

Voltage Noise Analysis with Ring Oscillator Clocks

Lucas Machado, Antoni Roca and Jordi Cortadella
Department of Computer Science
Universitat Politècnica de Catalunya
Barcelona, Spain

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). Ring Oscillator Clocks (ROCs) have been proposed as an alternative to mitigate the negative effects of voltage noise as technology scales down and power density increases. However, their effectiveness highly depends on the design parameters of the PDN, power consumption patterns of the system and spatial locality of the ROCs within the clock domains. This paper analyzes the impact of the PDN parameters and ROC location on the robustness to voltage noise. The capability of reacting instantaneously to unpredictable voltage droops makes ROCs an attractive solution, which allows to reduce the amount of decoupling capacitance without downgrading performance. Tolerance to voltage noise and related benefits can be increased by using multiple ROCs and reducing the size of the clock domains. The analysis shows that up to 83% of the margins for voltage noise and up to 27% of the leakage power can be reduced by using local ROCs.

Keywords—ring oscillators; adaptive clocking; voltage noise;

I. INTRODUCTION

Power integrity is one of the major challenges in the design of high-performance circuits. All components of the power delivery network (PDN) have a direct influence on the voltage fluctuations observed by the on-chip devices. Mitigating this noise is an arduous task that may have a significant impact on the design parameters: 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 are produced by the parasitic resistance of the PDN, whereas inductive noise is mainly caused by large current differences, associated with the switching activity of the chip.

Static voltage offsets can be estimated at design time. However, the dynamic variations are hard to predict and this is the reason why overly conservative margins are often added to prevent unexpected failures. Unfortunately, voltage droops that exceed the defined margins cannot be fully eliminated.

Clock and power gating are typical 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.

A common strategy to mitigate voltage noise is to increase the amount of on/off-chip capacitance by adding decoupling capacitors (decaps) [1]. Unfortunately, the additional decaps imply an increase in area and leakage power consumption. In [2], different amounts of integrated voltage regulators are investigated, analyzing the penalties in area and power for the voltage noise reductions obtained. Other proposals include improving the chip-package impedance, static and dynamic voltage margining, and performance throttling

and stalling using voltage sensors [3], [4]. All these have important overhead in design cost, area, power or performance.

Adaptive clocking [5]–[9] seems to be a promising solution with low overhead, but its efficiency is limited by the characteristics of the clock generators. Ring Oscillator Clocks (ROCs) [10], [11] can be considered an adaptive clocking proposal, which takes into account all sources of variability. If the ROC is correctly designed, 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 to scale the global and local parts of voltage noise?
- 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?

Voltage noise analysis has always been focused on estimating the global worst-case and deriving the timing margins required [12]. For ROCs, the key value is the largest *differential voltage* between the ROC and the critical path. This work presents an analysis on voltage locality for a design using ROCs as clock source. Voltage locality is introduced by multiple activity patterns using an on-chip power distribution model. A trade-off between the number of ROC domains and performance is presented. Also, modifications in the PDN are evaluated, such as removing on-chip decoupling capacitance and changing the placement of the power bumps.

II. VOLTAGE NOISE

The PDN is responsible for delivering the power and ground voltages to all devices of the design. Fig. 1 depicts the PDN model with its components: voltage regulator, board (PCB), package (PKG), the connection bumps and the on-chip power networks [1].

The parasitics of the PDN form LC circuits with different resonance frequencies, which are responsible for the *voltage droops*. The LC circuit composed of the die capacitance and the bumps inductance generates the *first droop*, which typically produces the largest voltage noise and has a resonance frequency of around 100-400MHz [13]. The second droop is controlled by C_{pkg} and L_{pkg} , and the third droop is dominated by C_{pcb} and L_{pcb} .

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 behaviour illustrated in Fig. 2(b) is observed when a single current spike is requested for this PDN: the first droop causes fast and large voltage swings in the order of ns ; then the voltage continues to fluctuate due to second and third droops, until it becomes stable after a few μs .

