15.6 A Power-Efficient 32b ARM ISA Processor Using Timing-Error Detection and Correction for Transient-Error Tolerance and Adaptation to PVT Variation

David Bull[1], Shidhartha Das[1], Karthik Shivshankar[1], Ganesh Dasika[2], Krisztian Flautner[1], David Blaauw[2]

[1]ARM, Cambridge, United Kingdom
[2]University of Michigan, Ann Arbor, MI

Razor [1-3] is a hybrid technique for dynamic detection and correction of timing errors. A combination of error-detecting circuits and micro-architectural recovery mechanisms creates a system which is robust in the face of timing errors, and can be tuned to an efficient operating point by dynamically eliminating unused guardbands.

Canary or tracking circuits [4-5] can compensate for certain manifestations of PVT variation. However they still require substantial margining to account for fast-moving or localized events, such as Ldi/dt, local IR drop, capacitive coupling, or PLL jitter. These events are often transient, and while the pathological case of all occurring simultaneously is extremely unlikely, it cannot be ruled out. A Razor system can survive both fast-moving and transient events, and adapt itself to the prevailing conditions, allowing excess margins to be reclaimed. The savings from margin reclamation can be realized either as a per device power efficiency (higher throughput same VDD, same throughput lower power), or as parametric yield improvement for a batch of devices.

Error-detection in Razor is performed by specific circuits which check for late-arriving signals. Error-correction is performed by the system using either stall mechanisms, or by instruction/transaction-replay. Measurements on a simplified Alpha pipeline [2] showed 33% energy savings. In [3], the authors evaluated error-detection circuits on a 3-stage pipeline, using artificially induced Vcc droops showing 32% throughput (TP) gain at same Vcc, or 17% Vcc reduction at equal TP.

This paper presents Razor applied to a processor with timing paths representative of an industrial design, running at frequencies over 1GHz, where fast-moving and transient timing-related events are significant. The processor implements a subset of the ARM ISA, with a micro-architecture design that has balanced pipeline stages resulting in critical memory access, and clock-gating enable paths. The design has been fabricated on a UMC [6] 65nm process, using industry standard EDA tools, with STA signoff frequency of 724MHz at the worst-case corner (0.9V/SS/125C). Silicon measurements on 63 samples, including split lots, show a 52% power reduction of the overall distribution for 1GHz operation. Error-rate driven dynamic voltage (DVS) and frequency scaling (DFS) schemes have been evaluated.

The micro-architecture is shown in Fig. 15.6.1. The pipeline is balanced using a combination of micro-architecture design and path-equalization performed by backend tools, such that all stages have similar critical-path delay. The pipeline includes forwarding and interlock logic, which contributes to both data and control critical paths, including clock gate enables, and memory access paths. Error recovery consists of flushing the pipeline and restarting execution from the next un-committed instruction. Razor stabilization stages, S0 and S1, delay instruction commit by two cycles. This allows synchronizing the potentially metastable error signal from the ME stage. Forwarding paths prevent any impact on IPC due to S0 and S1, which add 2.4% extra power overhead.

The Transition-Detector (TD) (Fig. 15.6.2), detects errors by generating a pulse in response to a transition at the D input of a flip-flop (FF) and capturing this pulse it within a window defined by a clock-pulse (CP) generated from the rising-edge. The sizing of the devices in the inverter and AND gates in the pulse-generators determines the width of the data pulse (DP). A delay on CK defines the width (T_{CK}) of the implicit CP, which is active when N1 and N2 are both on. Detection begins (ends) when the trailing (leading) edge of DP overlaps with the leading (trailing) edge of CP. The error-detection window is $T_D+T_{CK}-2T_{OV}$, where T_{OV} is the minimum overlap required. The min-delay constraint is $T_{CK}-T_{OV}$ which is less than the high clock-phase of previous designs [2-3]. The trade-off is increased pessimism, as the point at which transitions are flagged as errors is moved earlier. For 1GHz operation, this pessimism corresponds to ~5% of the cycle time, compared to when incorrect state starts to be latched.

In contrast to the RazorII FF [2] design, the TD can operate with conventional 50% duty-cycle clocks by integrating the CP generation with error-detection. Monitoring the input D, instead of the latch node, precludes the need for extra circuitry to suppress spurious error-detection for genuine transitions.

