

Ring Oscillator Clocks and Margins

Jordi Cortadella*, Marc Lupon*, Alberto Moreno*, Antoni Roca* and Sachin S. Sapatnekar[†]

*Department of Computer Science, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain

[†]Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA

Abstract—How much margin do we have to add to the delay lines of a bundled-data circuit? This paper is an attempt to give a methodical answer to this question, taking into account all sources of variability and the existing EDA machinery for timing analysis and sign-off. The paper is based on the study of the margins of a ring oscillator that substitutes a PLL as clock generator. A timing model is proposed that shows that a 12% margin for delay lines can be sufficient to cover variability in a 65nm technology. In a typical scenario, performance and energy improvements between 15% and 35% can be obtained by using a ring oscillator instead of a PLL. The paper concludes that a synchronous circuit with a ring oscillator clock shows similar benefits in performance and energy as those of bundled-data asynchronous circuits.

Index Terms—ring oscillator; on-chip variability; reactive clock;

I. INTRODUCTION

Asynchronous bundled-data circuits offer an attractive trade-off between synchronous and quasi delay-insensitive (QDI) circuits in terms of area, performance, and power. QDI provides robustness at the expense of an important cost in area and power for the use of delay-insensitive encoding techniques (e.g., dual rail). On the other hand, synchronous circuits are based on the generation of high-quality clock signals that provide reliable timing references. Phase-locked loops (PLLs) are commonly used to generate low-jitter clocks that are agnostic to the variability experienced by the circuit at runtime.

The way variability is handled in synchronous circuits is by adding guardband margins to the clock period that can accommodate the static and dynamic delay fluctuations. Static timing analysis (STA) with different process, voltage, and temperature (PVT) corners and on-chip variability (OCV) derating factors are typically used to estimate conservative bounds on delay variability [1].

The datapath of a bundled-data circuit is similar to the one of a synchronous circuit (see Fig. 1). The main difference lies on the clock signal generated by a set of distributed oscillators (delay lines) synchronized with other oscillators via handshake controllers (req/ack signals), such as the ones presented in [2].

Fig. 2 depicts a system with two clock domains, one driven by a ring oscillator (RO) and the other by a PLL. From the functionality point of view, both generators are interchangeable. It is even possible to design clock domains in which the clock generator can be dynamically selected (via multiplexers) at each time instant.

In the last few years, we have observed a proliferation of proposals for clock generation based on ROs [3]–[7]. The main motivation is the capability of tracking PVT variability, thus reducing the guardband margins and providing tangible improvements in power and performance. Bundled-data circuits also share the same motivation given that the role of the ROs and the delay lines is similar. The question we would like to answer in this paper is:

Which are the benefits of using an RO clock instead of a PLL?

*This work has been partially supported by funds from the Spanish Ministry for Economy and Competitiveness and the European Union (FEDER funds) under grant TIN2013-46181-C2-1-R, the Generalitat de Catalunya (2014 SGR 1034 and FI-DGR 2015) and a Fulbright award.




Fig. 1: Bundled-data asynchronous pipeline




Fig. 2: GALS system with two clock domains.

The answer to this question will implicitly give an estimation of the margins that must be applied to delay lines in bundled-data circuits.

The paper concludes that substituting PLLs by ROs is a practical alternative that inherits most of the benefits of asynchronous bundled-data circuits in terms of tolerance to variability. Moreover, the actual commercial EDA tools for STA can be used in RO clocks with minimum changes on the scripts required for sign-off.

II. WHY RING OSCILLATOR CLOCKS?

Designers usually advocate for robust timing references during STA. Low-jitter clocks and near-zero-skew clock trees are mechanisms that contribute to reduce the guardband margins required to achieve a target performance. For this reason, ROs have been generally disregarded as clock sources because of their jitter instability under the presence of variations, thus making STA either difficult or over-conservative (adding margins to cover a large clock jitter).

But looking at this problem from another angle, we can observe that the jitter generated by an RO is highly correlated with delay variability of the circuit [7].

This phenomenon is illustrated in Fig. 3. The horizontal bars represent delays. The critical paths of the circuit and the RO (or delay line in a bundled-data circuit) are competing paths. At different time instants (t_1 , t_2 , and t_3) the paths may experiment different delays due to the operating conditions (e.g., voltage and temperature fluctuations). However, and given that all paths are in the same neighbourhood, their delays are highly correlated. Therefore, the margins required for the RO clock only have to protect its *differential variability* with regard to the critical paths of the circuit (represented as Δd in the picture).

If we would use a PLL (agnostic to variability), the margins for timing correctness should cover the full variability range of the




Fig. 3: Guardband margins for a PLL and an RO clock.




Fig. 4: Period for a PLL and an RO with varying operating conditions.

critical paths. In conclusion, the main differences between a clock signal generated by a PLL or by an RO clock are:

- The margins for an RO clock are smaller and only have to cover the differential variability with regard to the critical paths.
- An RO clock generates a high-jitter clock whose average frequency is better than the one generated by a PLL.

Fig. 4 emphasizes the potential benefits of using an RO as a clock source. The period generated by a PLL must conservatively account for all possible (global and local) variations produced by the fluctuating operating conditions of the circuit, whereas an RO must only account for local variations. Overall, the average period of the RO is smaller than the one required for a PLL.

A key aspect of the RO is that it can *instantaneously* adapt the cycle period to dynamic variations (e.g., unexpected voltage droops). This particular aspect is crucial to save margins and bring significant improvements in power and performance. This immediate reaction allows the RO clock to outperform other techniques as dynamic voltage and frequency scaling [7] which reaction time is typically a few or hundred clock cycles [8].

