

A Self-Tuning Dynamic Voltage Scaled Processor Using Delay-Error Detection and Correction

Shidhartha Das, David Roberts, Seokwoo Lee, Sanjay Pant, David Blaauw, Todd Austin, Trevor Mudge, Krisztian Flautner*
 Department of EECS, University of Michigan, Ann Arbor, MI
 *ARM Inc., Cambridge, UK

Abstract

In this paper, we present the implementation and silicon measurements results of a 64bit processor fabricated in 0.18 μ m technology. The processor employs a delay-error detection and correction scheme called Razor to eliminate voltage safety margins and scale voltage 120mV below the first failure point. It achieves 44% energy savings over the worst case operating conditions for a 0.1% targeted error rate at a fixed frequency of 120MHz.

1. Introduction

Recently, we proposed a new voltage management concept for Dynamic Voltage Scaled (DVS) processors, called Razor [1], along with initial simulation based results. In this paper, we present the first silicon implementation of a Razor design. We discuss the circuit structures used in this new implementation and present silicon measurements for 33 tested dies. The chip implements a subset of the Alpha instruction set and was fabricated with MOSIS[7] in tsmc 0.18 μ m technology.

Traditional DVS techniques [2-6] use a delay chain or a lookup table to determine the minimum voltage necessary for error-free operation at a particular frequency. Hence they require voltage margins to ensure correct operation over process variation, temperature fluctuations and voltage drop. In contrast, the Razor processor uses a delay-error tolerant flip-flop on critical paths to detect when voltage is scaled to the point of first failure for a given frequency. Voltage control is based on the observed error rate and power savings are achieved by 1) eliminating the above margins under nominal operating and silicon conditions and 2) scaling voltage 120mV below the first failure point to achieve a 0.1% targeted error rate. The total measured energy savings over the worst case was 44% at 120MHz under nominal conditions.

Figure 1a shows the delay-error tolerant Razor flip-flop in concept. The standard positive edge triggered DFF is augmented with a *shadow latch* which samples at the negative clock edge. Timing errors are detected by comparing the main flip-flop data with that of the shadow latch. An additional detector flags the occurrence of metastability at the main flip-flop output. Error signals of individual RFFs are OR-ed together to generate the pipeline *restore*




Figure 1a. Abstract View of the Razor Flip-Flop

signal which overwrites the shadow latch data into the errant flip-flop. A distributed pipeline recovery mechanism [1] is implemented to recover correct pipeline state (figure 1b). The minimum allowed supply voltage is set to ensure that the shadow latch data is guaranteed correct and can be used for error recovery. The duration of the positive clock phase, when the shadow latch is transparent, determines the sampling delay of the shadow latch. The hold time constraint imposed by the shadow latch was met by inserting delay buffers.

The rest of the paper is organized as follows. Section 2 describes the transistor level design of the Razor flip-flop and the metastability detector. Section 3 describes Razor processor implementation details and measurement results are presented in Section 4. The Razor energy savings are quantified in Section 5. We give the details on Razor voltage control in Section 6 and draw our conclusions in Section 7.

2. Razor Flip-Flop Circuit Level Implementation

Figure 2 shows the transistor level schematic of the RFF. The error comparator evaluates in the negative phase when the data latched by the slave differs from the shadow. The metastability detector which shares the dynamic node *Err_dyn* with the comparator evaluates in the positive phase of the clock when the slave output could become metastable. Thus, the RFF error signal is flagged when either evaluate. This, in turn, evaluates the dynamic gate to generate the *restore* signal by OR-ing error signals of individual RFFs. The *restore* overwrites the master with the shadow latch data such that the slave gets the correct data at the next positive edge. In this positive phase, it also disables the shadow to protect state. The *rbar_latched* signal precharges the *Err_dyn* node for the next errant cycle. Compared to a regular DFF of the same drive strength and delay, the RFF consumes 22% extra (60fJ/49fJ) energy when sampled data is static and 65% extra (205fJ/124fJ) energy when sampled data switches. However, in the processor only 207 flip-flops out of 2388 flip-flops, or 9%, could become critical and needed to be RFFs. Hence, the net Razor power overhead, including the delay buffer power for short paths, was computed to be 3% of nominal chip power.




