

Design of Low power & Reliable Networks on Chip through joint crosstalk avoidance and forward error correction coding

Partha Pratim Pande¹, Amlan Ganguly¹, Brett Feero¹, Benjamin Belzer¹, Cristian Grecu²

¹*School of Electrical Engineering
& Computer Science
Washington State University
PO BOX 642752
Pullman, WA, USA
(pande, ganguly, bfeero, belzer)@eecs.wsu.edu*

²*Department of Electrical
& Computer Engineering
University of British Columbia
2332 Main Mall Vancouver,
BC, V6T 1Z4 Canada
grecuc@ece.ubc.ca*

Abstract

With the ever-increasing degrees of integration, design of communication architectures for big Systems on Chip (SoCs) is a challenge. The communication requirements of these large Multi Processor SoCs (MP-SoCs) are convened by the emerging network-on-a-chip (NoC) paradigm. To become a viable alternative IC design methodology, the NoC paradigm must address system-level reliability issues, which are among the dominant concerns for SoC design. The basic operations of NoCs are governed by on-chip packet switched networks. On the other hand, incorporation of different coding schemes in SoC design is being investigated as a means to increase system reliability. As NoCs are built on packet-switching, it is very natural to modify the data packets by adding extra bits of coded information to protect against any transient malfunction. By incorporating joint crosstalk avoidance coding (CAC) and forward error correction (FEC) schemes in the NoC data stream we are able to enhance the system reliability and at the same time reduce communication energy.

1. Introduction and motivation

The Network-on-Chip (NoC) design paradigm is viewed as an enabling solution [1] [2] for the integration of exceedingly high number of computational and storage blocks in a single chip. Widespread adoption of NoC paradigm will be possible if it addresses system level reliability issues in addition to easing the design process. The common characteristic of NoC interconnect architectures is that the functional blocks communicate with each other with the help of intelligent switches. Crosstalk between adjacent wires in the inter-switch link is an issue in NoC communication fabrics; it can cause timing violations and extra energy dissipation. Crosstalk avoidance codes (CACs) can be used to reduce the mutual inter-wire coupling capacitance and hence the energy dissipation of wire segments. By incorporating CAC schemes in NoC data streams we expect to enhance the system reliability and at the same time reduce communication energy, which will help to decrease the overall energy dissipation. CACs reduce the worst-case switching capacitance of a wire by ensuring that a transition from one codeword to another codeword does not cause adjacent wires to switch in opposite directions [3], but do not help to protect against any transient malfunctions like EM noise, alpha particle hit, ground bounce, etc. To make the system robust, in addition to CAC we need to incorporate FEC into the NoC data stream. CACs help to reduce the data dependent crosstalk and FECs will make the system robust.

Our aim in this paper is to explore the energy savings capability of joint crosstalk avoidance and forward error correction codes in different NoC architectures. Although energy dissipation of wire segments is inherently reduced in tandem with crosstalk, two additional energy dissipation sources must be considered. These include coder-decoder (codec) energy and the extra energy dissipation incurred with additional wire segments.

2. Related work

In recent years, there has been an evolving effort in developing on-chip networks to integrate increasingly large number of functional cores in a single die [1]. The role of different coding schemes in bus-based systems has been explored previously. In [4] the authors presented a unified framework for applying coding for systems on chips (SoCs). In addition, different low-power coding (LPC) techniques have been proposed to reduce power consumption of on-chip buses [5]. But these LPCs try to reduce only the self-transition in a wire. According to [6], the principal limitation of the applicability of the LPCs is that due to higher power dissipation in the codec blocks, the codes are energy efficient only if the length of the wire segment exceeds a certain limit. Error resiliency in NoC fabrics and the trade-offs involved in various error recovery schemes are discussed in [7]. In this work, the authors have investigated simple error detection codes like parity or cyclic redundancy check codes and single error-correcting, multiple error-detecting Hamming codes. The main drawback of this work is that these codes do not have any crosstalk avoidance characteristics, which is absolutely necessary in the deep submicron (DSM) era. The role of communication infrastructure on energy dissipation is discussed in [8]. Different strategies for power management of NoCs, such as power-aware on-off networks [9], and dynamic voltage scaling [10] have been addressed previously.

