

Applicability of Energy Efficient Coding Methodology to Address Signal Integrity in 3D NoC Fabrics

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

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

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

Abstract

Three dimensional (3D) Network on Chip (NoC) has attracted researchers' attention recently. 3D NoCs are capable of achieving better system throughput and lower latency compared to the corresponding 2D implementations. To fully exploit the performance benefits of 3D architectures, it is imperative to address signal integrity issues in the design phase and its implications on energy dissipation. In this work we show that by incorporating joint crosstalk avoidance and multiple error correction schemes it is possible to enhance the robustness and reduce the energy dissipation simultaneously for both the 3D and more conventional planar NoC architectures. The achievable energy savings in 3D NoCs is significantly more than that in 2D structures.

1. Introduction

With shrinking geometries, global interconnects are becoming the principal performance bottleneck for high-performance Systems-on-Chip (SoCs) [1] [2] in terms of communication latency and power. The Network-on-Chip (NoC) model is emerging as a revolutionary methodology in solving the performance limitations arising out of long interconnects. In addition to providing a solution for the global wire delay problem, the NoC paradigm also eases integration of high numbers of intellectual property (IP) cores in a single SoC. However, the conventional two dimensional (2D) IC has limited floor-planning choices, and consequently it limits the performance enhancements arising out of NoC architectures. On the other hand, preliminary research on emerging three-dimensional silicon structures has shown these to be advantageous in terms of power and speed [3]. With this in mind, the performance improvement arising from the architectural advantages of NoCs will be significantly enhanced if 3D ICs are adopted as the basic fabrication standard. The amalgamation of these two emerging paradigms allows for the creation of new structures that enable significant performance enhancements in terms of energy, latency, and throughput over more traditional planar solutions. Previous research [4] has shown that 3D NoC architectures can offer significant performance

enhancement over their 2D counterparts. Consequently, 3D Networks on Chip will be viable solutions for building high performance SoCs in the near future.

As we continue to scale down the feature size in ultra deep sub-micron (UDSM) regime, the signal integrity of 3D NoC architectures will be affected by varied sources of transient errors arising due to temporary conditions of the SoC and environmental factors. Among the transient failure mechanisms, we can enumerate factors such as crosstalk, electromagnetic interference, alpha particle hits, cosmic radiation, etc. [5] [6]. As NoCs are built on packet-switching, it is easy to modify the data packets by adding extra bits of coded information to protect against transient malfunctions. Crosstalk avoidance coding (CAC) schemes are effective ways of reducing the worst-case switching capacitance of a wire by ensuring that a transition from one codeword to another does not cause adjacent wires to switch in opposite directions. Though CACs are effective in reducing mutual inter-wire coupling capacitance, they do not protect against any other transient errors. To make the system robust, in addition to CAC we need to incorporate forward error-correction coding (FEC) into the NoC data stream. Among different FECs, single error correcting codes (SECs) are the simplest to implement. Performance of joint CAC with single-error correction, such as Duplicate Add Parity (DAP) [9], Modified Dual Rail (MDR) [10] and Boundary Shift Code (BSC) [11] have been previously studied for more conventional 2D NoC fabrics. However, this may not suffice for future implementations. Aggressive supply-voltage scaling and increasing deep sub-micron noise in future-generation NoCs will prevent SECs from satisfying reliability requirements. Hence, in addition to already existing SECs, we propose to investigate joint CAC and multiple error correcting codes (MECs) and their performance in 3D NoC architectures.

2. Related Work

In [4], 3D ICs were proposed to improve performance of chip multi-processors. Drawing upon 3D IC research, they chose a hybridization of busses and networks to provide the interconnect fabric between CPUs and L2 caches. This fusion of NoC and bus architectures was evaluated for performance using




Figure 1: NoC Architectures (a) 2D MESH, (b) 3D MESH, (c) Stacked MESH




Figure 2: Switch Architectures (a) 3D MESH, (b) Stacked MESH

standard CPU benchmarks. However, this analysis pertains only to chip multi-processors and does not consider the use of three-dimensional network structures for application-specific SoCs. Some of these concerns are addressed in [12], which compares 2D MESH structures with their 3D counterparts by analyzing the zero-load latency of each architecture. This work shows some of the advantages of 3D NoCs, but it neither applies any real traffic pattern, nor does it measure other relevant performance metrics.

