

Adaptive Performance Compensation with In-Situ Timing Error Prediction for Subthreshold Circuits

Hiroshi Fuketa^{†*}, Masanori Hashimoto^{†*}, Yukio Mitsuyama^{†*}, and Takao Onoye^{†*}

[†]Dept. Information Systems Engineering, Osaka University, JAPAN

^{*}JST, CREST

{fuketa.hiroshi,hasimoto,mituyama,onoye}@ist.osaka-u.ac.jp

Abstract—This paper presents an adaptive technique for compensating manufacturing and environmental variability in subthreshold circuits using “canary Flip-Flop” that can predict timing errors. A 32-bit Kogge-Stone adder whose performance was controlled by body-biasing was fabricated in a 65 nm CMOS process. Measurement results show that the adaptive control can compensate PVT variations and improve energy-efficiency of subthreshold circuits significantly compared to worst-case design and operation with guardbanding.

I. INTRODUCTION

Subthreshold circuits are attracting attention of designers implementing ultra-low power applications, such as a sensor-node processor [1], [2]. However, there is a problem that prevents subthreshold circuits from being widely used; the performances are extremely sensitive to manufacturing and environmental variability due to an exponential dependence of subthreshold current on threshold voltage (V_{th}) and supply voltage (V_{DD}). Therefore, a traditional “worst-case” design with guardbanding is totally inefficient, and an adaptive performance control is indispensable for subthreshold circuits.

Traditionally, replica circuits have been used for performance monitoring. Reference [3] presented a circuit for adaptive control with critical path replica. However, the critical path replica is inadequate for subthreshold circuits, since the delay mismatch between the replica and the actual critical path is remarkably large due to within-die random V_{th} variation. To overcome this mismatch, in-situ techniques have been proposed [4]–[7]. “Razor I” [4] and “Razor II” [5] detect timing errors in actual paths and correct the errors. However, applying “Razor” to general sequential circuits and simple processors such as subthreshold processors is impracticable, since “Razor” exploits re-execution mechanism integrated in high-performance processors. In contrast, a self-adaptive circuit with “Canary Flip-Flop” was proposed in [6], [7]. Canary FF cannot detect the timing errors but predict them. As long as the prediction is appropriate, the circuit is adaptively-controlled well without error recovery mechanism.

This paper presents the first work that applied the adaptive speed control with canary FF to subthreshold circuits and measured it on silicon. We will demonstrate that fabricated chips operate in subthreshold region with the adaptive speed and they compensate manufacturing and environmental variability. We will also reveal that the adaptive control provides much more energy-efficient operation in comparison with the worst-case operation with guardbanding, and refer to relations

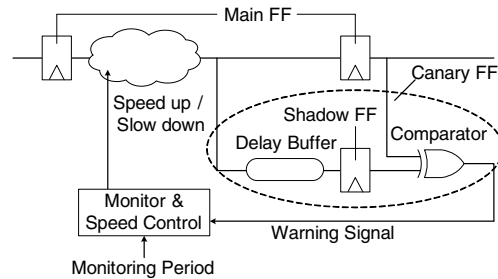


Fig. 1. Self-adaptive speed control with canary FF.

between timing error rate and design parameters, such as the location and the buffer delay of canary FF.

II. SELF-ADAPTIVE SPEED CONTROL WITH CANARY FF

A. Overview

Figure 1 shows an overview of the self-adaptive speed control with canary FF. Canary FF, which consists of a normal FF (we call it “shadow FF”), a delay buffer and a comparator, generates a warning signal to predict the occurrence of timing errors. The warning signal is monitored during a specified period. Once the warning signal is detected, the circuit is controlled to speed up. If no warning signals are generated during the monitoring period, the circuit is controlled to slow down. The circuit speed is controlled digitally and we use a term “speed level” to describe how fast or slow the circuit is controlled to be. Higher speed level means the circuit is controlled for faster operation.

When this adaptive speed control is applied to normal (non-test) operations, the occurrence of timing errors can not be completely eliminated because the circuit might be controlled to slow down excessively when the paths where canary FFs are inserted have not been activated for a long time [7]. We thus evaluate MTBF (Mean Time Between Failure) in addition to performance.