An error history (EHIST) diagnostic bit was added to each TD using an RS-latch, set whenever an error occurs. Reading out the EHIST allows identification of each TD that triggered over the course of a test. Simulation of a typical workload (WTYP) shows power overhead due to TD was 5.7% of the overall power with 1.3% overhead due to min-delay buffers.

Figure 15.6.3 shows TP and number of failing TDs versus frequency, as well as the EHIST map for WTYP at 1.1GHz and 1.2GHz. The TP linearly increases with frequency until the Point of First Failure (PoFF) at 1.1GHz, a 50% TP increase compared to the design point of 724MHz. Thereafter multiple errors occur due to the balanced nature of the pipeline and the TP degrades exponentially. Execution is correct until 1.6GHz, after which recovery fails.

DFS experiments use an on-die Adaptive Frequency Controller (AFC) which adapts to the dynamic workload variation by changing frequency in response to error-rate. Figure 15.6.4 shows the AFC response for a workload with 3 phases – a NOP loop, a combined critical path/power virus loop (PV), and typical workload (WTYP), running at a fixed 1V VDD. Highest frequency is measured in the NOP phase (1.2GHz) and the lowest in the PV phase (1GHz). In the TYP phase, there are 4 distinct frequencies (1143 - 1068MHz) due to a wider range of paths being exercised compared to the synthetic test cases.

Figure 15.6.5 shows the same 3-phase workload using an adaptive voltage controller at 1GHz frequency for 3 samples. Using Razor with the worst-case PV code on the slowest (SS6) part requires 1.17V, while WTYP requires 1.07V, which is below the 1.1V overdrive limit of the process. Considering parametric yield implications, conventional margining without Razor requires operation above 1.2V (3% VDD margin over PoFF) to achieve 100% yield at 1GHz, for reliable WC operation of SS6. This is unsustainable due to power and wear-out implications of excessive overdrive. Figure 15.6.6 shows the comparison between a baseline of 1.2V and Razor-tuned voltages. The max-power for the 1.2V distribution is due to the FF5 part, and is 52% higher than the Razor distribution, with a spread of 37mW compared to10mW.

An alternative to dynamic adaptation is to discard slower parts or reduce the max-frequency specification. As 6 out of 22 of our TT lot samples require more than 1.1V for the PV, discarding slower parts would almost certainly impact yield. Reducing the clock frequency to a point where yield was not impacted would limit the operation frequency to 800MHz. For the same distribution Razor provides potential for an effective 100% yield point at 1GHz, with supply voltage kept at or below 1.1V for all devices, except for extremely rare use cases equivalent to the pathological WC PV code. The die photograph and implementation details are shown in Fig. 15.6.7.

Acknowledgements:
We would like to thank staff at UMC (United Microelectronics Corporation) for providing, integrating and fabricating the silicon, as well as David Flynn, Sachin Idgunji and John Biggs at ARM for developing the "Ulterior" technology demonstrator chip that hosts the Razor subsystem.

References:

[1] S. Das, D. Roberts, S. Lee, S. Pant, et al., "A Self-Tuning DVS Processor Using Delay-Error Detection and Correction", *IEEE J. Solid-State Circuits*, vol. 41, pp.792-804, Apr. 2006.
[2] D. Blaauw, S. Kalaiselvan, K. Lai, et al., "RazorII: In situ Error Detection and Correction for PVT and SER Tolerance", *ISSCC Dig. Tech. Papers*, pp. 292-293, Feb. 2008.
[3] K. Bowman, J. Tschanz, N. S. Kim, et al., "Energy-Efficient and Metastability-Immune Timing-Error Detection and Instruction Replay-Based Recovery Circuits for Dynamic Variation Tolerance", *ISSCC Dig. Tech. Papers*, pp.402-403, Feb. 2008.
[4] A. Drake, R. Senger, H. Deogun, et al., "A Distributed Critical-Path Timing Monitor for a 65nm High-Performance Microprocessor", *ISSCC Dig. Tech. Papers*, Feb. 2007.
[5] J. Tschanz, N. S. Kim, S. Dighe, et al., "Adaptive Frequency and Biasing Techniques for Tolerance to Dynamic Temperature-Voltage Variations and Aging", *ISSCC Dig. Tech. Papers*, pp. 292-293, Feb. 2007.
[6] UMC, United Microelectronics Corporation, http://www.umc.com/

