

Robustness to Voltage Noise with Ring Oscillator Clocks

Lucas Machado, Antoni Roca, and Jordi Cortadella, *Fellow, IEEE*

Abstract—Voltage noise is the main source of dynamic variability in integrated circuits and a major concern for the design of Power Delivery Networks (PDNs). Lower supply voltages were made possible with technology scaling, but power density was also increased. Consequently, power integrity became a key factor in the design of reliable high performance circuits.

Ring Oscillator Clocks (ROCs) have been proposed as an alternative to mitigate the negative effects of voltage noise. However, the effectiveness highly depends on the design parameters of the PDN, power consumption patterns and spatial locality of the ROC within the clock domain. This paper analyzes the impact of the PDN parameters and ROC location on the voltage noise and the robustness achieved by using ROCs. The capability of reacting instantaneously to large voltage droops makes ROCs an attractive solution, which also allows to relax the constraints required for the PDN design. The experiments show that up to 83% of the margins for voltage noise and up to 27% of the total leakage power can be reduced by using ROCs. Also, PDN simplifications are possible, with fewer power interconnections or package decaps of lower quality. Tolerance to voltage noise and related benefits can be increased with multiple ROCs.

Index Terms—ring oscillators; voltage noise; adaptive clocking; power delivery network; multiple clock domains

I. INTRODUCTION

THE ESTIMATION of the path delays and their variability is critical for the reliability of digital circuits. In order to define a robust clock period, it is necessary to consider all conditions that may shift and affect the delay of every circuit path, such as the manufacturing process, the supply voltage, and the temperature (PVT). Static offsets of these conditions are estimated at design time and taken into account by adding guard band margins to the nominal clock period. Nevertheless, dynamic shifts are hard to predict and excessively conservative margins are often added to prevent failures.

Voltage noise is the main source of dynamic variability, and mitigating this noise is an arduous task that may have an significant impact on power, performance, and area. The main components of voltage drops are resistive and inductive [1]:

$$\Delta V = R \cdot i(t) + L \cdot \frac{di}{dt}. \quad (1)$$

IR drops (static and dynamic) are produced by the parasitic resistance of the Power Delivery Network (PDN).

This work was performed with the support of CNPq, Conselho Nacional de Desenvolvimento Científico e Tecnológico - Brasil, and has been partially supported by funds from the Spanish Ministry for Economy and Competitiveness and the European Union (FEDER funds) under grant TIN2017-86727-C2-1-R, and the Generalitat de Catalunya (2017 SGR 786).

The authors are with the Computer Science Department, Universitat Politècnica de Catalunya, Barcelona 08034, Spain (e-mail: lmachado@cs.upc.edu; aroca@cs.upc.edu; jordi.cortadella@upc.edu).

Inductive noise is mainly caused by current differences, associated with the switching activity of the chip. Clock and power gating are low-power techniques that can unintentionally produce large voltage droops. When many devices are simultaneously activated, a large di/dt is originated. If that situation is periodically repeated and aligned with a resonant frequency of the PDN, large voltage swings may appear, exceeding the ones tolerated by the system.

Augmenting the clock period offers more robustness against these changes in the operating conditions, but this comes at the expense of reducing performance. Another solution is to increase the amount of decoupling capacitors (decaps) [1], [2]. Voltage noise is mitigated when the system has a larger on-chip and off-chip capacitance. Unfortunately, the additional decaps imply an increase in area and leakage power, and variations that exceed the defined margins cannot be fully eliminated.

In [3], the use of integrated voltage regulators is investigated, quantifying the penalties in area and power for the voltage noise reductions obtained. Other proposals include: improving the PDN impedance, which requires adaptations for each particular circuit; static and dynamic voltage margining, which result in higher power consumption; and performance throttling and stalling [4], [5], which require high-quality voltage sensors, with additional area and power. Adaptive clocking [6]–[10] seems to be a promising solution with low overhead, but with its efficiency limited by the characteristics of its voltage sensors and clock generators.

Ring Oscillator Clocks (ROCs) [11], [12] can be considered an adaptive clocking proposal, which takes into account all sources of variability, voltage noise included. If the ROC is correctly designed, then a strong correlation can be achieved between the clock period and the delay of the critical paths. Considering that the ROC and the critical paths are exposed to the same sources of variability, the clock generator adapts immediately to the circuit demands.

Unfortunately, voltage fluctuations are not uniform across the die. Two distant points in the same die may have different voltage levels. This unsteady behavior raises some questions:

- How the global and local portions of voltage noise affect the performance when using ROCs?
- Is it possible to relax the PDN design by using ROCs?
- What is the relation between the required timing margins for an ROC and the size of its clock domain?
- Where to locate the ROC within a clock domain?

Voltage noise analysis has been focused on estimating the global worst-case and deriving the timing margins required [13]. For example, if the nominal voltage is 1V and the minimum voltage estimated is 0.85V, then a circuit with a rigid clock must consider a variation of 150mV for the clock period.




Fig. 1. Power delivery network (PDN) model with (a) off-chip and (b)(c) on-chip parasitics.

For ROCs, the key value is the largest *differential voltage* between the ROC and the critical path [14]. If the voltage at the ROC is 0.9V when it is 0.85V at the critical path, then the clock period margins should cover only the difference: 50mV.

This paper presents the following contributions:

- A conservative analysis on the benefits of using ROCs when facing problems related to voltage noise.
- Different activity patterns and ROC locations are explored to understand the effects on voltage variations.
- Voltage noise is generated using an activity frequency with very high impedance, in order to differentiate the global and local voltage variations.
- Simplifications of the PDN are investigated, reducing its cost and relaxing the design parameters to provide a stable supply voltage.

Notice that the number of ROCs and their placement have no significant influence on voltage noise. The goal is not to mitigate voltage noise, but to reduce voltage variations between the critical paths and the clock source, thus reducing the guard band margins required. Also, the results presented are based on simulations using models. This approach provides a flexible and easily reproducible method, but also illustrative enough to cover a broad range of potential applications.