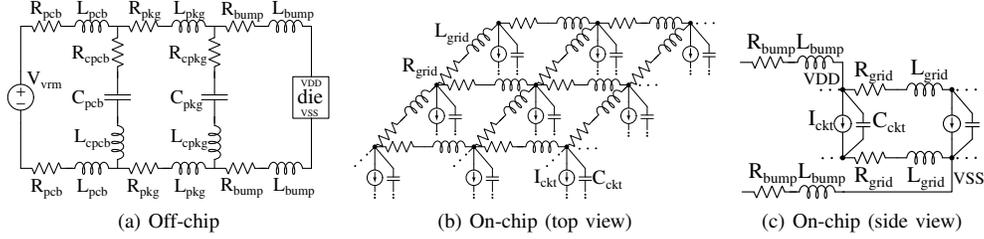


Figure 1. PDN model with off-chip and on-chip parasitics.

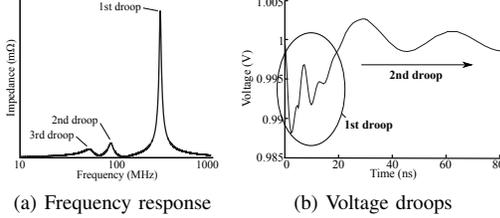


Figure 2. (a) The frequency response of a typical PDN, and (b) the voltage droops generated by a single current spike

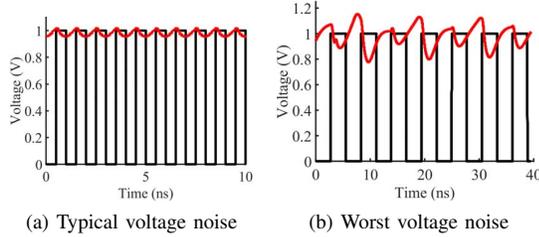


Figure 3. Voltage droops generated by periodical current differences at (a) low and (b) high impedance frequencies.

Voltage noise is minimized when the activity occurs at low impedance frequencies. Fig. 3(a) shows the supply noise and the clock signal for a circuit operating at 1GHz with this PDN, with voltage swings of $\pm 10\%$. The clock can be set to the frequency of the first droop to emulate the worst voltage noise, as seen in Fig. 3(b). In this case, the voltage noise amplitude goes up to 20%.

III. RING OSCILLATOR CLOCKS

The use of ring oscillators as the clock source has been disregarded due to the jitter, caused by the sensitivity of ROCs to the various sources of variability. Jitter and other clock uncertainties are generally handled in STA by adding timing margins, degrading the performance of the circuit. Therefore, clock generators with low-jitter, such as PLLs, became the *de facto* clock source paradigm.

Consider a synchronous circuit fed by a PLL or by an ROC, depending on the selection of a multiplexer. Fig. 4 illustrates the clock signals generated by the PLL and the ROC when a voltage droop occurs. The clock period of the PLL is not affected by the voltage swings, as it is designed to support them. However, the circuit paths have a different behavior: their delay increases when voltage decreases. If the PLL is the clock source, timing failure is prevented by adding margins considering the *minimum voltage*.

Differently, the ROC period is modified by the voltage variation, as seen in Fig. 4. Recent studies [10], [11] demonstrate that the jitter of ROCs is highly correlated with the delay variability of the circuit paths. In other words, as the ROC and the circuit paths are composed of similar gates, the PVT variations affect them similarly. This correlation enables the reduction of timing margins [10], [11].

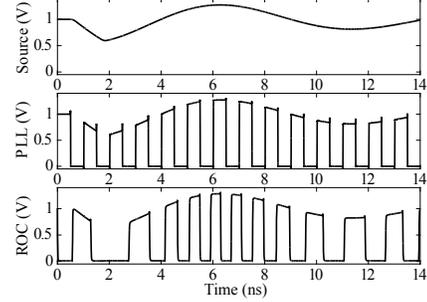


Figure 4. Clock signal generation in the presence of voltage noise.

