

# A Self-Aware Processor SoC using Energy Monitors Integrated into Power Converters for Self-Adaptation

Yildiz Sinangil<sup>1</sup>, Sabrina M. Neuman<sup>1</sup>, Mahmut E. Sinangil<sup>2</sup>, Nathan Ickes<sup>1</sup>, George Bezerra<sup>1</sup>, Eric Lau<sup>1</sup>, Jason E. Miller<sup>1</sup>, Henry C. Hoffmann<sup>3</sup>, Sridi Devadas<sup>1</sup> and Anantha P. Chandraksan<sup>1</sup>

<sup>1</sup>Massachusetts Institute of Technology, Cambridge, MA, USA, <sup>2</sup>NVIDIA, Westford, MA, USA, <sup>3</sup>University of Chicago, IL, USA

## Abstract

This paper presents a self-aware processor with energy monitoring circuits that can measure actual energy consumption of the key blocks. The monitors are embedded into on-chip DC/DC converters and generate results within 10% of accuracy with minimal power (<0.1%) and area (<1%) overhead. Our system, which is implemented in 0.18 $\mu$ m technology, is designed to be voltage scalable from 1.8V down to 0.6V. Low-voltage SRAM operation is made possible through the use of 8T bit-cells and write-assists. The d-caches are designed to be re-configurable in associativity and size to adapt to compute- versus cache-bound phases of applications. Cache configuration is performed in < 3 clock cycles including tag invalidation. These hardware features enable a software self-aware computation engine (SEEC) to dynamically adapt the processor to meet performance and energy goals. Measurement results show that up to 8.4 $\times$  energy savings can be achieved with DVFS and self-adaptation.

## Introduction

Modern processor systems must balance multiple and often competing design goals such as maximizing performance while minimizing energy. Furthermore, they have to work optimally under dynamic operating conditions such as temperature and voltage fluctuations, process variations, aging, and with a wide variety of applications with different phases. To cope with the complexity of this problem, recent systems leverage power management engines that use modeling to improve energy-efficiency [1][2]. However, power models cannot fully represent the actual profile of a complex processor system. Absolute energy monitoring circuits are demonstrated in [3], but additional benefits can be obtained by integrating them within the DC/DC converters. Recent work illustrates an energy monitoring circuit embedded into a DC/DC converter [4]. However, it only achieves 20% accuracy and requires a calibration process. This paper presents a self-aware processor SoC with energy monitoring circuits that can measure actual energy consumption on the fly. The monitors are embedded into DC/DC converters and do not require any extra off-chip components.

## System Description

Fig. 1 shows the block diagram of the self-aware processor SoC. The design is based on a LEON3 single-core processor. Efficient power conversion for the two power domains of the system is provided by two on-chip DC/DC converters that deliver variable load voltages from 0.6V to 1.8V. One of the DC/DC converters powers the core and i-cache while the other powers the d-cache. Two energy monitors allow the system to distinguish between energy spent on computational operations versus energy spent on data storage. Performance counters are included to track dynamic performance changes. The instruction and data caches are constructed from custom-design SRAMs.

The work in [5] introduced a Self-aware Computational (SEEC) model conceptually. In this work, the processor uses




Fig. 1 Self-aware processor block diagram.




Fig. 2 System simulation running four phases of a multi-media application using 1- self-aware adaptation, 2- static configuration with race-to-idle operation.

the SEEC engine to complement hardware adaptations at the software level (Fig. 2). For a system target, the SEEC engine uses absolute energy as well as performance counter data to make decisions for voltage and frequency, and d-cache size and associativity.

Fig. 2 shows SEEC optimizing the system for a multimedia application with four distinct phases: FFT, transpose, FFT, and histogram. While meeting the same performance goal, the self-aware design achieves an almost 2 $\times$  reduction in energy compared to a design with a static configuration and conventional race-to-idle operation. For a given performance target, each phase has a different optimal configuration and our proposed system is capable of finding it.