B. Circuit structure of test chip

We designed and fabricated a test circuit to demonstrate the adaptive speed control with canary FF in a 65 nm CMOS process. The structure of the test circuit is depicted in Fig. 2, and the micrograph is shown in Fig. 3. A 32-bit Kogge-Stone adder (KSA) was adopted as a circuit whose performance was controlled adaptively. S[32]–S[0] denote the outputs of the KSA, and S[32] is the most significant bit.

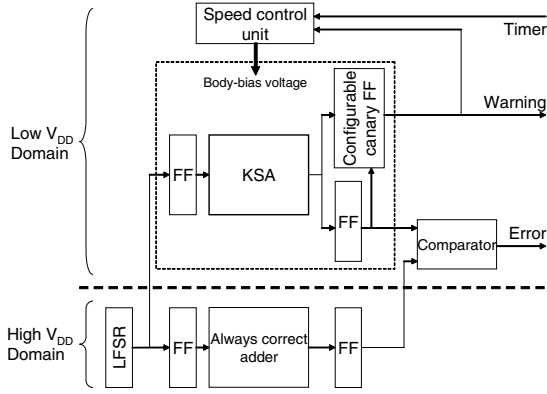


Fig. 2. Block diagram of the test circuit. 32-bit Kogge-Stone adder (KSA) is controlled adaptively with configurable canary FF.




Fig. 3. Micrograph of test chip.

Input patterns are generated by a linear feedback shift register (LFSR). The KSA outputs are compared to the answer to check if a timing error occurs. The answer is generated by “always correct” adder operating at higher supply voltage.

A timer signal is asserted when the monitoring period of the warning signal is elapsed. The monitoring period is counted off-chip. When implementing the timer circuit on-chip, ultra-low power timers such as a circuit using gate leakage currents [8] would be desirable.

The speed control unit alters by body-biasing the speed of the KSA, main FFs and canary FFs at inputs and outputs of the KSA. Figure 4 shows the schematic of the speed control unit. VPW/VNW denotes p-/n-well body-bias voltage of the KSA, main FFs, and canary FFs. Four speed levels can be provided by applying four pairs of body-bias voltage (VPW0–3 and VNW0–3), and each body-bias voltage is supplied by external DC voltage sources. VPW and VNW are selected from VPW0–3 and VNW0–3 according to the speed level stored in a two-bit register. The circuit operation starts at the maximum speed level. When the timer signal is asserted, the speed control unit immediately decrements the speed level by one and the circuit is controlled to slow down. In contrast, when the warning signal is asserted, the speed control unit increments the speed level by one.

We implemented the “configurable” canary FF such that the inserted location and the buffer delay can be configured. Figure 5 illustrates the configurable canary FF. The configurable canary FF is composed of 16 canary FFs with variable delay buffer. Each canary FF inserted at S[17] – S[32] can be enabled or disabled individually.




Fig. 4. Schematic of the speed control unit. VPW and VNW denote body-bias voltages of the KSA, main FFs, and canary FFs.

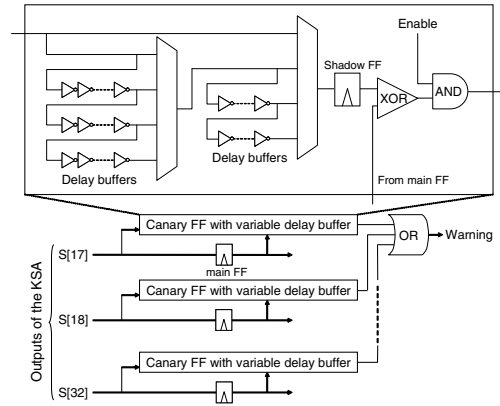


Fig. 5. Schematic of the configurable canary FF.

III. MEASUREMENT RESULTS

A. Operation example

Figure 6 shows an operation example with measured timing error, warning signals, and speed level transitions when the circuit was controlled adaptively with canary FF. The operation frequency and V_{DD} were 2 MHz and 350 mV. The step of body-biasing levels was set to 30 mV, which means speed level 1 corresponds to 30 mV forward body-bias (FBB) when speed level 0 is zero body-bias (ZBB). The speed level was altered according to the warning signal. A timing error occurred in this example.