The main contribution of this paper is to quantify the benefits of RO clocks and define a methodology for their timing sign-off.

III. STATIC TIMING ANALYSIS FOR A RING OSCILLATOR CLOCK

The purpose of Static Timing Analysis (STA) is to check whether the circuit meets a set of timing constraints that guarantee a proper propagation of data across the sequential elements. Two constraints are usually checked: setup and hold. The former is the one that determines the clock period and will be the center of our attention.

A timing constraint is specified as an inequality of two competing paths: the *launch* and the *capture* paths. For a setup constraint, the launch path usually starts at the clock generator, then it goes through the clock tree, the launch flip-flop, and the critical path, and ends at the capture flip-flop (red path in Fig. 5). The capture path (blue) starts from the clock generator, and ends at the capture flip-flop. It is important to note that the launching and capturing clock pulses are separated by a clock period. Thus, the setup constraint can be defined as¹:

$$LaunchPath < CapturePath + Period \quad (1)$$

¹For simplicity, we assume the setup time of the capturing flip-flop to be included in the delay of the launch path.




Fig. 5: Paths involved in a setup constraint with an RO clock.

The previous inequality must also take into account variability. Given that timing analysis cannot be performed under all possible operating conditions, the conventional approach for modern STA is to analyze the circuit in a discrete set of corners. Each corner defines the values for a set of parameters that model process and environmental variations (voltage and temperature).

For a given subset of global PVT operating conditions, the components of the circuit also suffer local (on-chip) variations. To cover on-chip variability (OCV for short), corner-based sign-off applies some derating factors to the timing paths of the circuit that scale the delays with regard to other competing paths in the timing constraints.

Finally, clock jitter and any pessimism derived from the inaccuracies and uncertainties of STA must also be modeled. Typically, they are modeled as a fixed margin in the timing constraints. In summary, in modern STA, variability is modeled using:

- **library corners** to model global variability.
- **derating factors** to model on-chip variability.
- **clock uncertainty** to model clock jitter and other inaccuracies.

Timing constraints must hold for all paths and corners under consideration. Given a library corner, the derating factors and clock uncertainty must be incorporated in the setup constraint:

$$\delta_L \cdot LaunchPath < \delta_C \cdot CapturePath + Period - Jitter \quad (2)$$

where $\delta_L \geq 1$ and $\delta_C \leq 1$ are the derating factors applied to the launch and capture paths, respectively. Clock jitter must be conservatively subtracted from the period.

A simplification of the model consists of making the derating factors symmetric and reducing the analysis to some ϵ such that:

$$\delta_L = 1 + \epsilon, \quad \delta_C = 1 - \epsilon. \quad (3)$$

The accuracy on how these derating factors model on-chip variability is crucial. Foundries usually provide conservative values, but more aggressive values can be used if designers have additional knowledge about the behavior and operating conditions of the circuit.

A. Timing sign-off for a ring oscillator clock

When using an RO as a clock source, the term *Period - Jitter* must be substituted by the delay of the RO in constraint (2):

$$\delta'_L \cdot LaunchPath < \delta'_C \cdot (CapturePath + RO) \quad (4)$$

with $\delta'_L = 1 + \epsilon'$ and $\delta'_C = 1 - \epsilon'$ being new derating factors.

Notice that the derating factor δ'_C is also applied to the delay of the RO (green path in Fig. 5). This is because the RO must be treated as a conventional timing path that experiments the same sources of variability as the other components of the circuit. The *J* term disappears in (4) as jitter is included in RO delay. The derating

factors in (2) can be different from those in (4) since δ'_C and δ'_L must also take into account the spatial correlation between the critical paths and the RO.

The key question is: which is the main difference between (2) and (4)? The answer is that these inequalities must hold for all PVT conditions (corners). While the PLL-based period is fixed for all conditions, the RO adapts its period “on-the-fly” and instantaneously, thus reducing the required guardband margins.

For simplicity, and without loss of generality, we will consider that the launch and capture paths are disjoint. In general, they share some part of the clock tree for which no derating factors should be applied. The technique of not applying derating factors for common paths is called *Common Path Pessimism Removal* (CPPR) [1]. In this paper we will assume that the common paths have already been cancelled out in the timing constraints. In the next sections, we will present a simple model to quantify the benefits of using an RO clock.

B. Multicorner static timing analysis

For timing closure, STA is performed on multiple corners that cover a spectrum of PVT variations. When using a PLL, the period is set to a frequency that guarantees the circuit to operate correctly under all the specified variations.

However, most dies may run faster than the specified frequency. To mitigate this pessimism, a process known as binning may be used to classify dies in such a way that each one can run at a different speed. Unfortunately, binning is a complex and expensive technique which is not always affordable.

Nevertheless, the RO clock period is determined by the PVT conditions that each die experiments at each time instant, which in practice is similar to performing binning if the RO is properly designed [7]. In this scenario, multicorner STA is necessary to guarantee the correct functionality of the RO at any available corner.

IV. DEFINING MARGINS AND DERATING FACTORS IN A RING OSCILLATOR CLOCK

This section presents a statistical model for PVT variations using an RO as a clock source. To ease comprehension, some of the details are presented in the appendix.