3. Data coding in NoC links

A few NoC interconnect architectures have been proposed by different research groups [4]. Figure 1 shows two of the most commonly used NoC architectures. Data exchange between the functional blocks takes place in the form of packets. Generally, wormhole switching is adopted for NoCs [8]. This scheme divides packets into fixed-length flow control units (flits), with I/O buffers storing only a few flits. The first flit, i.e., *header flit*, of a packet contains routing information. Header flit decoding enables the switches to establish the path and subsequent flits simply follow this path in a pipelined fashion.

The delay of an inter-switch wire in the NoC link depends on the transitions on the wire and wires adjacent to it. The worst case delay of a wire is $(1 + 4\lambda)\tau$, where τ is the delay of a crosstalk-free wire and λ is the ratio of the coupling capacitance to the bulk capacitance [11].

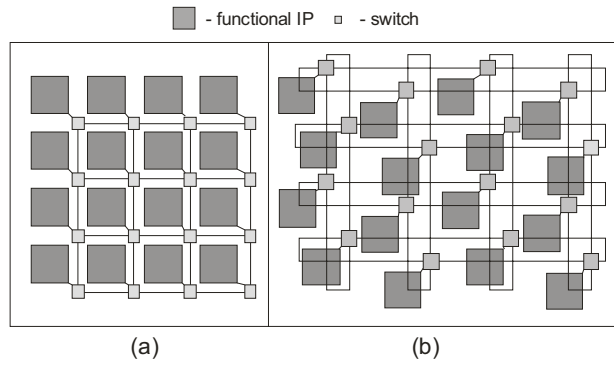


Figure 1: (a) Mesh; (b) Torus-based NoCs

The purpose of the crosstalk avoidance code is to reduce the delay of the line to $(1 + p\lambda)\tau$, where $p=1, 2$, or 3 is called the maximum coupling. As a result of reduction in the coupling capacitance, the CACs will reduce the energy dissipation per line in a NoC link. A number of crosstalk avoidance codes [3] were proposed in literature. In Forbidden Overlap Condition (FOC) codes the maximum coupling can be reduced to $p=3$ by ensuring that codewords with the bit pattern 010 do not transition to codewords with 101 at the same bit positions. Similarly in Forbidden Transition Condition (FTC) code and Forbidden Pattern Code (FPC) the maximum coupling can be reduced to $p=2$. To make the system robust, in addition to CAC we need to incorporate Forward Error Correcting Codes (FEC) into the NoC data stream. There are a few joint CAC and single error correction (SEC) codes among which Duplicate-add-parity (DAP) [4], Boundary Shift Code (BSC) [3] and Modified Dual Rail (MDR) code provide maximum coupling of $p=2$. Here we consider these three codes to study their effectiveness in reducing the communication energy in different NoC architectures.

3.1 DAP and MDR

The duplicate-add-parity (DAP) coding scheme uses duplication to reduce crosstalk [4]. By duplication, we can achieve Hamming distance of two, and with the addition of a single parity bit, the Hamming distance increases to three. The DAP encoder and decoder are shown in Figure 2. Encoding involves duplicating the bits of the link and calculating the parity of the link. Similarly, in decoding, the parity bit is recreated from one set of the data flit. As shown in Figure 2 bit y_8 is the previously calculated parity, and the other signal entering the xor gate is the parity of the more significant set (bits y_1, y_3, y_5 , and y_7). The new parity is compared with the original parity calculated in the encoder, and the error-free set is chosen. For example, in case of an error in the more significant set, the parities will differ, and the less significant set will be chosen as the decoded flit. On the other hand, if the error occurs in less significant set, the more significant set will be chosen.