Figure 1b. Distributed Pipeline Recovery Mechanism




Figure 2a. Razor Flip-Flop Circuit Schematic

The metastability detector consists of p- and n-skewed inverters which switch to opposite power rails under a meta-stable input voltage. The detector evaluates when input node *QS* can be ambiguously interpreted by its fan-out, inverter *G1* and the error comparator. The DC transfer curve (figure 3) of inverter *G1*, the error comparator and the metastability detector show that the “detection” band is contained well within the ambiguously interpreted voltage band. Table 1 gives the error detection and ambiguous interpretation bands for different corners.




Figure 3. DC Transfer Characteristics

Proc	Corner		Ambiguous Band	Detection Band
	VDD	TEMP		
Slow	1.2V	85C	0.57-0.60	0.53-0.64
Typ.	1.2V	40C	0.52-0.58	0.48-0.61
Fast	1.2V	27C	0.48-0.56	0.40-0.61
Slow	1.8V	85C	0.77-0.87	0.67-0.93
Typ.	1.8V	40C	0.71-0.83	0.65-0.90
Fast	1.8V	27C	0.64-0.81	0.58-0.89

Table 1. Metastability Detector Corner Analysis




Figure 2b. Restore Generation Circuitry

The probability that metastability propagates through the error detection logic and causes metastability of the *restore* signal itself was computed to be below $2e-30$. Such an event is flagged by the *fail* signal generated using double skewed flip-flops. In the rare event of a *fail*, the pipeline is flushed and the supply voltage is immediately increased. During 4 months of chip testing, this event was never detected.

3. Processor Implementation Details

The die photograph of the Razor processor and its details are shown in figure 4. To verify correct operation, the dcache/register file contents were scanned and compared with a PC emulating the same code. A 64b register records the number of errant cycles and was sampled to compute the error rate. An internal clock unit generates an asymmetric clock with a range between 60 MHz to 400 MHz in steps of 20MHz. The shadow latch sampling delay, defined by the positive clock phase, is configurable from 0ps to 3.5ns in steps of 500ps. The clock unit has a separate voltage domain that is not voltage scaled. Energy savings from Razor DVS were measured at 140 and 120MHz for 33 chips from 2 fabrication runs.




Figure 4. Die Photograph of the Chip

Technology Node	0.18 μ m
Max. Clock Frequency	140MHz
DVS Supply Voltage Range	1.2-1.8V
Total Number of Transistors	1.58million
Die Size	3.3mm*3.6mm
Measured Chip Power at 1.8V	130mW
Icache Size	8KB
Dcache Size	8KB
Total Number of Flip-Flops	2801
Total Number of Razor Flip-Flops	207
Number of Delay Buffers Added	2388
Error Free Operation (Simulation Results)	
Standard FF Energy (Static/Switching)	49fJ/124fJ
RFF Energy (Static/Switching)	60fJ/205fJ
% Total Chip Power Overhead	2.9%
Error Correction and Recovery Overhead	
Energy of a RFF per error event	260fJ

Table 2. Processor Implementation Details

4. Measurement Results

Figure 5 shows the error rates and energy gains versus supply voltage at 120 and 140 MHz for two chips. Energy at a particular voltage is normalized with respect to the energy at point of first failure. For all plotted points, correct program execution with Razor error correction was verified.

From the figure, we note that the error rate at the point of first failure is very low because only a few of the critical paths fail to meet the setup requirements. As voltage is scaled further into the sub-critical regime the error rate increases exponentially. The IPC penalty due to the error recovery cycles is negligible for error rates below 0.1%. Under such low error rates, the recovery overhead energy is also negligible and the total processor energy shows a quadratic reduction with the supply voltage. At error rates exceeding 0.1%, the recovery energy rapidly starts to dominate,




Figure 5. Error Rate and Normalized Energy Measurements for Chips 1 and 2

offsetting the quadratic savings due to voltage scaling. For the measured chips, the energy optimal error rate fell at approximately 0.1%. Table 3 shows the measured power at the point of first failure and the energy per instruction for both the chips at the point of first failure and at the point of 0.1% error rate. At 120MHz, chip 1 consumes 104.5mW at the first failure point and 89.7mW at an optimal 0.1% error rate, leading to 15% energy savings with negligible IPC hit. The energy gains for chip 2 are 18%. These gains are in addition to the energy saved by eliminating voltage margins.

