may impact the overall NoC latency.

Table 2 Primary characteristics of NoC component energy optimisation techniques

	<i>Increased leakage current</i>	<i>Reduced latency</i>	<i>Low power consumption</i>	<i>Good SNR</i>
DVS	√	×	√	×
DLS	×	×	√	√
Express cube	×	√	√	√
Segmented crossbar	×	√	√	×
Cut-through crossbar	×	√	√	×
Write-through crossbar	×	×	√	√

4 Circuit-level energy optimisation

Circuit-level tools deal mostly with encoding and synchronisation issues at the datagram (flit) level. To optimise power consumption at the circuit-level of NoC various schemes were proposed, with the most popular ones being low-swing voltage driver and reduced supply voltage technique. Voltage swing is the maximum peak voltage that the output circuit can produce before it starts clipping (Zhang et al., 2000). In Dumitras and Marculescu (2003), a low-swing (low output) voltage driver using dynamic diode-connected driver (DDCD) architecture is proposed. The resulting DDCD circuit had a simple inverter as a receiver and met the desired goals of low complexity single-wire low-swing driver, albeit at the cost of low SNR values. Lee et al. (2000) proposed an input-multiplexed transmitter to reduce clock load. The transmitter multiplexes the signal and then the multiplexed signal is fed to the clock. To reduce jitter, precision timing circuits based on delay locked loop (DLL) were used. DLL is a digital circuit that compares the phase of one of its outputs to the input clock, in order to generate an error signal which is then integrated and fed back as the control of all delay elements. A sensitive capacitive trimmed receiver also was proposed to ensure reliable operation at very low signal levels needed to reduce power consumption. Another way to minimise power consumption is to reduce the input voltage swing by using a multi-phase low-frequency clock instead of a high-frequency one as proposed in Lu (2008).

The most widely accepted technique is to reduce the voltage swing of signals on wires. To understand the relationship between voltage swing and energy consumption in wires, Wei and Horowitz (2000) proposed the following relationship:

$$E_{wire} = \alpha \times C \times V_{DD} \times V_{swing}, \quad (3)$$

where α is the switching activity per clock cycle of the signal that is being transmitted, C is the physical capacitance switched during signal transitions, V_{DD} is the supply voltage, and V_{swing} is the voltage swing across the wire. According to Equation (3), reduced voltage swing results in lower wire power consumption, when all the other variables are held constant. Various schemes, such as static driver with reduced supply (Nakagome et al., 1993), differential interconnect (Burd, 1998), dynamically enabled drivers (Colshan and Jaroun, 1994), and low-swing bus techniques (Yamauchi et al., 1995) are based on Equation (3). Zhang et al. (2000) measured the performance of these schemes using a worst-case noise analysis method. This method uses a driver circuit that converts a full-swing input into a reduced-swing interconnect signal that is converted back to a

full-swing output by the receiver. Results in Zhang et al. (2000) showed that differential interconnect (Burd, 1998) achieves energy savings up to four-folds and exhibits the highest SNR compared to all other schemes mentioned above.

To reduce both system area and power consumption, Srinivasan and Adve (2003) proposed a novel hybrid NoC structure and a dynamic job distribution algorithm that can reduce system area and power consumption. This optimisation technique reduces packet drop rate in a variety of applications, such as MPEG4 and MP3 decoders, global positioning system (GPS), and orthogonal frequency division multiplexing (OFDM) demodulators.

Table 3 Features of circuit-level optimisation methods

<i>Feature</i>	<i>No extra power supplies</i>	<i>No multiple threshold voltage process</i>	<i>Low power consumption of the circuits</i>	<i>Small area needed for the I/O circuits</i>	<i>Good SNR</i>
Power analysis of link level and end-to-end data protection in NoCing	×	×	√	×	×
Digitalised I/O techniques for high speed serial link system	×	√	√	√	√
Low-swing signalling using a dynamic diode-connected driver	√	√	√	√	√
A power efficient flit-admission scheme for wormhole-switched networks-on-chip	×	√	√	√	√
Low-power area efficient high-speed I/O circuit techniques	×	×	√	×	×

