

# A Multi-layer Approach to Green Computing: Designing Energy-efficient Digital Circuits and Manycore Architectures

Ajay Joshi, Chao Chen, Zafar Takhirov and Bobak Nazer

*ECE Department, Boston University,*

*Boston, MA, USA*

*Email: {joshi, chen9810, zafar, bobak}@bu.edu*

(Invited Abstract)

## I. INTRODUCTION

The performance of current and future VLSI systems is and will be highly constrained by energy making it imperative to explore design-time and run-time techniques for improving their energy-efficiency i.e. performance per unit energy. There are several efforts currently in place that are exploring techniques at device level, circuit level, architecture level, OS level and application level, to improve the energy efficiency. We present a run-time circuit-level technique that uses a feedback equalizer with Schmitt trigger circuit to improve the energy-efficiency of digital logic circuits and a run-time architecture-level technique that dynamically modulates the offered bandwidth, and hence the power consumption of a silicon-photonics network-on-chip (NoC) in a manycore system. Both techniques track the overlying application behavior to harness opportunities for lowering power dissipation at run time.

## II. ENERGY-EFFICIENT CIRCUIT DESIGN USING FEEDBACK EQUALIZER AND SCHMITT TRIGGER

Several circuit techniques involving voltage supply scaling, frequency scaling, switching capacitance reduction and threshold voltage management have been proposed to reduce the data-dependent and/or fixed energy consumption in digital CMOS circuits, while trading off performance and/or area. Of late, with the inherent unreliability in aggressively scaled CMOS devices, there is a big push towards developing circuits that trade-off reliability for low power [1]–[3]. The motivation for this approach is that not all overlying applications need 100% accuracy in the functionality of the underlying circuits.

We are exploring communications-inspired techniques for designing low-power digital circuits with error mitigation capabilities. The key idea here is that the combinational logic blocks in digital logic circuit can be treated as a communication channel, and standard error mitigation techniques like equalization can be used at the input and output flip-flops. As a case study we describe how a feedback equalizer circuit in combination with a Schmitt trigger (FEST) can be

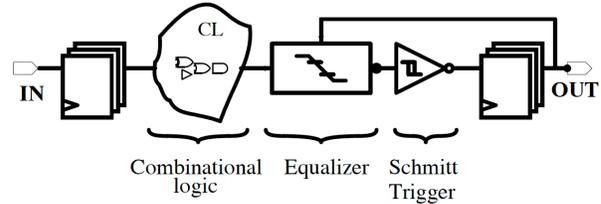


Figure 1: Combinational logic with feedback equalizer and Schmitt trigger (FEST) circuit.

used to mitigate timing errors resulting from voltage scaling in a 3-tap digital FIR filter, while maintaining performance [4]. The FEST circuit can be designed to achieve a target error rate that can be tolerated by the overlying application.

Figure 1 shows the use of the FEST circuit in between the combinational logic block output and the succeeding flip-flops. The feedback circuit stage reduces the inter-symbol interference caused by voltage overscaling while the Schmitt trigger stage smooths out any glitches created by the feedback circuit. On their own, the feedback circuit and the Schmitt trigger invert their inputs so the FEST circuit is non-inverting. Overall, the FEST circuit can be viewed as a variable threshold buffer that is robust to ISI and glitches.

We designed a 500 MHz 4-bit 3-tap FIR filter using 22 nm PTM [5] and compared the nominal filter design with a filter design having only the feedback equalizer circuit and a filter design with both feedback equalizer and Schmitt trigger. In Figure 2, we have plotted the WER for each of the three designs operating at 500 MHz. For the nominal design, design with only the feedback equalizer circuit, and design with FEST circuit, the errors start to appear below 680 mV, 610 mV, and 510 mV, respectively. Figure 2 also shows the energy consumption of the 3-tap FIR filter for the three designs. To keep the WER near zero, the energy consumed by the nominal design, design with only the feedback equalizer circuit, and design with FEST circuit is 420 fJ/op, 350 fJ/op, and 250 fJ/op, respectively. Thus, the design with only the feedback equalizer circuit