A. Static and dynamic on-chip variability

Process, voltage, and temperature are the main sources of on-chip variability. There exist other sources of variability as aging. In this paper we omit any other variability sources except PVT as there is no consensus in the community on how they can be quantified. We can classify variability sources depending on their behavior. On the one hand, *static* process (P) variability does not fluctuate along time, and it is caused by the uncertainties of manufacturing. Process variations are mainly due to voltage threshold variations, which can be accommodated in statistical models. In this paper, we assume that cell delays follow a normal distribution due to random process variations on cell voltage threshold [9], [10].

On the other hand, voltage and temperature (VT) variations are *dynamic*. Unflattened thermal distribution, unbalances in IR droops, computational activity, and floorplan power density are examples of VT variability [11]–[13]. As a result, VT changes locally in time depending on the chip activity and the fluctuations of the voltage source. The local component of dynamic variability can be quantified or bounded with enough detail by analyzing the target circuit with EDA tools. Note that local dynamic variability is usually marginal. Moreover, on-chip voltage has a large global component which is more pronounced when cells are close—more on this in Section VI.

Let us now model D_π as the delay of path π at time t :

$$D_\pi(t) = D_\pi^0 + D^P + D^V(t) + D^T(t)$$

where D_π^0 is the nominal critical path delay in a given corner, and D^P , $D^V(t)$, $D^T(t)$, represent the delay variation due to process, voltage, and temperature, respectively². Process delay variation is modeled as a normal distribution. Hence, D^P will be also modeled as a normal distribution. If $D^P \sim N(0, \sigma_\pi)$, $D_\pi(t)$, then:

$$D_\pi(t) \sim N(D_\pi^0 + D^V(t) + D^T(t), \sigma_\pi^2) \quad (5)$$

Simplifying, let us represent the delay of path π at time t as:

$$D_\pi(t) \sim N(D_\pi^0 + D_\pi^{VT}(t), \sigma_\pi^2)$$

where D_π^0 is the nominal delay in a specific corner, $D_\pi^{VT}(t)$ accounts for the dynamic variability produced by the fluctuations of voltage and temperature, and σ_π models the static (process) variations affecting the critical path.

B. On-chip variability margins in a PLL clock

Let $D_L(t)$ and $D_C(t)$ be the delay of the launch and capture paths at time t for a given corner. Both delays follow a normal distribution as defined in (5). That is, $D_L(t) \sim N(D_L^0 + D_L^{VT}(t), \sigma_L^2)$ and $D_C(t) \sim N(D_C^0 + D_C^{VT}(t), \sigma_C^2)$. Let P and J be the clock period and jitter of a rigid clock such as the PLL. The setup constraint (2) can be defined as:

$$P - J + D_C(t) - D_L(t) > 0,$$

which can be rewritten in:

$$P > (D_L^0 - D_C^0) + J + M_{PLL} \quad (6)$$

where $(D_L^0 - D_C^0)$ is its nominal corner value and M_{PLL} is the guardband margin required to cover on-chip variations. M_{PLL} can be decomposed into two different terms: a margin due to P (M_{pll}^P) variations and a margin due to VT (M_{pll}^{VT}) variations. Both margins are defined as

$$M_{pll}^P = \Psi \cdot \sigma_{PLL}; \quad M_{pll}^{VT} = \max_t |D_L^{VT}(t) - D_C^{VT}(t)| \quad (7)$$

where Ψ is a constant that depends on the number of critical paths and the desired yield (see (17) in the appendix), and $\sigma_{PLL}^2 = \sigma_L^2 + \sigma_C^2$. Note that the maximum difference along time between $D_L^{VT}(t)$ and $D_C^{VT}(t)$ is required to obtain a conservative clock period.

C. Guardband margins in an RO clock

Similar to the previous analyses, it is possible to estimate the margin required by an RO clock. Let $D_{RO}(t)$ be the delay of the RO defined as in (5), i.e., $D_{RO}(t) \sim N(D_{RO}^0 + D_{RO}^{VT}(t), \sigma_{RO}^2)$. Here, the setup constraint (4) for any timing path can be defined as:

$$D_{RO}(t) + D_C(t) - D_L(t) > 0,$$

and by developing the previous expression, we can obtain:

$$D_{RO}^0 \geq (D_L^0 - D_C^0) + M_{RO} \quad (8)$$

where $(D_L^0 - D_C^0)$ is its nominal value and M_{RO} is the margin required to absorb local PVT variations. As in the previous section, M_{RO} can be decomposed into a margin to cover P variations (M_{RO}^P) and a margin to cover VT variations (M_{RO}^{VT}). Both margins are defined as:

$$M_{RO}^P = \Psi \cdot \sigma_R; \quad M_{RO}^{VT} = \max_t |D_L^{VT}(t) - D_{RO}^{VT}(t) - D_C^{VT}(t)|$$

²Note that D_P models static variability and does not depend on t .




Fig. 6: Voltage received in the RO clock of 10-AES circuit.

where $\sigma_R^2 = \sigma_{RO}^2 + \sigma_L^2 + \sigma_C^2$. It is important to highlight that σ_R is larger than σ_{PLL} as it needs to account for the on-chip process variability suffered at the RO, represented by σ_{RO}^2 .

D. Comparing PLL and RO guardband margins

There exist some differences between the guardband margins used to deal with variability when comparing PLL with RO clocks. First, RO clocks do not consider any jitter margin. In addition, as σ_R is larger than σ_{PLL} , M_{RO}^P margin is larger than M_{pll}^P . This is not surprising as the RO, like any other element in the circuit, suffers process variability, while PLL does not. Nonetheless, M_{RO}^{VT} is typically smaller than M_{pll}^{VT} . This is a consequence of the correlation between the RO and critical paths when referring to voltage and temperature. Considering all sources of on-chip variability, the RO guardband margin should be larger than the PLL on-chip margin on normal PVT and jitter values.

