

Order is Power: Selective Packet Interleaving for Energy Efficient Networks-on-Chip

Amit Berman

Ran Ginosar

Idit Keidar

Department of Electrical Engineering, Technion – Israel Institute of Technology
Technion City, Haifa 32000, Israel

{bermanam@tx, ran@ee, idish@ee}.technion.ac.il

Abstract—Network-on-Chip (NoC) links consume a significant fraction of the total NoC power. We present Selective Packet Interleaving (SPI), a flit transmission scheme that reduces power consumption in NoC links. SPI decreases the number of bit transitions in the links by exploiting the multiplicity of virtual channels in a NoC router. SPI multiplexes flits to the router's output link so as to minimize the number of bit transitions from the previously transmitted flit. Analysis and simulations demonstrate a reduction of up to 55% in the number of bit transitions and up to 40% savings in power consumed on the link. SPI benefits grow with the number of virtual channels. SPI works better for links with a small number of bits in parallel. While SPI compares favorably against bus inversion, combining both schemes helps to further reduce bit transitions.

Keywords—network-on-chip; low-power design techniques; VLSI interconnects; routing;

I. INTRODUCTION

Power consumption is becoming a crucial factor in the design of high-speed digital systems, [1, 3, 4, 7, 10, 16, 17, 28, 31]. Whereas static power consumption is due to leakage and short-circuit currents, dynamic power consumption stems from switching activity. Interconnects consume the lion's share of dynamic power in modern chips. For example, studies show that interconnect links consume up to 60% of the dynamic power in NoCs [1, 29], more than 60% of the dynamic power in a modern microprocessor [18], and more than 90% in FPGA [14]. This portion is apparently growing [1, 8, 10, 23, 28, 31]. In this paper, we address the challenge of reducing dynamic power consumption in the context of NoC links.

A conventional NoC consists of a packet-switched network with a two-dimensional mesh topology [6, 26]. NoCs typically employ wormhole routing, i.e., each packet is divided into smaller units called flits, which are forwarded individually on links [15, 26]. NoC routers typically employ multiple virtual channels (VCs) [21, 26, 31], which allows them to transmit several flows in parallel by interleaving their flits on the outgoing link. Currently proposed NoCs employ between two and four VCs [15, 21], but studies argue that this number

should increase in future NoCs in order to supply higher throughput demands [26]. Each VC holds a buffer of flits pertaining to one flow. A NoC router with three VCs is depicted in Fig. 1.

We present Selective Packet Interleaving (SPI), a new flit transmission scheme for NoC-based systems. SPI reduces the dynamic power consumption on NoC links by reducing the number of bit transitions. SPI exploits the multiplicity of VCs in the NoC router, and selects the next flit to be transmitted so as to minimize the number of bit transitions with respect to the bus state (the previously transmitted flit). We illustrate this concept in Fig. 2, with two VCs of a certain output port. Each VC holds a packet consisting of four flits, and each flit consists of four bits. The VCs are multiplexed by the output port. Initially, the output is 0110. The output port can then select either flit 1001 of VC1 or flit 1110 from VC2. Selecting 1001 results in four bit transitions on the output link relative to the previously transmitted flit, whereas selecting of 1110 results in only one bit transition. Therefore, the output port selects the 1110 flit, according to SPI. Fig. 1 shows where SPI is inserted in the router.

SPI complements low-power design techniques such as voltage scaling, and can be implemented on top of such methods. In contrast to these approaches, SPI is technology-independent. Similarly, SPI is orthogonal to system-level design optimizations such as power-aware module placement, and unlike them, SPI does not require any a priori knowledge of the interconnect traffic patterns. Thus, SPI is broadly applicable, and can co-exist with a range of additional power optimizations. Our simulation results show that SPI can reduce the number of bit transitions by up to 55%. SPI benefits grow with the number of virtual channels. We simulate SPI with benchmarks as in [22].

We synthesize SPI with VLSI design tools, and derive the resulting power reduction using place and route design automation tools mapped to the Tower Semiconductor 0.18 μ m process library. Power analysis results show that SPI can yield up to 40% power savings in power consumed by the link.