There have been few efforts by different research groups to address signal integrity in NoC fabrics by employing error correction coding. Error resiliency in NoC fabrics and the trade-offs involved in various error recovery schemes are discussed in [14]. In this work, the authors investigated simple error detection codes like parity or cyclic redundancy check codes and single error-correcting, multiple error-detecting Hamming codes. The basic principle of this work is that the receiver corrects only single bit error in a flow-control-unit (flit), but for multiple bit errors, it requests retransmission. In [7] [8] performance of different CAC and joint CAC/SEC schemes in NoC fabrics has been investigated. All these works primarily targets conventional 2D NoC architectures. In this paper we investigate the efficiency of joint crosstalk avoidance

and multiple error correction schemes in 3D NoC architectures.

3. 3D NoC Architectures

One of the well known 2D NoC architectures is the MESH as shown in Figure 1(a). A straightforward extension of the popular 2D MESH is the 3D MESH, shown in Figure 1(b). Another 3D structure, introduced by [4], is shown in Figure 1(c). This is the Stacked MESH architecture. It features multiple layers of 2D MESH structures stacked one upon another, with communication in the third dimension taking place through busses instead of a multi-hop network. This is desirable due to the short distance between the 2D MESH layers. In the case of a 64-IP system, these 3D NoCs consist of four blocks in each dimension (4x4x4) with interconnect lengths in the z-dimension on the order of tens of microns. A 3D MESH NoC will employ 7-port switches as shown in Figure 2(a) (1 each for North, South, East, West, up, down, and the IP). On the other hand a typical switch for the stacked MESH architecture will contain 6 ports: one link to the IP block, 2 links each in the x and y dimensions, and one link to the bus, as shown in Figure 2(b). This eliminates one of the ports for the z-direction compared to a standard 3D MESH.

4. Data Coding in NoC Links

The common characteristic of NoC architectures is that the functional intellectual property (IP) blocks communicate with each other via intelligent switches. The data communication between IP's in a NoC takes place in the form of packets routed through a wormhole switching mechanism. The packets are broken down into fixed length *flow control units* or *flits*. The switch blocks need to store only a few flits. The header flits carry the relevant routing information. Consequently header decoding enables the establishment of a path that the subsequent payload flits simply follow in a pipelined

fashion. The transmitted flits are encoded to guard against possible transient errors and crosstalk. The joint crosstalk avoidance and single error correction (CAC/SEC) schemes like DAP, MDR, and BSC reduce the switching capacitance associated with an inter-switch wire segments from $(1+4\lambda)C_L$ to $(1+2\lambda)C_L$ [13] in addition to achieving single error correction. Here, λ is the ratio of the coupling capacitance to the bulk capacitance and C_L is the load capacitance, which includes the self capacitance of the wire. We propose a novel, simple joint crosstalk avoidance and double error correction scheme called Crosstalk Avoidance Double Error Correction code (CADEC), which achieves the same crosstalk reduction like DAP, MDR or BSC.

4.1 CADEC Scheme

In the current nanometer technologies, single error correction capability of the coding schemes discussed above is not enough to ensure reliable functionality of the NoC's. We propose a novel joint Crosstalk Avoidance and Double Error Correction (CADEC) coding scheme in this paper and evaluate its performance in a 3D NoC architecture as compared to a 2D NoC structure.

The encoder consists of a simple (n, k) Hamming coding followed by DAP duplication into 2 copies 'a' and 'b'. After the Hamming coding, an overall parity is computed on the Hamming coded bits. All the bits except, the parity, are then duplicated thus making the total number of bits in the encoded flit to be $2n+1$. This is shown in Figure 3, for an uncoded flit length of 32-bit. This increases the Hamming distance between the adjacent code words to enable double error correction. At the same time duplication provides with crosstalk avoidance capability.