The Modified Dual Rail (MDR) code is very similar to the DAP [12]. In the Dual Rail (DR) code, considering a link of k information bits, $m = k + 1$ check bits are added, leading to a code word length of $n = k + m = 2k + 1$.

We define the $k + 1$ check bits with the following equations:

$$c_i = d_i, \text{ for } i = 0 \text{ to } k - 1$$

$$c_k = d_0 \oplus d_1 \oplus \dots \oplus d_{k-1}$$

In the MDR two copies of parity bit C_k are placed adjacent to the other codeword bits, to reduce crosstalk.

3.2 The boundary shift code (BSC) scheme

The BSC coding scheme attempts to reduce crosstalk-induced delay by avoiding a shared boundary between successive codewords [13]. It is very similar to DAP in that it uses duplication and one parity bit to achieve crosstalk avoidance and single-error correction. However, the fundamental difference is that at each clock cycle, the parity bit is placed on the opposite side of the encoded flit. As can be seen in Figure 3, this coding scheme does not allow dependent boundaries in subsequent codewords. Encoding is achieved by duplicating bits and completing a parity calculation as in DAP. However, every second clock cycle will result in a one-bit shift. Similarly, the decoding structure is equivalent to that of DAP with the addition of a one-bit shift every other clock cycle before the parity check. Table 1 shows examples of different code words. The parity bits are indicated in bold.