The paper is organized as follows. A PDN description and a review of voltage noise sources is presented in Section II. Section III provides an overview about ROCs, with details about the jitter and ROC design. Section IV depicts the PDN and delay models, and the performance metrics used to compare ROCs and PLLs. Voltage locality is introduced by multiple activity patterns in Section V, analyzing the impact of voltage difference for ROCs, and the positive impact of increasing the number of clock domains. PDN modifications are evaluated in Section VI, regarding the on-chip capacitance, the placement of the power interconnections (bumps), and the package parasitics. Section VII provides a discussion on the experiments, and the advantages and disadvantages of using ROCs as the clock source. Section VIII concludes the paper.

II. VOLTAGE NOISE

The PDN delivers the power and ground voltages to all devices of a design. Fig. 1 depicts the PDN model with its components: voltage regulator (VRM), board (PCB), package (PKG), the connection bumps and the on-chip power grid [1]. The power distribution has parasitic inductances, resistances and capacitances, which can be modeled as depicted in Fig. 1. Decaps are placed at all levels of the PDN in order to




Fig. 2. (a) The frequency response of a typical PDN, and (b) the voltage droops generated by a single current spike.




Fig. 3. Voltage droops generated by periodical current differences at (a) low and (b) high impedance frequencies.

reduce voltage fluctuations. The parasitics of the capacitors are also known as equivalent series inductance (ESL) and equivalent series resistance (ESR). The PDN parasitics interact with each other, forming LC circuits with different resonance frequencies, which are responsible for the *voltage droops*.

The circuit composed by the on-chip capacitance and the power bumps inductance (L_{bump}) generates the *first droop*, which typically produces the largest voltage noise and has a resonance frequency of 100-400MHz [15]. The second and third droops usually have much lower resonance frequencies and amplitudes than the first droop.

Fig. 2(a) depicts the frequency response of a typical PDN, showing the impedance and the resonance frequency for the first, second and third droops. The supply voltage behavior illustrated in Fig. 2(b) is observed when a current spike is requested for this PDN: the first droop causes fast and large voltage swings in the order of nanoseconds; then the voltage continues to fluctuate due to second and third droops, until it becomes stable after a few microseconds.

Voltage noise is minimized when the activity takes place at frequencies with low impedance associated. Fig. 3(a) shows the supply noise and the clock signal for a circuit operating

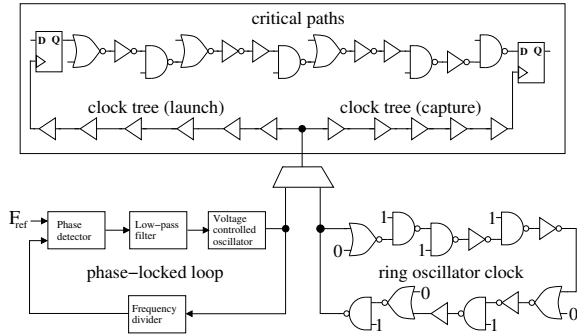


Fig. 4. Synchronous circuit with a PLL or an ROC as the clock source.




Fig. 5. Clock signal generation in the presence of voltage noise.

at 1GHz (low impedance), with voltage swings of $\pm 10\%$. The clock can be set to the frequency of the first droop to emulate the worst-case voltage noise, as seen in Fig. 3(b). In this case, the voltage noise amplitude goes up to 20%. Such large fluctuations are also known as *voltage emergencies*.

For designers, it is difficult to anticipate whether voltage emergencies will actually show up in their designs. Very often, they just *conjecture* that these events will not happen, without a full guarantee of safety. Note that a circuit designed for an application can be used for other purposes, with changes in the operating frequency, the submodules activated, the firmware, and the packaging. In this context, it is very difficult to predict the presence of such large voltage fluctuations. Still, if a voltage emergency occurs, then a timing failure may be originated and the circuit operation becomes unpredictable.

III. RING OSCILLATOR CLOCKS

Jitter and other clock uncertainties are typically covered by increasing the timing margins of the clock period, degrading circuit performance. For that reason, the use of ROCs as clock sources has been discarded, as they have a high jitter caused by their sensitivity to the various sources of variability. Therefore, rigid clock generators with low-jitter, such as Phase-Locked Loops (PLLs), became the *de facto* clock source paradigm.

Fig. 4 shows a synchronous circuit¹ fed by a PLL or by an ROC, depending on the selection of a mux. Fig. 5 illustrates the clock signals generated by the PLL and the ROC when

¹ The unconnected inputs of the circuit are connected to the rest of the design. This is a typical representation of critical paths in EDA tools, in which only the gates belonging to the path of largest delay are shown, and the rest of the circuit is omitted.

a voltage droop occurs. The clock period of the PLL is not affected by the voltage variations, as it is designed to support these fluctuations and deliver a low-jitter clock. Still, the circuit paths have a different behavior: their delay increases when voltage decreases. If the PLL is selected as the clock source, then timing failures are avoided by adding margins that consider the delay of the critical paths at the *minimum voltage*.

Differently, the period of the ROC is affected by the voltage, as seen in Fig. 5. In [16], it is shown that the power supply is dominant source of jitter for ring oscillators. Recent studies demonstrate that the jitter of ROCs is highly correlated with the delay variability of the circuit paths [11], [12].

In other words, as the ROC and the circuit paths are composed of similar gates, the PVT variations affect them likewise. If the circuit becomes slower due to a voltage droop or a temperature increase, then the frequency of the ROC slows down as well. This correlation enables the reduction of timing margins, and hence improve circuit average performance or reduce power [11], [12]. Note that the PLL jitter does not have similar correlation with circuit delay.

Obviously, there is not an exact match between the delay of the critical paths and the period of an ROC. Standard cells have different responses to PVT variations. Additionally, there are voltage and temperature differences across the chip, and process variability is not identical throughout the die [17]. All these factors must be considered in the design of an ROC.

In this work, a rigid clock source (PLL) is compared with an ROC, that is implemented using to the guidelines described in [11]. In summary, the design of an ROC consists of:

- Delay extraction of the critical paths of the circuit.
- Use the extracted delay to generate a path of library gates with similar delay behavior, considering all corners.
- Assemble these gates in a ring to create a clock.

The delay extraction is performed for all PVT corners available in the technology, using Static Timing Analysis (STA) tools. The extracted delays are the input to a path synthesizer tool, which produces a single chain of standard cells that is able to produce an oscillating signal, i.e. a clock. Note that the design of an ROC depends only on the manufacturing technology and the variability behavior. Hence, it is agnostic to the characteristics of the chip or the package.