In the example shown in Section VI, the RO and PLL on-chip margins are 11.7% and 9.4% of the nominal delay at the worst-case corner. Nonetheless, the RO clock does not require additional margins to handle global PVT variability, being more effective than PLLs in the typical and average case. Fig. 6 shows the amount of voltage received in the RO when a voltage droop periodically occurs because of circuit activity. In the example of Section VI, the RO has been designed with an on-chip margin of 158ps, whereas the PLL introduces an on-chip margin of 120ps.

In practice, that means that PLL period is 1515ps, as it needs to cover the worst-case scenario all the time. In contrast, the RO period fluctuates along time, being 1553ps in the worst-case scenario. As it can be seen in the figure, voltage fluctuates from 0.860mV up to 0.985mV in the RO path. Consequently, the RO period varies when voltage changes, ranging from 1553ps (worst-case) to 1130ps (best-case), without causing a timing failure. In average, the RO period is 1260ps, which is an 18% faster than the PLL period.

E. Applying derating factors in PLL and RO clocks

Derating factors are used to define the limits of guardband on-chip margins. It is possible to measure the PLL derating factor by rewriting (2):

$$P \geq (D_L^0 - D_C^0) + J + \epsilon(D_L^0 + D_C^0) \quad (9)$$

Note that it is possible to associate the on-chip margin M_{PLL} with the PLL derating factor ϵ by introducing (6) in (9). Thus, the derating factor ϵ for a PLL can be expressed as:

$$\epsilon = \frac{M_{PLL}}{D_L^0 + D_C^0}$$

TABLE I: RO and PLL margins and derating factors.

	Worst	Typical	Best
CP delay (ps)	1427	794	562
RO OCV Margin (ps)	162	73	47
PLL OCV Margin (ps)	155	69	44
ϵ (%)	10.8	8.7	7.8
ϵ' (%)	5.4	4.2	3.8




Fig. 7: Schematic floorplan of 10 AES.

Similarly, we could define ϵ' as the derating factor required when performing timing sign-off in an RO clock. By introducing (8) in (4), we can obtain:

$$\epsilon' = \frac{M_{RO}}{2D_L^0 + M_{RO}}$$

Table I shows the on-chip margins and derating factors used in the circuit implemented in Section VI. As it can be seen, the RO margin is larger for any corner than the PLL on-chip margin, but the derating factors applied at RO timing sign-off are smaller than the ones required by the PLL. Note that the PLL derating factors are different from those used in RO timing closure, as RO clocks are more sensitive to on-chip variability. Because of that, they require a throughout analysis at timing sign-off of the circuit. Nonetheless, if it is not possible at design time to perform such analysis, it is valid to apply the PLL derating factors that are provided by the foundry, as they are conservative from RO's point of view.

V. EXPERIMENTAL METHODOLOGY

In order to validate the timing models presented in previous sections, we have analyzed the variability of a digital circuit using commercial EDA tools. Our experimental circuit comprises 10 instances of an AES encryption module [14] operating in the same clock domain. The AES encryptors were synthesized, placed, and routed using the Synopsys Design Compiler®, the Synopsys IC Compiler®, and a 65nm commercial library. Additional logic has been introduced to enable independently each individual AES encryptor. Fig. 7 shows the layout of the digital circuit evaluated in this analysis. The AES encryptors were organized in a 5×2 matrix. Each AES encryptor occupied $350 \times 350 \mu m^2$, requiring the circuit a total die area of $2488 \times 1285 \mu m^2$.

Metal layers ranging from 2 to 6 were used for routing purposes, while metal layers 9 and 10 distributed power and ground. The power delivery network was designed to keep the IR drop below a 5% of the nominal voltage. To this end, ten power pads and ten ground pads (dotted IO pads) were equidistantly placed through the IO ring. Current flip chips allow IO pads to be placed across the die, and thus, flip chip voltage droops are more uniform than the ones presented in this paper. Despite some of the results enclosed in this paper are related to the usage of peripheral IO pads, we believe that the main conclusions of this study are still valid.

Synopsys VCS[®] was used to run multiple simulations and extract the switching activity of each standard cell of the circuit. Simulations covered different scenarios in which one or several AES encryptors computed in isolation, or other scenarios where all instances operated at the same time.

Synopsys Primerail[®] was used to perform static and dynamic IR drop analysis. The switching activity obtained from simulations allowed to estimate the voltage at any power/ground pin of the standard cells. For a complete voltage noise analysis, package [15] and voltage source models were incorporated for simulation. Voltage source variability has been set to 3% of its nominal value.

Performing an accurate temperature analysis is a difficult task and requires the usage of technological parameters that are typically not available at design time. In previous work, the temperature difference between any pair of nodes was quantified in real silicon for 65nm and 32nm technologies [16], [17]. A maximum temperature gradient of 4°C and 50°C were obtained, respectively. Due to the technology proximity, we took 4°C as a reference—the same difference reported in [16]³. Finally, we also assumed that the maximum temperature reachable in our circuit was 125°C (i.e., the maximum temperature at the worst corner).

We performed SPICE simulations over the most representative critical paths of the circuit using Synopsys HSPICE[®]. The critical paths were extracted using Synopsys Primetime[®]. The parameters of the transistors of each logic gate were modified to model random process variations by adding a random component to the nominal value, according to a Gaussian distribution and using linear interpolation between the values of the parameters at different library corners ($\pm 3\sigma$ was assumed for the best/worst corners).