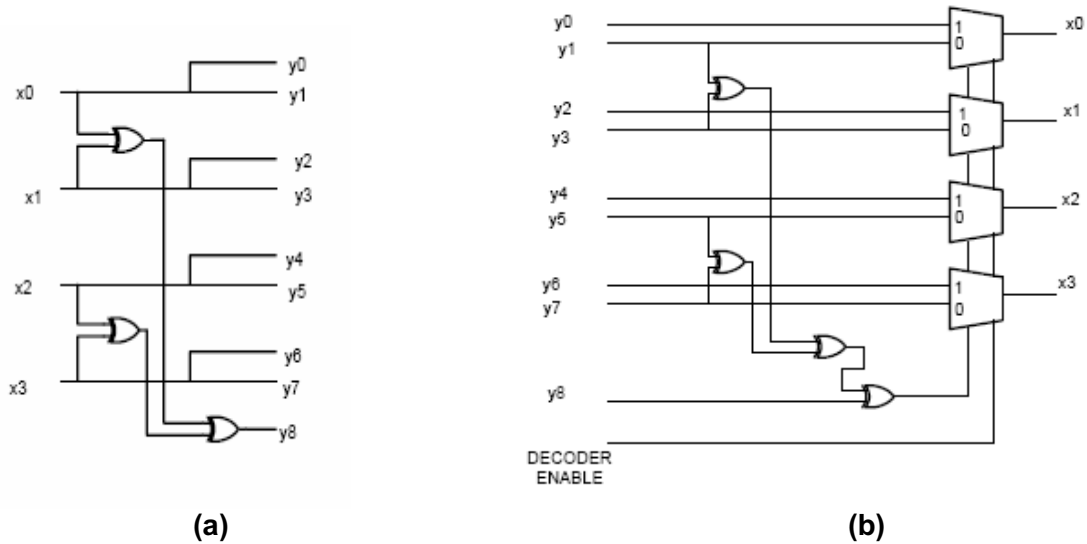


Figure 2: Duplicate-Add-Parity (DAP) coding (a) encoder; (b) decoder

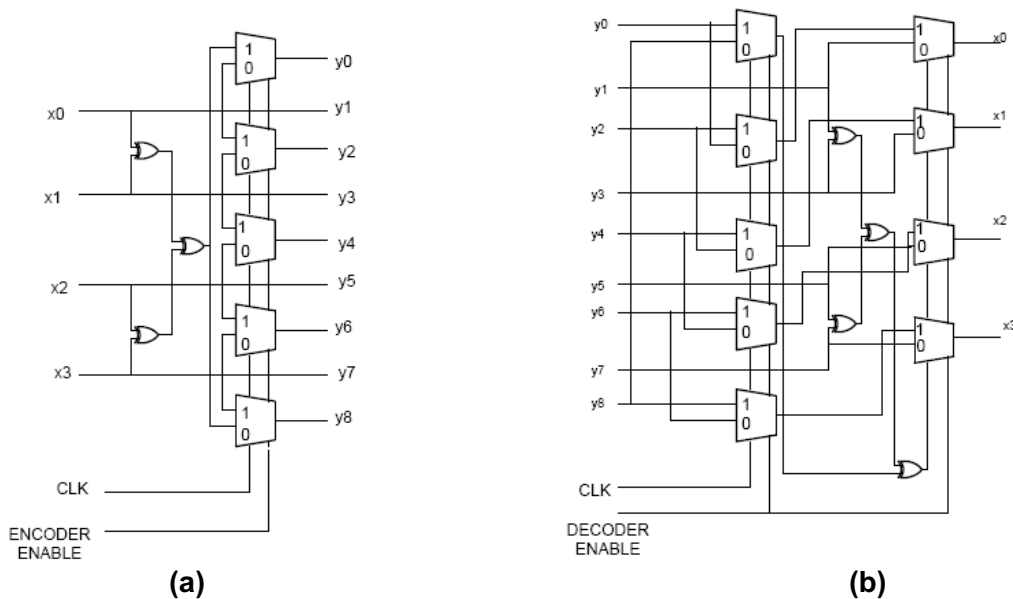


Figure 3: Boundary Shift Coding (BSC) (a) encoder; (b) decoder

Table 1: Coded flit structure for different coding schemes

Clock Cycle	Flit	BSC	DAP	MDR
1	0010	100001100	100001100	1100001100
2	0010	000011001	100001100	1100001100
3	1100	011110000	011110000	0011110000
4	1010	110011001	111001100	1111001100
5	0100	100110000	100110000	1100110000
6	0011	000011110	000001111	0000001111

4. Energy dissipation in a NoC-based SoC

When flits travel on the interconnection network, the inter-switch wires and the logic gates in the switches will toggle, resulting in energy dissipation. The flits from the source nodes need to traverse multiple hops consisting of switches and wires to reach destinations. Consequently, we determine the energy dissipated in each interconnect and switch hop. The energy per flit per hop for a NoC with uncoded bits is given by

$$E_{hop} = E_{switch} + E_{interconnect} \quad (4.1)$$

In presence of coding E_{hop} will be modified as

$$E_{hop} = E_{switch} + E_{codec} + E_{interconnect} \quad (4.2)$$

The energy dissipated in transporting a flit consisting of n bits through h hops can be calculated as

$$E_{flit} = \sum_{j=1}^h E_{hop,j} \quad (4.3)$$

Let P be the total number of flits transported, n be the number of bits in a flit and E_{flit_i} be the energy dissipated by the i^{th} flit, where i ranges from 1 to P . The average energy per bit $\overline{E_{bit}}$ is then calculated according to the following equation:

$$\overline{E_{bit}} = \frac{\sum_{i=1}^P E_{flit_i}}{P.n} = \frac{\sum_{i=1}^P \left(\sum_{j=1}^{h_i} E_{hop,j} \right)}{P.n} \quad (4.4)$$

In order to quantify the energy dissipation profile for a NoC interconnect architecture, we determine the energy dissipated in each switch, E_{switch} , and each codec, E_{codec} by running SynopsysTM Prime Power on the gate-level netlist of the switch and codec blocks and feeding a large set of data patterns. The experimental data set included long sequences of 1s and 0s to account for the possible cases where low transition activity data were to be transported. To determine the interconnect energy $E_{interconnect}$, the capacitance of each interconnect stage, $C_{interconnect}$ is calculated by taking into account the specific layout of each topology [14]. According to [15], the worst case capacitance of a wire is given by the following equation, where AR is the aspect ratio of the wire and C_{fringe} is the fringe capacitance.