Another technique to minimise energy consumption is to reduce supply voltage as proposed in Kim and Hwang (2008). However, reducing supply voltage results in circuit delays. To circumvent such an anomaly, Raghunathan et al. (2003) proposed the shortening of critical data link paths. However, reducing supply voltage to the data link paths introduces new problems, such as increase in the routing cost and also exposes the network to lower SNR (Bierrgaard and Mahadevan, 2006). Therefore, the network becomes more susceptible to interwire crosstalk, supply power noise, and radiation induced defects. Such a noisy interconnect behaves as an unreliable transport medium introducing errors at the transmitted signals. In Lauter et al. (2005), channel coding was used to restore network's reliability. This method is based on adding redundancy to the information symbol vector resulting in a longer coded vector that is nevertheless distinguishable at the output of the channel (Bierrgaard and Mahadevan, 2006). Besides reducing voltage swing and supply voltage, Fu and Ampadu (2008) also proposed an energy-efficient multiwire error control scheme using Hamming product codes. To

design a reliable and energy efficient NoC, Lauter et al. (2005) advocated that packet retransmission schemes are more energy-efficient than error protection schemes for long wires and strong codes. Strong error correcting codes have higher fault tolerance; however, coding schemes tend to use heavy compute power and as a consequence are energy inefficient. The same work also showed that error control schemes should be implemented at network level (end-to-end from source-to-destination) rather than link level (switch-to-switch).

Table 3 lists the pros and cons of the aforementioned circuit-level optimisation methods. The methods discussed in Lee et al. (2000) and Lu (2008) require small system area, have good SNR, but require extra power supply. On the other hand the methods presented in Lauter et al. (2005) and Zhang et al. (2000) do not need extra power supply but suffer from low SNR. Last, Ferretti and Beerel (2001) combine the benefits of the aforementioned works to some extent; however, the representative model is complex.

5 Architecture-level power modelling tools

Power models have been proposed for a variety of network types, such as on-chip FPGA networks (Zhang et al., 2000) and IP (Ye et al., 2000). All of these studies focused on investigating the power aspect of different network topologies. Their goal was to build tools for estimating average power consumption based on transistor count (Patel et al., 1997) and switch width (Zhang et al., 2000). However, the previous studies did not explore architectural-level information. Therefore, worst-case scenarios cannot be efficiently analysed.

Circuit-level power estimation tools, such as DVS (Shang et al., 2002) and DLS (Taylor et al., 2002a, 2002b) provide excellent accuracy but at the expense of long execution time. Improving simulation time has motivated architectural-level power simulators for processors and memories, such as Wattch (Brooks et al., 2000) and Orion (Wang et al., 2002). Wattch (Brooks et al., 2000) is a framework for analysing and optimising microprocessor power dissipation at the architectural-level. Orion (Wang et al., 2002) is a power-performance interconnection network simulator that is capable of providing detailed power characteristics. In addition, it also explores power-performance trade-offs at the architectural-level.

As an architecture-level tool for achieving energy efficiency in NoC routers, Shang et al. (2006) proposed a dynamic power management (DPM) scheme, termed PowerHerd (Shang et al., 2006). PowerHerd is a distributed scheme for dynamically satisfying peak-power constraints in interconnection networks. In PowerHerd, each router dynamically maintains a local power budget, controls its local power dissipation, and exchanges spare power resources with its neighbouring routers to optimise network performance. Simulations demonstrate that PowerHerd can effectively regulate network power consumption and can meet peak-power constraints with negligible network-performance penalty.

6 Achieving thermal efficiency in NoC

Chip temperature is an accumulated effect of both processing and communication components. Because networks consume a significant portion of the chip power budget

(Benini and Micheli, 2002; Wang et al., 2003) the power consumption induces a substantial thermal impact. Understanding the joint thermal behaviour of all on-chip components, both processors and networks, is key to achieve efficient thermal design. Authors in Brooks and Martonoshi (2001a) and Chen and Peh (2003) were among the first to consider thermal issues in microprocessors. Unlike centralised microprocessors, networks are distributed in nature imposing unique requirements on thermal modelling and management.

Different dynamic thermal management (DTM) schemes have been proposed in Brooks and Martonoshi (2001b) and Huang et al. (2001) for high performance microprocessors. Shang et al. (2002) proposed a model, termed ThermalHerd. ThermalHerd is a distributed run-time mechanism that dynamically regulates network temperature. Evaluation using NoC traffic traces from the UT-Austin TRIPS CMP (Hu and Marculescu, 2003), demonstrated that ThermalHerd can effectively regulate network temperature and eliminate thermal emergencies. Moreover, ThermalHerd proactively adjusts and balances the network thermal profile to achieve a lower junction temperature. Furthermore, Borkar (1999) proposed to reduce the die (defined in aforementioned text) size. However, reducing the die size results in increased supply current and decreased performance. Therefore, Gonzalez et al. (1997) suggested the term, energy delay product to be considered to overcome the issue. The energy delay product is the product of packet-switching delay of the router and energy consumed by the NoC. However, minimising energy delay product requires scaling of threshold voltage, which results in increased leakage current.

