A Comparison of Digital Droop Detection Techniques in ASAP7 FinFET

Kevin Zheng

***Abstract*—This report compares current techniques for digital supply droop detection. While there are many droop detection circuits in use, they have been thus far studied in isolation; it is not clear what the design merits and trade-offs of different topologies are. This project presents design methodologies for two single-cycle digital droop detectors, which are used to design and compare droop detector sensitivity and selectivity for two representative example specifications in ASAP7.**

I. Introduction

Digital circuits and systems, particularly processors, are sensitive to transient droops in supply voltage. These droops occur as a result of transient current spikes due to digital switching action or large digital blocks coming out of sleep. These spikes cause local supply voltage droops due to resistive or inductive impedances in the on and off-chip power supply network. These droops are problematic because logic delays increase with lower supply voltage, so logic and memory cells require greater timing margins. This problem is exacerbated by increased logic density, power gating, and design rule limitations of modern processors that complicate the design of effective power supply networks.

In order to ensure performance and reliability, processors employ droop detection and mitigation techniques. Droop mitigation strategies fall roughly into two categories: reducing droop using a local supply regulator (for example, using a digital low-dropout-regulator), or dynamically reducing frequency. Both techniques require a fast droop detector that works reliably across process and temperature variations to detect when a supply droop is occurring.

It is also desirable, particularly for digital designers and system integrators, that droop detectors are implementable entirely using digital standard cells. Digital droop detectors can be integrated into a digital design using the same synthesis and place-and-route flow without requiring custom analog/mixed-signal (AMS) design and layout. It is not yet clear whether custom "generators" that automate AMS design [1] will obviate the desire for fully-digital droop detectors.

This report targets single-cycle detection of the shortest and deepest so called "first droop" events that are dominated by package inductance and on-die capacitance [2]. In practice, it may take several more cycles for a real system to propagate a droop signal and then react appropriately. Third droop events can be readily mitigated by bulk capacitors on the motherboard. Second droop events, while important, are not nearly as deep and as fast as first droop events.

The rest of this report is organized as follows: Section II reviews and explains the operation of several droop detection techniques found in recent publications. Section III presents a comparison of two particular droop detection techniques two example implementations in ASAP7. Section IV summarizes and interprets the results.

II. Droop Detection Techniques

Droop detectors are effectively fast, coarse-grained voltage sensors. Most droop detectors fall into two broad categories: voltage and time-based sensors. Voltage-based droop detectors compare the local digital supply voltage *DVDD* to some generated voltage threshold. This technique generally requires some externally-routed voltage reference. Time-based droop detectors do not directly measure voltage, but rather some supply voltage-dependent parameter like inverter delay. This timing parameter is compared with an external frequency reference and converted into a measurement of voltage. This technique replaces the need for a reference voltage by a reference frequency, but fortunately frequency references are typically readily available on clocked digital circuits.

Presented here are a few droop detectors that have been found in recent publications. Where appropriate, simplified schematics and simulated operation waveforms in ASAP7 are shown. Some of these droop detectors are part of digital regulator (DLDO) designs; others are used stand-alone in processor systems.

A. NAND Divider Detector

One theme of voltage-based digital droop detectors is the emulation of analog functions using only digital standard cells. One example of this is the droop detector presented in [3] (Fig. 1), which from here we will call the *NAND divider detector*.

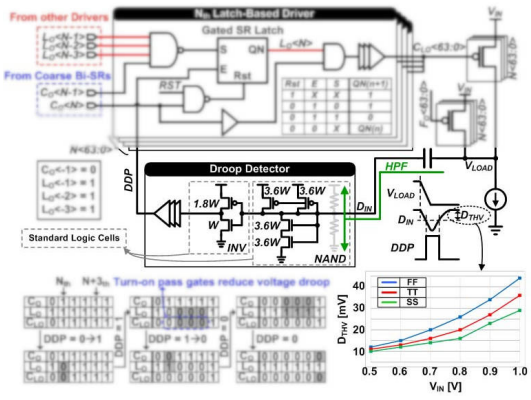

Fig. 1. Droop detector presented in [3]