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_{cpkg}	0.17 m Ω	L_{cpkg}	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	-	-

Obviously, there is not an exact match between the delay of the critical paths and the ROC period. 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 [14]. All these factors must be considered in the design of an ROC.

IV. PRELIMINARIES

A. PDN model

The chip-grid presented in [15] is used as PDN model, which represents an SoC with four cores of Pentium-4. The PDN components, illustrated in Fig. 1, are described in SPICE netlists using the values of Table I. As external regulators typically do not regulate high frequency variations, the voltage regulator (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, as seen in Fig. 1(b). Both the power and ground networks are considered in the model, with a V_{DD} or a V_{SS} bump connected at each point. A grid point models a section 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. 5(a)).

Additionally, a decoupling capacitor is added at each grid 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]. The PDN model, with 200nF of on-chip decoupling capacitance, has the frequency response of Fig. 5(b) observed at the power bumps.

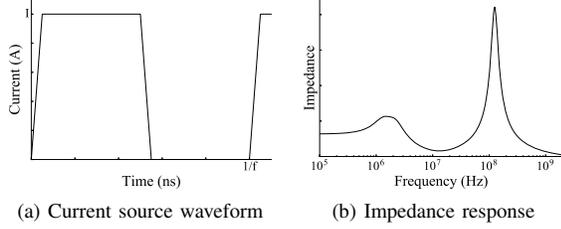


Figure 5. Waveform and impedance response with 200nF of decaps.

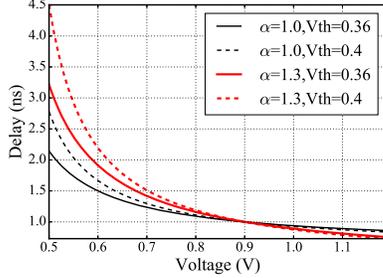


Figure 6. Path delay given by (2), with $td = 1ns$, $V_{DD}=0.9V$.

B. Delay model

A simplification of the gate delay formulation was proposed in [16], which is still widely accepted. By this model, the variation of the delay due to the supply voltage depends on the threshold voltage and α , a technology fitting value in the range of 1-2. This model was defined for a single gate, but the relationship between delay and voltage holds for a path composed of several 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 voltage.

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

A 65nm commercial library with 1V nominal voltage is used as reference. The average V_{th} of all combinational cells is 0.36V for 75°C and 0.4V for 125°C. A typical value of α is 1.3 [16], and this parameter is closer to 1 for current technologies. Typically, $\pm 10\%$ offsets are defined for the voltage swings during STA. Considering a clock source of 1GHz, the critical path at $V_{DD} = 0.9V$ must have a maximum delay of 1ns. Fig. 6 shows the path delay curves with the k values calculated using (2), with $V_{DD} = 0.9V$, $td = 1ns$, $\alpha = [1.0, 1.3]$ 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 and at the minimum voltage, as shown in (3).

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

The correct design of an ROC must consider the delay behaviour of Fig. 6, keeping the ROC period larger than the critical path delay for any voltage. For example, if the ROC has a larger V_{th} than the critical path, then margins might be smaller. However, for simplification, the delay behaviour of the ROC and the critical path are both given by (2) with the same parameters.

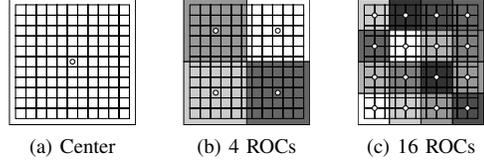


Figure 7. ROC placement strategies and number of clock domains.