 978-1-4244-6034-2/10/$26.00

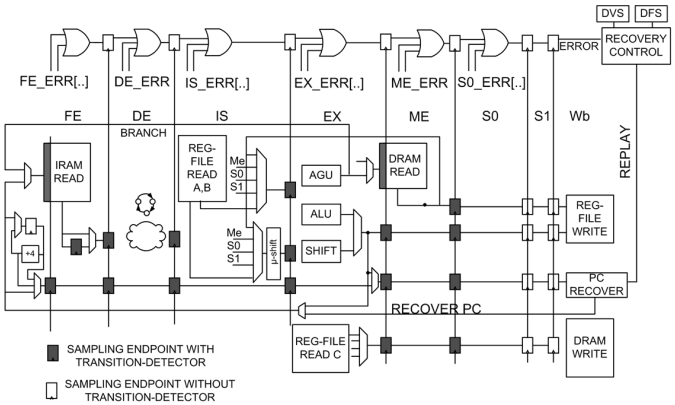

Figure 15.6.1: Pipeline diagram of the ARM ISA processor showing error-detecting TD and recovery control.

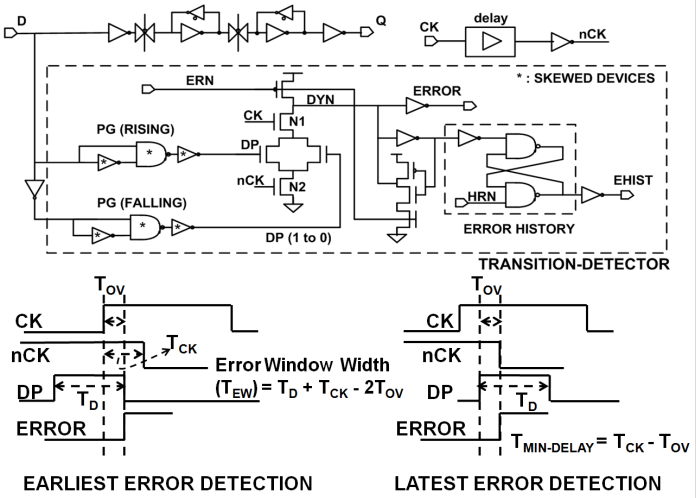

Figure 15.6.2: Transition-Detector circuit schematic and conceptual timing diagrams showing principle of operation.

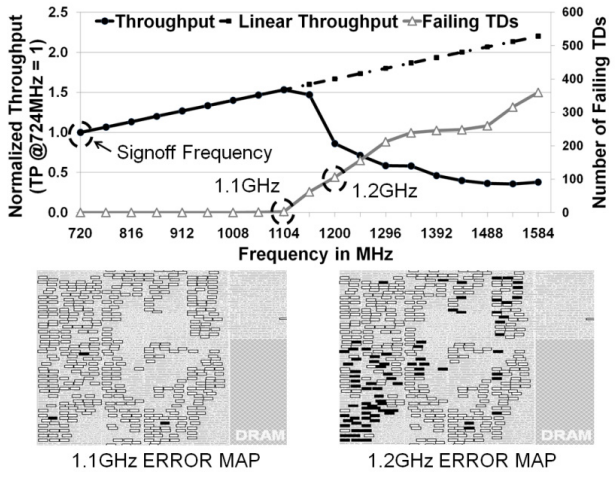

Figure 15.6.3: Measured throughput (TP) versus frequency for a typical workload (WTYP). At 1.1GHz, maximum TP gain occurs as only 3 TDs fail. At 1.2GHz, 122 TDs fail and TP degrades drastically.

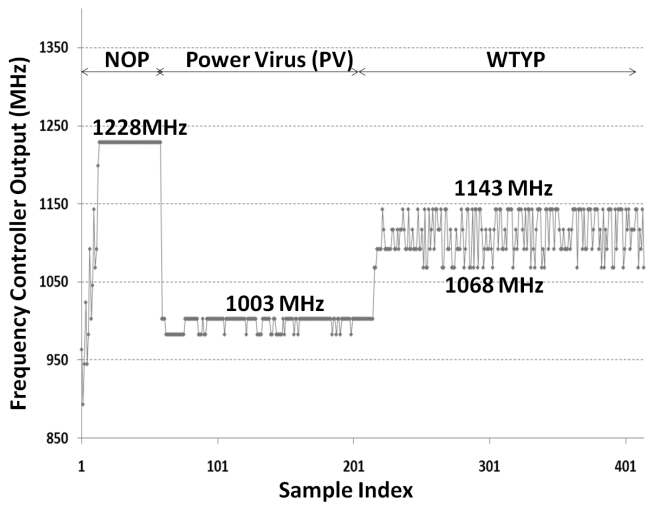

Figure 15.6.4: The Adaptive Frequency Controller response for a 3-phase code consisting of loops of NOP, Power Virus and typical workloads.

