



Comparison of in-situ delay monitors for use in Adaptive Voltage Scaling

N. Pour Aryan¹, L. Heiß¹, D. Schmitt-Landsiedel¹, G. Georgakos², and M. Wirnshofer¹

¹Lehrstuhl für Technische Elektronik, Technische Universität München, Germany

²Infiniteon Technologies AG, München-Neubiberg, Germany

Correspondence to: N. Pour Aryan (n.aryan@tum.de)

Abstract. In Adaptive Voltage Scaling (AVS) the supply voltage of digital circuits is tuned according to the circuit's actual operating condition, which enables dynamic compensation to PVT variations. By exploiting the excessive safety margins added in state-of-the-art worst-case designs considerable power saving is achieved. In our approach, the operating condition of the circuit is monitored by in-situ delay monitors. This paper presents different designs to implement the in-situ delay monitors capable of detecting late but still non-erroneous transitions, called Pre-Errors. The developed Pre-Error monitors are integrated in a 16 bit multiplier test circuit and the resulting Pre-Error AVS system is modeled by a Markov chain in order to determine the power saving potential of each Pre-Error detection approach.

1 Introduction

In today's advanced integrated circuits, with ever increasing performance demands, methods and schemes aiming to minimize power consumption in digital circuits are becoming more of a concern. The state-of-the-art worst-case guardbanding approach adds several safety margins considering Process, Voltage and Temperature variability and also Aging (PVT) to the supply voltage required for correct operation under nominal condition. This fixed supply voltage approach results in unnecessary power dissipation in non-worst-case scenarios. Adaptive Voltage Scaling (AVS) controls the supply voltage according to the operating condition of the circuit by exploiting unused safety margins.

The operating condition of the circuit is commonly monitored by measuring the timing of the circuit. The global speed monitor approach in Drake et al. (2007) uses delay lines (replica paths), aiming to track the effect of temperature changes or supply voltage reduction on the critical paths.

Delay lines are only replicas of the real circuit and thus local variations, process gradients or aging mechanisms are not covered by these global monitors.

In contrast to the global speed monitor, in-situ delay monitors measure the timing inside the real circuit and thus provide reliable timing information considering both local and global variations. The in-situ delay monitors are enhanced flip-flops providing information about the timing of the circuit. The AVS approach in Bowman et al. (2009) and Das et al. (2009) uses in-situ delay monitors capable of detecting timing errors. Additional complexity is introduced in these approaches for error recovery. Moreover, repetition of erroneous computation in case of an error occurrence excludes the application of error detection approaches to real time systems.

Therefore, Wirnshofer et al. (2011) use in-situ delay monitors that are able to detect critical transitions instead of error detectors. These in-situ delay monitors, also called Pre-Error flip-flops, are inserted at the end of critical paths in the circuit. The supply voltage is adjusted during normal operation of the circuit based on the timing information obtained by Pre-Error flip-flops. As the voltage adaptation is based on Pre-Errors instead of errors, no additional circuitry and clock cycles are needed for data recovery. Therefore, the Pre-Error AVS is applicable to real time systems.

In this paper, different designs to implement the Pre-Error flip-flops are developed and investigated. The rest of this paper is organized as follows. In Sect. 2, an overview of Pre-Error AVS is given. Section 3 discusses the design of the developed Pre-Error flip-flops. In Sect. 4, the Markov model representing the Pre-Error AVS to estimate the power saving potential is introduced. Comparisons between the different Pre-Error flip-flops are presented in Sect. 5 and the paper is concluded in Sect. 6.