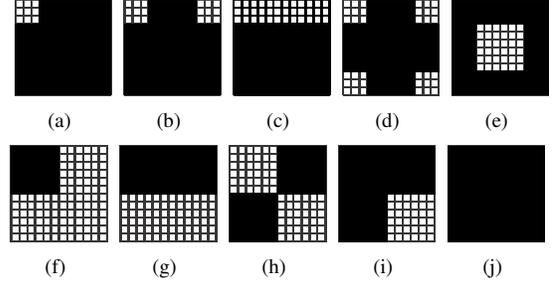


Figure 8. Patterns defining the grid points that are active.

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

- The voltage at the ROC is 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 (delay variations are larger for lower voltages).
- Positive effects due to the clock distribution are not taken into account, such as clock-data compensation [13].

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} . But the ROC margin varies with its location, as the voltage difference is smaller between points that are closer to each other. Fig. 7 depicts the 3 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, one at the center of each 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 overhead in performance and power. Therefore, this case is reported but its results are not compared with the PLL.

V. VOLTAGE LOCALITY ANALYSIS

Different patterns of active logic on top of the grid model are proposed, which are depicted in Fig. 8. The dark areas represent the regions of the chip that are active. The grid points that are not active are modeled with constant current sources.

A decoupling capacitance of 200nF is necessary to keep the voltage swings within $\pm 10\%$ when all current sources are operating at 1GHz. Using a grid model with this configuration, the activity patterns of Fig. 8 are simulated with Synopsys HSPICE[®] for 50 clock cycles at 125°C, saving the minimum voltage (V_{min}) of all grid points, and the maximum voltage difference between any two points in the grid (ΔV_{DD}).

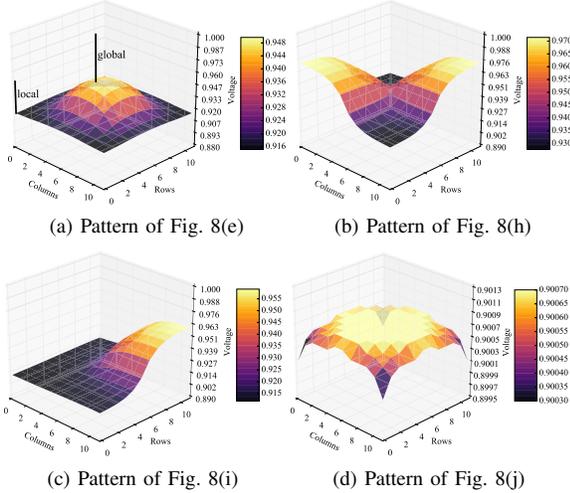


Figure 9. Voltage distribution for some activity patterns of Fig. 8.

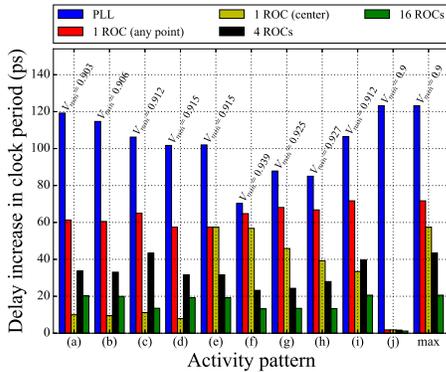


Figure 10. Delay increase in the clock period for each activity pattern (200nF decaps, activity at 1GHz).

Fig. 9 shows the global and local effects due to some of the patterns. These images show the voltage levels at each grid point when the minimum voltage is reached. The pattern in Fig. 8(j) generates the lowest voltage, reaching a maximum current of 28A.

A. Typical voltage noise

Fig. 10 is generated with the voltage data gathered, using (3) and (4). 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 point, and 57ps if it is placed at the center of the die.

In Fig. 11, the activity patterns of Fig. 8(e) and Fig. 8(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 center. Fig. 11(a) depicts the worst case for the ROC, while Fig. 11(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 ROC similarly. Therefore, ROCs enable a higher performance for the same level of voltage noise robustness.

