Ultra-Fine Grain Power Management at Datapath-Level: Fact or Fiction

Mehmet E. Belviranli
University of California, Riverside
belviram@cs.ucr.edu

Weize Yu
University of South Florida
weizeyu@mail.usf.edu

Selçuk Köse
University of South Florida
kose@usf.edu

1. INTRODUCTION

The continuous struggle to consume less power while increasing the system performance propels the research community to experiment with circuit, architecture, and system level innovations. The recent advancements in the semiconductor technology permit multi-billion transistors on a single die, paving the way to significantly increasing functionality and hence the performance. To compensate for the severe power requirements, a naive solution is to not use high performance processors for low-intensity tasks and instead use low power replacement processors.

Intel's QuickAssist [1] technology allows servers to employ both high-performance Xeon and low-power Atom processors on the same front-side-bus (FSB) for interchangeable use based on the application demand. Similarly, ARM has recently proposed a big.LITTLE [2] architecture which uses two multi-core processors – the bigger and more powerful *big* Cortex-A15 processor is active during demanding tasks whereas the smaller and more power-efficient *LITTLE* Cortex-A7 processor is used when there is no need for high performance. In both cases, in spite of the significant differences between the performance and power-efficiency of the coupled processors, ISAs (*i.e.*, x64 and armv7) remain the same for the coupled processors, enabling applications with diverse requirements to be run on the same platform.

The primary difference between the coupled processors is the micro-architecture. For example, LITTLE processor is an in-order, non-symmetric dual issue processor with 8 to 10 pipeline stages whereas *big* processor is an an out-of-order sustained triple issue processor with 15 to 24 pipeline stages. The major cost for such architectures is the added area and hardware resources plus the additional scheduling complexity to be introduced by the O/S.

If we have the ability to temporarily deactivate the unused elements in the datapath of the big processor at runtime, the power consumption of the big processor can be reduced significantly. The ability to adaptively deactivate certain system components would eliminate the need for the LITTLE processor and the *big* processor can behave like the LITTLE processor when the utilization is low, thereby saving significant chip area while reducing the cost of the system. Moreover, transparent handling of the power-mode switching by the chip will relieve the O/S from the burden of load-driven task scheduling.

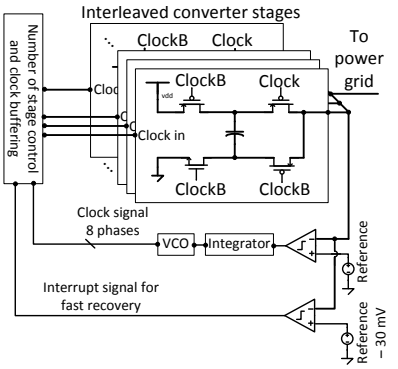

Figure 1: Proposed on-chip power voltage converter architecture.

2. ULTRA-FINE GRAIN POWER MANAGEMENT

Historically, one or two large board or package level off-chip voltage converters down-convert the 12-V power supply voltage to a 3-V to 1-V. This voltage is then delivered to on-chip circuitry via dedicated power/ground pins and off-chip interconnection network. Having the voltage regulator off-chip has several disadvantages such as slower transient response time, higher power/ground noise, huge number of dedicated I/O pins (*i.e.*, ∼60% are dedicated to power and ground), and significantly slower voltage scaling [3].

Up until recent advancements in semiconductor technology and circuit-level innovations, fully integrated distributed on-chip voltage regulation had been infeasible. SInce then, there is a considerable amount of on-going work to improve the stability, speed, and power efficiency of on-chip distributed power delivery systems. Intel Haswell processor has 13 different multi-phase buck type fully on-chip voltage converters to increase the power efficiency over 90%. The recent IBM POWER8 processor has 64 distributed LDO type voltage regulators to provide over 10 A current. The motivation for distributed on-chip power delivery is primarily to i) reduce the I/O pin count dedicated to power/ground, ii) minimize the transient response time, iii) increase the power conversion efficiency, and iv) enable faster DVS. Although on-chip power conversion has enabled core-level DVS, to the best of authors' knowledge, there is no work on datapath-

This research is supported in part by the National Science Foundation CAREER Award under contract No. CCF-1350451.