The rest of the paper is organized as follows: in Section 2 we survey related work on low-power design techniques for NoC. In Section 3 we present the SPI scheme in detail. In Section 4, we simulate SPI with benchmarks. In Section 5, we present a conceptual VLSI implementation of SPI together with

power, area and latency analysis. Finally, in Section 6 we conclude the paper and propose some ideas for future research.

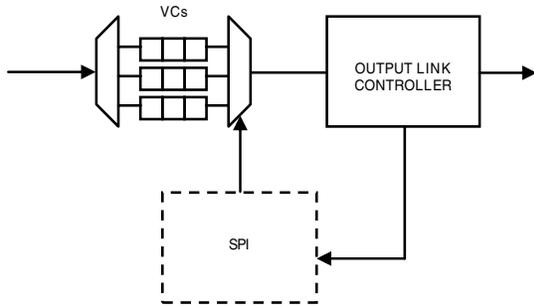


Figure 1. Virtual channels (VCs) in a NoC router's output buffer.

II. RELATED WORK

System-level power design approaches include synthesis algorithms to increase the power efficiency in interconnection networks via better module placement [11] or improved application design [19]. In such methods, the traffic patterns among the cores need to be known a-priori. In contrast, the approach we present in this paper does not require a-priori knowledge of the interconnect usage.

Embedded power design approaches include techniques to minimize the number of bit transitions sent in each packet. For example, in [2] a routing algorithm is used to reduce the redundant bits transmission implied by error protection code to enable fault tolerant communication. That method is applicable only when parity bits are used. SPI can complement it, and combining both schemes can help to further reduce power.

Data encoding is often employed to decrease the number of bit transitions over interconnects. Popular methods include bus-invert (BI) [28], adaptive coding [12], gray coding [20] and transition method [25]. Of these, we elaborate only on bus-invert, (BI), which was shown to be the most effective in NoCs [22]. Bus-invert compares the data to be transmitted with the current data on link. If the hamming distance (the number of bits in which the data patterns differ) between the new information and the link state is larger than half the number of bits (wires) on the link, then the data pattern is inverted before transmission. To enable restoring the original data pattern, an extra control wire is added to the link, in which a transmission of 1 indicates data inversion. Analysis [27] shows that on link widths of more than 8 bits, the savings are insufficient to justify the overhead of encoding circuits, and therefore wider links are segmented.

Previous work [22] has investigated the reduction of NoC power consumption achieved using the various data encoding schemes. Experiments in 0.35μ technology showed that BI achieves the best results. We therefore compare SPI to and combine SPI with BI in this paper. Nevertheless, the achieved power gain is offset by the overhead required to implement the BI encoding scheme. In contrast, the power savings achieved by SPI are higher than the power consumed by the required overhead.

Power may also be reduced using low-power device and circuit design techniques, such as dynamic voltage and frequency scaling (DVFS) [13], which adjust the supply voltage and clock rate dynamically according to circuit parameters. The energy efficiency of DVFS is highly dependant on the slack of the circuit. Another approach uses low-swing signaling techniques [9, 14, 30], the efficiency of which depends on circuit layout and manufacturing parameters. In contrast, the approach we present in this paper does not require knowledge about the circuit layout or manufacturing parameters.

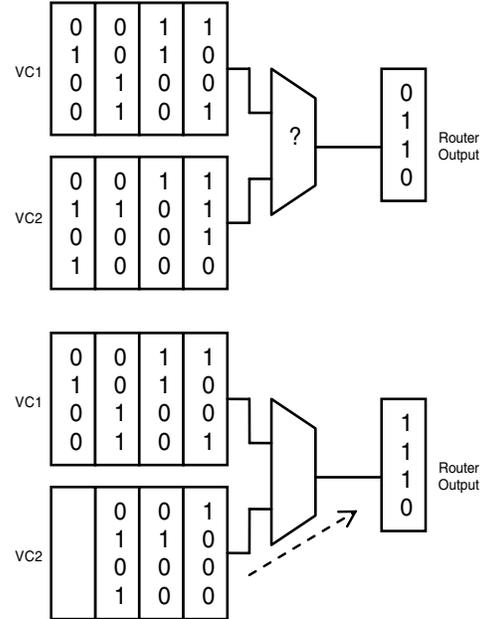


Figure 2. Example of SPI flit transmission in routing two packets using two virtual channels (VCs). Each packet consists of four 4-bit flits, marked as separate vertical blocks. SPI selects to transmit the 1110 flit from VC2 as it entails fewer bit transitions than 1001 with respect to the link's previous state (0110).