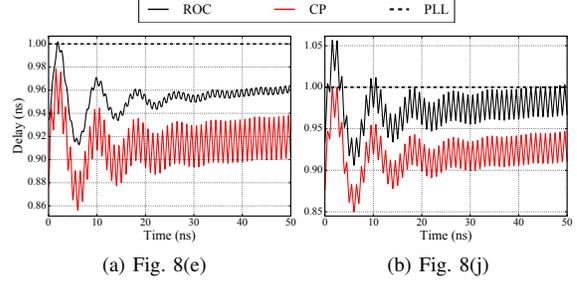


Figure 11. PLL and ROC period, and critical path (CP) delay for the activity patterns of Fig. 8(e) and Fig. 8(j).

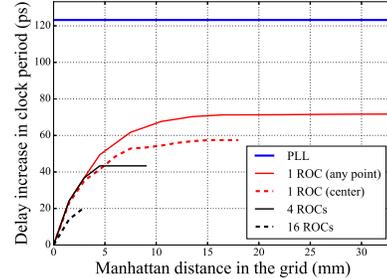


Figure 12. Largest delay increase vs. the distance between the ROC and the critical path (200nF decaps, activity at 1GHz).

Fig. 12 depicts the largest delay increase for each distance between two grid points, considering all activity patterns. As expected, the delay is smaller if the critical path is closer to the ROC. So, it is possible to perform a *trade-off* between performance and the number of ROC domains. For example, the required delay is reduced to 43ps with 4 ROCs and to 20ps with 16 ROCs.

B. Worst voltage noise

The delay increase shown in Fig. 10 is required for a typical voltage noise, but larger voltage droops may happen if current differences take place at frequencies with high impedance, as seen in Sect. II. The first droop frequency of the grid model with 200nF of on-chip decaps is 125MHz. This means that 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, but with the current sources operating at 125MHz, instead of 1GHz.

Fig. 13 depicts the delay increase for all activity patterns in this case. As expected, the voltage noise is boosted due to the high impedance, and the delay increase required for the clock period of the PLL is 1535ps. Therefore, if worst voltage noise is considered, then 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, comparable to the delay increase of the PLL for typical voltage noise.

VI. ON-CHIP DECOUPLING CAPACITANCE

The voltage noise reduction obtained by increasing the on-chip decoupling capacitance has a direct impact on performance. This effect is illustrated in Fig. 14(a), varying the amount of on-chip decaps, with activity at 1GHz.

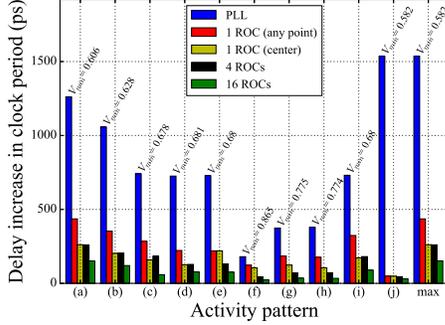
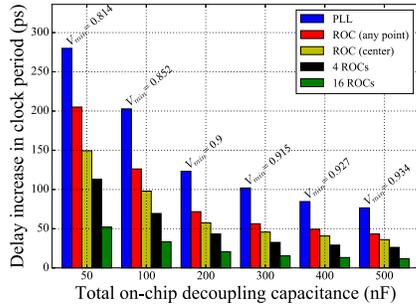
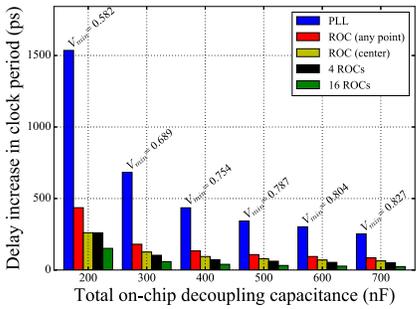


Figure 13. Delay increase in the clock period for each activity pattern (200nF decaps, activity at first droop (125MHz)).