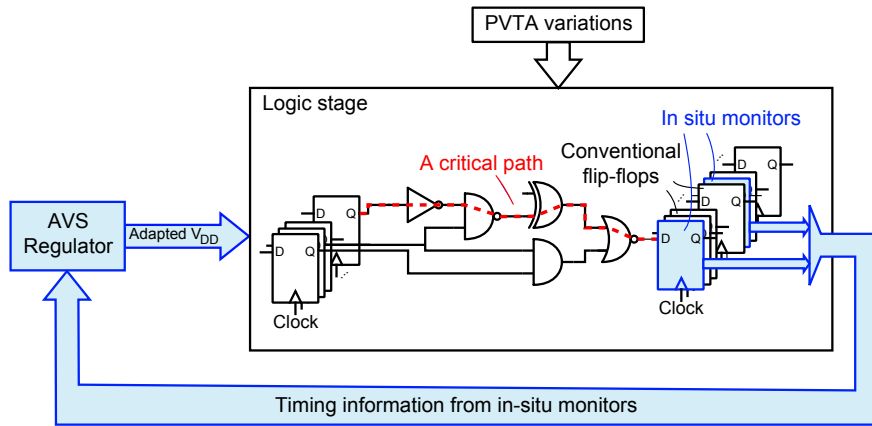


Fig. 1. Block diagram of the Pre-Error AVS system: timing information extracted by in-situ delay monitors is used to tune the supply voltage in a closed-loop configuration.

2 Overview of Pre-Error Adaptive Voltage Scaling

Pre-Error flip-flops are in-situ delay monitors with the ability to distinguish between relaxed and critical operation of the circuit. Timing information extracted by Pre-Error flip-flops inserted at the end of critical paths is used for the closed-loop voltage control, shown in Fig. 1. Data transitions during a certain time interval before clock rising edge, the Pre-Error detection window, result in a Pre-Error signal.

For different input patterns the delay of the digital circuit varies. Figure 2 shows the delay histogram of the 16 bit multiplier's critical output for random input patterns. The part of the distribution lying right of the dashed line denoting the detection window will result in Pre-Errors. Note that for the same input pattern, different operating conditions result in different delays for the digital circuit. Reducing the supply voltage increases the delay of the circuit and moves the histogram to larger delays, resulting in more Pre-Errors.

The supply voltage is adjusted based on the count of Pre-Errors during an observation interval, i.e. a defined number of clock cycles. To adapt the supply voltage, two decision limits for the Pre-Error count after an observation interval are defined. If n_{pre} , the count of Pre-Error occurrence in the previous observation interval, is under a lower threshold of $n_{limit\downarrow}$ the voltage is decreased. If n_{pre} is above an upper threshold of $n_{limit\uparrow}$ the voltage is increased. For Pre-Error counts between $n_{limit\downarrow}$ and $n_{limit\uparrow}$ the voltage is maintained. As the voltage adaptation is based on the statistics of Pre-error occurrence, an overcritical voltage reduction, i.e. the risk of timing errors, can never be completely ruled out. However, the error rate can be adjusted in a user-defined controlled manner by adjusting the number of clock cycles for the observation interval, the detection window length T_{Pre} or the decision limits $n_{limit\downarrow}$ and $n_{limit\uparrow}$.

With varying activity rates in digital circuits, regulating the voltage based on the number of occurred Pre-Error pulses during a *fixed* time interval would result in aggressive volt-

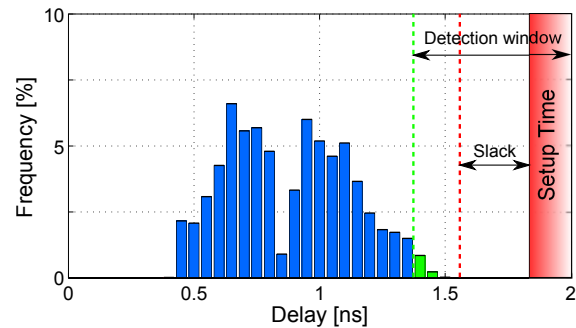


Fig. 2. Delay histogram of the multiplier's critical output for different input patterns at a fixed operating condition.

age reduction in phases with low activity. This would lead to unpredictable risk of timing errors. Therefore, the observation interval should only consist of active clock cycles. To distinguish between active and inactive clock cycles, we introduce a transition detector monitoring all data transitions, either relaxed or critical.

