DOE Award Number: Early CAREER Program (DE-SC0003936)

Institution: Oregon State University

Title: CAREER: Sustainable Silicon – Energy-Efficient VLSI Interconnect for Extreme-Scale Computing

PI: Patrick Yin Chiang

Date: Period covering 2010-2013
[NOTE: Accelerated spending, due to ARRA money]

THEME-1: Off-Chip VLSI Interconnect Links

Short-Range Links: Achieving Sub-1mW/Gbps Energy Efficiency

Task 1: Optimizing power consumption in global and local clock distribution to parallel links

We published numerous papers in this area over the past 3-years, working with our industry mentors Intel, LSI, and IBM. Our first work in 2010 was a 6.4Gbps serial-link RX, operating at an energy-efficiency of 0.6mW/Gbps [1a]. This was one of the 1st-papers using injection-locked oscillators (ILO) for both phase-locking and phase-shifting, and has been cited over 40 times.

In 2011/2012, we improved on this energy efficiency by 5x by lowering the supply to near-threshold operation, achieving an 8Gbps RX-data-rate with an energy-efficiency of 0.16mW/Gbps [1b,1c,1d], which is still one of the best energy-efficiencies to-date.

In 2013/2014, we increased the data-rate from 8Gbps to 16Gbps for the TX, while still achieving a sub-1mW/Gbps energy-efficiency, in the top conference in our field [1e]. Furthermore, we also demonstrated excellent energy-efficiency for the RX-based equalizer, a 0.25mW/Gbps decision-feedback equalizer (DFE), again at our top conference [1f]. Finally, we have already successfully measured a 0.6mW/Gbps 16Gbps-RX operating in the near-threshold region, and will submit that for publication in Jan-2014 [1g].

In summary, we demonstrated with our many techniques that achieving sub-1pJ/bit (1mW/Gbps) energy-efficiency for off-chip serial links is possible, both at 8Gbps and 16Gbps data-rates. Our future work in this field is to extend this level of energy-efficiency to 25Gbps serial links.

Task 2: Theoretical and practical limits to fine-grain, instantaneous power gating of on/off links

We fabricated a test-chip in May-2013, that exhibits a 0.7V, 500fJ/bit, 16Gbps transmitter for short-range interconnects. The chip incorporates fast wakeup, using distributed injection-locked oscillators to achieve nanosecond resolution wakeup. This chip has already been fabricated, was successfully tested, and has been accepted for publication at the International Solid-State Circuits Conference, 2014, our most prestigious conference [1d]. In this work, we demonstrated measured wakeup times of the TX of less than 2ns, using burst-mode injection-locking synchronization. Our future work in this area will be to understand how to use this same fast-wakeup mechanism of injection-locking for the clock/data-recovery system in the RX.

Long-Range Links: Energy-Efficient, ADC-based Links

Task 3: Reducing jitter limitations of gigahertz ADCs by using front-end, sine-wave sampling

In 2011, we demonstrated the improvement in SNR if we move to non-binary signaling, showing that duobinary signaling transmission is better than NRZ for a 25Gbps data-rate [2d]. However, in order to achieve multi-level signaling, the system requires an ADC-based serial link. Conventional multi-gigasample ADCs exhibit a problem for timing uncertainty, both process-induced static phase offsets and dynamic random jitter. Because the static phase offsets are a function of process variations, they can be measured during operation and calibrated away. We demonstrated this approach in [2f], showing that the residual jitter can be as little as 0.6ps, after calibration. Dynamic random jitter is more difficult to counteract, since they are a function of the oscillator phase noise and power-supply induced noise. We demonstrated in [2e] that using sine-wave sampling, we could minimize the sampling jitter of the initial front-end sample/hold, thereby improving the jitter accuracy, since a pure sine-wave clock generator will introduce the least amount of random jitter. Finally, we were able to transfer this technology into commercial use, by joint licensing this work with LSI Corporation, one of our corporate sponsors [2g].