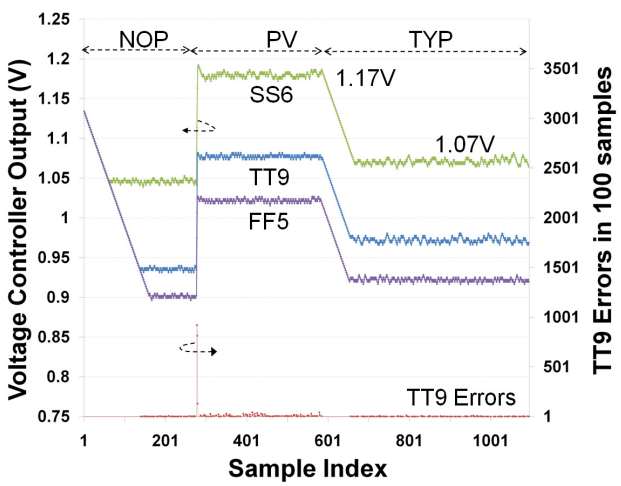

Figure 15.6.5: Dynamic Voltage Controller Response. A Proportional Controller adjusts voltage according to measured error-rates. Error-rate spike going from NOP to PV phase results in a sharp VDD increase.

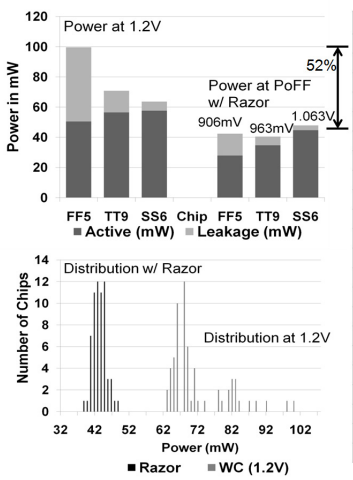

TT Chip	PV PoFF	1GHz/1.1V
TT5	1.061	Pass
TT7	1.062	Pass
TT19	1.065	Pass
TT17	1.068	Pass
TT8	1.068	Pass
TT9	1.071	Pass
TT31	1.071	Pass
TT47	1.072	Pass
TT34	1.079	Pass
TT3	1.08	Pass
TT18	1.08	Pass
TT32	1.084	Pass
TT16	1.084	Pass
TT10	1.087	Pass
TT33	1.09	Pass
TT45	1.09	Pass
TT30	1.102	Fail
TT15	1.11	Fail
TT26	1.114	Fail
TT2	1.122	Fail
TT27	1.126	Fail
TT28	1.144	Fail

Figure 15.6.6: Measured energy savings due to Razor-enabled operation for SS6, TT9 and FF5 chips. Without Razor, limiting voltage overdrive to 1.1V impacts parametric yield.

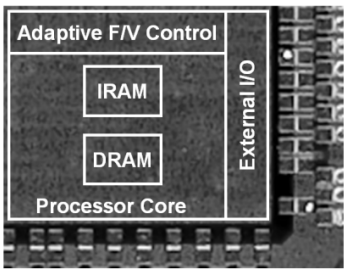

Flip-flops	2976
Flip-flops with TD	503
ICGs	149
ICGs with TD	27
TD for RAMs	20
TD Power Overhead	5.7%
Power Overhead of Min-delay buffers	1.3%
Process Technology	UMC 65SP
Nominal VDD range	1.0V-1.1V
IRAM and DRAM size	2KB
STA signoff frequency @ 0.9V/SS/125C	724MHz
Total die (FF/TT/SS)	63 (20/22/21)

Figure 15.6.7: Die Photograph and Implementation Details.

978-1-4244-6034-2/10/$26.00