A simplified schematic and simulated operation waveform are shown in Fig. 2. DVDD (a) is capacitively coupled into a resistive divider biased slightly above mid-rail (b). The divider drives the gate of one or more inverters. When DVDD is constant, node (b) remains at its bias point and node (c) stays low. If DVDD changes, it is capacitively coupled into node (b). A sufficiently large change will cause node (b) to fall below the crossing point of the inverter, and node (c) will rise. The output of node (c) is captured in an asynchronous latch, whose output node (d) is the droop detector output. In order for this technique to work, the resistive divider must be operated from a separate, stable reference supply *AVDD*.

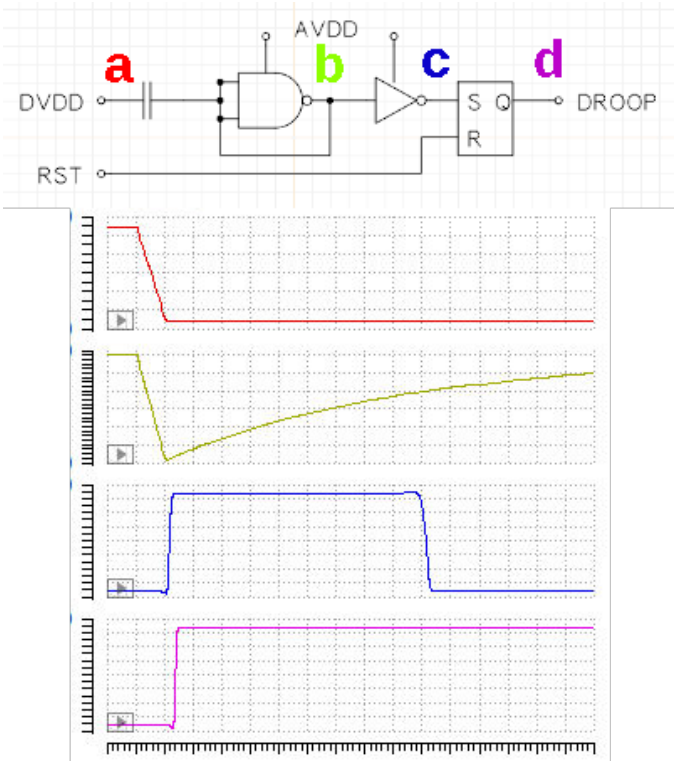

Fig. 2. Simplified schematic and simulated operatoin waveform of NAND droop detector in ASAP7

Resistors are often unavailable in advanced digital integrated circuit process technologies, not to mention digital place-and-route flows. A resistive voltage divider can be synthesized using a standard NAND cell (Fig. 3). NAND gates are typically sized for equal pull-up and pull-down strength compared to an inverter. If a unit inverter's PMOS width is k times wider than the NMOS width, then all the gates in a unit 2-NAND gate are designed to be k wide. When all of the inputs and outputs of a 2-NAND are connected together, it functions like an inverter with input and output shorted. If the equivalent pull-up resistance is R, then the equivalent pull-down resistance is $2R$. This way, the input/output node forms a voltage divider biased at approximately 2/3 VDD.

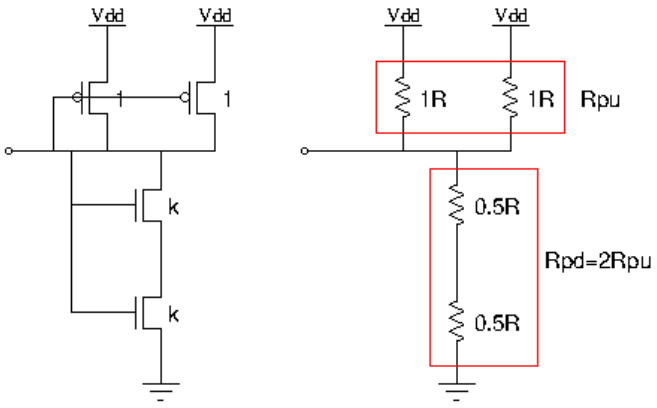

Fig. 3. Synthesizing a resisitive divider using a 2-input NAND gate

This technique also requires a 2-terminal AC-coupling capacitor. Since bypass capacitor ("decap") standard cells are usually implemented using 1-terminal MOS capacitors, they are not suitable for AC coupling. 2-terminal capacitors can be implemented using MOM (see the design example in Section III) or MIM capacitors when available.