The AVS Regulator shown in Fig. 1 consists of an AVS control unit and a voltage regulator. The digital AVS control unit counts the Pre-Errors in an observation interval of N active clock cycles and decides whether to change the supply voltage. The communication to the voltage regulator is done via a binary control word representing the voltage level.

3 Implementation of Pre-Error flip-flops

A Pre-Error flip-flop is a conventional flip-flop with additional circuitry enabling it to detect Pre-Errors and data transitions. For realizing the Pre-Error detection window, either the duty-cycle of the clock signal or a delay element can be used. All Pre-Error flip-flops are designed in a 65 nm low power CMOS technology.

3.1 Duty-cycle based Pre-Error approach

In this approach the Pre-Error flip-flop is implemented in dynamic and static design style, respectively.

3.1.1 Dynamic Pre-Error flip flop

The Dynamic Pre-Error flip flop is designed in dynamic logic, shown in Fig. 3a. The low phase of the clock signal is exploited as the Pre-Error detection window. The transition pulse generator, which is an XOR gate with two inputs of data and delayed data, generates a pulse at its output X for every data transition. Data transitions occurring during the detection window are assigned as Pre-Errors by transistor M_2 . All data transitions from one to the following clock rising edge are assigned as Transition pulses by transistor M_4 .

For the AVS scheme it is important to define a detection window as accurate and robust as possible. Therefore, the variations of the detection window length are investigated in the following. Since the duty cycle of the clock signal is exploited as the detection window, variations caused by the Pre-Error flip-flop itself and the clock tree are considered. Thus, the clock tree as an H-tree with three levels of buffers is included in the design. The H-tree is the basic clock topology for many clock distribution systems (Tam et al., 2004).

Figure 3b illustrates the deviations from the ideal detection window length in the nominal and corner cases¹ over V_{DD} . The ideal detection window length is considered the length at the supply voltage of $V_{DD} = 1.0$ V, nominal process and a temperature of $T = 27^\circ\text{C}$. Moreover, Monte Carlo simulations are performed to determine the uncertainties of the detection window length due to local variations. For the supply voltage range of $V_{DD} = 1.2$ V down to $V_{DD} = 0.8$ V, the 3σ interval due to local variations is derived from Monte Carlo simulations and depicted in Fig. 3b. In this design, global variations have a minor impact compared to local ones.

3.1.2 Static Pre-Error flip flop

For the Static Pre-Error flip-flop, illustrated in Fig. 4a, the falling edge of the clock signal is used as the starting point for the detection window.

The Static Pre-Error flip-flop is implemented in static design style with standard library elements, requiring less design effort than the Dynamic Pre-Error flip-flop. The Pre-Error detector is comprised of a flip-flop with inverted clock as its clock input. The data input and the output Q 2 are compared by the XOR1 gate. The flip-flop with inverted clock

¹Note that typically the fast corner is at fast process and low temperature. In this example, however, the voltage is scaled to a point where the circuit is operated at temperature inversion. Here, the effect of decreasing threshold voltage with temperature exceeds the mobility degradation. Consequently, the circuit exhibits an inverted temperature characteristic, as it speeds up with increased temperature and vice versa.