IV. MODELS AND METRICS

A. PDN model

The chip-grid presented in [18] represents an SoC with four cores of Pentium 4 and it is used as the PDN model in this work. The PDN components illustrated in Fig. 1 are described in SPICE netlists using the values of Table I. In Fig. 1(a), the values R_{die} , C_{die} and I_{die} represent equivalent values of the on-chip PDN. For example, C_{die} depends on the amount of on-chip decaps, and I_{die} depends on the activity pattern. As external regulators typically do not regulate high frequency variations, the voltage regulator module (VRM) is modeled as a fixed voltage source delivering 1V at the power bumps.

The on-chip power distribution is modeled with a 12×12 grid [13], as seen in Fig. 1(b). Both the power and ground

TABLE I
PDN PARAMETERS

Param.	Value	Param.	Value	Param.	Value
R_{pcb}	0.094 m Ω	L_{pcb}	21 pH	V_{vrm}	1 V
R_{cpcb}	0.17 m Ω	L_{cpcb}	1 pH	C_{pcb}	240 μ F
R_{pkg}	1 m Ω	L_{pkg}	120 pH	C_{pkg}	26 μ F
R_{cpkg}	0.54 m Ω	L_{cpkg}	5.61 pH	C_{ckt}	120 pF
R_{bump}	40 m Ω	L_{bump}	72 pH	I_{ckt}	195 mA
R_{grid}	50 m Ω	L_{grid}	5.6 fH	-	-




Fig. 6. Current waveform and impedance response (200nF of on-chip decaps).

networks are considered in the model, with a V_{DD} or a V_{SS} bump connected at every grid point.

Each point in the grid models a portion of the circuit, with an *intrinsic decoupling capacitance* and a current source emulating the circuit operation, with rise, high and fall times set to 5%, 45%, and 5%, respectively (see Fig. 6(a)).

Additionally, a decoupling capacitor is added at each point. Note that spreading the decaps uniformly is the best placement in order to reduce voltage fluctuations, considering a similar power consumption throughout the die [1], [13]. The frequency response of Fig. 6(b) is observed at any point of grid, considering a total of 200nF of on-chip decoupling capacitance.

B. Delay model

A simplification of the gate delay formulation was proposed in [19], which is still widely accepted. This model proposes the delay variation with the voltage based on the threshold voltage (V_{th}) and a technology fitting value α in the range of 1-2. Details on how to calculate the α for different technologies can be obtained in [19]. Notice that the model was defined for a single gate, but the relationship between delay and voltage holds for a path composed of multiple gates. Considering that V_{th} , α and k have small variation with the voltage, then it is possible to calculate the constant k in (2) and have the path delay based on the supply voltage.

$$td(V_{DD}) = \frac{k \cdot V_{DD}}{(V_{DD} - V_{th})^\alpha} \quad (2)$$

A 65nm commercial library with nominal voltage of 1V is used as reference. The average V_{th} of all combinational cells of the library is 0.36V for 75°C, and 0.4V for 125°C. A typical value of α is 1.3 [20], and this parameter is closer to 1 for more advanced technologies. Generally, $\pm 10\%$ offsets are defined for the voltage swings during STA. Therefore, the critical path at $V_{DD} = 0.9V$ must have a maximum delay of 1ns, considering a clock source of 1GHz.

Fig. 7 shows the path delay curves with the k values calculated using (2), with $V_{DD} = 0.9V$, $td = 1ns$, $\alpha = [1.0, 1.3]$




Fig. 7. Path delay given by (2), with $td = 1ns$ and $V_{DD}=0.9V$.




Fig. 8. ROC placement strategies with different number of clock domains.

and $V_{th} = [0.36, 0.4]$. For a *conservative analysis*, $V_{th} = 0.4V$ and $\alpha = 1.3$ are selected, indicating larger delay variations for smaller voltage differences, with $k = 0.45$.

C. Performance Metric

In this work, the required timing margin is used to compare the performance of the ROC and the PLL. For the PLL, the margin is the difference between the critical path delay at the nominal voltage (V_{nom}) and at the minimum voltage (V_{min}):

$$margin_{PLL} \geq td(V_{nom}) - td(V_{min}). \quad (3)$$

The design of an ROC must consider the delay behavior of Fig. 7 in order to keep the clock period larger than the delay of the critical paths for any given voltage. For simplification of the analysis, the delay behavior of the ROC and the critical paths are both given by (2) with the same parameters. Still, if the ROC has a larger V_{th} than the critical path, then margins may be smaller.

In order to perform a *conservative analysis* of the required timing margins for the ROC, the following claims are made:

- The voltage at the ROC is always higher than at the critical path.
- The critical path is placed at the point with the largest voltage difference with respect to the ROC.
- The largest voltage difference happens at the minimum voltage, as delay variations are larger for lower voltages.
- Positive effects due to the clock distribution are not taken into account, such as clock-data compensation [15].

Thus, the margin for the ROC is given by (4), which is the difference between the critical path delay at the minimum voltage and the ROC period at the largest voltage difference.

$$margin_{ROC} \geq td(V_{min} + \max(\Delta V_{DD})) - td(V_{min}) \quad (4)$$

The PLL margin is required regardless of its placement, as the clock period must consider the critical path delay at V_{min} .




Fig. 9. Patterns determining the grid points that are active.

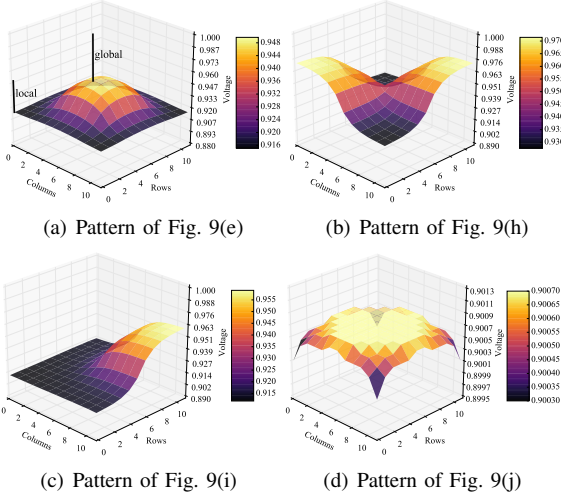


Fig. 10. Voltage distribution for some activity patterns of Fig. 9.