III. SELECTIVE PACKET INTERLEAVING

The goal of SPI is to reduce the number of bit transitions in NoC interconnect links. SPI achieves this goal by minimizing the Hamming distance between every pair of successive flits. SPI affects flit transmission: at every transmission time slot, it selects a flit from one of the m virtual channels and transmits it over the router's outgoing link. Only flits at the heads of virtual channel buffers can be selected. SPI can replace any existing interleaving scheme. Typical NoC router implements simple interleaving schemes like round-robin [16].

We first introduce some notations.

- $d_H(f_p, f_q)$ - The Hamming distance between flits f_p and f_q , i.e., the number of ones in the bit-wise xor of the flits;
- m - the number of virtual channels connected to an output port in a NoC router;

- m - the number of virtual channels connected to an output port in a NoC router;
- f_1, f_2, \dots, f_m – the flits at the head of the virtual channel buffers;
- f_{LINK} - the current logic state of the output link.

SPI selects a flit out of f_1, f_2, \dots, f_m for which the minimum number of bit transitions between successive flits is obtained. This rule can be expressed using the following equation:

$$SPI(f_1, f_2, \dots, f_m, f_{LINK}) = \arg \min_{1 \leq i \leq m} d_H(f_{LINK}, f_i)$$

Fig. 3 illustrates the SPI flit transmission mechanism. The SPI control logic is marked with dashed line.

In order to guarantee packet transmission latency, internal counters limit the use of SPI for each virtual channel. Without such counters, one or more VCs may be unable to transmit flits over a long time, because the first flit in their buffers incurs more bit transitions than flits in other VCs. This condition is known as starvation. In typical traffic patterns the probability for starvation is small. Nevertheless, one may wish to improve fairness in order to meet certain real-time requirements. It is easy to do so, and eliminate starvation altogether, by using counters. Since such an extension of SPI is straightforward, we do not elaborate on it here.

IV. BENCHMARK SIMULATIONS

In this section, we simulate flit transmission schemes with real workloads of the file type as used in [22]. We simulate the uncoded, bus-invert, SPI and the combined SPI+BI flit transmission schemes. In Section A we describe the simulated model and setup. In Section B we present the benchmark simulation results, which validate our analysis. In this section, we do not simulate the VC identification bits. Finally, in Section C, we add the VC identification bits to the transmissions, and examine their impact on the number of transitions.

A. Methodology

We use a cycle-accurate router simulator that models the output buffer with a given number of VCs, a given flit transmission scheme, and the router's output link. We assume that flit size and link width are equal, having the same number of bits n . The model is implemented in MATLAB. We count the number of bit transitions on the output link with different numbers of VCs, different benchmark streams, and the four studied flit transmission schemes.

We measure the average number of bit transitions over the link, i.e., the total number of bit transitions divided by (number of packets) × (flits per packet) in the streams. We assume that the flit size is equal to the link width and there are 128 flits per packet. In each simulation, streams of files of the same benchmark type are transmitted via all VCs of the simulated router, so that all VCs are fully utilized at all time. Each file is sent through a different VC. Results are shown in Fig. 4 with

(a) two VCs and 16b width links; and (b) eight VCs and 8b width links.

The workload file types are the same as used in [22], namely: jpg, pdf, mp3, bmp, tiff, wav, html, gcc, gzip, raw, and bytecode. For each file type, we download from the web a random collection of files of this type.

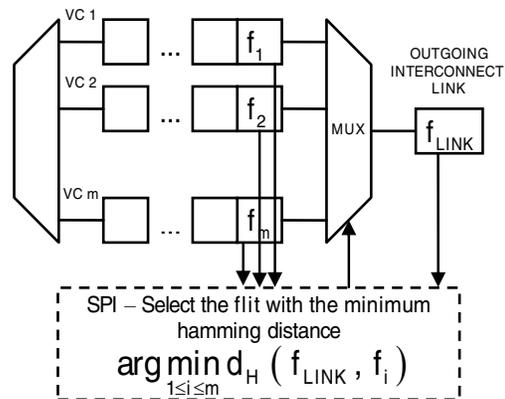


Figure 3. SPI flit selection with m virtual channels in a NoC router.