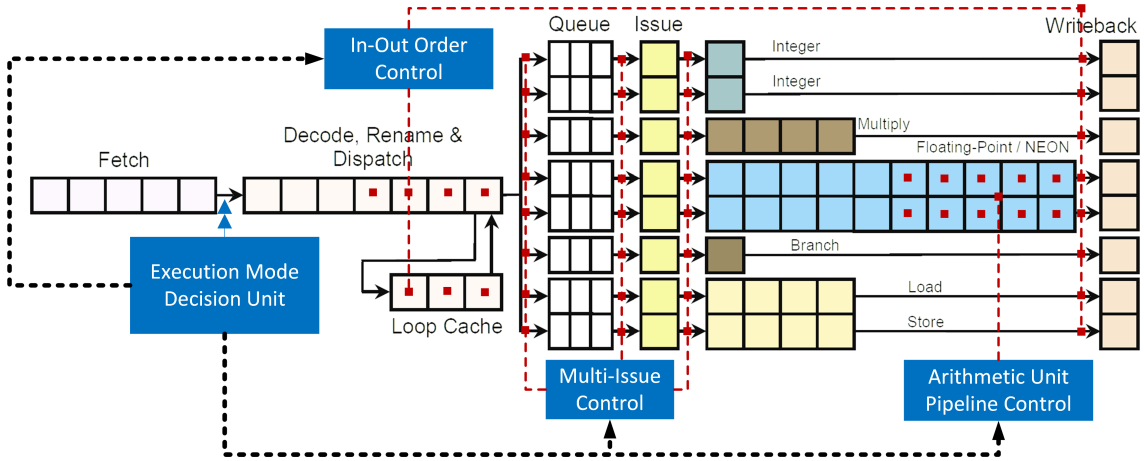

Figure 2: Proposed control units (in blue) for fine granular datapath component gating on the Cortex A-15 pipeline [2] (Base image is courtesy of ARM). Bold dashed lines (in black) indicate signaling between the execution mode decision unit (EMDU) and other component level control units. Red squares in front of each component illustrate regulated-gates and light dashed lines (in red) are the links between gates and corresponding control units.

level power management that would further improve power savings [4]. In this work, we propose to exploit on-chip multi-phase switched capacitor (SC) on-chip voltage converters to enable datapath-level power gating and DVFS to potentially boost the power efficiency of *big* processors under low intensity tasks. SC voltage converters utilize a flying capacitor which is charged by the input voltage and discharged to the load circuit at a frequency high enough to generate a DC output with a small ripple voltage. The power density of SC converters can be significantly increased to over 4.6W/mm^2 with the deep trench capacitors[1]. This power density surpasses the power density of distributed LDO regulators in POWER8 processor which is ∼2W/mm^2.

We will exploit our recent work on adaptive on-chip power management technique, converter-gating, (shown in Fig. 1) to turn-off the voltage converter stages that provide power to certain components in a datapath based on the workload.

3. ARCHITECTURAL IMPLICATIONS

Integrating per-component voltage converters and shut-off gates into datapaths of existing high performance processors requires several architectural design considerations. Figure2 depicts the components of the proposed pipeline modifications based on the Cortex-A15 architecture given in [2]:

Execution mode decision: Chip-wide power-mode detection and control mechanisms are the core components of our pipeline design. An execution mode determination unit (EMDU) will detect frequency patterns for each instruction type as well as the total instruction issue rate and will decide multi-level power mode for each of the control types described below. EMDU will send signals to proper converter stages based on the current level while considering the switching delay. The performance-mode may be implemented either as per-component basis or globally, depending on complexity of the overall design. EMDU will also support manual mode-set via a new instruction, hence enabling SW management of the power level.

Multiple issue: Dynamically changing the instruction count for multiple issues involves gating extra arithmetic units, register file and memory unit access ports, muxes and data buses.

Functional units: High-performance cores embed longer pipelines for ALUs and especially for FPUs to exploit a finer grain pipeline parallelism. A hybrid dual-design for such arithmetic unit pipelines may employ both finer and coarser pipelined units and may selectively prefer corresponding component based on the global execution mode flag.

In-order(I-O)/out-order(O-O) execution: The ability to switch from higher IPC O-O to I-O instruction execution requires hybrid design where O-O execution components such as re-order buffers and per-unit instruction queues can be selectively disabled.

Switching delay: The activation/deactivation for each converter stage takes ∼6-8 ns assuming a 150 MHz switching frequency for the SC voltage converter. For a processor running at 1GHz, the switching delay is estimated to be ∼6-8 clock cycles.

[1]IBM uses trench capacitors to increase the power density of the distributed LDO regulators in POWER8 processor.

4. FUTURE WORK AND CONCLUSION

Although there are potentially lots of challenges at the circuit and system level implementation, the proposed datapath level power management technique may significantly reduce the power consumption without significant area overhead under high intensity tasks.

5. REFERENCES

[1] Ian McCallum, “Intel quickassist technology accelerator abstraction layer (aal),” *Intel Corporation*, 2007.

[2] Peter Greenhalgh, “Big. little processing with arm cortex-a15 & cortex-a7,” *ARM White paper*, 2011.

[3] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, “System level analysis of fast, per-core dvfs using on-chip switching regulators,” February 2008, pp. 123–134.

[4] G. Yan, Y. Li, Y. Han, X. Li, M. Guo, and X. Liang, “Agileregulator: A hybrid voltage regulator scheme redeeming dark silicon for power efficiency in a multicore architecture,” February 2012, pp. 1–12.