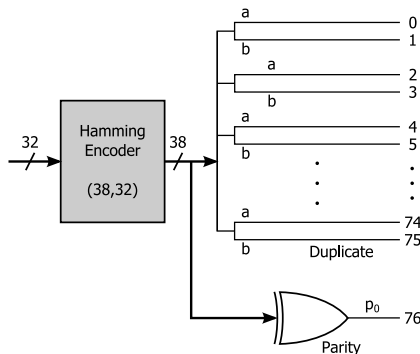


Figure 3: The CADEC encoder

At the decoder the parities are recreated locally from the individual Hamming copies and compared with each other. If the two local parities differ then the copy whose parity matches with the transmitted parity is selected for a final (n, k) single error correcting Hamming decoding. If the two local parities match, the syndrome of one copy is calculated and if the syndrome

is non-zero then that copy is selected, otherwise the other copy is selected for the final stage of (n, k) single error correction Hamming decoding. This algorithm is illustrated with the help of a flowchart in Figure 4.

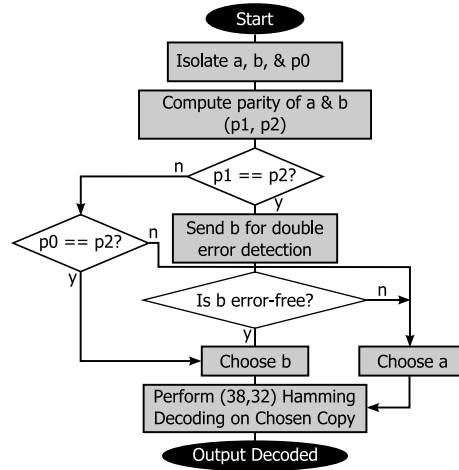


Figure 4: Decoding Algorithm for CADEC

5. Voltage swing reduction on interconnects

One of the significant outcomes of data coding in NoC links is that it increases reliability of the communication fabric. Consequently the noise margins and hence the voltage swings on the interconnect can be reduced without compromising system reliability. This reduction can be quantified by considering a white Gaussian noise voltage of magnitude V_N and variance or power of σ_N^2 that represents the cumulative effect of all the different sources of noise. This gives the probability of bit error, ϵ , also called the channel bit error rate (BER) as

$$\epsilon = Q\left(\frac{V_{dd}}{2\sigma_N}\right), \quad (1)$$

where, the Q -function is given by

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{y^2}{2}} dy \quad (2)$$

The word error probability is a function of the channel BER, ϵ . If $P_{unc}(\epsilon)$ is the probability of word error in the uncoded case and $P_{ecc}(\epsilon)$ is the residual probability of word error with error control coding, then it is desirable that, $P_{ecc}(\epsilon) < P_{unc}(\epsilon)$. Using (1), we can reduce the supply voltage in presence of coding to \hat{V}_{dd} , given by

$$\hat{V}_{dd} = V_{dd} \frac{Q^{-1}(\hat{\epsilon})}{Q^{-1}(\epsilon)} \quad (3)$$

In (3), V_{dd} is the nominal supply voltage in the absence of any coding and $\hat{\epsilon}$ is the modified BER such that $P_{ecc}(\hat{\epsilon}) = P_{unc}(\epsilon)$. Therefore, to compute the \hat{V}_{dd} for

various schemes we find the residual word error probability for each of the schemes used in this paper.

If (38,32) shortened Hamming code is used for error detection (ED) only, then the probability of undetected error in this case, for small values of BER ϵ , turns out to be

$$P_{ED}(\epsilon) = (n-k)\epsilon^2 \quad (4)$$

where $n=38$ and $k=32$ for the (38,32) shortened Hamming code.

The probability of an uncorrected error for DAP is shown in [9] to be

$$P_{DAP}(\epsilon) = \frac{3k(k+1)}{2}\epsilon^2. \quad (5)$$

where, k is the number of bits in the original flit which is 32 in this paper.

The probability is the same for BSC and MDR schemes as they are all single error correcting codes and the decoders are essentially similar in principle.

