Research Statement

Radu Teodorescu, The Ohio State University
http://arch.cse.ohio-state.edu

For many decades we have taken for granted the fact that computer performance has been doubling every one or two years. This extraordinary growth created an industry that has impacted almost every aspect of our lives – from the way we work to the way we play, and how we communicate or provide healthcare. This revolution was enabled in no small part by one of the computer's core technologies: the microprocessor. Over the last fifty years, microprocessors have benefited tremendously from technology innovations that have delivered more and faster transistors with every new generation. Unfortunately, that technology has reached an impasse in recent years, as transistors have approached low-nanometer dimensions. Even though semiconductor technology can continue to deliver a Moore's Law-like rate of growth in transistor integration, the power efficiency of those transistors is increasing at a much lower rate. Moreover, these transistors are less predictable and more likely to suffer various types of failures.

These technological challenges are happening at a time when the need for energy efficient computing is greater than ever. The recent explosive growth of ultra-portable computing devices like smartphones and tablets is reshaping the consumer computing landscape. Expectations of battery life for these devices are now measured in days rather than hours. At the same time, these ultra-portable devices are rapidly becoming gaming consoles and productivity platforms with high performance demands. Meeting these performance requirements, while keeping power consumption under control, requires dramatic improvements in the energy efficiency of computation.

Research Accomplishments

Our research group is taking on the challenge of designing faster and more energy efficient microprocessors under the most adverse technological challenges our industry has ever faced. Our mission is to enable a new class of energy-conscious microprocessors that deliver the performance of supercomputers to mobile form factors, and enable environmentally-responsible growth in computing.

The principal approach used in this undertaking is a new computing paradigm generally referred to as Near-Threshold Computing (NTC). This technique relies on lowering the supply voltage of a chip to a level only slightly higher than the threshold voltage – the level at which transistors begin conducting current. Voltage is the most powerful lever for improving energy efficiency, because it impacts both dynamic and static power superlinearly. Even though NTC significantly reduces chip speed, it allows for many cores to be powered on simultaneously for the same power cost. Multi-threaded workloads that

can take advantage of the increased parallelism can run much more efficiently at NTC. Experimental data shows these applications can attain 8× to 10× higher energy efficiency at NTC compared to conventional super-threshold computing. A recent prototype of a low-voltage chip from Intel Corp. is showing very promising results. Unfortunately, Near-Threshold Computing faces multiple challenges before it can become a mainstream technology. This is because NTC is less reliable than conventional technology, requiring additional protection against failures. NTC also amplifies the effects of process (post-manufacturing) and voltage (runtime) variability.

Generally speaking, our research addresses the limitations of low-voltage technology with the goal of speeding up its successful deployment in production systems. We developed energy-efficient solutions for improving reliability of processors [3], a very powerful algorithm for error correction that can be implemented in processor caches [4], solutions for reducing the effects of process variability on the performance and efficiency of microprocessors[6, 2], novel solutions for reducing voltage variability by rethinking thread synchronization [5] and a new firmware-based system for dynamically reducing voltage margins in server processors [1].

Process Variability

Process variability is caused by the relative imprecision in the manufacturing process of chips with very large numbers of minute transistors. This imprecision makes transistor dimensions and properties somewhat uncertain and with a high degree of variability. This variability affects crucial transistor parameters such as threshold voltage and leads to heterogeneity in transistor speed and power consumption. A microprocessor, which is generally homogeneous by design, will instead behave more like a heterogeneous system with cores that have different speeds and power consumption profiles. This heterogeneity is very large at NTC, with design-identical cores having 2-3× variation in top speed. This very large and unpredictable heterogeneity can be very detrimental to the performance and energy efficiency of synchronized parallel applications such as those used in scientific computing.

We developed mechanisms [6, 2] for addressing variation-induced heterogeneity in near-threshold chips. One such solution [2] relies on a microprocessor design that provides two power supply lines to each core set at two very low but different voltages. Each core in the CMP can be dynamically assigned to either of the two power rails using a gating circuit. This allows each core to periodically switch between two different maximum frequencies, on a predetermined schedule. The schedule is different for each core and is chosen so that core frequencies average to the same value over a finite interval. This means that cores that are inherently slow are scheduled to spend more time on the high voltage rail while those that are fast will spend more time on the low voltage rail. The result is a CMP that achieves performance homogeneity from an underlying heterogeneous fabric, improving both performance and energy efficiency.

Reliability

Low-voltage operation slows down transistors and makes them more likely to behave unpredictably which can lead to computation or memory retention errors. Large memory blocks such as those used in large on-chip caches are especially vulnerable. They are optimized for density and therefore built using the smallest transistors, which are the most sensitive to low-voltage operation.

Our group has developed a new error correction technique [4] designed for near-threshold caches. The basic idea is to apply a simple error correction code repeatedly, in multiple permutations, to a data block to achieve an exponential increase in error correction capability. The proposed solution trades off correction strength for decoding time. While this iterative decoding process takes longer than traditional schemes, it is a pay-as-you-go approach (depending on the severity of the errors and the correctness requirements). This means applying more decoding iterations to lines that suffer larger numbers of errors and fewer decoding iterations to the others, while keeping the storage overhead low. This technique allows near-threshold chips to operate reliably in the presence of high error rates. It is also one of the first practical hardware implementations of an iterative error correction solution.

Voltage Variability

Another significant challenge of NTV operation is the increased sensitivity to voltage fluctuations. These fluctuations are caused by abrupt changes in power demand triggered by processor activity variation with workload. If the voltage deviates too much from its nominal value, it can lead to so-called voltage emergencies, which can cause timing and memory retention errors.