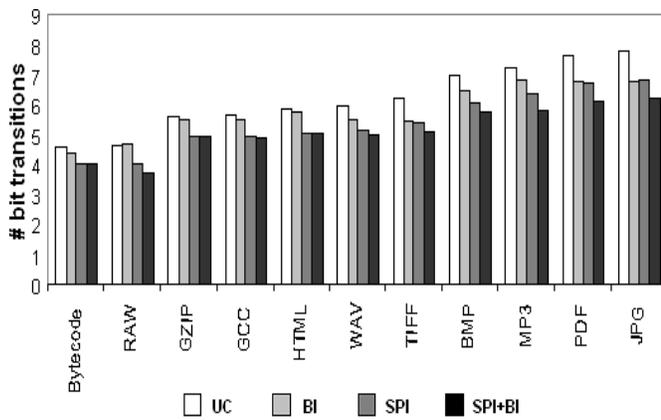
B. Simulation Results

The simulation results of all benchmarks are depicted in Fig. 4. We observe similar reductions in the number of bit transitions over all benchmarks. With two VCs and 16b links, the percentage of reduction in bit transitions of SPI relative to uncoded transmissions ranges from 10% to 13%, whereas with eight VCs and 8b links, we see reductions of 45%-55%. In general, the SPI+BI scheme presents the best results. However, the gap between SPI+BI and SPI in the tested configuration is small and they achieve relatively similar results.

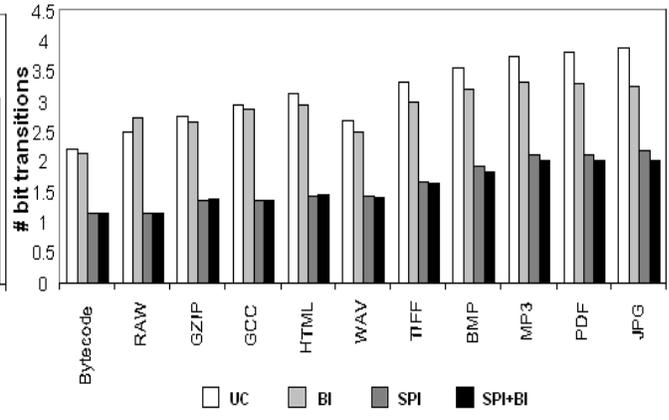
We next zoom in on one of the benchmarks, and experiment with it in a larger design space of parameter values. We simulate the specific MP3 benchmark with the four studied flit transmission schemes. Results are shown in Fig. 5 with (a) SPI and BI improvement over uncoded; and (b) SPI+BI improvement over uncoded. We observe that for 8b links, SPI outperforms BI starting from two VCs. For 16b and 32b links, SPI outperforms BI starting from three VCs.

C. Evaluation of Link and Interleaving Overhead

In practice, the VC identification is transmitted on separate lines. Hence, an additional $\log_2 m$ bits is required in order to identify the VC number. We now simulate the four studied bit transmission schemes with link widths of $n + \log_2 m$ bits. We use the same benchmarks as described in Section A, and add benchmark of random data patterns. In the uncoded case, we assume round-robin interleaving, where each transmission incurs one bit transition in the VC identification bits. Results are shown in Fig. 6. We observe a 22%-57% reduction in the number of bit transitions with SPI+VC identification compared to the uncoded+VC identification case. SPI+VC identification advantage increases with the number of VCs.