## Energy Monitoring Circuit

Fig. 3 shows the block diagram of the DC/DC converters and embedded energy monitoring circuits. Our buck converters use a PFM mode control. To cover a large voltage range of operation with high conversion efficiency, the pulse width of the control signals for M1 and M2 can be configured using the *Config* signal, resulting into up to 3% efficiency improvement.

During *normal operation*, depending on the desired output voltage, the necessary pulses are supplied to M1 and M2 (Fig. 4). When an energy monitoring period occurs, the monitoring

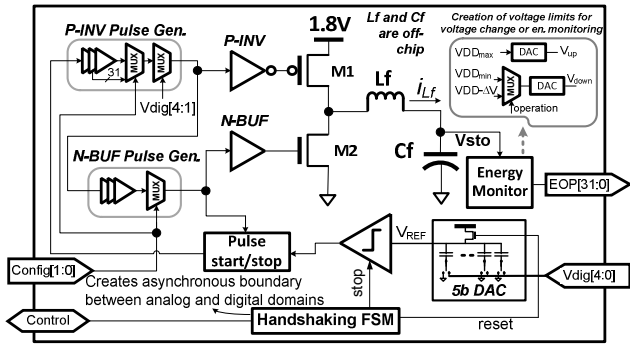


Fig. 3 DC/DC design with the embedded energy monitoring circuit.




Fig. 4 1- System demonstration during voltage change and energy monitoring cycles. 2- Energy monitoring cycles.

circuit uses a two-step process to generate energy per operation (EOP) information. Step 1 is the *discharge* phase where the M1 and M2 pulses are kept OFF. The number of clock cycles it takes for a  $\Delta V$  voltage drop across a known filtering capacitor,  $C_f$ , is observed. Then, in step 2, the voltage is restored and EOP is calculated by using  $EOP = CF \times VDD \times \Delta V / N$  assuming  $\Delta V$  is small [6].  $\Delta V$  can be set with  $\sim 50$  mV steps through an on-chip capacitive DAC to achieve high monitoring accuracy.

On the processor side, dedicated registers are used to adjust the operating voltage of the two domains and to issue energy monitoring operations. Fig. 4 shows the oscilloscope output of the system while performing energy monitoring operations and voltage changes. First, a voltage change from 1.8 V to 1.7 V is performed, then it is followed by an energy monitoring period. The system can decide to scale the operating voltage down or up, and this is communicated to the on-chip DC/DC converters. In our example, system decides to decrease the voltage to 1.6 V and performs a second energy monitoring period.

The DC/DC converter control circuits and energy monitors are designed to work with a fixed clock, whereas the core frequency is adjusted depending on the operating voltage. Hence, a four-phase handshake protocol is used between the core domain and energy monitors.

#### Adaptive Cache with Tag Invalidation

To enable operation across a large voltage range, custom SRAMs are designed using 8T bit-cells. To improve write-ability at low-voltages, peripheral row-drivers boost the word-line voltage up by  $\sim 200$  mV. Both the d-cache and i-cache are designed to be 16 kB and configured as 16 blocks as shown in Fig. 1. For each 1 kB SRAM block, there is a 128 B tag memory block. The d-cache memory is designed with dynamic associativity and size scalability. During runtime, our d-cache can be configured to be 1- to 4-way set associative and the size of each set can be configured to change from 1 kB to 4 kB. After certain reconfigurations, the tag memories need to be invalidated. To perform this quickly, the tag memories can be cleared in a single clock cycle. When asserted, a synchronous *CLR* input to these memories causes all 32 words of each block to be overwritten with '0's simultaneously (Fig. 1).