Experiments in this paper set the monitoring period to 10^7 cycles to observe the timing errors more frequently for demonstration purpose and measurement practicability. The monitoring period in practical use should be set to longer cycles. This issue will be discussed in III-D.

B. Adaptive compensation of environmental variability

Figure 7 shows the power dissipation at the various temperature conditions (25–70 °C) when the operation frequency was set to 3 MHz in the following cases;

- CT1: the circuit was controlled adaptively with canary FF,
- CT2: 200 mV-FBB, which was the minimum body-bias for 3 MHz operation at 25 °C, was fixedly applied,
- CT3: the minimum FBB voltage required for 3 MHz operation at each temperature was applied.




Fig. 6. Measured timing error, warning signals, and transitions of speed level (2 MHz @ $V_{DD} = 0.35$ V).

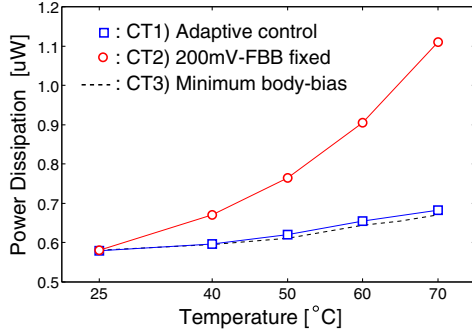


Fig. 7. The power dissipation at the various temperature conditions (3 MHz @ $V_{DD} = 0.35$ V). The circuit operates CT1) adaptively, CT2) with 200 mV-FBB fixedly, and CT3) with a minimum body-bias required for 3 MHz operation at each temperature.

In CT1, a canary FF at S[20] was enabled and its buffer delay was 130 ns at ZBB and 25°C. The power dissipation includes those of the KSA, the main FFs, the speed control unit, and canary FF. The power overhead of a canary FF was estimated to be around 2% by circuit simulation. In this measurement, we set four speed levels from seven speed levels (ZBB – 180 mV-FBB) at each temperature. No timing errors were observed during 1.8×10^9 cycles at all temperature conditions.

We can see that the power dissipation of CT1 is very close to that of CT3, which means optimal body-bias voltages were selected adaptively at each temperature. On the other hand, when the 200 mV-FBB was fixedly applied (CT2), the power dissipation at 75°C was 63% larger than that of CT1.

Figure 8 shows the power dissipation at the various supply voltages (0.33–0.38 V) when the operation frequency was set to 2 MHz in the following cases;

- CV1: the circuit was controlled adaptively with canary FF,
- CV2: 150 mV-FBB, which was the minimum body-bias required at $V_{DD} = 0.33$ V, was fixedly applied,
- CV3: the minimum FBB voltage required for 2 MHz operation at each supply voltage was applied.

No timing errors were observed during 1.2×10^9 cycles at each supply voltage. The power dissipation of the adaptive control (CV1) follows that with minimum body-bias at each supply voltage (CV3), which means the circuit was adaptively controlled appropriately.

C. Comparison to the operation considering the worst-case

We next demonstrate how inefficient the worst-case design and operation for PVT is for subthreshold circuits, and clarify how beneficial the adaptive performance control is.

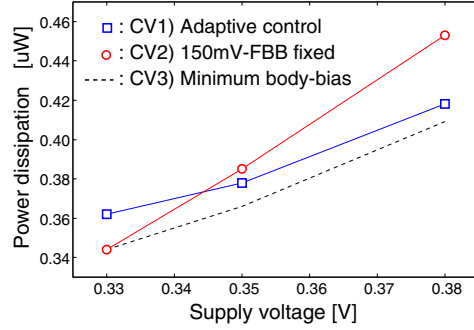


Fig. 8. The power dissipation at the various supply voltage (2 MHz). The circuit operates CV1) adaptively, CV2) with 150 mV-FBB fixedly, and CV3) with minimum body-bias required for 2 MHz operation at each voltage.