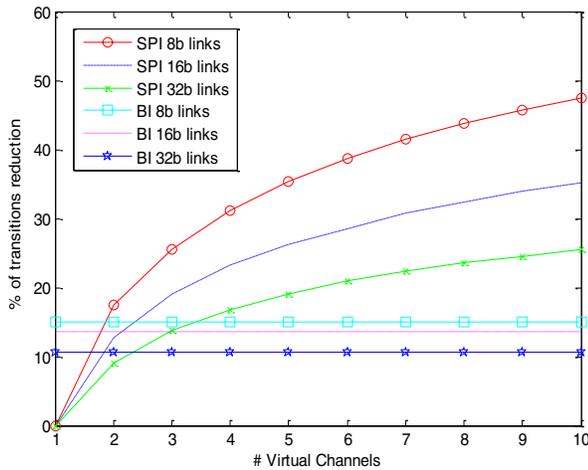


(a) Benchmark simulations: 2 VCs, 16b links.

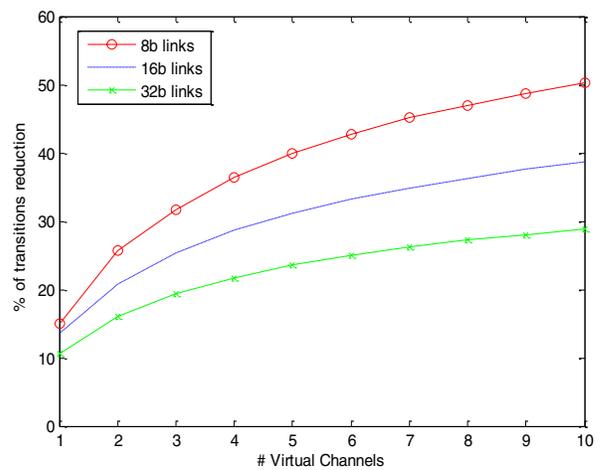


(b) Benchmark simulations: 8 VCs, 8b links.

Figure 4. Average number of bit transitions with various benchmarks. We observe that the percentage of improvement is similar for all benchmarks: with two VCs and 16b links, improvement rates of SPI and SPI+BI over uncoded are 10% to 13%, whereas with eight VCs and 8b links, the improvements are 45%-55%.



(a) SPI and BI improvement over uncoded



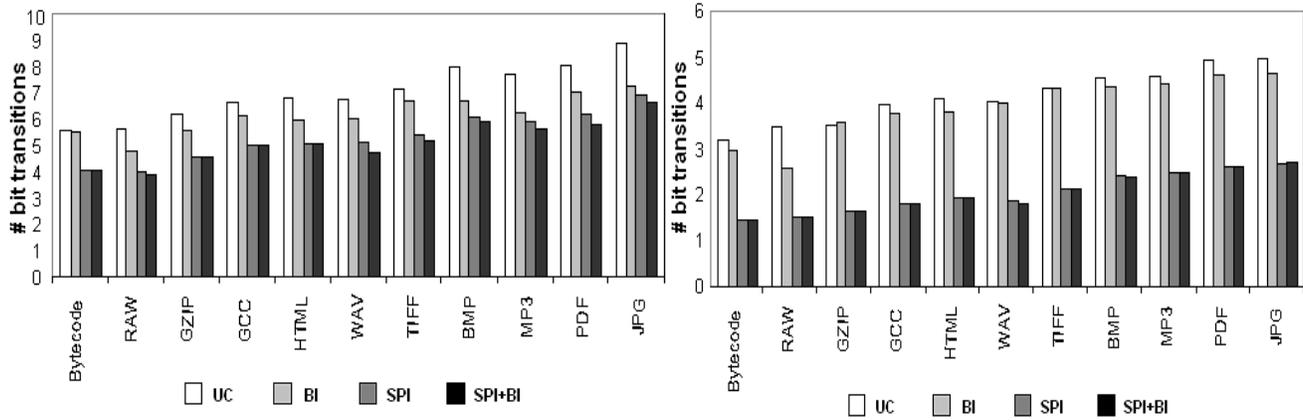
(b) SPI+BI improvement over uncoded

Figure 5. Percentage of reduction in the number of bit transitions between consecutive flits relative to uncoded flit transmission in the simulated MP3 benchmark. For 8b links, SPI outperforms BI starting from two VCs. For 16b and 32b links, SPI outperforms BI starting from three VCs.