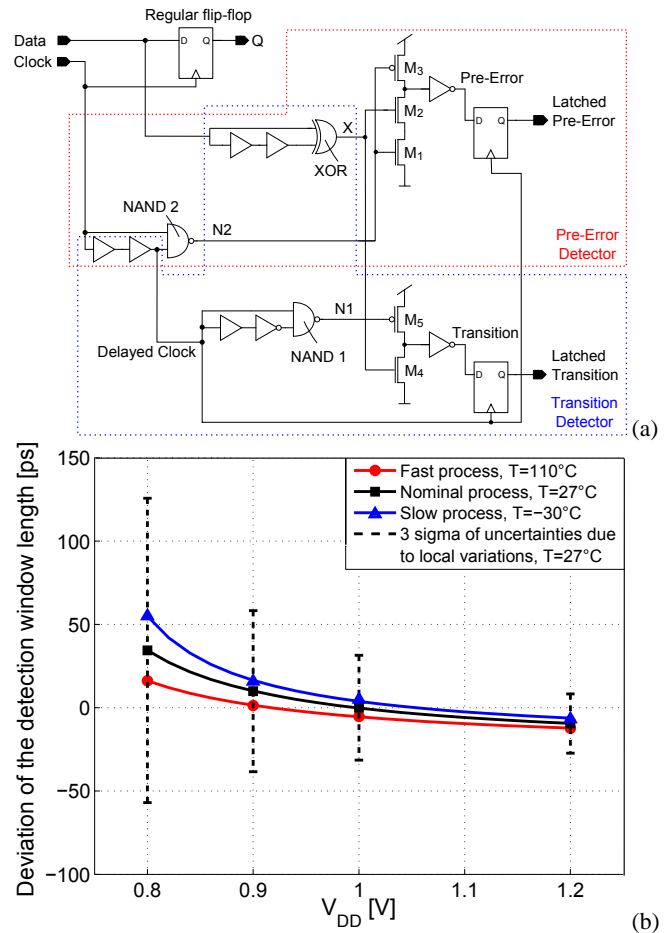


Fig. 3. Dynamic Pre-Error flip-flop, (a) Schematic, (b) Deviations of the detection window over V_{DD} under variations.

latches data at the beginning of the detection window. Therefore, if a data transition occurs during the detection window this flip-flop fails to detect valid data and a Pre-Error is generated by XOR1. For the Transition detector circuit, the inputs of XOR2 are the Data signal and the output Q of the regular flip-flop. In case of a data transition, the input signal Data will differ from its value in the previous clock cycle, stored as Q . Hence, a transition signal is generated by XOR2.

Deviations of the detection window over V_{DD} for corner cases and under local variations are shown in Fig. 4b. The Static Pre-Error flip-flop is more robust to global and local variations compared to the Dynamic Pre-Error flip-flop.

In duty-cycle based Pre-Error approaches, the deviations of the detection window over V_{DD} or due to variations are independent of the length of the nominal detection window.

3.2 Delay element based Pre-Error approach

In the Delay element based Pre-Error approach a shadow flip-flop with delayed data is added in parallel to the regular flip-flop (Eireiner et al., 2007). Figure 5a shows the schematic

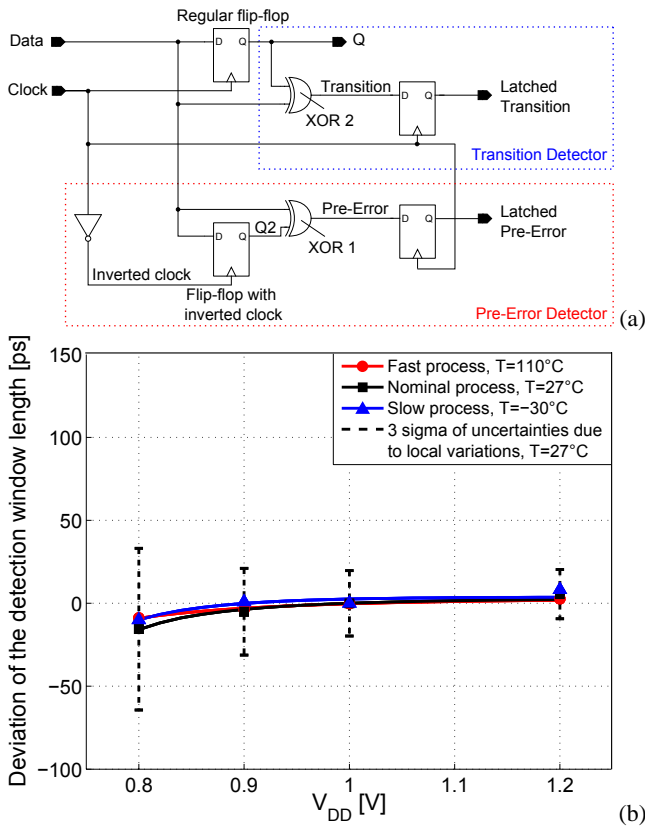


Fig. 4. Static Pre-Error flip-flop, (a) Schematic, (b) Deviations of the detection window over V_{DD} under variations.