But the ROC margin varies with its location, as the voltage difference is smaller between points closer to each other.

Fig. 8 depicts the three placement strategies analyzed, with circles at ROC locations and squares around the grid points on the same clock domain: one ROC at the center of the chip; 4 ROCs, with one at the center of each processor core; and 16 ROCs uniformly distributed. Additionally, one ROC placed at an arbitrary grid point is analyzed, reporting the placement that requires the largest margin.

Notice that 16 ROCs would require additional synchronization between the clock domains, with an overhead in performance and power not investigated. Therefore, this case is reported but its results are not compared with the PLL.

V. VOLTAGE LOCALITY ANALYSIS

The different patterns depicted in Fig. 9 are proposed to stimulate voltage variations across the die. The dark areas represent the portions of the chip that are active. The parts that are not active are modeled with constant current sources.

Fig. 10 shows the global and local effects due to some of the proposed patterns. These images show the voltage levels at each grid point when the minimum voltage is reached in the simulation. The pattern in Fig. 9(j) generates the lowest voltage, reaching a maximum current of 28A. An on-chip decoupling capacitance of 200nF is necessary to keep the voltage swings within $\pm 10\%$ for this activity pattern, considering an activity frequency of 1GHz.

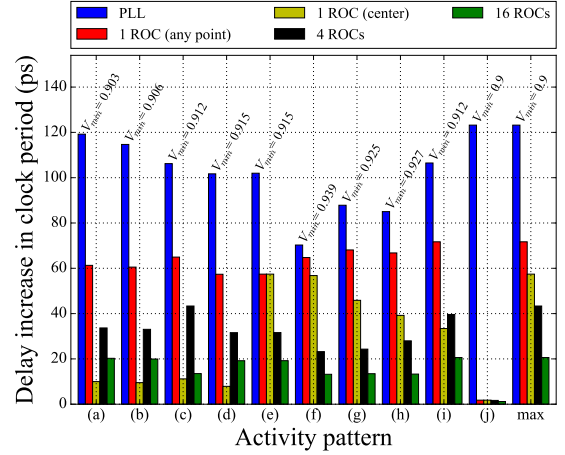


Fig. 11. Delay increase in the clock period for each activity pattern (200nF decaps, activity at 1GHz).

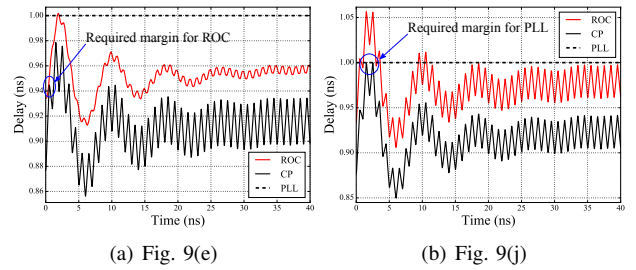


Fig. 12. Critical path (CP) delay, and the clock period of the PLL and the ROC, for the activity patterns of Fig. 9(e) and Fig. 9(j).

Using the grid model with 200nF of on-chip decoupling capacitance, the activity patterns of Fig. 9 are simulated with Synopsys HSPICE[®] for 50 clock cycles at 125°C, gathering the minimum voltage (V_{\min}) of all grid points, and the maximum voltage difference between any two points in the grid (ΔV_{DD}). Two cases are analyzed: a typical voltage noise, generated by the designed clock period, with low impedance (1GHz); and the worst-case voltage noise, caused by an activity frequency with very high impedance (first droop).

A. Typical voltage noise

Fig. 11 is generated with the voltage data gathered, using (3) and (4) to obtain the required margins. The delay increase for the PLL is proportional to the number of active points, which is related with the total current and the minimum voltage. In the worst case for the PLL, $V_{\min} = 0.9V$ and the delay increase is 123ps.

For the ROC, the delay increase is related with the voltage difference between the ROC and the critical path (CP). Considering all activity patterns, delay increase is 71ps if the ROC is placed at any grid point, and 57ps if it is at the center.

In Fig. 12, the activity patterns of Fig. 9(e) and Fig. 9(j) are simulated, keeping track of the voltage at the center of the grid and at the point with the largest voltage difference. A 57ps margin is added to the ROC period, as it is placed at the grid center. Fig. 12(a) depicts the worst case for the ROC, whereas Fig. 12(b) shows the largest delay of the critical path. Notice that the first and second voltage droops are present. As these effects are global, they affect the critical path and the

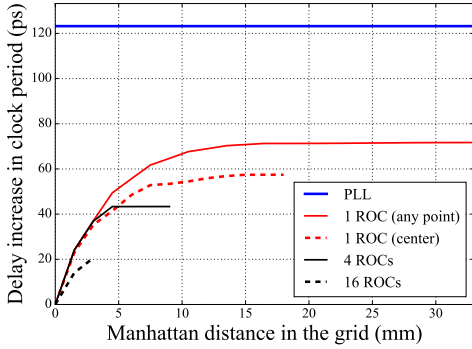


Fig. 13. Largest delay increase vs. the distance between the ROC and the critical path (200nF decaps, activity at 1GHz).

ROC similarly. Therefore, ROCs enable a 53% better average performance for the same level of voltage noise robustness.

Fig. 13 depicts the largest delay increase for each distance between any two grid points, considering all activity patterns. As expected, the delay is smaller if the critical path is closer to the ROC. This graph shows that a *trade-off* is possible between performance and the number of clock domains. The required delay is reduced to 43ps with 4 ROCs, and to 20ps with 16 ROC domains.

B. Worst-case voltage noise

The delay increase shown in Fig. 11 is required for a typical voltage noise, but larger voltage droops may happen if the activity frequency has a high impedance associated, as seen in Section II.

The first droop frequency of the grid model with 200nF of on-chip decaps is 125MHz. As a result, the voltage noise is amplified if a large current difference happens every 8 clock cycles, considering a clock source of 1GHz. In order to evaluate this phenomenon, the previous experiment is repeated with the current sources operating at 125MHz.

Fig. 14 depicts the delay increase for each activity pattern in this case. As expected, the voltage noise is boosted due to the high impedance, and the delay increase required for the PLL is 1.5ns. Therefore, if worst-case voltage noise is considered, a design with a PLL *cannot* operate at 1GHz with this PDN.