V. POWER, AREA AND LATENCY OF A VLSI IMPLEMENTATION

Having shown that SPI effectively reduces the number of bit transitions, we proceed to examine the impact this has on power reduction and whether this reduction justifies the overhead of implementing SPI. We have implemented SPI using verilog HDL, synthesized by Synopsis Design Compiler using the Tower Semiconductor 0.18 μ process library and placed and routed by Cadence Encounter EDA tools. For the

NoC link, we assume a wire length of 3mm, and derive its other parameters from the "global interconnect" setting described by the PTM models in [32]: width of 0.8 μ , spacing of 0.8 μ , thickness of 1.25 μ , height of 0.65 μ and k of 3.5. Input data is assumed to be random, and all virtual channels are assumed to be fully utilized. We assume the link width includes flit width and VC identification. SPI is implemented for two or more virtual channels. SPI hardware architecture is shown in Fig. 7.



(a) Benchmark simulations: 2 VCs, 16b links.

(b) Benchmark simulations: 8 VCs, 8b links.

Figure 6. Simulated values of the number of bit transitions between successive flits. Link width includes flit size and VC identification. We observe 22%-57% reduction in the number of bit transitions comparing the uncoded case. SPI's advantage increases with the number of VCs.

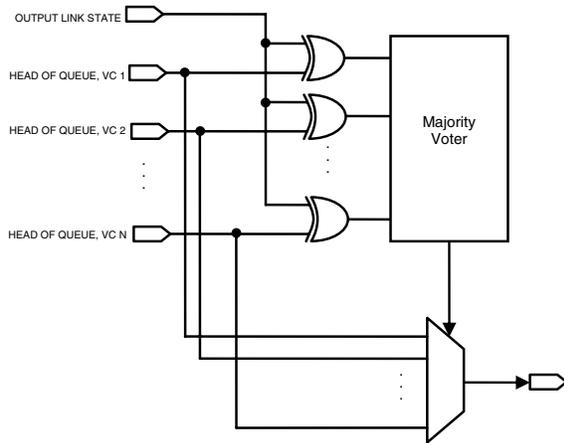


Figure 7. The circuit implementation of SPI.

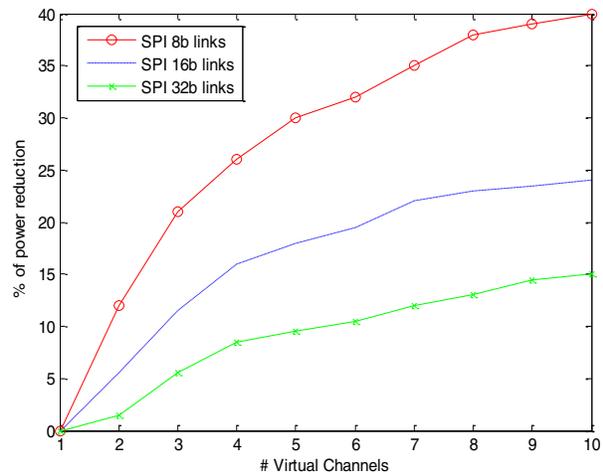


Figure 8. SPI power reduction. For 8b width links and four VCs, we observe more than 25% power reduction. For 16b width links and four VCs, we observe 15% power reduction and for 32b width links and four VCs, the reduction is about 10%.

Measurements of the dynamic power dissipated on the link, along with the power consumed by the SPI module, are jointly referred to as power consumption. The percentage of reduction in power consumption relative to uncoded flit transmission is shown in Fig. 8. As with bit transitions, we observe increasing reduction in power consumption with the growing number of VCs. For example, with 8b width links and four VCs, we observe a power reduction of more than 25%. With 16b width links and four VCs, we observe 15% power reduction, and with 32b width links and four VCs, a power reduction of about 10%. These findings are consistent with the reductions in bit transitions found in the benchmark simulations described in

Section 4. The latency of the SPI module is about 1ns in 0.18 μ technology, while the router latency at 200Mhz operation is about 1-2 clock cycles and is unaffected by the SPI module. The SPI circuit's area is 155 μ m², 0.0001% of a typical 144mm² die.

Note that the percentage of power reduction is smaller than the reduction in the number of bit transitions due to the power overhead of implementing SPI. Nevertheless, in all cases, SPI is cost-effective, and saves more power than it consumes.