of this so-called Crystal-ball Pre-Error flip-flop. The delay element comprising an inverter chain specifies the length of the detection window. When a data transition occurs closer to the clock rising edge than the delay of the delay element, the shadow flip-flop will miss to latch the input data and the Pre-Error pulse is generated. Transition detection is implemented the same way as in the Static Pre-Error flip-flop.

For a delay element comprising 44 inverters designed to implement a nominal detection window of 650 ps at $V_{DD} = 1.0$ V, deviations of the detection window over V_{DD} for corner cases and under local variations are shown in Fig. 5b. For the Crystal-ball Pre-Error approach, the detection window length is strongly dependent on the delay of the delay element, which has a large shift in corner cases. For an ideal detection window of $T_{Pre} = 650$ ps at a supply voltage of $V_{DD} = 1.0$ V, nominal process and temperature of $T = 27^\circ\text{C}$, the shifts in the detection window length are more than 10 times larger than the shifts for the duty-cycle based Pre-Error flip-flops. Remember that an enlarged detection window increases the probability for a voltage increment and vice versa. Thus, the strong voltage dependence of the Crystal-ball Pre-error flip-flop might seem critical at first, but it stabilizes the adapted voltage around one level.

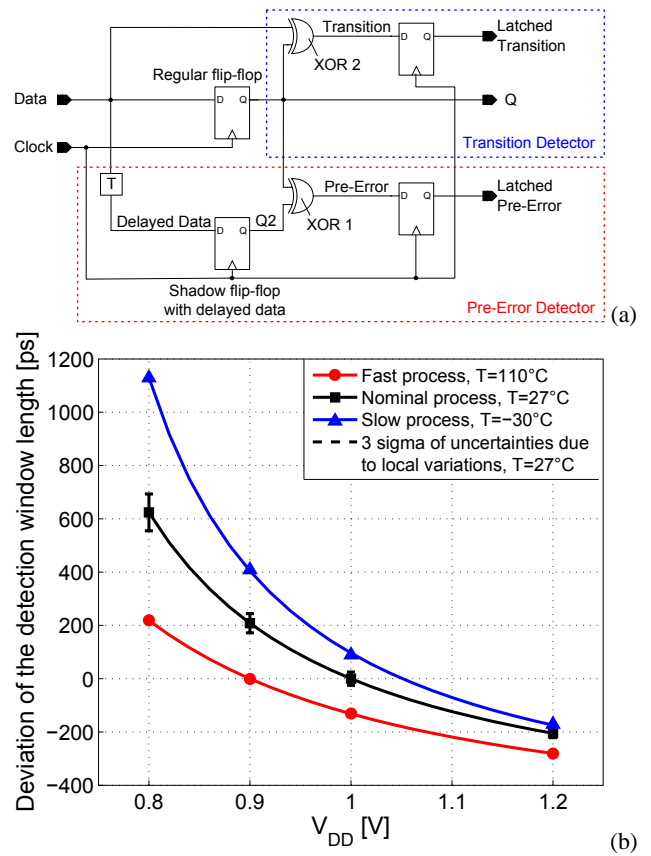


Fig. 5. Crystal-ball Pre-Error flip-flop, (a) Schematic, (b) Deviations of the detection window over V_{DD} under variations.

4 Analyzing the modeling of the Pre-Error AVS

As shown in Fig. 2, for different input patterns the output delays of the digital circuit vary due to changing signal propagation paths. The probability P_{pre} of a Pre-Error in one clock cycle is the number of critical delays (occurring during Pre-Error detection window) divided by the number of all delays. In the Pre-Error AVS, the supply voltage is regulated between a finite number of voltage levels, M , based on the number of Pre-Errors in the previous observation interval.