In case of CADEC, the probability of correct decoding can be found by considering each of the cases where the decoder can correctly decode flits despite errors. The various cases for which the CADEC scheme can correctly decode a received flit include the case where one Hamming copy is totally error-free while the other has anywhere from zero to all bits in error. This condition needs the sent parity bit to be error-free. The other possibility of correct decoding is when each of the two copies has a single error. The decoder is then able to correctly decode irrespective of the parity bit being in error or not. When the parity bit is in error there are several cases of correct decoding like, no error in either copy, zero or single error in one copy and any even number of errors and finally, the case where there is a single error in one copy and odd number of errors in the other. Except these cases all the other conditions will result in an erroneous decoding and the probability of those conditions give the residual probability of word error for the CADEC scheme. This probability can be now derived from the above conditions to be

$$P_{CADEC}(\epsilon) = n^2(n-4)\epsilon^3. \quad (6)$$

Where, n is the number of bits in a single Hamming coded copy which is 38 here, after the (38, 32) Hamming coding on the original flit of 32 bits.

Now, using the word error probabilities for the different coding scheme and equation (3) we can plot the maximum tolerable reduction in voltage swing on the interconnects as a function of varying word error rate. This plot is shown in Figure 5. The nominal voltage at the 90 nm technology node is assumed to be $V_{dd} = 1V$. As can be seen from Figure 5, the voltage swing is lower than the nominal voltage for all the coding schemes. The CADEC scheme provides

maximum voltage reduction as it can correct and also detect more errors than the others. For the purpose of simulations the voltage swings for different coding schemes corresponding to the word error probability of 10^{-20} is used later in the paper.

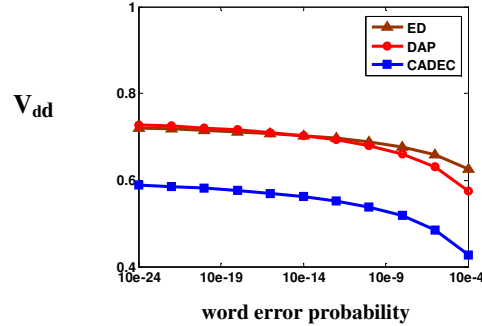


Figure 5: Possible lowering of voltage swing on interconnect wires in presence of coding

6. Simulation Methodologies

In order to simulate the performance of different NoC architectures under consideration, a flit-driven, cycle-accurate simulator using wormhole routing is used. The simulator models message injection using a self-similar distribution [15], since it has been shown to exhibit a more realistic traffic model than Poisson distributions. Self-similar traffic is characterized by modeling injections through a multitude of ON-OFF message sources according to a Pareto distribution. In order to calculate energy, the simulator traces each flit transmission, tracking energies associated with the switches, codecs, and inter-switch wires. It tracks energies in each hop since not all links are congruent. The measurements are made after initially running the simulations for long enough to allow transients to settle.

7. Experimental Results and Analysis

In order to quantify the performance of 3D NoC architectures in presence of various coding methodologies discussed here, we considered a system consisting of 64 IP blocks and mapped them onto 4x4x4 3D MESH and 8x8 2D MESH architectures. We analyze the performance of the 3DNoC architectures by incorporating the proposed CADEC scheme, as well as the previously proposed joint CAC/SEC code like DAP. At the same time to achieve more complete analysis, comparative results with respect to the sole error detection (ED) code were also obtained.

In order to characterize the performance of the 3D NoC architectures in presence of coding, we plotted the energy savings profile of the various architectures with varying injection load. All the schemes necessitate the use of different number of redundant bits and hence a fair comparison is possible only when we consider a full

coded flit in each case. Flits traverse multiple hops to reach their destinations, and we can model energy in the form of hop energy. To model savings, we must consider three parameters: uncoded inter-switch link energy $E_{link,uncoded}$, coded inter-switch link energy $E_{link,coded}$, and the codec energy E_{codec} . To determine the inter-switch link energy, the capacitance of each interconnect stage in the x-y dimension was calculated taking into account the specific layout of each topology [16]. Additionally the total interconnect length in the z-dimension was assumed to be 100 microns [3]. In presence of coding schemes incorporating crosstalk avoidance, the effective coupling component of this capacitance is reduced [13]. Thus, in addition to the voltage reduction due to increased reliability, the effective wire capacitance is also reduced. These two effects together contribute to reduce the energy dissipation of the inter-switch wire segments in presence of the joint codes. However, in addition to the codec energy the effect of the redundant wires due to incorporation of coding, which adds to the energy expenditure is also considered. The metric used for energy savings in this work for comparison thus takes into account the savings in energy due to the reduced crosstalk and reduced voltage level on the inter-switch wire segment, additional energy dissipated in the extra redundant wires and the codecs.