The ROC takes advantage of the global characteristic of voltage droops, and the delay increase is 435ps if it is placed at an arbitrary point, and 260ps if placed at the center. Hence, it is possible to reduce the delay in 83%, without increasing the number of clock domains. Also, it is possible to reduce margins by increasing ROC domains, with a delay increase of 151ps with 16 ROCs, which is comparable to the delay increase of the PLL for a typical voltage noise.

VI. RELAXING PDN PARAMETERS

The design of the PDN is a difficult task that must take into account the circuit specification, the decaps and parasitics. It is necessary to adjust the characteristics of the PDN in order to avoid undesired voltage droops, which may happen when the switching activity is aligned with a resonance frequency.

This section shows how the robustness of ROCs contributes to *relax* the PDN design constraints, given the tolerance to handle global voltage variations. Three parameters are analyzed:

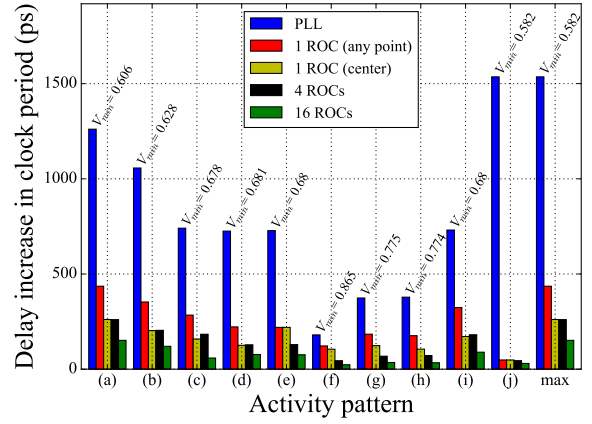


Fig. 14. Delay increase in the clock period for each activity pattern, for the PLL and the ROC (200nF decaps, activity at first droop frequency).




Fig. 15. Impedance response of the PDN with 200nF, 300nF, 400nF and 500nF of on-chip decoupling capacitance.

on-chip decoupling capacitance, the number and placement of power bumps, and the parasitics of the package decaps.

A. On-chip decoupling capacitance

Fig. 15 depicts the impedance response of the PDN with 200nF, 300nF, 400nF, and 500nF of on-chip decoupling capacitance. Notice that adding decaps to the chip has a linear increase in area and power, whereas the impedance reduction is important but not linear.

The voltage noise reduction obtained by increasing the on-chip decaps has a direct impact to the performance, as seen in Fig. 16(a). All activity patterns of Fig. 9 are simulated for the different amounts of on-chip decaps, with activity at 1GHz.

The behavior is similar with the activity aligned with the first droop frequency, with significant margin reductions shown in Fig. 16(b). Notice that the first droop frequency varies with the amount of on-chip capacitance (see Fig. 15). The lower impedance is one of the reasons for the performance improvements seen in Fig. 16. Still, there is a saturation on the positive effect of adding decaps.

Generally, on-chip decaps do not imply an increase in area, given that the core utilization for standard cells is typically 70-90%, and decaps are placed in the white space. Still, the leakage power consumption of the decaps is important. As ROCs support larger voltage fluctuations with lower margins than static clocks, it is possible to reduce the amount of decaps and leakage power without degrading performance.

Leakage power can be modeled by expression (5), where P_{std}^{sq} and P_{dec}^{sq} are the leakage power per area of the standard

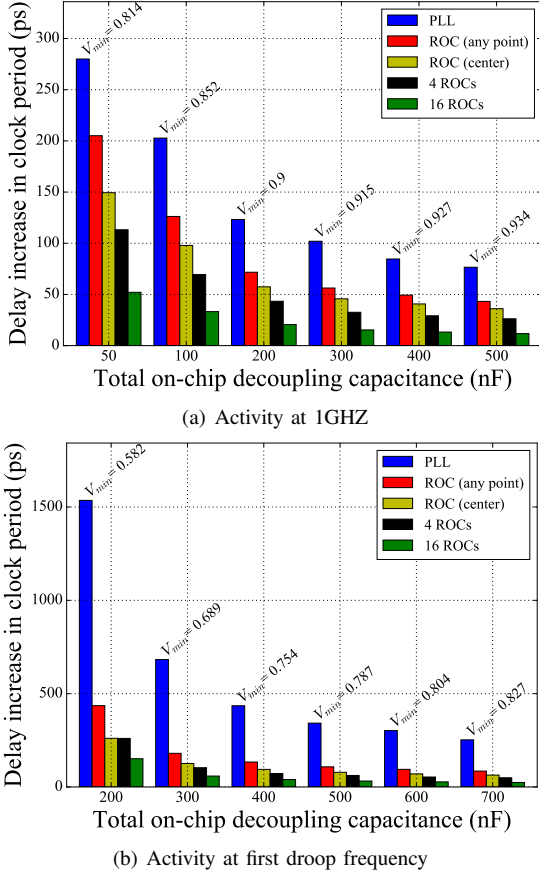


Fig. 16. Delay increase for the PLL and ROC, with different amounts of on-chip decoupling capacitance.

cells and the decaps, respectively. The area occupied by standard cells and decaps are A_{std} and A_{dec} , respectively.


$$P_{leak} = P_{std}^{sq} \cdot A_{std} + P_{dec}^{sq} \cdot A_{dec} \quad (5)$$

The leakage savings are estimated by using the parameters of a commercial 65nm library. The least leaky decap cell is selected, with a capacitance per area of 6nF/mm² and leakage power consumption of 2.5mW/nF. Hence, the leakage power per area of decaps is defined as 15mW/mm².

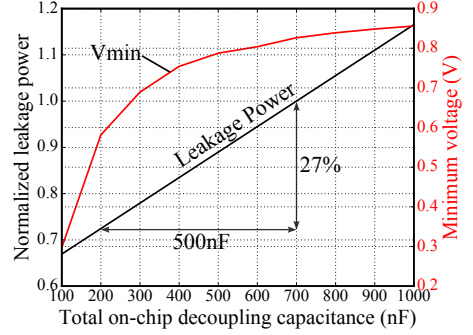
For standard cells, leakage per area is estimated based on a design with a representative mix of combinational gates and flip-flops [21], obtaining 20.9mW/mm². These values are *conservative*, as decaps typically have a larger average leakage power than standard cells. For the area ratio, it is assumed that 200nF represent 20% of the core area (utilization of 80%).