VI. CONCLUSIONS

Modern integrated circuits introduce low power design challenges. The lion's share of power consumption lies with the interconnect switching activity, and this share is expected to grow in years to come [1, 10, 23, 24, 28, 31]. Data encoding is commonly used to reduce the switching activity over interconnects, but it expends additional power in redundant circuits, which in some cases offsets the achieved power reduction.

In this paper, we presented SPI – selective packet interleaving, a flit transmission scheme for energy efficient NoCs. SPI exploits the multiplicity of virtual channels to transmit a dynamically chosen flit so as to minimize bit transitions between consecutive flits at the same time. SPI uses simple, low-complexity circuits. We have analyzed the savings achieved by SPI, and have shown that SPI yields significant improvement in power consumption, which outweighs the cost of implementing SPI. For example, with 8b width links and 4 VCs, SPI reduces the average number of bit transitions over the link by more than 35% and reduces the power consumption by 25%. Analysis and simulations demonstrate a reduction of up to 55% in the number of bit transitions and up to 40% savings in power consumed on the link. SPI benefits grow with the number of virtual channels.

REFERENCES

- [1] A. Banerjee, R. Mullins, S. Moore, "A Power and Energy Exploration of Network-on-Chip Architectures", International Symposium on Networks-on-Chip (NOCS), pp. 163-172, 2007.
- [2] A. Berman, I. Keidar, "Low Overhead Error Detection for Networks-on-Chip", International Conference on Computer Design (ICCD), pp. 219-225, 2009.
- [3] L. Benini, A. Macii, E. Macii, M. Poncino, R. Scarsi. "Architecture and Synthesis Algorithms for Power-Efficient Bus Interfaces", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (ICCAD), Vol. 19, pp. 969-980, Sept. 2000.
- [4] L. Benini, G. De Micheli, "Powering Networks on Chips", 14th international symposium on Systems synthesis (ISSS) pp.33-38, Oct. 2001.
- [5] V. Chandra, A. Xu, H. Schmit, "A Low Power Approach to System Level Pipelined Interconnect Design", International Workshop on System-Level Interconnect Prediction (SLIP), Feb. 2004.
- [6] W.J. Dally and B. Towles, "Route Packets, Not Wires: On-Chip Interconnection Networks, Proceedings of Design Automation Conference (DAC), pp. 684-689, 2001.
- [7] S. Devadas and S. Malik, "A Survey of Optimization Techniques Targeting Low Power VLSI Circuits", Proceedings of the 32nd annual ACM/IEEE Design Automation Conference (DAC), pp. 242-247, 1995.
- [8] R. Dobkin, A. Morgenshtein, A. Kolodny, R. Ginosar, "Parallel vs. Serial On-Chip Communication", International Workshop on System-Level Interconnect Prediction (SLIP), April. 2008.
- [9] R. Golshan, B. Haroun, "A novel reduced swing CMOS BUS interface circuit for high speed low power VLSI systems", IEEE International symposium on circuits and systems (ISCAS), pp. 351-354, Jun. 1994.
- [10] P. Grosse1, Y. Durand and P. Feautrier, "Power Modeling of a NoC Based Design for High Speed Telecommunication Systems", Proceedings of the 16th international workshop on Integrated Circuit and System Design. Power and Timing Modeling, Optimization and Simulation (PATMOS) Sept. 2006.
- [11] J. Hu, R. Marculescu, "Exploiting the Routing Flexibility for Energy/Performance Aware Mapping of Regular NoC Architectures", Proceedings of Design, Automation and Test in Europe Conference and Exhibition (DATE), Feb. 2004.
- [12] C. Jose et-al. , "Adaptive Coding in Networks-on-Chip: Transition Activity Reduction Versus Power Overhead of the Codec Circuitry, Proceedings of the 16th international workshop on Integrated Circuit and System Design. Power and Timing Modeling, Optimization and Simulation (PATMOS) Sept. 2006.