The energy dissipated in each codec, E_{codec} was determined by running SynopsysTM Prime Power on the gate-level netlist of the codec blocks using 90 nm standard cell libraries from CMP [17]. The energy dissipation due to the retransmission buffers and ARQ signals for the ED scheme were considered as well. With these parameters, the savings in energy per flit per hop, say the j^{th} hop, can be given by

$$E_{savings,j} = E_{link,uncoded} - (E_{link,coded} + E_{codec}) \quad (7)$$

From this, the energy savings in transporting a single flit, say the i^{th} flit, through h_i hops can be calculated as

$$E_{savings,i} = \sum_{j=1}^{h_i} E_{savings,j} \quad (8)$$

Therefore, the average energy savings in transporting a packet consisting of P such flits through h_i hops for each flit will be given as,

$$\bar{E}_{savings} = \frac{\sum_{i=1}^P \sum_{j=1}^{h_i} (E_{savings,j})}{P} \quad (9)$$

The plots in Figure 6 to 8 show the energy savings profile for 2D MESH, 3D MESH and 3D Stacked MESH architectures compared to the uncoded case for each of them, incorporating different coding schemes at the two representative values of λ for the 90nm technology node, namely 1 and 6 [18].

As seen from these figures incorporation of a coding scheme always results in savings in energy dissipation irrespective of the architecture used. The

savings are however, different depending on the particular architecture. It is evident that the CADEC is the most energy efficient scheme. Though CADEC adds more redundant wires than the other coding schemes considered here, it allows much larger reduction in voltage swing due to the double error correcting capability. The plots clearly show that the advantages of coding that were obtained in the 2D architectures are also equally possible in the 3D cases. Both the 3D NoC architectures achieve significant energy savings in presence of coding compared to the corresponding uncoded situation.




Figure 6: Energy savings profile in 2D MESH.




Figure 7: Energy savings profile in 3D MESH.




Figure 8: Energy savings profile in 3D Stacked MESH.

However, a better picture of the advantages of adopting a 3D architecture from the stand point of energy dissipation can be seen if the packet energies are investigated. The packet energy is the energy dissipated by a single packet as it travels from the source to its destination through multiple hops. Figure 9 shows the difference in energy savings per packet between the different architectures with respect to the basic 2D MESH without any coding.

To understand this, two factors must be taken into account. First, 3D structures have less number of hops between a pair of source and destination [12] compared to their 2D counterparts. Fewer hops clearly translate into significantly less energy per packet. Second, the wire lengths in the z direction are much smaller than that in x-y directions. This happens due to the fact that the planer layers of silicon are separated by very small distance, on the order of a few tens of microns. As a result of this, the energy dissipated by data transmission in z-dimension is very small. It is so small that in the z-dimension, the achievable energy savings due to crosstalk avoidance and voltage swing reduction is nullified by the energy addition from the codecs and redundant wires. Yet, savings are still prevalent in the x- and y-dimensions (interconnects are much longer in the 2D case). So, overall, the reduction in hop count, and significant savings in the x- and y-dimensions result in considerable overall energy savings.