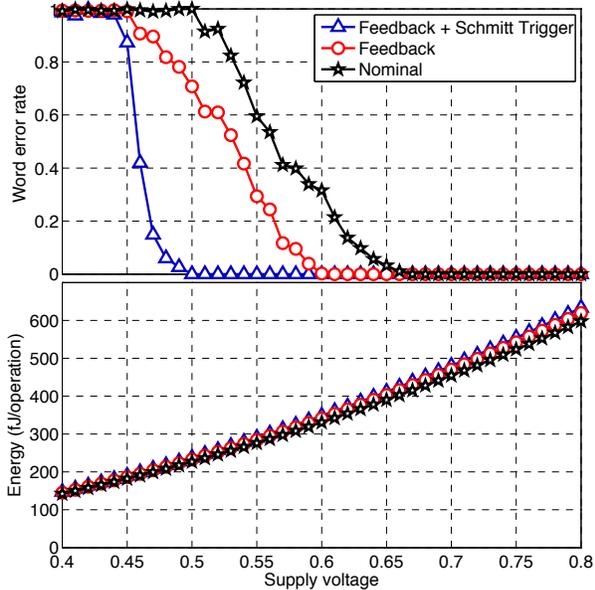


Figure 2: Word error rate and energy consumption for the nominal design, design with only the feedback circuit, and design with the FEST circuit for a 3-tap FIR filter.

and design with the FEST circuit can provide more than 16% and 40% energy savings, respectively. If the overlying application can tolerate upto 10% error rate, then compared to the nominal design with 100% accuracy we get more than 20% energy savings with only the feedback circuit and more than 50% energy savings with the FEST circuit. The area overhead for the FEST circuit is 15.8%.

### III. ENERGY-EFFICIENT SILICON-PHOTONIC MANYCORE ARCHITECTURE DESIGN

The ever-increasing core count (a core includes ALU, FPU and cache) and NoC size in manycore systems without any significant increase in the system power budget requires the development of novel techniques to improve system energy efficiency of computation and communication blocks. In particular, we target silicon-photonic networks that are expected to supplant electrical networks for intra- and inter-chip communication in manycore systems. Our focus is on improving the energy-efficiency of the NoC.

For on-chip communication, compared to the global electrical links, silicon-photonic links provide an order of magnitude higher bandwidth density and comparable data-dependent power. However, the power consumed in the laser source that powers the silicon-photonic links can more than offset these advantages [7]. Hence, it is imperative to develop techniques that will reduce

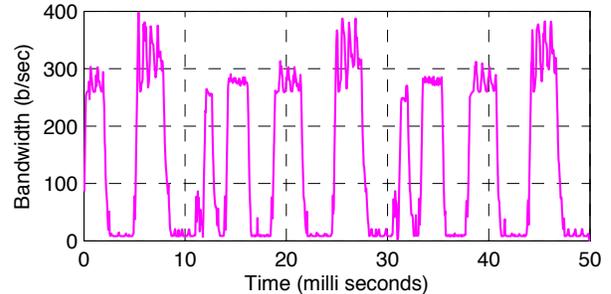


Figure 3: Temporal variations in the network bandwidth requirements for core-to-memory controller traffic for the ‘ft’ benchmark from the NAS parallel benchmark suite [6]. These temporal variations provide the opportunity to modulate the network bandwidth and save laser power.

the power consumed in the laser sources, while maintaining system performance. There are a few efforts in place that explore techniques for dynamically managing laser power. A shared optical channel in a crossbar architecture and dynamic allocation of board-to-board photonic network bandwidth have been proposed and evaluated using synthetic benchmarks and traces [8], [9]. However, a more thorough analysis that explicitly considers the temporal and spatial dynamics of many-core applications to identify opportunities for run-time laser power management is necessary. As an example, Figure 3 shows the temporal variations in the core-to-memory controller bandwidth requirements of the ‘ft’ benchmark from the NAS Parallel Benchmark Suite [6]. These variations present opportunities to reduce the provided network bandwidth and in turn reduce the total laser power.

For a multi-bus NoC architecture, we have developed a weighted time division multiplexing (TDM) technique that modulates the bandwidth of each bus depending on the temporal and spatial variations in the bandwidth requirements of the applications using those buses. The scaling of the TDM weights not only results in redistribution of the current bandwidth across all the buses but it also creates opportunities for switching OFF some laser sources, which in turn can be harnessed to save laser power.