A. Generating a timing path for an RO clock

Our evaluation methodology replaces a conventional clock source with an RO implemented according to the teachings described in [7]. The design process consists of extracting the delays of the circuit’s critical path in all the available corners using static sign-off tools (in our development framework, Synopsys Primetime[®]). These delays feed a path synthesizer tool that produces a single chain of gates that generates an oscillating signal. The delay of the RO almost fits the delay of the circuit’s critical path in all the analyzed corners, although it is a bit slower. The generated RO comprises an additional margin which goal is to avoid timing violations when the circuit operates in conditions not covered by the corners provided in commercial libraries. The more corners a library provides, the smaller this margin can be, making the RO more precise.

In our analysis the RO is connected to the 10-instance AES encryption module, and it is used as a clock source. Unlike other clock sources such as the PLLs, the clock signal delivered by the RO is strongly correlated with the global and local PVT variability that takes place in the circuit. The aim of the next section is to quantify the effects of such correlation, and to assess whether RO clocks are feasible, beneficial, and robust.

VI. VARIABILITY ANALYSIS IN RING OSCILLATOR CLOCKS


A. Adaptability to voltage noise

Voltage noise has two main components: global and local. Global fluctuations affect uniformly the whole die (or clock domain), whereas local fluctuations affect small regions. We will show that,

³The delay impact of temperature is usually smaller than other components such as voltage or process. For the commercial 65nm library used in this study, the maximum temperature difference barely affects the results presented in the paper.

TABLE II: Characterizing on-chip voltage variations

Nominal V_{dd}	Static IR drop	Max. V_{dd} droop	Avg. V_{dd}
1V	0.044V	0.168V	0.940V



(a) All (10) AES modules active. (b) Bottom-left AES module active.

Fig. 8: Die voltage map: two different scenarios.

despite voltage droops are highly dependent on switching activity, the global component dominates over the local, favoring the argument that ROs are in a privileged position when used as a clock source.

In this characterization, voltage noise has been divided into two components: static and dynamic. Static voltage droops are caused by energy losses in resistors when no circuit activity exists. On the contrary, dynamic voltage droops are produced when circuit current fluctuates due to switching activity and voltage source ripple. Hence, dynamic voltage droops depend on circuit resistance (dynamic IR drop) and inductance (Ldi/dt) [18], [19]. Commonly, the maximum voltage droop is used as a measure to determine the margin required in PLLs to cover voltage variability (see Sect. IV-D).

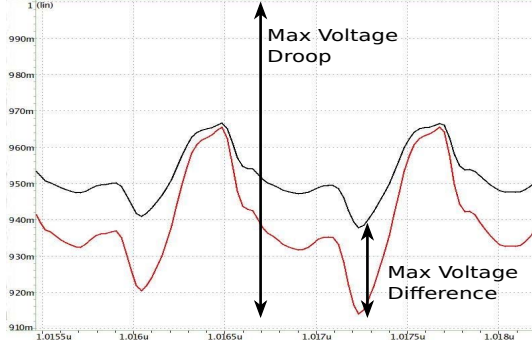
Table II summarizes the behavior of the voltage droop in the evaluated circuit. The largest component of the voltage droop is caused by dynamic variations produced by the switching activity and the voltage source ripple. Moreover, the maximum voltage droop occurs at the center of the die when all AES modules are active. However, in this scenario, the average voltage supply along time is 940mV. This value is important as it determines the average RO performance.

The next figures map the voltage of the circuit in two different situations. On the one hand, Fig. 8a shows the maximum voltage droop achieved at any cell of the circuit when all AES encryptors are active. The voltage droop is more pronounced for the cells that are far away from the IO pins, reaching the highest rates at the center of the die. On the other hand, Fig. 8b shows an extreme case where the switching activity is strongly localized. In that occasion, only the left-bottom AES encryptor is working. However, the maximum droop is much lower than in Fig. 8a, since the total switching activity is minimized. In particular, the maximum droop is reduced from 168mV (all encryptors active) to 59mV (a single AES encryptor active).


Unlike PLLs, RO clocks do not measure the maximum voltage droop to determine the required voltage margin. Instead, they compute the maximum voltage difference between the RO cells and any other cell in the clock domain. Fig. 9a depicts the voltage fluctuations of the two cells of the die that show the largest voltage difference during the experimental simulations. It can be observed that the maximum difference is smaller than the maximum voltage droop. Fig. 9b shows a voltage map of the die at the instant that the voltage difference between cells is the largest (i.e., 29mV).

It is important to understand that the maximum voltage droop does not necessarily generate the maximum voltage difference between cells, given that the voltage droop has a large global component. Note that the power delivery network propagates voltage droops across the die independently from the switching activity.

Table III shows the maximum voltage difference between any two



(a) Voltage difference between two cells.



(b) Die voltage map,

Fig. 9: Analysis of the maximum voltage droop.

TABLE III: Voltage droop analysis in three situations.

Scenario	Max. V droop	Max. V Diff. between cells
All active	168mV	25mV
Bottom-left active	59mV	18mV
Half active	85mV	29mV

cells in three complementary scenarios: (a) when all AES encryptors are active, (b) when only the bottom-left AES encryptor is active, and (c) when all the AES encryptors placed at the left half of the die are active. These three scenarios are a good representation of extreme situations that the circuit may experiment. For instance, the maximum voltage happens when all the AES encryptors are active. In that situation, the maximum droop is 168mV, but the maximum difference between circuit cells is 25mV. Nonetheless, the maximum difference between circuit cells (29mV) occurs when half of the AES encryptors are active. In that scenario, the maximum voltage droop is only 85mV.