$$C_{wire} = 2\epsilon_d \epsilon_0 \left(\frac{1 + 2AR^2}{AR} \right) + C_{fringe} \quad (4.5)$$

In the presence of coding the value of C_{wire} will be reduced according to the coding scheme and this will help in reducing $E_{interconnect}$. On the other hand incorporation of the codec

blocks will increase energy dissipation in the switches along with increased energy dissipation in the redundant wires as these extra wires will also dissipate energy now. Our aim is to study the effects of all these contributions on the overall energy dissipation in NoC communication infrastructures.

5. Experimental results and analysis

To evaluate the role of the joint coding schemes discussed above on the energy dissipation characteristics of NoC communication infrastructures, we considered a system consisting of 64 IP blocks and mapped them onto mesh and folded-torus based NoC architectures as shown in Figure 1. Messages were injected with a uniform traffic pattern, i.e., in each cycle, all IP cores can generate messages with the same probability. The routing mechanism used for all simulations was *e-cube* (dimension order) routing [8].

In NoC the inter-switch wire segments together with the intra-switch stages constitute a highly pipelined communication medium with the delay of each stage within the limit of one clock cycle [14]. By incorporating the coder/decoder we increase the number of pipelined stages in the communication path. We designed and synthesized the codec blocks using 130 nm standard cell libraries. We have verified through circuit level design and timing analysis that the delay of these blocks are well within this clock cycle limit. Consequently by adding the codec blocks we do not violate the pipelined communication architecture.

The energy dissipation of each inter-switch wire segment is a function of λ , the ratio of the coupling capacitance to the bulk capacitance. For a given interconnect geometry, the value of λ depends on the metal coverage in upper and lower metal layers [4]. We vary the value of λ from 1 to 4 [4]. Figure 4 shows the average bit energy dissipation as a function of the injection load with $\lambda=1$ and $\lambda=4$ respectively at the 130 nm technology node for an 8×8 mesh-based NoC. The injection load is expressed as the number of flits injected by each IP per cycle. By incorporating the error correction schemes in a NoC data stream, the reliability of the system is enhanced. Consequently, the supply voltage can be reduced without compromising system reliability. According to [4], the voltage level driving the inter-switch wire segments in presence of DAP, BSC, MDR was assumed to be 0.86V, which is lower than the nominal supply voltage of 1.2V. With this lowered operating voltage the residual probability of word error of a coded system is same as that of an uncoded system [4]. This is possible due to the enhanced robustness of the system. The average energy dissipation in transmitting a bit through the NoC is calculated according to equations (4.1)-(4.4).




Figure 4: Bit energy dissipation characteristics for a MESH-based NoC (a): $\lambda=1$; (b): $\lambda=4$.

It is evident from Figure 4 that all the coding schemes have potential for energy savings in a mesh-based NoC. Figure 5 shows the energy dissipation characteristics of a folded-torus-based

NoC, in which the energy savings are greater due to the fact that the inter-switch wire segments of the folded torus are longer than those of a mesh.




Figure 5: Bit energy dissipation characteristics for a Folded Torus-based NoC (a): $\lambda=1$; (b): $\lambda=4$.

Tables 2 and 3 quantify the energy dissipation characteristics of the coding schemes for two different values of λ at throughput saturation for both the MESH and Folded Torus networks. The relative amount of energy savings corresponding to each coding scheme is denoted by ΔES . It is evident from these tables that MDR and DAP are the most energy efficient.