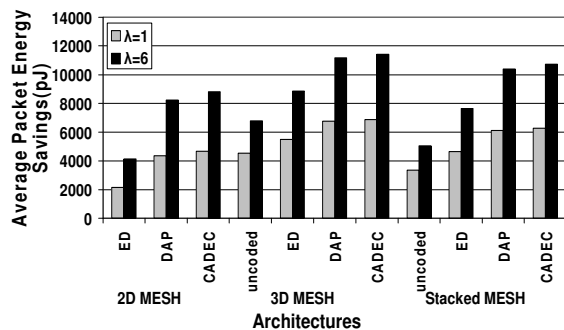


Figure 9: Savings in average packet energy over an uncoded 2D MESH architecture

8. Conclusions

3D NoC architectures are emerging as a promising solution for high-performance multi-core systems on chip (SoCs). They have clear advantages over more conventional planar counterparts in terms of system level performance metrics like throughput and latency. In the ultra deep submicron era one of the critical challenges is to address signal integrity in an energy efficient manner. The sources of transient errors that limit reliable inter-IP communication in conventional planer NoC's will continue to affect the 3D architectures. With device dimensions shrinking rapidly, existing single error correction mechanisms are not adequate to handle varied sources of transient noise affecting the signal integrity. Simple joint crosstalk avoiding multiple error correcting schemes need to be incorporated to enhance the robustness of the system. In this paper we have demonstrated that by incorporating simple joint crosstalk avoidance and double error correction codes it is possible to achieve significant amount of energy savings in 3D NoC architectures, especially when compared to the corresponding 2D counterparts.

9. References

- [1] L. Benini and G. De Micheli, "Networks on Chips: A New SoC Paradigm," *IEEE Computer*, Jan. 2002, pp. 70-78.
- [2] P. Magarshack and P.G. Paulin, "System-on-Chip beyond the Nanometer Wall," Proceedings of 40th Design Automation Conf. (DAC 03), ACM Press, 2003, pp. 419-424.
- [3] A. W. Topol et al., "Three-dimensional integrated circuits," IBM J. RES. & DEV. VOL. 50 NO. 4/5 JULY/SEPTEMBER 2006
- [4] F. Li et al., "Design and Management of 3D Chip Multiprocessors Using Network-in-Memory", Proceedings of the 33rd International Symposium on Computer Architecture (ISCA'06), pp. 130-141
- [5] E. Dupont, M. Nicolaidis, P. Rohr, "Embedded Robustness IPs for Transient-Error-Free ICs", *IEEE Design and Test of Computers*, Volume 19, Issue 3, May-June 2002 pp: 54 – 68.
- [6] S. Mitra, N. Seifert, M. Zhang, Q. Shi and K.S. Kim, "Robust System Design with Built-In Soft Error Resilience," *IEEE Computer*, Vol. 38, Number 2, Feb. 2005, pp. 43-52.
- [7] P. P. Pande, H. Zhu, A. Ganguly, C. Grecu, "Crosstalk-aware Energy Reduction in NoC Communication Fabrics", Proceedings of IEEE International SOC Conference, SOCC 2006, 24th-27th September, 2006, pp: 225-228.
- [8] P. P. Pande, A. Ganguly, B. Feero, B. Belzer, C. Grecu, "Design of Low power & Reliable Networks on Chip through joint crosstalk avoidance and forward error correction coding", Proceedings of 21st IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT 06), 4th-6th October, 2006.
- [9] S. R. Sridhara, and N. R. Shanbhag, "Coding for System-on-Chip Networks: A Unified Framework", *IEEE Transactions on Very Large Scale Integration (TVLSI) Systems*, vol. 13, no. 6, June 2005, pp. 655-667.
- [10] D. Rossi, C. Metra, A. K. Nieuwland and A. Katoch, "Exploiting ECC Redundancy to Minimize Crosstalk Impact", *IEEE Design & Test of Computers*, Volume 22, issue 1, Jan 2005 pp:59 – 70.
- [11] 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.
- [12] V. F. Pavlidis and E. G. Friedman, "3-D Topologies for Networks-on-Chip", Proceedings of IEEE SOC Conference, (SOCC 2006), pp. 285-288.
- [13] 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.
- [14] S. Murali, G. De Micheli, L. Benini, T. Theocharides, N. Vijaykrishnan, and M. Irwin, "Analysis of Error Recovery Schemes for Networks on Chips," *IEEE Design & Test of Computers*, vol. 22, no. 5, 2005, pp. 434-442.
- [15] K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*. John Wiley & Sons, 2000.
- [16] 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.
- [17] <http://cmp.imag.fr/index.php>
- [18] D. Sylvester and C. Hu, "Analytical modeling and characterization of deep-submicrometer interconnect," *Proc. IEEE*, vol. 89, no. 5, pp. 634-664, May 2001.