B. Finding the best location for an RO clock

The previous voltage analyses and the die voltage maps should help designers to locate a good region to place the RO clock. A near-optimal location (x, y) would be the one that minimizes the maximum difference of voltage along cells and time, i.e.,

$$\min_{(x,y)} [\max_{g,t} (V_{x,y}(t) - V_g(t))] \quad (10)$$

where g ranges over all cells of the circuit and t ranges over time, respectively. In other words, the best location to fit the RO clock




Fig. 10: Maximum voltage difference between an RO cell and any other cell in the die.

TABLE IV: Critical path delays and RO margins.

	Worst SS/0.9V/125°C	Typical TT/1.0V/25°C	Best FF/1.1V/0°C
CP delay (ps)	1427	794	562
RO margin (ps)	162	73	47
% margin	11.4%	9.2%	8.4%

would be the one that maintains the voltage of the RO cells as similar as any other cell of the circuit.

Theoretically, the optimal point to place the RO is usually close to the point that experiments the maximum voltage droop. At that location, the RO cells experiment a larger delay, increasing the clock period. This guarantees enough time slack for all critical paths. The simulations performed are aligned with the previous hypothesis, as they indicate that best location for the RO is the center of the die because it is the place that suffers the worst voltage droop in most scenarios.

Fig. 10 plots the maximum voltage difference for each location of the die according to (10). The figure shows that the best location is around the center of the die, where the maximum voltage difference between a cell placed in the center and any other cell is 14mV.

C. Validating margins for ring oscillator clocks

Table IV reports the critical path delays of our circuit at three different corners: worst, typical, and best. The analysis assumes no on-chip variability. The table also reports the margins introduced when designing the RO and their percentage with respect to each critical path delay.

To estimate the impact of variability on the yield, a set of SPICE simulations with the most significant critical paths of the circuit have been executed. On-chip variability has been modeled with different PVT parameters. For process variability, transistor models have been randomly generated with a probability distribution of V_{th} following a normal distribution $N(V_{th}, \sigma_{V_{th}}^2)$, where $3\sigma_{V_{th}} = 0.4V_{th}$ (ITRS [20]). The maximum temperature difference with respect to the RO has been set to 4°C as in [16] and the maximum voltage difference between the RO and other circuit cells has been set to 14mV, as estimated in Section VI-A. The margin of the RO clock has been calculated for a yield of 97% (see (13) in appendix) and set to 162ps.

Fig. 11 shows the minimum time slack when the circuit is driven by the RO clock. From a total of 500 different circuit configurations and simulations (each one with its own global and local PVT conditions), failures (negative slack) have only been produced in five of them. Note that the number of failing simulations is less than the desired yield, a result that validates our timing model.

D. PLL vs. RO clock performance

Fig. 12 shows the performance of circuit when the clock is driven by a PLL or by an RO clock. Two cases are considered when




Fig. 11: Critical path slack with an RO clock source.




Fig. 12: PLL vs. RO clock performance.

defining the PLL frequency: worst-case sign-off and speed binning. For worst-case sign-off, the frequency determined by the worst corner is assigned to all dies. For speed binning, each die is assigned the best possible frequency according to its process parameters (this is an optimistic assumption for binning).

Three process corners (best, typical, and worst) have been analyzed. When analyzing the PLL, worst local process variation has been assumed on each cell in order to determine the clock period. Similarly, voltage and temperature have been set to their worst value, which is determined by the library. Temperature is set to 125°C and voltage is set to 832mV, which is the minimum achievable voltage (considering global and local variations).

The RO clock period has been measured for the same three process corners. For each process corner, the average RO clock period is shown. To compute the average RO clock period, temperature has been set to 75°C, as it is the average estimated temperature along time. The voltage is set to 940mV, as it has been measured in Sec. VI-A to be the average voltage of the circuit when all AES modules are active.

Fig. 12 also shows the interval between the minimum and maximum RO achievable clock period. These values represent the theoretical minimum and maximum RO clock period if worst/best local PVT conditions remain unalterable along time. The maximum RO clock period is computed when the RO temperature and voltage are at 125°C and 832mV, respectively, and all the cells of the RO suffer the worst on-chip process variation. On the contrary, the minimum RO clock period is achieved when temperature and voltage are at 25°C (nominal voltage), and all cells of the RO have the best on-chip process variation.

Results show that the RO clearly outperforms PLL in all the analyzed corners. When no binning is applied, RO performance is 11.6%, 35.7%, and 53.1% larger than PLL for the worst, typical, and best corners, respectively. If perfect binning takes place for the PLL, the RO still outperforms the PLL. In this case, conservative improvements of 11.6%, 15.2%, and 22.7% are achieved. These benefits would be improved if realistic binning procedures were assumed.

VII. CONCLUSIONS

Guardband margins for covering static and dynamic variability have a significant impact on power and performance. This paper has presented a statistical analysis on the margins required for delay lines when used either for asynchronous circuits with bundled data or synchronous circuits with ring oscillator clocks.

In the future we envision more efforts towards techniques that can instantaneously adapt to dynamic variations at the expense of using clock generators with fluctuating frequencies that can better explore power/performance trade-offs at runtime.