After each observation interval comprising N active clock cycles, the supply voltage is decreased if the count of Pre-Error occurrence, n_{pre} , is below $n_{limit\downarrow}$. Therefore the probability of decreasing the supply voltage, $P_{V_{DD}\downarrow}$, for each voltage level, V_{DD} , is calculated by

$$P_{V_{DD}\downarrow} = P[n_{pre} < n_{limit\downarrow}] \quad (1)$$

$$= \sum_{n_{pre}=0}^{n_{limit\downarrow}-1} \binom{N}{n_{pre}} \cdot (P_{pre})^{n_{pre}} \cdot (1 - P_{pre})^{N-n_{pre}}$$

Similarly the probability of increasing and maintaining the supply voltage, $P_{V_{DD}\uparrow}$ and $P_{V_{DD}\rightarrow}$ are calculated. The transition probabilities automatically adapt to the actual operating condition of the circuit. For a fixed operating condition,

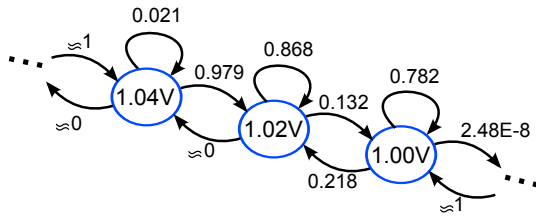


Fig. 6. Markov chain used to model the Pre-Error AVS. The values next to each arrow denote the transition probabilities.

the transition probabilities satisfy the Markov property, and thus we use a Markov chain to model the Pre-Error AVS. An example of the resulting Markov chain with a voltage granularity of 20 mV is shown in Fig. 6.

The corresponding Markov matrix is formed assuming a maximum voltage of $V_{DD} = 1.2\text{ V}$ and a right Markov matrix (each row summing up to 1):

$$\mathbf{P} = \begin{pmatrix} P_{1.20\text{V}\rightarrow} & P_{1.20\text{V}\downarrow} & 0 & 0 & 0 & \dots \\ P_{1.18\text{V}\uparrow} & P_{1.18\text{V}\rightarrow} & P_{1.18\text{V}\downarrow} & 0 & 0 & \dots \\ 0 & P_{1.16\text{V}\uparrow} & P_{1.16\text{V}\rightarrow} & P_{1.16\text{V}\downarrow} & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \ddots \end{pmatrix} \quad (2)$$

For a fixed operating condition, the voltage adaptation comes into the steady state. The corresponding steady state vector of probabilities π reads as

$$\pi = (P_{1.20\text{V}} \ P_{1.18\text{V}} \ P_{1.16\text{V}} \ \dots) \quad (3)$$

where $P_{V_{DD}}$ is the final probability of being at the voltage level V_{DD} . The steady state vector of probabilities satisfies the equation

$$\pi \mathbf{P} = \pi \quad (4)$$

which defines π as the left eigenvector of the Markov matrix, with the corresponding eigenvalue 1. For changing operating condition, π automatically adapts.

4.1 Power saving potential of the Pre-Error AVS approach

The designed Pre-Error flip-flops are integrated into a 16 bit multiplier test circuit in a 65 nm CMOS technology. The clock frequency is chosen as 500 MHz ($T = 2\text{ ns}$). Pre-Error flip-flops are placed at the end of the three most critical paths (9.4 % of all 32 outputs). An OR tree combines the outputs of three Pre-Error flip-flops to generate the overall Pre-Error. Similarly, the Transition signals are ORed together. The power saving potential is evaluated using the Markov Model for each Pre-Error approach.