Gwennap (1996) suggested that the overall chip energy consumption and thermal efficiency may limit not only what can be integrated onto a chip, but also how fast the chip can be clocked. Gowan et al. (1998) proposed the use of hierarchical clocking scheme to lower the clock power consumption and overall temperature of a chip. However, the proposed scheme in Gowan et al. (1998) reduces chip performance and increases delay (circuit-switching delay of the chip).

To lower the energy consumption and increase thermal efficiency, Reinman et al. (2002) suggested small size cache for the router memory. However, because of the limiting size of the cache, the proposed technique is not scalable. Buyuktosunoglu et al., (2002) presented two techniques that reduce queue power consumption at routers. Evaluation revealed that for a given NoC, the proposed techniques resulted in less power consumption and consequently reduced thermal dissipation.

Pering et al. (1998) proposed a model that includes a thermal temperature sensor in hardware and an interrupt capability to notify software when a threshold temperature has been reached. The model also can include an instruction cache throttling mechanism that allows the NoC's bandwidth to be reduced when the system reaches a preset temperature (Seng et al., 2000). However, these models were mainly geared towards improving battery life for portable machines. The Transmeta Crusoe processor includes 'LongRun' technology that dynamically adjusts supply voltage and frequency to reduce power consumption and increase thermal efficiency (Transmeta, 2000). Although voltage and frequency tuning are quite effective at reducing power consumption, the delay in triggering these responses is high. Ghiasi et al. (2000) advocated a model based on the advanced configuration and power interface (ACPI) specification, in which hardware and software cooperate with each other to manage the NoC temperature dynamically. ACPI involves actions, such as turning on or off input/output (I/O) devices and managing multiple batteries. Rohou and Smith (1999) considered temperature measurement

feedback to guide the system in controlling the temperature of NoC. Sanchez et al. (1997) performed simulations using Watch (Brooks et al., 2000) to correlate power dissipation with other NoC parameter statistics. The results showed that the model reduces NoC temperature significantly with a slight degradation of performance.

All the noise sources, such as coupling capacitances between two neighbouring wires collectively induce a noise voltage on the channel that follows a Gaussian distribution (Sylvester and Keutzer, 2000). This induced voltage results in increased temperature of the NoC (Wingard, 2001). Simulation results in Worm et al. (2002) showed that tangible savings in energy can be attained while achieving more robustness to large variations in actual workload and noise.

Several routing algorithms have been developed to achieve thermal efficiency in NoC. Schurgers and Srivastava (2001) proposed an algorithm that minimises communication energy consumption and also balances out the spatial distribution of energy consumption in the network. Hu and Marculescu (2003) suggested an algorithm that avoids local hot spots and thereby simplifies SoC thermal management. Brooks and Martonoshi (2001a, 2001b) proposed a speculation based algorithm that uses dynamic thermal management to reduce the cooling system costs of NoC.

7 Conclusions

Energy efficient NoC design is necessary to optimise power consumption at SoCs and CMPs. Energy efficiency can be achieved through careful optimisation at different levels, such as component, circuit, and architectural level. This paper reviewed different techniques and tools to achieve energy efficiency at each of the aforementioned levels. Adopting a combination of techniques each aiming at different level is presumably the most viable approach in designing energy efficient interconnection network. That said, one must also tackle thermal issues for the NoC to be truly implementable.

Acknowledgements

This research was supported by the FNR under grant C09/IS/05 and by the North Dakota State University Research, Creative Activities & Technology Transfer grant.

References

- Adler, V. and Friedman, E.G. (1997) 'Repeater design to reduce delay and power in resistive interconnect', *The 1997 IEEE International Symposium on Circuits and Systems*, pp.2148–2151.
- Benini, L. and Micheli, G. (2001) 'Powering NoCs', *The 14th International Symposium on System Synthesis (ISSS)*, IEEE, pp.33–38.
- Benini, L. and Micheli, G. (2002) 'Networks on chips: a new SoC paradigm', *IEEE Computer*, Vol. 35, No. 1, pp.70–78.
- Bierrgaard, T. and Mahadevan, S. (2006) 'A survey of research and practices of NoC', *ACM Computing Surveys*, Vol. 38, No. 1, pp.1–51.
- Borkar, S. (1999) 'Design challenges of technology scaling', *IEEE Micro*, Vol. 19, No. 4, pp.23–29.