Fig. 17(a) shows the leakage power and the minimum voltage for different amounts of on-chip decaps, for typical voltage noise. Leakage power is normalized with respect to 200nF. Considering the margins seen in Fig. 16(a), it is possible to reduce up to 150nF in decaps without degrading performance, by using ROCs. This reduction represents 11% of the total leakage power of the design.

Similarly, Fig. 17(b) depicts the leakage and minimum voltage, but for the worst-case voltage noise produced by activity aligned with first droop frequency. In this case, leakage power is normalized with respect to 700nF. Considering the data in Fig. 16(b), it is possible to have 200nF decaps instead



(a) Activity at 1GHz



(b) Activity at first droop frequency

Fig. 17. Normalized leakage power and minimum voltage for different amounts of on-chip decoupling capacitance.




Fig. 18. Different power bumps placement strategies (a VDD connection is a black circle, and VSS connection is a white circle).

of 700nF, without degrading average performance, with ROCs. Removing 500nF means a reduction of 27% in the total design leakage power consumption. Furthermore, if 200nF occupy all the white space, then 700nF entail a non-negligible area increase that can be simply avoided by using ROCs.

B. Power interconnections

The amount (and placement) of power bumps is another characteristic that influences voltage locality. The experiments in previous sections were performed with 72 pairs of VDD/VSS bumps uniformly distributed (see Fig. 18(a)). This placement minimizes the impedance between the chip and the package [1], and any grid point has practically the same impedance response. As seen in Fig. 10, such placement reduces significantly the voltage differences across the die.

This section considers different bump placements in the grid model with 200nF, for typical voltage noise (activity at 1GHz). Two additional placements are analyzed: 36 VDD/VSS pairs uniformly distributed, as in Fig. 18(b); and 40 VDD/VSS pairs placed in the borders (similar to wire bonding), depicted in Fig. 18(c). These placements affect the impedance response




Fig. 19. Impedance response of all grid points (200nF of on-chip capacitance) with (a) 36 bumps distributed and (b) 40 bumps in the borders.

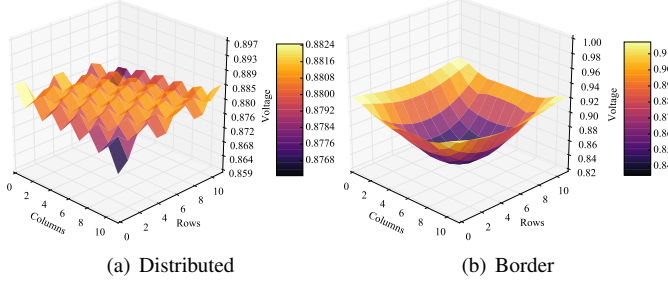


Fig. 20. Voltage distribution for activity pattern of Fig. 9(j) with (a) 36 VDD/VSS bumps distributed ($V_{\min} = 0.872V$) and (b) 40 VDD/VSS bumps in the borders ($V_{\min} = 0.837V$).

across the die, as observed in Fig. 19. Such configurations also have a huge impact in the power distribifon (see Fig. 20).

All activity patterns are simulated, producing the results of Fig. 21. As the impedance is higher, the minimum voltage is lower, indicating larger margins. Also, ROC margins have an important increase, due to larger voltage differences. Still, it is possible to reduce the bumps configuration using ROCs, with same or better performance of a PLL.

With bumps placed in the border, it is possible to take further advantage of ROC characteristics by placing it at the center. In this case, the ROC will typically have the lowest voltage in the die, enabling a higher average performance.

C. Package decoupling capacitance

The design of the board and the package is a key factor in the quality of the supply voltage at the chip devices. Small




Fig. 21. Required margins for the PLL and ROC with different bump placements (200nF of decoupling capacitance, activity at 1GHz).




Fig. 22. Impedance responses with 500nF of on-chip capacitance and different package decap parasitics.




Fig. 23. Required margins for the PLL and ROC with the different package decap parasitics (500nF of on-chip decaps, activity at first droop frequency).

parasitics in the off-chip PDN may have a great impact in the global voltage variations. This section proposes an analysis with different package decoupling capacitance parasitics:

- Package 1 (*PKG1*): the same used in previous sections, with typical ESL: $L_{\text{cpkg}} = 5.61pH$.
- Package 2 (*PKG2*): with almost ideal decoupling capacitance, maximizing voltage noise reduction: $L_{\text{cpkg}} = 2pH$.
- Package 3 (*PKG3*): using decaps with higher inductive parasitics, increasing the equivalent inductance that forms the LC circuit with the die capacitance: $L_{\text{cpkg}} = 12pH$.

In order to enforce a voltage variation of $>10\%$ for all cases and compare their impact on the reference performance, all current sources are active and aligned with the first droop frequency, which is different for each package (see Fig. 22), with a total on-chip decoupling capacitance of 500nF. This configuration generates voltage swings large enough to provoke a voltage emergency for *PKG2*, and to keep the delay increase less than 1ns for *PKG3*.

Fig. 22 depicts the impedance responses for the 3 packages. Notice that the ESL parasitics in the package decaps have a massive influence in the quality of the PDN. The very low inductance of the *PKG2* decaps results in a lower impedance at the first droop and a great voltage noise mitigation.

On the opposite side, the decaps with higher ESL of *PKG3* increase the equivalent inductance connected to the chip, resulting in a higher peak impedance. In practice, *PKG3* can be used as a reference in terms of impedance as if the flip chip interconnection would be replaced by a wire bonding, which it is known to have higher impedance and to be cheaper [22].

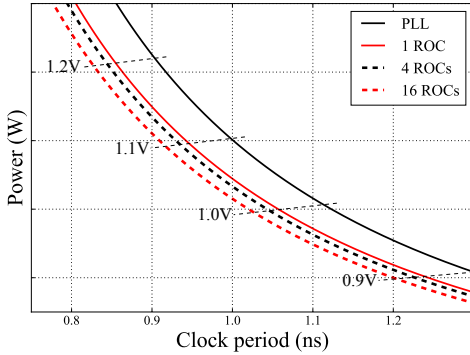


Fig. 24. Power/Performance trade-off for $\pm 10\%$ voltage noise.