Task 4: Energy-efficient circuit topologies for time-interleaved, successive-approximation ADCs

In 2010, we demonstrated a 1.3GSs, 6b, 6mW SAR-ADC structure in 40nm-CMOS [2a, 2b]. This is one of the fastest single-channel SAR-ADC designs, which demonstrating a competitive energy Figure-of-Merit (~200fJ/bit). A follow-up chip was fabricated in 40nm-TSMC process, operating at 14GSs, 5b, and 60mW, which we are currently submitting that design for publication [2c].

THEME-2: On-Chip VLSI Interconnect

Task 5: Fundamental and experimental limits to using low-voltage swing, parallel interconnect for core-core and crossbar links

In 2010, we demonstrated a 4-core network-on-a-chip prototype, utilizing low-voltage signaling for the core-core and crossbar links. This was one of the first demonstrations of low-swing signaling for NoC applications, demonstrating an energy-efficiency of 40fJ/bit/mm. This improves upon conventional on-chip transceiver wire energy by 4x. We extended on this work with a 4fJ/bit/mm on-chip transceiver for short-range on-chip wires, which was published in 2011 [3c]. By operating this link at the sub/near-threshold supply region, this chip demonstrated an energy-efficiency 40x better than conventional on-chip inverter-based wires.

Task 6: The design and usage of compact, low-power equalization for RC-limited wires

We submitted a paper to the IEEE Journal of Solid-State Circuits [3d], entitled: "A 0.2-1V, 1MHz-1.66GHz, 4-113fJ/bit/mm Capacitive Charge Sharing Transmitter for Short-Range On-chip Wires". This circuit demonstrated energy-scalable performance, for VDD=0.2V-1.0V, and energy scalable from 4fJ-113fJ per bit. This paper demonstrated not only some of the lowest reported figures-of-merit for energy-efficiency, but also introduced several concepts related to compact equalization. For example, embedded hysteresis enables low-power equation, using the intrinsic circuit hysteresis for equalization purposes. A patent was submitted to the USPTO, in regards to "LOW-VOLTAGE SWING CIRCUIT MODIFICATIONS", where we propose to insert low-swing transceiver circuits automatically using our proposed low-swing circuits [3e].

Other works: Due to the ARRA accelerated spending, my group pursued other research areas extremely relevant to the Exascale program, while simutaneously increasing the spending rate.

THEME-3: Silicon Photonics: Silicon photonics is an important research topic, as it enables distance-independent communications at very competitive energy-efficiencies. With our partner HP-Labs, we have demonstrated ring-resonator-based systems with CMOS/photonic hybrid interconnects [4a]. At 8Gbps data-rate, we achieved a competitive energy/bit of 1.5mW/Gbps [4b]. Furthermore, our prototype demonstrated the ability for incorporating on-chip thermal tuning compensation, which is critical for tracking resonance shifts due to temperature shifts. We are currently working on our 2nd-generation design, operating at 25Gbps x 4-channels, enabling 100Gbps bandwidth.

THEME-4: Reliability and Resiliency: Circuit reliability and resiliency are critical concerns for next-generation computing, especially as the supply voltage is pushed deep into the near-threshold regime. We began looking into trying to achieve both energy-efficiency and timing variation resiliency for parallel SIMD computation in [5a, 5b, 5d]. We demonstrated in 45nm-SOI a 10-core processor with Razor-based error detection, along with decoupling queues to enable limited asynchronicity between lanes.

Further investigation determined that one of the major problems with low-voltage operation is the detection of timing errors. In [5c], we proposed using analog-based sensing to detect timing faults versus conventional digital-based error detectors. We found analog-based sensing could increase throughput by more than 50%, due to the higher sensitivity and larger window for timing speculation.

Finally, the susceptibility to soft-errors (i.e. cosmic radiation) while under near-threshold operation is a phenomenon that has not been heavily researched. Working with our collaborators at Intel Circuits Research Labs, we have designed a circuit-testing platform that enables quick and accurate testing of soft-errors upsets. Our preliminary 65nm-CMOS chip, integrating both memory and logic structures, was tested under nuclear irradiation at Los Alamos National Labs. Experimental results showed that SEU rates increased ~5x when operating in the near-threshold regime, along with correlated increases in multi-bit upsets. These results are the first-of-their-kind to try to understand the behavior of near-threshold operation and soft-error upsets, and will be submitted for publication in Jan-2014 [5e, 5f].