When analyzing the power saving of each Pre-Error approach, the overhead of the AVS circuitry has to be considered. Additional power consists of the extra circuitry for the integrated Pre-Error flip-flops, OR trees and the AVS control unit. For all Pre-Error approaches integrated in the 16bit

multiplier test circuit, the power overhead of the AVS control unit is evaluated as 3.4 %. The power overhead corresponding to the OR trees is less than 1 %. Finally the power overhead due to the three Pre-Error flip flops is 9.8 %, 4.3 % and 10.1 %, resulting in the total power overhead of 13.3 %, 7.8 % and 13.7 % for the Dynamic, Static and Crystal-ball Pre-Error approaches ($T_{\text{Pre}} = 650\text{ ps}$), respectively. To evaluate the risk of timing errors, the error probability is determined for each supply voltage level with SPICE simulation. The overall error rate is determined using the simulated error rates at each voltage level, and the steady state vector of probabilities:

$$P_{\text{err}} = \sum_{V_{DD}} P_{V_{DD}} \cdot P_{\text{err}, V_{DD}} \quad (5)$$

where $P_{V_{DD}}$ is the final probability of being at a certain voltage level V_{DD} after voltage adaptation and $P_{\text{err}, V_{DD}}$ is the error probability of the corresponding supply voltage level of V_{DD} . Similarly, the total power consumption is evaluated as:

$$P_{\text{dyn}} = \sum_{V_{DD}} P_{V_{DD}} \cdot P_{\text{dyn}, V_{DD}} \quad (6)$$

in which $P_{\text{dyn}, V_{DD}}$ is the simulated dynamic power consumption of the supply voltage level of V_{DD} .

5 Results and comparison

Different variations can be included into the Markov model: the delay variations of the multiplier test circuit, the setup-time variations of all flip-flops and the Pre-Error detection window variations for each approach. To be able to compare between the proposed approaches, we excluded the local variations of the multiplier test circuit for the following results.

For the Static Pre-Error approach, Fig. 7 shows the resulting power saving potential considering detection windows with nominal lengths of 650 ps (lowest error rates), 550 ps, 400 ps and 250 ps (highest error rates). For all designs, the error rate deviation from the nominal values in fast and slow corners is small compared to the impact of local variations. To compare the robustness of the proposed Pre-Error designs, we inspected the uncertainty of an adjusted error rate. The chosen error rate is 10^{-9} , which requires a detection window of approximately 550 ps. The Crystal-ball Pre-Error flip-flop implementing the detection window of 550ps has a total power overhead of 12.9 %. The normalized error uncertainty in presence of local variations is obtained by dividing the error uncertainty shown in Fig. 7 by the nominal error rate. For an error rate of 10^{-9} , the resulting normalized error uncertainty is 2.3 and 4.8 for Static and Dynamic Pre-Error flip-flops. For the Crystal-ball Pre-Error flip-flop, it is only 1.6 as a result of the strong voltage dependence for the Pre-Error detection length, which reduces voltage alterations and provides even more stabilized voltage control. Table 1 shows the power saving with the nominal error rate of 10^{-9} .

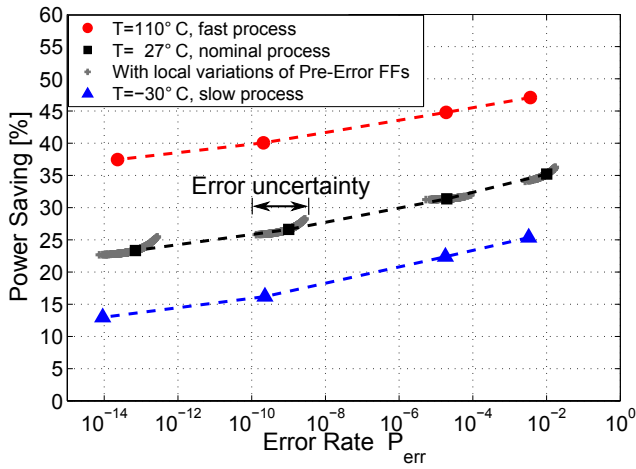


Fig. 7. Power saving potential in relation to the error rate for the Static Pre-Error approach.

As shown in Table 1, the power saving potential of the Static Pre-Error approach is larger than for the other two approaches as it has less power overhead. The power saving potential of the Static Pre-Error approach in the nominal case is 26.6% compared to the conventional fixed voltage approach. For the fast corner ($T = 110^\circ\text{C}$, fast process) lower voltages are more likely due to increased circuit speed, resulting in higher power savings. As the AVS also exploits the conservative timing margins produced by the synthesis tool, there is considerable power saving even in the slow corner.