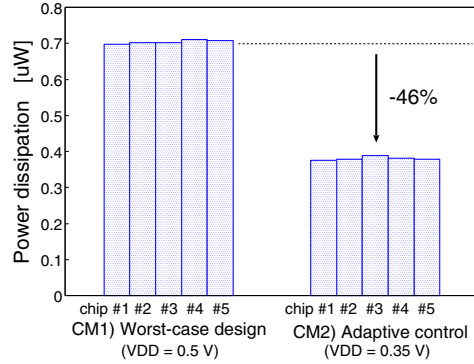


Fig. 9. The power dissipation when the operation frequency is 2 MHz in the following cases; CM1) all chips operate at $V_{DD} = 0.5$ V, CM2) all chips operate with the adaptive control at $V_{DD} = 0.35$ V.

We first focus on the worst-case design in terms of manufacturing variability. Supposing 2 MHz operation, the supply voltage must be 0.5 V or higher for a chip at SS device corner, for example. In this case, all chips should operate at $V_{DD} = 0.5$ V when the traditional worst-case design with guardbanding is adopted. Figure 9 shows the power dissipation of five chips in the following cases;

- CM1: all chips operated at $V_{DD} = 0.5$ V, which was the minimum V_{DD} for a chip at SS device corner,
- CM2: all chips operated with adaptive control at $V_{DD} = 0.35$ V.

One canary FF was enabled, and its location and the buffer delay were determined such that no timing errors occurred during 1.2×10^9 cycles (10 minutes). The power dissipation with the adaptive control (CM2) was smaller than that of the operation with guardbanding (CM1) by 46%, because of lower supply voltage.

Figure 10 shows the power dissipation when temperature is 60°C (3 MHz @ $V_{DD} = 0.35$ V) in the following cases;

- CVT1: body-bias voltage required for operation at the worst-case environmental condition (here, 25°C and $V_{DD} = 0.33$ V) was fixedly applied assuming that the body-bias voltage can be ideally obtained and given for each chip at a pre-shipment test,
- CVT2: the circuit was controlled adaptively with canary FF.

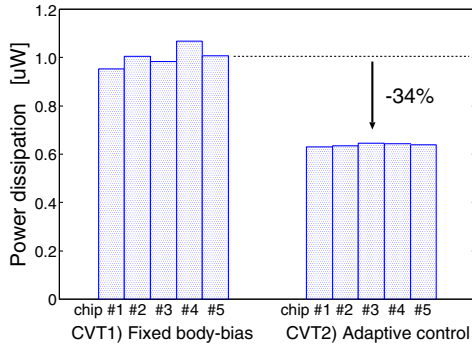


Fig. 10. Power dissipation when temperature is 60°C (3 MHz @ $V_{DD} = 0.35$ V). “Fixed body-bias” denotes body-bias voltage required to operate at the worst-case of the environmental condition (in this example, 25°C and $V_{DD} = 0.33$ V) is fixedly applied assuming that the body-bias voltage can be ideally obtained at a pre-shipment test.

The power of CVT2 is 34% smaller than that of CVT1.

Even if an optimal body-bias could be given for each chip through expensive delay testing and manufacturing variability unique to each chip could be eliminated, we have to assure correct operation at the worst-case environmental condition. The body-bias selected for the worst-case is higher than needed at other environmental conditions. The adaptive speed control can select the appropriate body-bias according to the present environmental condition in addition to manufacturing variability, and hence the design with the adaptive control is much more efficient in power dissipation than the worst-case design.

D. Dependence on design parameters

We will discuss the measured dependences on the design parameters such as the inserted location of canary FF and the buffer delay. Figure 11 shows the timing error rate (MTBF) when the enabled canary FF was varied while keeping the buffer delay constant. The timing error rate at each location was calculated by counting the number of timing errors during 10 minutes (1.2×10^9 cycles). This figure indicates that the timing error rate depends on the location of canary FF. Please note that we here chose the design parameters to make it easier to measure MTBF, i.e. we chose the parameters that caused frequent errors. For a practical use, the design parameters should be decided to satisfy the required timing error rate (MTBF).