As a case study we explored a 256-core system, with the multi-bus network providing connectivity between private L2 caches and on-chip memory controllers. We assumed an aggressive laser power switch ON/OFF time of 1  $\mu$ sec with a network reconfiguration every 50  $\mu$ sec. Figure 4 shows the performance and laser power savings for each application from the NAS Benchmark Suite running on a 32-core group in a 256-core system. The runtime management reduces the laser power costs

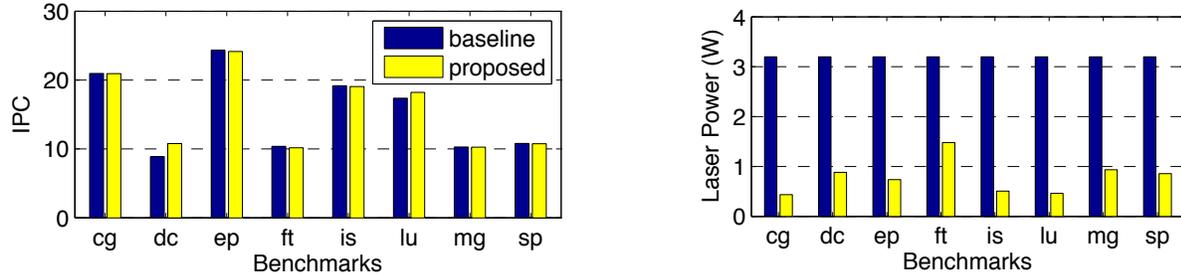


Figure 4: Impact of using run-time laser power management – Laser power management can save an average of 2.42 W of laser power per benchmark across the 8 NAS benchmarks with each benchmark running on a 32-core group in a 256-core system with a multi-bus NoC architecture. There is minimal change in the power dissipated in the cores and the other network components (transceivers and thermal tuning circuits).

by 19.3 W (more than 75%) with minimal impact on the system performance (up to 1.9%). The core power dissipation does not change significantly as there is minimal degradation in IPC. However, there is a 3.1 W power overhead for network reconfiguration. The overall system power, including cores and silicon-photonics networks can be reduced by 16.2 W.

#### IV. SUMMARY

We presented a communications-inspired circuit-level technique that can take advantage of the error tolerance of the overlying applications to reduce power dissipation, and an architecture technique that takes advantage of the spatial and temporal variations in NoC bandwidth requirements to reduce power dissipation. We believe that such an integrated approach, where we explore the opportunities for improving the energy efficiency at each level in the design hierarchy based on the constraints/specifications at other levels in the hierarchy, needs to be adopted to design highly energy-efficient VLSI systems.

#### ACKNOWLEDGMENT

Part of the work presented in this abstract was supported by the BU CoE Dean’s Catalyst Award.

#### REFERENCES

- [1] R. Abdallah and N. Shanbhag, “Error-resilient low-power Viterbi decoder architectures,” *Signal Processing, IEEE Transactions on*, vol. 57, no. 12, pp. 4906–4917, Dec. 2009.
- [2] A. Abdallah and N. Shanbhag, “Minimum-energy operation via error resiliency,” *Embedded Systems Letters, IEEE*, vol. 2, no. 4, pp. 115–118, Dec. 2010.
- [3] A. Kahng, S. Kang, R. Kumar, and J. Sartori, “Slack redistribution for graceful degradation under voltage overscaling,” in *Design Automation Conference*

- (*ASP-DAC*), 2010 15th Asia and South Pacific, Jan. 2010, pp. 825–831.
- [4] Z. Takhirov, B. Nazer, and A. Joshi, “Error mitigation in digital logic using feedback equalization with schmitt trigger (fest) circuit,” in *International Society on Quality Electronic Design Symposium*, March 2012.
- [5] (2011) Predictive technology model. [Online]. Available: <http://ptm.asu.edu/>
- [6] D. Bailey *et al.*, “The NAS parallel benchmarks,” Tech. Rep. RNR-94-007, 1994.
- [7] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, “Silicon-photonics crosstalk networks for global on-chip communication,” in *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, ser. NOCS ’09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 124–133.
- [8] Y. Pan, J. Kim, and G. Memik, “Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar,” in *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*, Jan 2010, pp. 1–12.
- [9] A. Kodi and A. Louri, “Energy-efficient and bandwidth-reconfigurable photonic networks for high-performance computing (HPC) systems,” *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, no. 2, pp. 384–395, March-April 2011.