Table 2: Average bit energy dissipation (pJ) for a MESH-based NoC at saturation

λ	Uncoded	BSC	ΔES (%)	DAP	ΔES (%)	MDR	ΔES (%)
1	215.06	177.18	17.61	168.75	21.53	168.90	21.46
4	253.92	184.04	27.52	175.61	30.84	175.73	30.79

Table 3: Average bit energy dissipation (pJ) for a Folded Torus-based NoC at saturation

λ	Uncoded	BSC	ΔES (%)	DAP	ΔES (%)	MDR	ΔES (%)
1	286.73	206.32	28.04	197.88	30.98	198.03	30.93
4	389.91	220.03	43.56	211.60	45.73	211.75	45.69

This trend in energy dissipation for different coding schemes can be explained by considering the effects of the redundant wires the coding schemes add on the energy dissipation profile. Though by incorporating coding the capacitance of each wire is reduced, extra wires are added to the system. We consider an uncoded 32-bit-wide inter-switch link and denote the energy of the whole link by $E_{uncoded}$ and the energy of the coded link including the redundant wires

by E_{coded} . The energy dissipated by the codec is denoted by E_{codec} . The energy savings in the interconnect segment arising out of incorporation of coding schemes is denoted by η , where $\eta = (E_{uncoded} - E_{coded})$. We are interested in the situations where $E_{coded} \leq E_{uncoded}$ and η is equal to or greater than E_{codec} as this is the case in which we

can achieve net energy savings per link. We plot the ratio $\gamma = \frac{\eta}{E_{codec}}$ as a function of λ for all

the coding schemes under consideration. For those values of λ where $\gamma \geq 1$, we have net energy savings per link. In these plots we assume that the uncoded and coded link have the same frequency of operation.

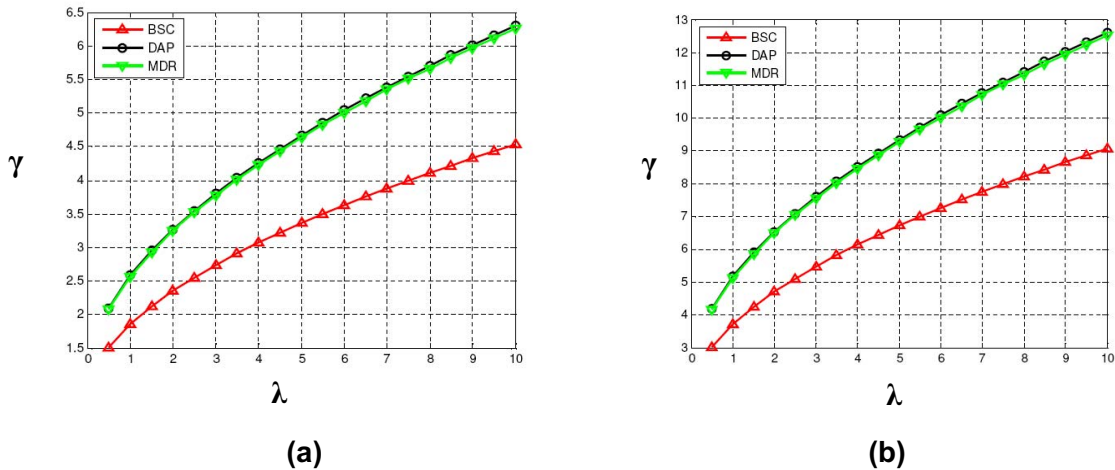


Figure 6: Energy savings characteristics of the coding schemes (a) Mesh (b) Folded Torus

It is clear from Figure 6 that the coding schemes under consideration provide energy savings for $\lambda \geq 1$ at 130 nm technology node for both the Mesh-based and Folded Torus-based NoC.

6. Conclusions and future work

Network-on-chip is emerging as a revolutionary method to integrate numerous cores in a single SoC. Widespread adoption of the NoC paradigm will be possible if it addresses system level reliability issues in addition to easing the design process. By incorporating joint crosstalk avoidance and single error correction codes into NoC data streams, it is possible to reduce the coupling capacitance of inter-switch wire segments, and at the same time, the communication infrastructure becomes robust against transient malfunctions. Our experiments show that the joint codes are capable of reducing the communication energy of the NoC infrastructure by around 31% for a 64-IP Mesh-based NoC and around 45% for a 64-IP folded torus-based NoC.