Prior work that addressed this issue has focused primarily on single-core or low core count systems. In recent work [5] we showed that, as the number of cores in future CMPs increases, the effects of chip-wide activity variation will overshadow the effects of within-core workload variability. As a result, core-local activity is less likely to cause large power fluctuations that lead to emergencies. However, chip-wide coordinated activity such as that forced by global synchronization in multithreaded applications leads to much larger power fluctuations. For instance, barrier synchronization causes blocked threads to idle with very low power consumption. When all idle threads are released from a barrier the associated jump in activity across participating cores leads to very large power spikes that can lead to voltage emergencies. Going forward, as the number of cores integrated on a chip increases, coordinated activity across multiple cores will dominate within-core activity variation. As a result, new mechanisms for eliminating voltage emergencies will be needed. To address this issue we developed a new, emergency-aware synchronization library [5] that smooths-out power spikes and eliminates synchronization-related voltage emergencies. The solution allowed a reduction in voltage guardbands that lowered energy consumption by over 30%.

Dynamic Self-Tuning of Voltage Margins

High variability in transistor parameters coupled with increased sensitivity to fluctuations in the supply voltage are forcing more conservative design choices. To ensure correct execution, high guardbands are added to the supply voltage to account for the worst case behavior of a chip. These guardbands need to consider various factors, such as temperature, circuit aging, process variability and workload characteristics. As a result, V_{dd} margins are often unnecessarily conservative and lead to wasted energy.

To address this inefficiency, we developed a new mechanism [1] for dynamically reducing voltage margins and lowering supply voltage while maintaining the chip operating frequency constant – a technique generally refered-to as "voltage speculation." Unlike previous approaches that rely on dedicated hardware to avoid or recover from timing violations, our solution does not require additional hardware. Instead, it maintains low, but safe operating voltage margins by leveraging error correction (ECC) support already available in modern processors. A key observation made in this work is that, as supply voltage is lowered, correctable errors in ECC-protected functional units manifest before uncorrectable errors or data corruption. Starting from this observation, we implement a voltage speculation system that uses the rate and type of runtime correctable errors as indicators for finding the lowest safe voltage point at which each core can operate. The solution adapts to on-chip core-to-core variability by tailoring the supply voltage to each core's safe operating level. It also adapts to variability in workload vulnerability to low voltage execution.

The system was prototyped on an HP Integrity Server that uses Intel's Itanium 9560 processors. A control system that dynamically adjusts the V_{dd} of four core pairs in the 8-core chip was implemented in System Firmware. The control system monitors the rate of correctable errors posted by the hardware and makes voltage assignment decisions. The appropriate voltage level is selected to keep the processor operating close, but still above the safety margin to ensure correct operation. Measurements performed on real hardware demonstrated power savings ranging from 18% to 23%, with minimal performance impact.

Technology Transfer and Industrial Partnerships

The applied nature of this research benefits tremendously from close collaboration with industry partners interested in NTC such as Intel and Mentor Graphics. As part of an GOALI (Grant Opportunities for Academic Liaison with Industry) project funded by the National Science Foundation our team is working with Mentor Graphics researchers to design and enhance the capabilities of modeling and CAD tools to accurately support NTC development.

Our research is contributing to and is supported in part by the Defense Advanced Research Projects Agency (DARPA) through the Power Efficiency Revolution for Embedded Computing Technologies (PERFECT) program, an ambitious 5.5 years project that seeks to build an embedded computing system for military applications with a computation efficiency of 75 gigaflops/watt. By comparison, a typical system today achieves a com-

putation efficiency of about 1 gigaflop/watt, almost two orders of magnitude lower. The driving technology behind this embedded system is Near-Threshold Computing. Our team together with teams from University of Illinois and University of Wisconsin are responsible for designing circuits, architecture and runtime components that mitigate and tolerate parameter variations at Near Threshold.

Long-term Research Vision

Fundamental changes are happening in the way we design, program and use microprocessors. These changes have far-reaching effects throughout our industry. In this context, I have the following long-term research objectives: (1) Invent technologies that improve the energy efficiency of computation by an order of magnitude. (2) Develop novel, integrated software, architectural and circuit design solutions for addressing the reliability and variability challenges brought by transistor scaling into low nanometer sizes. (3) Broaden the Computer Architecture fields understanding of the manufacturing and process technologys impact on microprocessor design.

References

[1] Anys Bacha and Radu Teodorescu. Dynamic reduction of voltage margins by leveraging on-chip ECC in Itanium II processors. In *International Symposium on Computer Architecture (ISCA)*, pages 1–11, June 2013.

[2] Timothy N. Miller, Xiang Pan, Renji Thomas, Naser Sedaghati, and Radu Teodorescu. Booster: Reactive core acceleration for mitigating the effects of process variation and application imbalance in low-voltage chips. In *International Symposium on High-Performance Computer Architecture (HPCA)*, pages 27–38, February 2012.

[3] Timothy N. Miller, Nagarjuna Surapaneni, and Radu Teodorescu. Flexible error protection for energy efficient reliable architectures. In *International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pages 1–8, October 2010.

[4] Timothy N. Miller, Renji Thomas, James Dinan, Bruce Adcock, and Radu Teodorescu. Parichute: Generalized turbocode-based error correction for near-threshold caches. In *International Symposium on Microarchitecture (MICRO)*, pages 351–362, December 2010.

[5] Timothy N. Miller, Renji Thomas, Xiang Pan, and Radu Teodorescu. VRSync: Characterizing and eliminating synchronization-induced voltage emergencies in many-core processors. In *International Symposium on Computer Architecture (ISCA)*, pages 249–260, June 2012.

[6] Timothy N. Miller, Renji Thomas, and Radu Teodorescu. Mitigating the effects of process variation in ultra-low voltage chip multiprocessors using dual supply voltages and half-speed stages. *IEEE Computer Architecture Letters*, 11(2), 2012.