Fig. 23 shows the delay increase for the PLL and the different ROC configurations, taking into account all activity patterns of Fig. 9. For the PLL, it is necessary to cover the deepest droops and to ensure that the delay of the critical paths are always shorter than the clock period. The largest generated droop was -329mV , leading to a performance degradation of up to 84% with *PKG3*, comparing the delay increase of the PLL (773ns) with the ROC in the center of the grid (123ns). As seen in Section V-B, ROCs take advantage of the global characteristics of voltage droops, requiring smaller margins and achieving higher average performance.

VII. DISCUSSION

Voltage droops have a great impact in the performance when using rigid clocks. For this reason, a significant effort must be invested in designing high-quality PDNs: adding decaps at all levels, reducing the impedance at each interconnection, considering the frequency response w.r.t. the activity of the circuit, and using elements with low parasitics and *low variability*.

Section VI presented different and illustrative configurations of the PDN, demonstrating how harmful low-quality PDNs can be for PLLs. ROCs provide a better alternative to tackle power integrity problems without degrading performance. This section presents a summary of advantages and disadvantages of using ROCs as the clock source.

A. Simpler voltage scaling

ROCs offer instantaneous adaptation to static and dynamic variability. Such characteristic can be used for a simpler version of dynamic voltage/frequency scaling (DVFS). Differently from the DVFS techniques currently used, in which both frequency and voltage must be controlled, with ROCs it is possible to define the performance only with the voltage.

Furthermore, voltage scaling can be used for an improved power/performance trade-off [11]. Fig. 24 depicts the trade-off between power and performance for the PLL and the different ROC strategies, with iso-voltage curves. Notice that ROCs naturally adapt to the process variability, and voltage scaling can be used after fabrication in order to find the *minimum energy point* for the performance required.

B. EMI reduction

Electromagnetic interference (EMI) is an aspect that must be considered to comply with the regulations in the application

domain. In digital systems, EMI is mostly produced by the periodic current differences around clock edges.

A known approach to mitigate electromagnetic radiations is the use of spread-spectrum clock generators, that outspread the energy over a wider bandwidth, reducing peak amplitude [23]. This technique consists of inserting intentional jitter to the clock generator, which implies additional timing margins.

The presence of dynamic variations implicitly injects jitter to the clock period of ROCs. Fortunately, this jitter does not need to be margined since the period variability is correlated with the circuit delays. Therefore, a *natural* spread-spectrum effect is produced without affecting performance.

C. Benefits of multiple ROC domains

In a Globally Asynchronous Locally Synchronous (GALS) design methodology with multiple ROC domains, the period of each ROC is defined based on the critical path within the local clock domain, and not on the global worst-case. Thus, EMI reduction benefits can be boosted [24], while side-channel security is also improved. In addition, clock tree synthesis is simpler with smaller clock domains, whereas power consumption can be minimized with lower clock frequencies.

D. Disadvantages

ROCs can *surf* over deep voltage fluctuations while sustaining an average performance. This comes at the expense of a clock period with high jitter and potentially large frequency variations. Systems operating with ROCs must tolerate these characteristics along the executing time of the applications.

It is difficult to design an ROC with a stable duty cycle, and the duty cycle cannot be guaranteed. Therefore, this may be a limitation for applications that require both clock edges, such as DDR SDRAMs. However, a simple solution is to use more than one clock source, e.g., a PLL with 50% duty cycle for the DDR memory interface, and ROCs for the random logic.

The GALS methodology has an important characteristic: it requires cross-domain crossing (CDC) techniques to be applied between the different ROC regions. There are several known techniques that perform CDC [24]. Each technique has its pros and cons, but there is an overhead in area, power and throughput, independently of the approach defined. Still, for multi-core or very large chips, the use of multiple clocks is already required [25], and the use of multiple ROC domains could be applied without additional costs.

VIII. CONCLUSIONS


Power integrity is a major concern due to low supply voltages and high power density in high-performance circuits. ROCs provide a robust scheme that tolerates large fluctuations in the supply voltages, and have been shown to be a competitive alternative to the rigid clocks generated by PLLs, with reductions of up to 83% in performance margins and up to 27% in leakage power. The design of the PDN is a difficult task that must consider the characteristics of the circuit and deliver a high-quality supply voltage. With an ROC, it was shown that the PDN design constraints can be relaxed, without performance loss.

We are facing a future in which many devices will have to operate in environments with scarce energy in which scavenging mechanisms will be essential to survive. Providing reliable DC voltages under these scenarios may be difficult and costly. ROCs emerge as a potential solution to operate robustly in hostile environments with low-cost PDNs. Furthermore, considering the use of integrated circuits in safety critical applications, the ROCs characteristic of adapting to undesirable operating conditions may be crucial to support situations of scarce energy or large voltage noise.

As future work, some directions might be explored. Additional positive effects could be considered, such as clock-data compensation [15], and voltage noise reduction, as the ROC adaptation to the supply voltage acts as negative feedback, reducing the amplitude of the voltage variations. Also, we do believe that the measurements in manufactured circuits would give a significant value and would confirm the conclusions presented in this work.