THEME-1: Task1 and Task2

[1a] K. Hu, T. Jiang, J.G. Wang, F. O'Mahony, and P. Chiang, "A 0.6mW/Gbps, 6.4-7.2Gbps Serial Link Receiver Using Local, Injection-Locked Ring Oscillators in 90nm CMOS", Journal of Solid-State Circuits, April 2010.
[1b] K. Hu, T. Jiang, S. Palermo, and P. Chiang, "Low-Power 8Gb/s Near-Threshold Serial Link Receivers Using Super-Harmonic Injection Locking in 65nm CMOS", IEEE Custom Integrated Circuits Conference, San Jose, CA, Sep. 2011.
[1c] K. Hu, R. Bai, T. Jiang, C. Ma, A. Ragab, S. Palermo, and P. Chiang, "0.16-0.25pJ/bit, 8Gb/s Near-Threshold Serial Link Receiver with Super-Harmonic Injection-Locking" IEEE Journal of Solid-State Circuits, August 2012.
[1d] Y. Song, R. Bai, N. Yang, K. Hu, P. Chiang, and S. Palermo, "A 0.47-0.66pJ/bit, 4.8-8Gb/s I/O Transceiver in 65nm-CMOS", IEEE Journal of Solid-State Circuits, May 2013.
[1e] Y-H. Song, H-W. Yang, H. Li, P. Y. Chiang, S. Palermo, "An 8-to-16Gb/s 0.65-to-1.05pJ/b 2-Tap Impedance-Modulated Voltage-Mode Transmitter with Fast Power-State Transitioning in 65nm CMOS", accepted, International Solid-State Circuits Conference, Feb. 2014.
[1f] R. Bai, S. Palermo, and P. Y. Chiang, "A 0.25pJ/b 0.7V 16Gb/s 3-Tap Decision-Feedback Equalizer in 65nm CMOS", accepted, International Solid-State Circuits Conference, Feb. 2014.
[1g] L. Hao, S. Chen, L. Yang, R. Bai, W. Hu, S. Palermo, P. Chiang, “A 0.8V, 625fJ/bit, 16Gbps Injection-locked Receiver with 4+1 Rotating Digital CDR”, in submission, VLSI Circuits Symposium, Jun 2014.

THEME-1: Task3 and Task4

[2a] T. Jiang, W. Liu, C. Zhong, F. Zhong, P. Chiang, "Single-Channel, 1.25-GS/s, 6-bit, Loop-Unrolled Asynchronous SAR-ADC in 40nm-CMOS", IEEE Custom Integrated Circuits Conference, Sep. 2010.
[2b] T. Jiang, K. Hu, W. Liu, F. Zhong, C. Zhong, and P. Chiang, "A Single-Channel, 1.25-GS/s, 6-bit, 6.08-mW Asynchronous Successive-Approximation (SAR) ADC with Improved Feedback Delay in 40nm-CMOS", accepted, IEEE Journal of Solid-State Circuits, 2012.
[2c] 14.4-GS/s, 5-bit, 50mW Time-Interleaved ADC with Distributed Track-and-Hold Circuits and Sampling Instant Synchronization for ADC-Based Serial Links, in submission.
[2d] K. Hu, L. Wu, and P. Chiang, "A Comparative Study of 20-Gb/s NRZ and Duobinary Signaling Using Statistical Analysis", IEEE Transactions on VLSI Systems, May 2011
[2e] R. Bai, J. Wang, L. Xia, F. Zhang, Z. Yang, W. Hu, P. Chiang, "Sinusoidal Clock Sampling for Multi-Gigahertz ADCs", IEEE Transactions on Circuits and System-I, July 2011
[2f] Lingli Xia, Jingguang Wang, Will Beattie, Jacob Postman, and Patrick Yin Chiang, "Sub-2ps, Static Phase Error Calibration Technique Incorporating Measurement Uncertainty Cancellation for Multi-Gigahertz Time-Interleaved T/H Circuits", IEEE Transactions on Circuits and System-I, August 2011.
[2g] T. Jiang, P. Chiang, F. Zhong, “Time-interleaved track-and-hold circuit using distributed global sine-wave clock”, Patent No. 8,487,795.