- Brooks, D. and Martonoshi, M. (2001a) 'Dynamic thermal management for high-performance microprocessors', *The Seventh International Symposium on High Performance Microprocessors*, pp.171–182.
- Brooks, D. and Martonoshi, M. (2001b) 'Adaptive thermal management for high-performance microprocessors', *Workshop on Complexity Effective Design*, pp.117–121.
- Brooks, D., Tiwari, V. and Martonoshi, M. (2000) 'Wattch: a framework for architectural level power analysis and optimization', *The 27th Annual International Symposium on Computer Architecture*, pp.83–94.
- Burd, T.D. (1998) *Energy Efficient Processor System Design*, Doctoral thesis, University of California, Berkeley.
- Burd, T.D. and Brodersen, R.W. (2000) 'Design issues for dynamic voltage scaling', *International Symposium on Low Power Electronics and Design*, pp.9–14.
- Burd, T.D., Pering, T.A., Stratakos, A.J. and Brodersen, R.W. (2000) 'A dynamic voltage scaled microprocessor system', *IEEE Journal of Solid-state Circuits*, Vol. 35, pp.1571–1580.
- Buyuktosunoglu, A., Albonesi, D., Schuster, S., Brooks, D., Bose, P., and Cook, P. (2002) 'Power-efficient issue queue design', in Buyuktosunoglu, A., Albonesi, D., Schuster, S., Brooks, D., Bose, P. and Cook, P. (Eds.): *Power Aware Computing*, pp.35–58, Kluwer Academic Publishers, Norwell, MA.
- Chen, T.C., Chien, S.Y., Huang, Y.W., Tsai, C.H., Chen, C.Y., Chen, T.W. and Chen, L.G. (2006) 'Analysis and architecture design of an HDTV720p 30 frames/s H.264/AVC encoder', *IEEE Transaction on Circuits and System for Video Technology*, Vol. 16, No. 6, pp.673–688.
- Chen, X. and Peh, L.S. (2003) 'Leakage power modeling and optimization in interconnection network', *The 2003 International Symposium on Low Power Electronics and Design*, pp.90–95.
- Colshan, R. and Jaroun, B. (1994) 'A novel reduced swing CMOS BUS interface circuit for high speed low power VLSI systems', *IEEE International Symposium on Circuits and Systems*, pp.351–354.
- Dally, W.J. and Peh, L.S. (2000) 'Flit-reservation flow control', *IEEE Transactions on Parallel and Distributed Systems*, Vol. 3, No. 3, pp.194–205.
- Dally, W.J. and Poulton, W.J. (1998) *Digital Systems Engineering*, Cambridge University Press.
- Dally, W.J. and Towles, B. (2001) 'Route packets, not wires: on-chip interconnection networks', *Design Automation Conference*, pp.684–689.
- Dumitras, T. and Marculescu, R. (2003) 'On-chip stochastic communication', *IEEE Design Automation and Test in Europe*, pp.790–795.
- Ferretti, M. and Beerel, P.A. (2001) 'Low-swing signaling using a dynamic diode-connected driver', *The 27th European Solid State Circuits Conference*, pp.369–372.
- Fu, B. and Ampadu, P. (2008) 'An energy-efficient multiwire error control scheme for reliable on-chip interconnects using hamming product codes', research article, *VLSI Design*, Vol. 2008, Article ID 109490, 14pp., Hindawi Publishing Corporation.
- Ghiasi, S., Casmira, J. and Grunwald, D. (2000) 'Using IPC variation in workloads with externally specified rates to reduce power consumption', *Complexity Effective Design at ISCA27*, pp.294–296.
- Ghosh, P., Sen, A., and Hall, A. (2009) 'Energy efficient application mapping to NoC processing elements operating at multiple voltage levels', *3rd ACM/IEEE International Symposium on Networks-on-Chip*, pp.80–85.
- Gonzalez, R., Gordon, B., and Horowitz, M. (1997) 'Supply and threshold voltage scaling for low-power CMOS', *IEEE J. Solid State Circuits*, Vol. 32, No. 8, pp.1210–1216.
- Gowan, M., Biro, L. and Jackson, D. (1998) 'Power considerations in the design of the alpha 21264 microprocessor', *35th Annual Design of Automation Conference*, pp.726–731.
- Gwennap, L. (1996) *Power Issues May Limit Future CPUs*. Microprocessor Report, 10pp., Intel Corporation.