REFERENCES

- [1] M. Popovich, A. V. Mezhiba, and E. G. Friedman, *Power Distribution Networks with On-Chip Decoupling Capacitors*, 1st ed. Springer, 2008.
- [2] M. D. Pant, P. Pant, and D. S. Wills, "On-chip decoupling capacitor optimization using architectural level prediction," *IEEE Transactions on VLSI Systems*, vol. 10, no. 3, pp. 319–326, 2002.
- [3] Z. Zeng, X. Ye, Z. Feng, and P. Li, "Tradeoff analysis and optimization of power delivery networks with on-chip voltage regulation," *Proc. of DAC*, 2010, pp. 831–836.
- [4] R. Joseph, D. Brooks, and M. Martonosi, "Control techniques to eliminate voltage emergencies in high performance processors," *Proc. of HPCA*, 2003, pp. 79–90.
- [5] K. A. Bowman, C. Tokunaga, T. Karnik, V. K. De, and J. W. Tschanz, "A 22 nm all-digital dynamically adaptive clock distribution for supply voltage droop tolerance," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 4, pp. 907–916, 2013.
- [6] J. Tschanz, N. S. Kim, S. Dighe, J. Howard, G. Ruhl, S. Vangal, S. Narendra, Y. Hoskote, H. Wilson, C. Lam *et al.*, "Adaptive frequency and biasing techniques for tolerance to dynamic temperature-voltage variations and aging," in *Proc. of ISSCC*, 2007, pp. 292–604.
- [7] N. Kurd, P. Mosalikanti, M. Neidengard, J. Douglas, and R. Kumar, "Next generation Intel core micro-architecture clocking," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 4, pp. 1121–1129, 2009.
- [8] K. Wilcox, R. Cole, H. R. Fair III, K. Gillespie, A. Grenat, C. Henrion, R. Jotwani, S. Kosonocky, B. Munger, S. Naffziger *et al.*, "Steamroller module and adaptive clocking system in 28 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 1, pp. 24–34, 2015.
- [9] S. Nasir, S. Gangopadhyay, and A. Raychowdhury, "All-digital low-dropout regulator with adaptive control and reduced dynamic stability for digital load circuits," *IEEE Transactions on Power Electronics*, vol. 31, no. 12, 2016.
- [10] D. Kamakshi, M. Fojtik, B. Khailany, S. Kudva, Y. Zhou, and B. Calhoun, "Modeling and analysis of power supply noise tolerance with fine-grained GALS adaptive clocks," in *Proc. of ASYNC*, 2016, pp. 75–82.
- [11] J. Cortadella, L. Lavagno, P. López, M. Lupon, A. Moreno, A. Roca, and S. Sapatnekar, "Reactive clocks with variability-tracking jitter," in *Proc. of ICCD*, 2015, pp. 511–518.
- [12] J. Cortadella, M. Lupon, A. Moreno, A. Roca, and S. Sapatnekar, "Ring oscillator clocks and margins," in *Proc. of ASYNC*, 2016, pp. 19–26.
- [13] S. Pant and E. Chiprout, "Power grid physics and implications for CAD," in *Proc. of DAC*. ACM, 2006, pp. 199–204.
- [14] L. Machado, A. R. Perez, and J. Cortadella, "Voltage noise analysis with ring oscillator clocks," in *Proc. of ISVLSI*. IEEE, 2017, pp. 1–6.
- [15] K. Wong, T. Rahal-Arabi, M. Ma, and G. Taylor, "Enhancing microprocessor immunity to power supply noise with clock-data compensation," *IEEE J. of Solid-State Circuits*, vol. 41, no. 4, pp. 749–758, 2006.
- [16] T. Pialis and K. Phang, "Analysis of timing jitter in ring oscillators due to power supply noise," in *Proc. of ISCAS*. IEEE, 2003, pp. 685, 688.
- [17] K. Agarwal and S. Nassif, "Characterizing process variation in nanometer CMOS," in *Proc. of DAC*, 2007, pp. 396–399.
- [18] M. S. Gupta, J. L. Oatley, R. Joseph, G.-Y. Wei, and D. M. Brooks, "Understanding voltage variations in chip multiprocessors using a distributed power-delivery network," in *Proc. of DATE*, 2007, pp. 1–6.
- [19] T. Sakurai and A. R. Newton, "A simple MOSFET model for circuit analysis," *IEEE Transactions on Electron Devices*, vol. 38, no. 4, pp. 887–894, 1991.
- [20] T. Sakurai, "A JSSC classic paper: the simple model of CMOS drain current," *IEEE Solid State Circuits Society Newsletter*, vol. 9, no. 4, pp. 4–5, 2004.
- [21] M. Litochevski and L. Dongjun, "High throughput and low area AES," 2012. [Online]. Available: http://opencores.org/project,aes_highthroughput_lowarea
- [22] A. Fontanelli, "System-in-package technology: opportunities and challenges," in *Proc. of ISQED*. IEEE, 2008, pp. 589–593.
- [23] J. Kim, D. G. Kam, P. J. Jun, and J. Kim, "Spread spectrum clock generator with delay cell array to reduce electromagnetic interference," *IEEE Transactions on Electromagnetic Compatibility*, vol. 47, no. 4, pp. 908–920, 2005.
- [24] M. Krstic, E. Grass, F. K. Gürkaynak, and P. Vivet, "Globally asynchronous, locally synchronous circuits: Overview and outlook," *IEEE Design & Test of Computers*, vol. 24, no. 5, pp. 430–441, 2007.
- [25] K.-D. Schubert, W. Roesner, J. M. Ludden, J. Jackson, J. Buchert, V. Paruthi, M. Behm, A. Ziv, J. Schumann, C. Meissner *et al.*, "Functional verification of the IBM POWER7 microprocessor and POWER7 multiprocessor systems," *IBM Journal of Research and Development*, vol. 55, no. 3, pp. 10–1, 2011.



Lucas Machado received the B.S. and the M.S. degrees in computer engineering from the Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil, in 2010 and 2013, respectively. He is currently pursuing the Ph.D. degree in computer science at the Universitat Politècnica de Catalunya, Barcelona, Spain. His current research interests include computer-aided design of integrated circuits, with special interest on logic synthesis, security, reliability and asynchronous circuits.



Antoni Roca received the M.S. degree in telecommunications and PhD in computer science from Universitat Politècnica de València, Spain, in 2006 and 2012, respectively. He was a post-doc researcher at the Universitat Politècnica de Catalunya, Barcelona, from 2013 to 2016. His current research interests include network-on chip, integrated circuits design, and chip variability.



Jordi Cortadella (S'87-M'89-F'15) is currently a Professor with the Computer Science Department, Universitat Politècnica de Catalunya, Barcelona, Spain. His current research interests include formal methods and computer-aided design of VLSI systems with a special emphasis on asynchronous circuits, concurrent systems, and logic synthesis. Prof. Cortadella is a member of Academia Europaea. He received best paper awards at the International Symposium on Advanced Research in Asynchronous Circuits and Systems in 2004 and 2016, the Design Automation Conference in 2004, and the International Conference on Application of Concurrency to System Design in 2009. He has served on the technical committees of several international conferences in the field of design automation and concurrent systems, and is an Associate Editor of the *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*.