THEME-2: Task5 and Task6

[3a] T. Krishna, J. Postman, L.- S. Peh, P. Chiang, "SWIFT: A SWing-reduced Interconnect For a Token-based Network-on-Chip in 90nm CMOS", International Conference on Computer Design (ICCD), Amsterdam, Netherlands, October 2010.
[3b] J. Postman, T. Krishna, C. Edmonds, L.S. Peh, and P. Y. Chiang, "SWIFT: A Low-Power Network-On-Chip implementing the Token Flow Control Router Architecture with Swing-Reduced Interconnects", accepted, IEEE Transactions on VLSI Systems, 2013.
[3c] J. Postman and P. Chiang, "Energy-Efficient Transceiver Circuits for Short-Range On-chip Interconnects", IEEE Custom Integrated Circuits Conference, San Jose, CA, Sep. 2011.
[3d] J. Postman and P. Chiang, “A 0.2-1V, 1MHz-1.66GHz, 4-113fJ/bit/mm Capacitive Charge Sharing Transmitter for Short-Range On-chip Wires”, submitted, IEEE Journal of Solid-State Circuits, 2013.
[3e] J. Postman and P. Chiang, “LOW-VOLTAGE SWING CIRCUIT MODIFICATIONS”, USPTO patent application filed.

THEME-3: Silicon Photonics

[4a] Chin-Hui Chen, Cheng Li, Rui Bai, Ayman Shafik, Marco Fiorentino, Zhen Peng, Patrick Chiang, Samuel Palermo, and Ray Beausoleil, "Hybrid Integrated DWDM Silicon Photonic Transceiver with Self-Adaptive CMOS Circuits", accepted, Optical Interconnects Conference, May 2013
[4b] C. Li, R. Bai, A. Shafik, E. Tabasy, G. Tang, C. Ma, C-H. Chen, Z. Peng, M. Fiorentino, P. Chiang, S. Palermo, "A Ring-Resonator-Based Silicon Photonics Transceiver with Bias-Based Wavelength Stabilization and Adaptive Power-Sensitivity Receiver", ISSCC, Feb. 2013.
[4c] C. Li, et. al., “Silicon Photonic Transceiver Circuits with Microring Resonator Bias-Based Wavelength Stabilization in 65-nm CMOS”, in submission, Journal of Solid-State Circuits, 2014.

THEME-4: Near-Threshold Operation and Error Resiliency

[5a] Joseph Crop, Evgeni Krimer, Nariman Moezzi-madani, Thomas Ruggeri, Robert Pawlowski, Patrick Chiang, Mattan Erez, "Error Detection and Recovery Techniques for Variability-Aware CMOS Computing: A Comprehensive Review", Special Issue on Low Power Design and Methodologies and Applications, Journal of Low Power Electronics and Applications, 2011.
[5b] Robert Pawlowski, Evgeni Krimer, Joseph Crop, Jacob Postman, Nariman Moezzi-Madani, Mattan Erez, Patrick Chiang, "A 530mV 10-Lane SIMD Processor With Variation Resiliency in 45nm SOI", Feb., ISSCC-2012.
[5c] J. Crop, R. Pawlowski, and P. Chiang, "Regaining Throughput Using Completion Detection for Error-Resilient, Near-Threshold Logic", Design Automation Conference (DAC), 2012.
[5d] E. Krimer, P. Chiang, and M. Erez, "Lane Decoupling for Improving the Timing-Error Resiliency of Wide-SIMD Architectures", International Symposium on Computer Architecture (ISCA), Portland, OR, June 2012.
[5e] J. Crop, E. Krimer, P. Chiang, M. Erez, "Replication-Free Single-Event Upset (SEU) Detection for Eliminating Silent Data Corruption in CMOS Logic", accepted, SELSE Workshop, 2013.
[5f] R. Pawlowski, J. Crop, and P. Chiang, “Characterization of radiation-induced soft errors in on-chip memory structures and logic from 0.33-1.0V in 65nm CMOS”, in submission, VLSI Circuits Symposium, 2014.