(a) Activity at 1GHZ



(b) Activity at first droop frequency

Figure 14. Delay increase for the PLL and ROC, with different amounts of on-chip decoupling capacitance.

The behavior is similar with the activity aligned with the first droop frequency, with significant margin reductions due to lower impedance, as seen in Fig. 14(b). Notice that Fig. 14 is generated considering all activity patterns of Fig. 8, and that the first droop frequency varies with the amount of on-chip capacitance.

Generally, on-chip decaps do not increase 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 of decaps is important. As ROCs support larger voltage fluctuations with lower margins than the PLL, it is possible to reduce the amount of decaps and leakage power without degrading performance.

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

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

The leakage savings are estimated by using the parameters

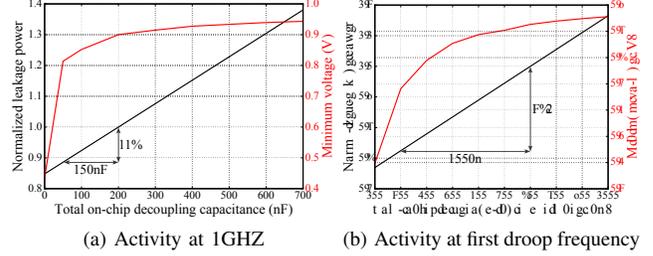


Figure 15. Normalized leakage power and minimum voltage for different amounts of on-chip decoupling capacitance.

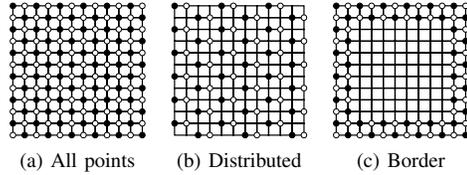


Figure 16. Different power bumps placement strategies.

of a commercial 65nm library. The least leaky decap cell is selected, with a capacitance per area of 6nF/mm² and leakage power of 2.5mW/nF. Hence, the leakage power per area of decaps is 15mW/mm². For standard cells, leakage per area is estimated based on a design with a representative mix of combinational gates and flip-flops [17], 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 (for an utilization of 80%).

Fig. 15(a) shows the leakage power and the minimum voltage for different amounts of decaps, for typical voltage noise. Leakage power is normalized with respect to 200nF. Considering the margins in Fig. 14(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, for typical conditions.

Similarly, Fig. 15(b) depicts the leakage and minimum voltage, but for the worst 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. 14(b), it is possible to have 200nF decaps instead of 700nF, without degrading average performance, with ROCs. Removing 500nF means a reduction of 27% in total leakage. Furthermore, if 200nF occupy all the white space, then 700nF imply a non-negligible area increase that can be simply avoided by using ROCs.

VII. POWER INTERCONNECTIONS

The amount (and placement) of power bumps is another PDN characteristic that influences voltage locality. The experiments in previous sections were performed with 72 pairs of VDD/VSS bumps uniformly distributed (see Fig. 16(a)), in order to minimize the impedance between the chip and the package [1]. As seen in Fig. 9, such placement reduces the voltage differences.

This section considers the placement of bumps in the grid model with 200nF, for typical voltage noise (activity at 1GHz). Two strategies are analyzed: 36 VDD/VSS pairs uniformly distributed, as in Fig. 16(b); and 40 VDD/VSS pairs placed in the borders (similar to wire bonding), depicted in Fig. 16(c). Fig. 17 shows that the bump placement has a huge impact in the power distribution, affecting voltage throughout the die.

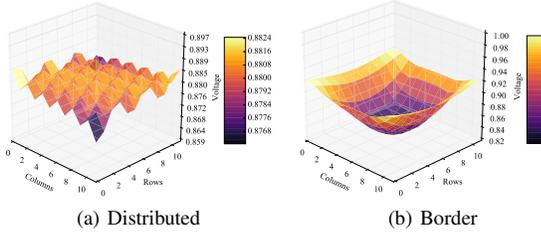


Figure 17. Voltage distribution for activity pattern of Fig. 8(j) with (a) 36 VDD/VSS bumps distributed and (b) 40 VDD/VSS bumps in the borders.