To attain longer MTBF, we introduce the following two ways – to lengthen the monitoring period and the buffer delay of canary FF. For example, if the monitoring period is changed from 10^7 used in the measurement to 10^9 cycles (about 5 minutes at 3 MHz operation), MTBF will become at least 100 times larger. In addition, MTBF depends on the buffer delay. Figure 12 shows the timing error rate (MTBF) as a function of the buffer delay when the inserted location was fixed. The longer buffer delay achieved longer MTBF.

IV. CONCLUSION

In this paper, we presented the self-adaptive compensation technique using canary FF for subthreshold circuits. Measure-

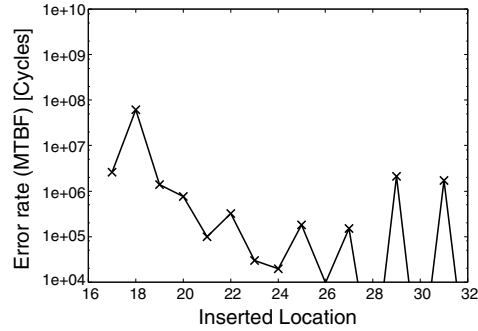


Fig. 11. Timing error rate (MTBF) at each inserted location when the buffer delay is constant (2 MHz @ $V_{DD} = 0.35$ V). The buffer delay is 100 ns at ZBB and 25°C.

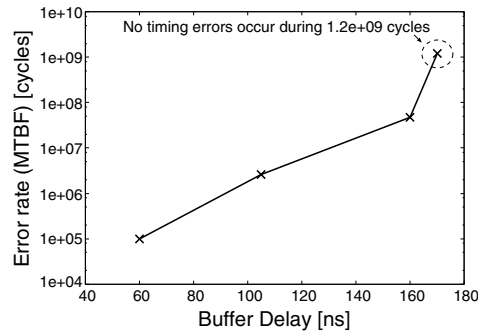


Fig. 12. Timing error rate (MTBF) as a function of the buffer delay when the inserted location is fixed. The X-axis represents the average value of rise and fall delay.

ment results showed the adaptive control compensated manufacturing and environmental variability and reduced power dissipation by 46% compared to traditional worst-case design.

ACKNOWLEDGMENTS

This research was performed by the authors for STARC as part of the Japanese Ministry of Economy, Trade and Industry sponsored “Next-Generation Circuit Architecture Technical Development” program. The VLSI chip in this study has been fabricated in the chip fabrication program of VLSI Design and Education Center (VDEC), the University of Tokyo in collaboration with STARC, e-Shuttle, Inc., and Fujitsu Ltd. This study was partly supported by NEDO.

REFERENCES

- [1] M. Seok, et al., “The Phoenix Processor: A 30pW Platform for Sensor Applications,” in *Symp. VLSI Circuits*, pp.188–189, 2008.
- [2] N. Ickes, et al., “A 10-pJ/instruction, 4-MIPS Micropower DSP for Sensor Applications,” in *Proc. ASSCC*, pp.289–292, 2008.
- [3] J. Kao, et al., “A 175-mV Multiply-Accumulate Unit Using an Adaptive Supply Voltage and Body Bias Architecture,” *JSSC*, pp.1545–1554, Nov. 2002.
- [4] S. Das, et al., “A Self-Tuning DVS Processor Using Delay-Error Detection and Correction,” *JSSC*, pp.792–804, Apr. 2006.
- [5] S. Das, et al., “RazorII: In Situ Error Detection and Correction for PVT and SER Tolerance,” *JSSC*, pp.32–48, Jan. 2009.
- [6] T. Sato and Y. Kunitake, “A Simple Flip-Flop Circuit for Typical-Case Designs for DFM,” in *Proc. ISQED*, pp.539–544, 2007.
- [7] H. Fuketa, et al., “Trade-off Analysis between Timing Error Rate and Power Dissipation for Adaptive Speed Control with Timing Error Prediction,” in *Proc. ASP-DAC*, pp.266–271, 2009.
- [8] L. Shiang, et al., “A sub-pW timer using gate leakage for ultra low-power sub-Hz monitoring systems,” in *Proc. CICC*, pp.397–400, 2007.