APPENDIX

Given a PVT corner, the delay D of a gate can be modeled by a well-accepted expression:

$$D \propto \frac{V}{\mu(T)(V - V_{th})^\alpha} \quad (11)$$

where V , μ , V_{th} , and T , represent the supply voltage, carrier mobility, voltage threshold, and temperature, respectively, at a given corner [21], [22]. α ($\in [1, 2]$) is a technological parameter related to transistor velocity saturation. When considering on-chip variability, the delay of a gate at time t , denoted as $D(t)$, can be decomposed as follows:

$$D(t) = D^0 + D_{ocv}(t)$$

where $D_{ocv}(t)$ represents the on-chip variations. $D_{ocv}(t)$ can be decomposed into three different terms representing the PVT variations:

$$D_{ocv}(t) = D^P + D^V(t) + D^T(t)$$

Each source of variability can be approximated linearly using a first-order Taylor expansion around the nominal value of the corner:

$$\begin{aligned} D^P &= \left. \frac{\partial D}{\partial V_{th}} \right|_0 (V_{th} - V_{th0}) \\ D^V(t) &= \left. \frac{\partial D}{\partial V} \right|_0 (V(t) - V) \\ D^T(t) &= \left. \frac{\partial D}{\partial T} \right|_0 (T(t) - T) \end{aligned}$$

where $\left. \frac{\partial D}{\partial X} \right|_0$ represents the first derivative of the delay (11) with respect to parameter X , measured at the nominal value of the corner. Under the assumption that V_{th} is a statistical random variable that follows a Gaussian distribution $N(V_{th0}, \sigma_{V_{th}}^2)$, D^P can be modeled as another Gaussian distribution $N(0, \sigma_P^2)$, where:

$$\sigma_P = \left. \frac{\partial D}{\partial V_{th}} \right|_0 \sigma_{V_{th}}$$

Let's define $D_{L_i}(t)$, $D_{C_i}(t)$ as the delay of the launching and capturing paths of the i th critical path with L and C gates, respectively. Both delays follows a normal distribution as defined in (5). That is, $D_L(t) \sim N(D_L^0 + D_L^{VT}(t), \sigma_L^2)$ and $D_C(t) \sim N(D_C^0 + D_C^{VT}(t), \sigma_C^2)$. We assume that capturing and launching paths have a similar variance-to-mean ratio (index of dispersion). Then:

$$\frac{\sigma_L^2}{D_L} \approx \frac{\sigma_C^2}{D_C} \quad (12)$$

Being P and J the PLL clock period and its jitter, the setup constraint for a critical path at a given corner is defined as:

$$S_i = P - J + D_{C_i}(t) - D_{L_i}(t) > 0$$

Given the statistical nature of launching and capturing paths, circuit timing correctness is achieved when the clock period is larger than

any critical path delay with a probability Y (yield). In a circuit with N critical paths,

$$P(\cap_{i=1}^N S_i > 0) \geq Y$$

Assuming that the N critical paths are identical ($S = S_i, \forall i$) and uncorrelated⁴, we could redefine the previous expression as:

$$P(S > 0)^N = \left(1 - P(S < 0)\right)^N \geq Y \quad (13)$$

where $P(S \leq 0)$ is the well-known cumulative distribution function (CDF) of the normal distribution $S \equiv N(\mu_{PLL}, \sigma_{PLL}^2)$ such as:

$$\mu_{PLL} = P - J + D_C^0 + D_C^{VT}(t) - D_L^0 + D_L^{VT}(t) \quad (14)$$

Using (12), we can quantify σ_{PLL}^2 as:

$$\sigma_{PLL}^2 = \sigma_L^2 + \sigma_C^2 = \left(1 + \frac{D_C^0}{D_L^0}\right) \sigma_L^2 \quad (15)$$

After introducing the normal distribution CDF into (13), we obtain:

$$\left(1 - \frac{1}{2} \left(1 + \operatorname{erf}\left(\frac{-\mu_{PLL}}{\sqrt{2}\sigma_{PLL}}\right)\right)\right)^N \geq Y.$$

where erf is the error function (and its complementary is erfc). Last expression can be reordered as:

$$\mu_{PLL} \geq \Psi \sigma_{PLL} \quad (16)$$

where Ψ is a design parameter that depends on the number of critical paths and the desired yield (Y):

$$\Psi = \sqrt{2} \cdot \operatorname{erfc}(2(1 - Y^{1/N})) \quad (17)$$

When introducing (14) and (15) into (16), we obtain the PLL clock period accounting for on-chip variability:

$$P \geq (D_L^0 - D_C^0) + J + \Psi \sigma_{PLL} + D_L^{VT}(t) - D_C^{VT}(t)$$

As it can be seen, the clock period of the PLL can be defined as its nominal period in a given corner ($D_L^0 - D_C^0$) plus some additional margins to cover on-chip PVT and jitter variations.

Similar to previous analyses, it is possible to compute the required RO clock period to cover on-chip variations. Let's define $D_{RO}(t)$ as the delay of the RO defined as in (5). That is, $D_{RO}(t) \sim N(D_{RO}^0 + D_{RO}^{VT}(t), \sigma_{RO}^2)$. In this case, the setup constraint for any critical path at a given corner is:

$$R_i = D_{RO}(t) + D_{C_i}(t) - D_{L_i}(t) > 0 \quad (18)$$

Analogous to the PLL, we can assume that all critical paths are identical and independent ($R = R_i, \forall i$). Thus, we can define R as a normal distribution $R \sim N(\mu_R, \sigma_R^2)$, where:

$$\mu_R = D_{RO}^0 + D_{RO}^{VT}(t) + D_C^0 + D_C^{VT}(t) - D_L^0 - D_L^{VT}(t) \quad (19)$$

Using (12), we can quantify σ_R^2 as:

$$\sigma_R^2 = \sigma_L^2 + \sigma_C^2 + \sigma_{RO}^2 = \left(1 + \frac{D_C^0}{D_L^0} + \frac{D_{RO}^0}{D_L^0}\right) \sigma_L^2 \quad (20)$$

By performing the same mathematical analysis as in the PLL approach, we can obtain the RO clock period considering on-chip variability:

$$D_{RO}^0 \geq (D_L^0 - D_C^0) + \Psi \sigma_R + D_L^{VT}(t) - D_{RO}^{VT}(t) - D_C^{VT}(t) \quad (21)$$

Like in PLLs, the RO clock period can be defined as its nominal period in a given corner ($D_L^0 - D_C^0$) plus some additional margins used to only cover PVT variations.