With aggressive supply voltage scaling and increase in deep sub micron noise, single error correcting codes cannot satisfy the reliability requirements. In the next phase of this work we propose to investigate joint crosstalk avoidance and multiple error correcting codes and their performance in NoC fabrics.

7. References

- [1] L. Benini and G. De Micheli, "Networks on Chips: A New SoC Paradigm," *Computer*, vol. 35, no. 1, Jan. 2002, pp. 70-78
- [2] P. Magarshack and P.G. Paulin, "System-on-Chip beyond the Nanometer Wall," *Proc. 40th Design Automation Conf. (DAC 03)*, ACM Press, 2003, pp. 419-424.
- [3] S. R. Sridhara, and N. R. Shanbhag, "Coding for reliable on-chip buses: fundamental limits and practical codes", *Proceedings of IEEE International Conference on VLSI Design*, 3-7 Jan 2005, pp. 417-422
- [4] S. R. Sridhara, and N. R. Shanbhag, "Coding for System-on-Chip Networks: A Unified Framework", *IEEE transactions on very large scale integration (VLSI) systems*, vol. 13, no. 6, June 2005, pp. 655-667
- [5] M. R. Stan and W. P. Bursleson, "Low-power encodings for global communication in CMOS VLSI", *IEEE transactions on very large scale integration (VLSI) systems*, Volume 5, Issue 4, Dec. 1997 pp:444 – 455
- [6] C. Kretzschmar, A. K. Nieuwland, D. Muller, "Why Transition Coding for Power Minimization of on-Chip Buses does not work", *Proceedings of the Design, Automation and Test in Europe Conference and Exhibition, (DATE)*, 16-20 Feb. 2004 pp. 512-517
- [7] S. Murali, G. De Micheli, L. Benini, T. Theodorides, N. Vijaykrishnan, and M. Irwin, "Analysis of Error Recovery Schemes for Networks on Chips," *IEEE Design & Test of Computers*, vol. 22, no. 5, pp. 434-442, 2005.

- [8] P. P. Pande, C. Grecu, M. Jones, A. Ivanov, R. Saleh, "Performance Evaluation and Design Trade-offs for Network on Chip Interconnect Architectures", *IEEE Transactions on Computers*, vol. 54, no. 8, pp. 1025-1040, August 2005.
- [9] V. Soteriou and L. S. Peh, "Design-space exploration of power-aware on/off interconnection networks", Proceedings of IEEE International Conference on Computer Design (ICCD), 11-13 Oct. 2004, pp:510 - 517
- [10] L. Shang, L. S. Peh and N. K. Jha, "Dynamic voltage scaling with links for power optimization of interconnection networks", Proceedings of the 9th International Symposium on High Performance Computer Architecture (HPCA-9), 8-12 Feb. 2003, pp. 91 – 102
- [11] P. P. Sotiriadis and A. P. Chandrakasan, "A bus energy model for deep submicron technology", *IEEE transactions on very large scale integration (VLSI) systems*, Volume 10, Issue 3, June 2002, pp. 341-350
- [12] D. Rossi, et al. "New ECC for Crosstalk Impact Minimization," *IEEE Design and Test of Computers*, 2005, pp. 340-348.
- [13] K. N. Patel, and I.L. Markov, "Error-Correction and Crosstalk Avoidance in DSM Busses," *IEEE transactions on very large scale integration (VLSI) systems*, Special Issue for System Level Interconnect Prediction (SLIP), 2003, pp. 1-5.
- [14] C. Grecu, P. P. Pande, A. Ivanov, R. Saleh, "Timing Analysis of Network on Chip Architectures for MP-SoC Platforms", *Microelectronics Journal, Elsevier*, Vol. 36, issue 9, pp. 833-845.
- [15] K.C. Saraswat, et al., Technology and reliability constrained future copper interconnects—part II: performance implications, *IEEE Transactions on Electron Devices*, Vol. 49, issue 4, 2002 pp. 598–604.