- Hedge, R. and Shanbhag, N.R. (2000) 'Toward achieving energy efficiency in presence of deep submicron noise', *IEEE Transactions on VLSI systems*, Vol. 8, No. 4, pp.379–391.
- Hu, J. and Marculescu, R. (2003) 'Exploiting the routing flexibility for energy/performance aware mapping of regular NoC architectures', *IEEE Design Automation and Test in Europe*, pp.688–693.
- Huang, M., Renau, J., Yoo, S.M., and Torrellas, J. (2001) 'The design of DEETM: a framework for dynamic energy efficiency and temperature management', *Journal of Instruction-Level Parallelism*, Vol. 3, No. 9, pp.327–336.
- Irani, S., Shukla, S. and Gupta, R. (2003) 'Algorithms for power savings', *The 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp.37–46.
- Jayaseelan, R. and Mitra, T. (2009) 'Dynamic thermal management via architectural adaptation', *46th Annual Design Automation Conference*, pp.484–489.
- Jejuriker, R., Pereira, C. and Gupta, R.K. (2004) 'Leakage aware dynamic voltage scaling for real time embedded system', *The 41st Annual Design Automation Conference*, pp.275–280.
- Kim, W.J. and Hwang, S.Y. (2008) 'Design of an area-efficient and low-power NoC architecture using a hybrid network topology', *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, Vol. E91-A, No. 11, pp.3297–3303.
- Koziri, M., Dadaliaris, A., Stamoulis, G. and Katsavounidis, I. (2007) 'A novel low-power motion estimation design for H.264', *IEEE Conf. on Application-Specific Systems, Architectures and Processors (ASAP)*, pp.247–252.
- Lauter, R., Jantsch, A. and Vitkowski, A. (2005) *Power Analysis of Link Level and End-to-end Data Protection in NoCing*, Royal Institute of Technology Press, Stockholm, Sweden.
- Lee, M.J.E., Dally, W.J. and Chiang, P. (2000) 'Low-power area efficient high-speed I/O circuit techniques', *IEEE Journal of Solid-State Circuits*, Vol. 35, No. 11, pp.1591–1599.
- Lu, H.W. (2008) *Digitalized I/O Techniques for High Speed Serial Link System*, Doctoral Thesis, 146 pp., National Central University, Taiwan, Republic of China.
- Maheshwari, A. and Burleson, W. (2001) 'Current sensing techniques for global interconnects in very deep submicron (VDSM) CMOS', *IEEE Computer Society Workshop on VLSI*, p.66.
- Nakagome, Y., Itoh, K., Isoda, M., Takeuchi, K. and Aoki, M. (1993) 'Sub-1-V swing internal bus architecture for future low-power ULSI's', *IEEE Journal of Solid-State Circuits*, Vol. 28, No. 4, pp.414–419.
- Patel, C.S., Chai, C.S., Yalamanchili, S. and Schimmel, D.S. (1997) 'Power constrained design of multiprocessor interconnection networks', *International Conference on Computer Design*, pp.408–416.
- Pering, T., Burd, T. and Brodersen, R. (1998) 'The simulation and evaluation of dynamic voltage scaling algorithms', *International Symposium on Low Power Electronics and Design*, pp.23–27.
- Raghunathan, V., Srivastava, V. and Gupta, R.K. (2003) 'A survey of techniques for energy efficient on-chip communication', *Design Automation Conference*, pp.900–905.
- Reinman, G., Calder, B. and Austin, T. (2002) 'High performance and energy efficient serial prefetch architecture', *4th International Symposium on High Performance Computing*, pp.146–159.
- Rohou, E. and Smith, M. (1999) 'Dynamically managing processor temperature and power', *2nd Workshop on Feedback Directed Optimization*, pp.325–327.
- Sanchez, H., Kuttanna, B., Olson, T., Alexander, M., Gerosa, G., Philip, R. and Alvarez, J. (1997) 'Thermal management system for high performance power PCTM processors', *42nd IEEE Computer Conference*, p.325.
- Sankaralingam, K., Nagarajan, R., Liu, H., Kim, C., Huh, J., Burger, D., Keckler, S.W. and Moore, S.R. (2003) 'Exploiting ILP, TLP and DLP with the polymorphous TRIPS architecture', *ACM SIGARCH Computer Architecture News*, Vol. 31, No. 2, pp.422–433.