- [13] W. Kim, J. Kim, S. Min, "A Dynamic Voltage Scaling Algorithm for Dynamic-Priority Hard Real-Time Systems Using Slack Time Analysis", Design, Automation and Test in Europe Conference and Exhibition (DATE) 2002.
- [14] R. Krishnan, J. Gyvez, H. Veendrick, "Encoded-Low Swing Technique for Ultra Low Power Interconnect", Field Programmable Logic and Applications, pp. 240-251, Springer Publishers, 2003.
- [15] S. Kumar, A. Jantsch, J.P. Soininen, M. Forsell, M. Millberg, J. Berg, K. Tiensy and A. Hemani, "A Network on Chip Architecture and Design Methodology", IEEE Computer Society Annual Symposium on VLSI (ISVLSI), pp. 105-112, Apr. 2002.
- [16] K. Lee, S.J. Lee and H.J. Yoo, "Low-Power Network-on-Chip for High-Performance SoC Design", IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 14, pp. 148-160, Feb. 2006.
- [17] E. Macii, "Ultra low power electronics and design", Kluwer Academic Publishers, Chp. 12, pp. 214-233, 2004.
- [18] N. Magen, A. Kolodny, U. Weiser and N. Shamir, "Interconnect-Power Dissipation in a Microprocessor", Proceedings of the 6th International Workshop on System-Level Interconnect Prediction (SLIP), Feb. 2004.
- [19] C. Marcon1, N. Calazans, F. Moraes, A. Susin, I. Reis and F. Hessel, "Exploring NoC Mapping Strategies: An Energy and Timing Aware Technique", Proceedings of Design, Automation and Test in Europe Conference (DATE), pp. 502-507, 2005.
- [20] H. Mehta, R. Owens, M. J. Irwin. "Some Issues in Gray Code Addressing", GLS-VLSI-96, pp. 178-180, Mar. 1996.
- [21] F. Moraes, N. Calazans, A. Mello, L. Möller and L. Ost, "HERMES: an infrastructure for low area overhead packet-switching networks on chip", Integration, the VLSI Journal, Vol. 38, pp. 69-93, Oct. 2004.
- [22] J. Palma, L. Indrusiak, F.G. Moraes, A. Garcia Ortiz, M. Glesner, R. Reis, "Inserting Data Encoding Techniques into NoC-Based Systems" IEEE Computer Society Annual Symposium on VLSI (ISVLSI), pp. 299-304, March 2007.
- [23] V. Raghunathan, M.B. Srivastava and R.K. Gupta, "A survey of techniques for energy efficient on-chip communication", Proceedings of Design Automation Conference (DAC), pp. 900-905, 2003.
- [24] P. Ramos, A. Oliveira, "Low Overhead Encodings For Reduced Activity in Data And Address Buses", IEEE International Symposium on Signals, Circuits and Systems, 1999.
- [25] S. Ramprasad, N.R. Shanbhag and I.N. Hajj, "A Coding Framework for Low-Power Address and Data Busses", IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 7, pp. 212-221, Jun. 1999.
- [26] E. Salminen, A. Kulmala and T.D. Hamalainen, "Survey of Network-on-chip proposals", white paper, OCP-IP, March 2008.
- [27] M. Stan and W. Burleson, "Bus-Invert Coding for Low-Power I/O", IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Vol. 3, pp. 49-58, 1995.
- [28] K. Sundaresan and N.R. Mahapatra, "Accurate Energy Dissipation and Thermal Modeling for Nanometer-Scale Buses", Proceedings of the 11th Int'l Symposium on High-Performance Computer Architecture (HPCA) 2005.
- [29] H. Wang, L.S. Peh, S. Malik, "Power-driven Design of Router Microarchitectures in On-chip Networks", 36th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), Dec. 2003.
- [30] H. Zhang, V. George, J. Rabaey, "Low-Swing On-Chip Signaling Techniques: Effectiveness and Robustness", IEEE Transactions on very large scale integration (VLSI) systems, Vol. 8, No. 3, June 2000.
- [31] L. Zhong, N.K. Jha, "Interconnect-aware High-level Synthesis for Low Power", International Conference on Computer Aided Design (ICCAD), 2002.
- [32] Arizona State University, Predictive Technology Model [Online]. Available at <http://www.eas.asu.edu/~ptm/>