⁴This assumption is not completely true as part of the OCV is common to all critical paths. However, this assumption is conservative as it generates an upperbound for the clock source period

REFERENCES

- [1] J. Bhasker and R. Chadha, *Static Timing Analysis for Nanometer Designs*. Springer-Verlag, 2009.
- [2] G. Birtwistle and K. Stevens, "Modelling mixed 4phase pipelines: Structures and patterns," in *20th IEEE Int. Symp. on Asynchronous Circuits and Systems (ASYNC)*, 2014, pp. 27–36.
- [3] A. E. Sjogren and C. J. Myers, "Interfacing synchronous and asynchronous modules within a high-speed pipeline," *IEEE Transactions on VLSI Systems*, vol. 8, no. 5, pp. 573–583, 2000.
- [4] N. Toosizadeh, S. G. Zaky, and J. Zhu, "Varipipe: low-overhead variable-clock synchronous pipelines," in *Proc. International Conf. Computer Design (ICCD)*, 2009, pp. 117–124.
- [5] M. Garg, C. Chai, and J. Bridges, "Adaptive clock generators systems and methods," US Patent application 2011/0 140 752 A1, Jun. 16, 2011.
- [6] K. Bollapalli and T. Raja, "Clock generation circuit that tracks critical path across process, voltage and temperature variation," US Patent application 2015/0 008 987 A1, Jan. 8, 2015.
- [7] J. Cortadella, L. Lavagno, P. López, M. Lupon, A. Moreno, A. Roca, and S. S. Sapatnekar, "Reactive clocks with variability-tracking jitter," in *Proc. International Conf. Computer Design (ICCD)*, Oct. 2015, pp. 540–547.
- [8] K. Chae and S. Mukhopadhyay, "All-digital adaptive clocking to tolerate transient supply noise in a low-voltage operation," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 59, no. 12, pp. 893–897, Dec. 2012.
- [9] C. Hernandez, A. Roca, F. Silla, J. Flich, and J. Duato, "On the impact of within-die process variation in GALS-Based NoC performance," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 31, no. 2, pp. 294–307, Feb. 2012.
- [10] Y. Cao and L. Clark, "Mapping statistical process variations toward circuit performance variability: An analytical modeling approach," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 26, no. 10, pp. 1866–1873, Oct. 2007.
- [11] M. Shafique, S. Garg, J. Henkel, and D. Marculescu, "The EDA challenges in the dark silicon era: Temperature, reliability, and variability perspectives," in *Proc. of the 51st Annual Design Automation Conference (DAC)*, 2014, pp. 185–185.
- [12] P. Chaparro, J. Gonzalez, G. Magklis, Q. Cai, and A. Gonzalez, "Understanding the thermal implications of multi-core architectures," *IEEE Transactions on Parallel and Distributed Systems*, vol. 18, no. 8, pp. 1055–1065, Aug. 2007.
- [13] T. Sato, J. Ichimiya, N. Ono, K. Hachiya, and M. Hashimoto, "On-chip thermal gradient analysis and temperature flattening for SoC design," in *Proc. ACM/IEEE Design Automation Conference*, 2005, pp. 1074–1077.
- [14] M. Litochevski and L. Dongjun, "High throughput and low area AES," 2012. [Online]. Available: http://opencores.org/project,aes_highthroughput_lowarea
- [15] N. Devnani and E. Murray, "Power Supply Control from PCB to Chip Core," Avago Technologies, White Paper, Mar. 2010.
- [16] S. Sharifi and T. S. Rosing, "Accurate direct and indirect on-chip temperature sensing for efficient dynamic thermal management," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 29, no. 10, pp. 1586–1599, 2010.
- [17] E. K. Ardestani, F. J. Mesa-Martinez, G. Southern, E. Ebrahimi, and J. Renau, "Sampling in thermal simulation of processors: Measurement, characterization, and evaluation," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 32, no. 8, pp. 1187–1200, 2013.
- [18] S. K. Nithin, G. Shanmugam, and S. Chandrasekar, "Dynamic voltage (IR) drop analysis and design closure: Issues and challenges," in *International Symposium on Quality Electronic Design (ISQED)*, 2010, pp. 611–617.
- [19] M. Gupta, J. Oatley, R. Joseph, G.-Y. Wei, and D. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *Design, Automation Test in Europe Conference Exhibition, 2007. DATE '07*, April 2007, pp. 1–6.
- [20] "International technology roadmap for semiconductors (2009)," Tech. Rep. [Online]. Available: <http://www.itrs.net/reports.html>
- [21] A. Srivastava, D. Sylvester, and D. Blaauw, *Statistical Analysis and Optimization for VLSI: Timing and Power*. Springer US, 2005.
- [22] T. Sakurai and A. Newton, "Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas," *IEEE Journal of Solid-State Circuits*, vol. 25, no. 2, pp. 584–594, Apr 1990.