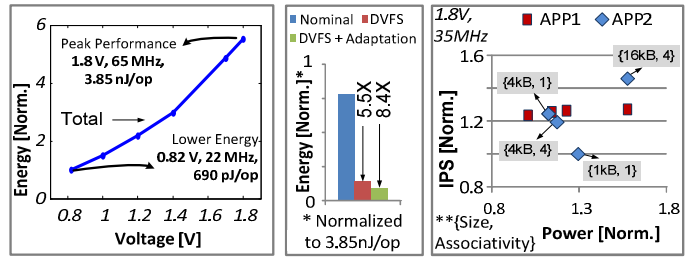


Fig. 5 Measurement results show up to 8.4× energy savings. APP1 is matrix transpose (16×16) and APP2 is matrix transpose (32×32).

#### Measurement Results

Fig. 5 shows the energy consumption of the system running a matrix transpose benchmark, APP1. Total energy per operation scales from 3.85 nJ to 690 pJ with DVFS only. Up to 8.4× lower energy is achieved through both DVFS (5.5×) and dynamic adjustments of d-cache size and associativity (1.54×) compared to operating the system at 1.8 V and full memory.

Fig. 5 also shows a scatter plot of measured power and performance trade-offs for two applications. Each point represents an operation using a different cache configuration at the same voltage (1.8 V) and clock frequency (CLK=35 MHz). The power numbers represent the total power of the core, caches and I/O power. APP1 is not cache-bound and therefore does not benefit from increasing the cache size. On the other hand, APP2 has a larger working set and benefits from increasing the cache size. For example, a 1 kB cache size results in an increase in cache misses and total power consumption rises due to the larger I/O power. For minimum energy consumption, our system can choose to work with an intermediate memory configuration (set = 4 kB, way = 1).

The test-chip specifications are summarized in Fig. 6 alongside the chip micrograph. Our system is implemented in 0.18  $\mu\text{m}$  technology for proof-of-concept and it would scale well with technology. The energy monitors impose around 1% area overhead. The proposed self-aware processor can adapt itself based on the application it is running during its operation to achieve up to 8.4× lower energy consumption.

0.18 $\mu\text{m}$ CMOS		Core DC/DC & En. Monitor	Cache DC/DC & En. Monitor	Work	[1]	[3]	[4]	[6]	This work
8T w. write assists									
SRAM	Data Memories	16 (256×32b blocks)		Adaptation	Core, L3 cache, DVFS	DVFS	no adapt.	DVFS	DVFS, d-cache
	Tag Memories	1 (32×32b blocks)							
	Adaptation	Size 1-16kB, 1-4 Way							
DC/DCs	Resolution	bit	5	VDD [V]	0.7 – 1.15	0.37 – 1.2	3.3	0.25 – 0.7	0.6 – 1.8
	Modulation	PFM			Accuracy	model based	10%	20%	N/A
En. Monitor	Input / Output VDD	V	1.8 / 0.6 – 1.8	Area overhead	N/A	16%	N/A	21%	1%
	$\Delta V$ Step Size	mV	$\sim 50$		Requirement	thermal sensor	capacit. (off-chip)	PFM	capacit. (shared)
En. Monitor	Dyn. Power Overhead	%	< 0.1	En. Savings (due to DVFS)	N/A	4.8X	no DVFS	4.1X	5.5X
					En. Savings (due to adaptation)	N/A	no adapt.	no adapt.	no adapt.

Fig. 6 Summary with micrograph of the test-chip and comparison with previous work.

#### Acknowledgements:

This work was funded by the U.S. Government under the DARPA UHPC program. The views and conclusions contained herein are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

#### References

- [1] M. Yuffe, et al., *ISSCC*, pp.264-266, Feb. 2011.
- [2] S. Damaraju, et al., *ISSCC*, pp.56-57, Feb. 2012.
- [3] Y. Sinangil and A.Chandrakasan, *A-SSCC*, pp.69-72, June 2012.
- [4] P. Dutta, et al., *IPSN*, pp.283-294, Apr. 2008.
- [5] H. Hoffmann, et al., *DAC*, pp.259-264, June, 2012.
- [6] Y. Ramadass and A.Chandrakasan, *JSSC*, pp.256-265, Jan. 2008.