6 Conclusions

Three Pre-Error detection approaches for use in Pre-Error AVS were designed and optimized. The Dynamic Pre-Error flip-flop requires a lot of design effort to achieve a robust and reliable design. In contrast, both Static and Crystal-ball Pre-Error flip-flops are designed in static design style using standard library elements, requiring less design effort. Monte-Carlo simulations and corner analysis were performed to evaluate the variations of the detection window length. The maximum 3σ of the detection window length deviation occurring at the lowest evaluated supply voltage ($V_{DD} = 0.8\text{ V}$) is 50 ps, 70 ps and 90 ps for Static, Crystal-ball and Dynamic Pre-Error designs, respectively. To determine the power saving and robustness, the Pre-Error AVS was modeled by a Markov chain. The variations of detection window length were included in the Markov model. For each Pre-Error design, the power saving potential and error rate uncertainties were determined for local and global variations. For a nominal error rate of 10^{-9} at a clock frequency of 500 MHz, the Static Pre-Error design has the largest power saving of 26.6% and Crystal-ball and Dynamic Pre-error designs have power savings of 21.6% and 21.1%, respectively. The Crystal-ball Pre-Error design has the smallest normal-

Table 1. Power saving potential of Pre-Error AVS scheme applied to the 16 bit multiplier test circuit (in %).

Pre-Error detection approach	Fast process $T = 110^\circ\text{C}$	Nominal process $T = 27^\circ\text{C}$	Slow process $T = -30^\circ\text{C}$
Static	40.1	26.6	16.9
Dynamic	34.6	21.1	10.7
Crystal-ball	34.9	21.6	11.1

ized error uncertainty of 1.6 and Static and Dynamic Pre-Error designs have the normalized error uncertainties of 2.3 and 4.8.

Acknowledgements. This work was supported by the German Research Foundation (DFG) within Grant SCHM 1478/8-1.

References

- Bowman, K. A., Tschanz, J. W., Nam Sung Kim, Lee, J. C., Wilkerson, C. B., Lu, S.-L. L., Karnik, T., and De, V. K.: Energy-efficient, metastability-immune resilient circuits for dynamic variation tolerance, *IEEE J. Solid-State Circuits*, 44, 49–63, 2009.
- Das, S., Tokunaga, C., Pant, S., Ma, W.-H., Kalaiselvan, S., Lai, K., Bull, D. M., and Blaauw, D. T.: RazorII: In situ error detection and correction for PVT and SER Tolerance, *IEEE J. Solid-State Circuits*, 44, 32–48, 2009.
- Drake, A., Senger, R., Deogun, H., Carpenter, G., Ghiasi, S., Nguyen, T., James, N., Floyd, M., and Pokala, V.: A Distributed critical-path timing monitor for a 65nm high-performance micro-processor, *IEEE International Solid-State Circuits Conference (ISSCC)*, 11–15 February 2007, 398–399, 2007.
- Eireiner, M., Henzler, S., Georgakos, G., Berthold, J., and Schmitt-Landsiedel, D.: In-Situ delay characterization and local supply voltage adjustment for compensation of local parametric variations, *IEEE J. Solid-State Circuits*, 42, 1583–1592, 2007.
- Tam, S., Limaye, R. D., and Desai, U. N.: Clock generation and distribution for the 130-nm Itanium^{REG} 2 processor with 6-MB on-die L3 cache, *IEEE J. Solid-State Circuits*, 39, 636–642, 2004.
- Wirnshofer, M., Heiß, L., Georgakos, G., and Schmitt-Landsiedel, D.: A variation-aware Adaptive Voltage Scaling technique based on in-situ delay monitoring, *IEEE 14th International Symposium on Design and Diagnostics of Electronic Circuits Systems (DDECS)*, 261–266, 2011.