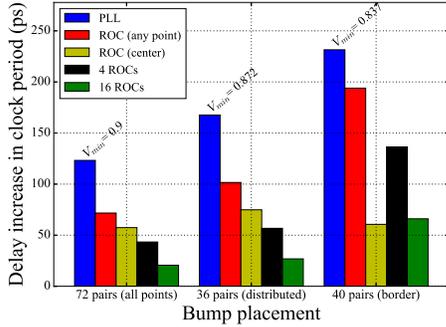


Figure 18. Required margins for the PLL and ROC with different bump placements (200nF of decoupling capacitance, activity at 1GHz).

All activity patterns are simulated, producing the results of Fig. 18. As the impedance is higher, the minimum voltage is lower, indicating more margins. Also, ROC margins have a larger increase, due to higher voltage differences. Still, it is possible to reduce bumps 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.

VIII. CONCLUSIONS

Power integrity is a major concern nowadays due to low supply voltages and high power density in high-performance circuits. ROCs have been shown to be a competitive alternative to the classical rigid clock paradigm, with reductions of up to 83% in margins and up to 27% in leakage power. ROCs do not only provide significant advantages in performance and power, but a robust scheme that tolerates large fluctuations in the supply voltages, and live with low-quality PDNs.

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.

ACKNOWLEDGMENT

The present work was performed with the support of CNPq, Conselho Nacional de Desenvolvimento Científico e Tecnológico - Brasil. 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, and the Generalitat de Catalunya (2014 SGR 1034).

REFERENCES

- [1] M. Popovich, A. V. Mezhiba, and E. G. Friedman, *Power Distribution Networks with On-Chip Decoupling Capacitors*, 1st ed. Springer, 2008.
- [2] Z. Zeng, X. Ye, Z. Feng, and P. Li, "Tradeoff analysis and optimization of power delivery networks with on-chip voltage regulation," in *Proc. of DAC*, 2010, pp. 831–836.
- [3] R. Joseph, D. Brooks, and M. Martonosi, "Control techniques to eliminate voltage emergencies in high performance processors," in *Proc. of HPCA*, 2003, pp. 79–90.
- [4] 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 J. of Solid-State Circuits*, vol. 48, no. 4, pp. 907–916, 2013.
- [5] 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.
- [6] N. Kurd, P. Mosalikanti, M. Neidengard, J. Douglas, and R. Kumar, "Next generation Intel core micro-architecture clocking," *IEEE J. of Solid-State Circuits*, vol. 44, no. 4, pp. 1121–1129, 2009.
- [7] 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 J. of Solid-State Circuits*, vol. 50, no. 1, pp. 24–34, 2015.
- [8] S. Nasir, S. Gangopadhyay, and A. Raychowdhury, "All-digital low-dropout regulator with adaptive control and reduced dynamic stability for digital load circuits," *IEEE Trans. on Power Electronics*, vol. 31, no. 12, 2016.
- [9] 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. ASYNC*, 2016, pp. 75–82.
- [10] 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.
- [11] J. Cortadella, M. Lupon, A. Moreno, A. Roca, and S. Sapatnekar, "Ring oscillator clocks and margins," in *Proc. of ASYNC*, 2016, pp. 19–26.
- [12] S. Pant, E. Chiprout, and D. Blaauw, "Power grid physics and implications for CAD," vol. 24, no. 3, pp. 246–254, 2007.
- [13] 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.
- [14] K. Agarwal and S. Nassif, "Characterizing process variation in nanometer CMOS," in *Proc. of DAC*, 2007, pp. 396–399.
- [15] 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.
- [16] 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.
- [17] M. Litochevski and L. Dongjun, "High throughput and low area AES," 2012. [Online]. Available: http://opencores.org/project_aes_highthroughput_lowarea