- Schurgers, C. and Srivastava, M. (2001) 'Energy efficiency routing in sensor networks', *Milcom*, pp.357–361.
- Seng, J., Tullsen, D. and Kai, G. (2000) 'Power-sensitive multi threaded architecture', *International Conference on Computer Design*, pp.123–129.
- Sgroi, M., Sheets, M., Mihal, A., Keutzer, K., Malik, S., Rabaey, J. and Sangiovanni, A. (2001) 'Addressing the system-on-a-chip interconnect woes through communication-based design', *The 38th Annual Design Automation Conference*, pp.667–672.
- Shang, L., Peh, L.S. and Jha, N.K. (2002) 'Power-efficient interconnection networks: dynamic voltage scaling with links', *Computer Architecture Letters*, Vol. 1, No. 1, p.6.
- Shang, L., Peh, L.S., and Jha, N.K. (2006) 'PowerHerd: dynamically satisfying peak power constraints in interconnection networks', *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 25, No. 1, pp.92–110.
- Shang, L., Peh, L.S., Kumar, A. and Jha, N.K. (2004) 'Thermal modeling, characterization and management of NoCs', *The International Symposium on Microarchitecture*, pp.67–78.
- Srinivasan, J. and Adve, S.V. (2003) 'Predictive dynamic thermal management for multimedia applications', *International Conference on Supercomputing*, pp.109–120.
- Sylvester, D. and Keutzer, K. (2000) 'A global wiring paradigm for deep submicron design', *IEEE Transaction CAD*, pp.242–252.
- Taylor, M.B., Agarwal, A., Lee, W. and Amarasinghe, S. (2002a) 'Scalar operand networks: on-chip interconnect for ILP in partitioned architectures', *International Symposium on High Performance Computer Architecture*, pp.341–353.
- Taylor, M.B., Kim, J., Miller, J., Wentzlaff, D., Ghodrati, F., Greenwald, B., Hoffman, H., Johnson, P., Lee, J.W., Ma, A., Saraf, A., Seneski, M., Shnidman, N., Strumpfen, V., Frank, S., Amarasinghe, S. and Agarwal, A. (2002b) 'The raw microprocessor: a computational fabric for software circuits and general-purpose programs', *IEEE Micro*, Vol. 22, No. 2, pp.25–35.
- Tenhunen, H. and Jantsch, A. (2003) 'Will networks on chip close productivity gap?', in Tenhunen, H. and Jantsch, A. (Eds.): *Networks on Chip*, Ch. 1, pp.3–18, Kluwer Academic Publishers, Hingham, MA.
- Tiwari, V., Singh, D., Rajgopal, S., Mehta, G., Patel, R. and Baez, F. (1998) 'Reducing power in high-performance micro-processors', *Design Automation Conference*, pp.732–737.
- Transmeta Corporation (2000) *The Technology Behind The Crusoe Processor Whitepaper*.
- Wang, H.S., Peh, L.S. and Malik, S. (2003) 'Power-driven design of router microarchitectures in NoCs', *International Symposium of Microarchitecture*, pp.105–116.
- Wang, H.S., Peh, L.S. and Malik, S. (2002) 'Orion: a power performance simulator for interconnection networks', *International Symposium of Microarchitecture*, pp.294–305.
- Wei, G.Y. and Horowitz, M. (2002) 'Energy-efficient design of high-speed links', in Pedram, M. and Rabaey, J. (Eds.): *Power Aware Design Methodologies*, Kluwer Academic Publishers, Norwell, MA, Ch. 8.
- Wingard, D. (2001) 'Micro-network based integration for SOCs', *ACM/IEEE Design Automation Conference*, pp.673–677.
- Worm, F., Lenne, P., Thiran, P. and Micheli, G. (2002) 'An adaptive low power transmission scheme for on-chip networks', *15th International Symposium on System Synthesis*, pp.92–100.
- Yamauchi, H., Akamatsu, H. and Fujita, T. (1995) 'An asymptotically zero power charge-recycling bus architecture for battery operated ultra-high data rate ULSI's', *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 4, pp.423–431.
- Ye, T.T., Benini, L. and Micheli, G.D. (2000) 'Analysis of power consumption on switch fabrics in network routers', *Design Automation Conference*, pp.524–529.
- Zhang, H., George, V. and Rabaey, J. (2000) 'Low-swing on-chip signaling techniques: effectiveness and robustness', *IEEE Transactions on VLSI systems*, Vol. 8, No. 3, pp.262–272.