120MHz 27C	Point of First Failure		Point of 0.1% Error Rate	
	Power	Energy per Instruction (Power/IPC/Freq)	Power	Energy per Instruction (Power/IPC/Freq)
Chip1	104.5mW	870pJ	89.7mW	740pJ
Chip2	119.4mW	990pJ	99.6mW	830pJ

Table 3. Error Rate and Energy/Instruction at Point of First Failure and Point of 0.1% Error Rate

Figure 6 shows the distribution of the first failure voltage for the 33 measured chips. The first failure voltage for chips 1 and 2 in figure 5 are 1.63V and 1.72V, respectively and hence represent typical and worst case process conditions. The scatter plot shows the correlation between the first failure voltage and the 0.1% error rate voltage. The relative “flatness” of the linear fit indicates less sensitivity to process variation when running at a 0.1% error rate than at point of first failure. The distribution of energy savings from running at 0.1% error rate at 120MHz and 27C is shown for all chips and ranges from 5% to 23%. The measured error rates at different operating temperatures shows 80mV shift in the point of first failure from 1.46V to 1.54V for a temperature increase from 45 to 95C as shown in figure 7.

5. Razor Energy Savings

The bar graph in figure 8 shows the energy for the two chips in figure 4 when operating at 120MHz and 45C. The first set of bars




Figure 6. Distribution of Point of First Failure, Point of 0.1% Error Rate and Normalized Energy across 33 measured chips




Figure 7. Temperature Dependence of Error Rate

shows the energy when Razor is turned off and worst-case margins are added to ensure correct operation. For chip 1, 80mV temperature margin, 130mV process margin (compared to the worst-case chip out of the 33 chips) and an estimated 180mV power supply margin (10% nominal Vdd) were added. The second and third sets of bars show the energy when operating with Razor at first point of failure and at 0.1% error rate.

Total energy gains for chip 1 (71mW, 44%) and chip 2 (63mW, 39%) are comparable because greater process margin in chip 1 (100mV greater) is compensated by increased savings for chip 2 when scaling below first failure point. The distribution of the total energy savings over the worst case at 120 and 140MHz for all 33 chips is shown in figure 9 and shows an average savings of approximately 50% for 120MHz and 45% for 140MHz.




Figure 8. Razor Energy Savings @120Mhz




Figure 9. Distribution of Net Energy Savings

6. Razor Voltage Control

Figure 10 shows the voltage controller, implemented in software that regulates the supply voltage by reacting to error rates. The controller samples the error register and adjusts the supply voltage to achieve a targeted error rate. The response of the voltage controller for a 0.5% targeted error rate for a test code with alternating high and low error rate phases is shown. The controller settles at 1.52V at high error rate phases and at 1.45V at low error rate phases.




Figure 10. Razor Voltage Controller Response

7. Conclusion

In this paper, we present a self tuning DVS processor using delay error tolerant flip-flops. We obtained 44% energy savings by eliminating voltage margins and operating at a 0.1% error rate.

References

- [1] D. Ernst, et. al., pp.7-18, *MICRO 36*.
- [2] K. Suzuki, et. al., pp. 587-590, *CICC 97*
- [3] S. Sakiyama, et. al., pp. 99-100, *Symposium on VLSI circuits, 1997*
- [4] T. Burd, et. al., pp.294-295, *ISSCC, 2000*.
- [5] S. Akui, et. al., pp.64-65, *ISSCC 2004*.
- [6] Nowka, et. al., pp.340-341, *ISSCC 2002*
- [7] www.mosis.org