B. Ring Oscillator Detector

The output frequency of a ring oscillator changes almost instantaneously with its supply voltage. Lower supply voltages increase gate delay, which lowers the oscillation frequency. The time-domain simulation output of a pre-extraction 3-stage ring oscillator implemented in ASAP7 is shown in Fig. 4.

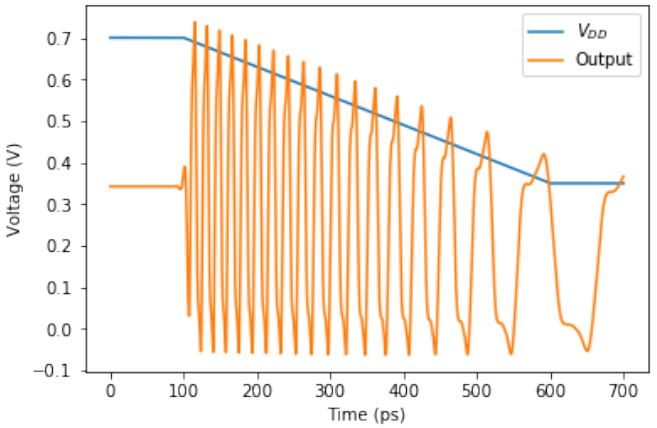

Fig. 4. Simulated pre-extraction 3-stage ring oscillator response to DVDD in ASAP7

The time-based droop detector presented in [4] makes use of the ring oscillator's sensitivity to supply voltage. A simplified schematic of a modified implementation for our design examples in ASAP7 (Section III) is shown in Fig. 5:

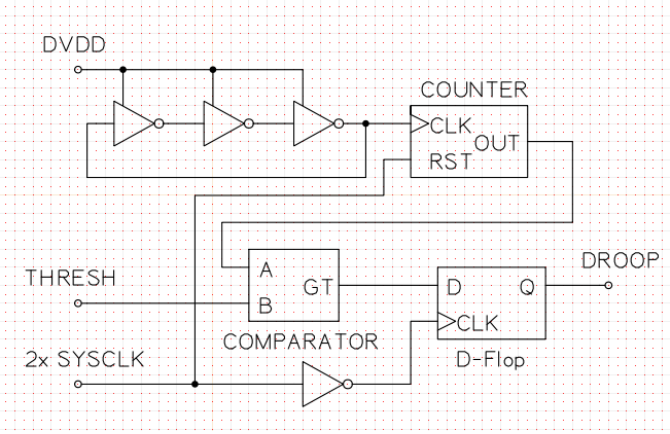

Fig. 5. Simplified schematic of a ring oscillator droop detector

During half of the system clock period, the counter counts the number of oscillations in the ring oscillator. At the end of the half clock cycle, the count is compared to a programmable threshold and captured in a D-flip-flop, whose output is the droop detection signal for the current cycle. When DVDD drops enough to cause the count to drop below the threshold, a droop event is signaled.

This droop detector is readily implementable with standard digital synthesis and place-and-route flows. It requires an external frequency reference, but clocks are readily available to digital circuits. Our particular implementation requires a two-phase clock detection in half a cycle.

C. Other Detectors

Several other droop detectors can be found in recent literature. These droop detectors are not compared in this report, but are still worth mentioning.

1) Delay Chain (TDC) Detector: Another way to directly measure gate delay is with a delay chain. The time-based droop detector presented in [5] (Fig. 6) uses flip-flops to sample the successive outputs of a delay chain, typical of a digital time-to-digital converter (TDC). In this detector, the detection signal is a sampled thermometer code output which represents a delay proportional to supply droop.

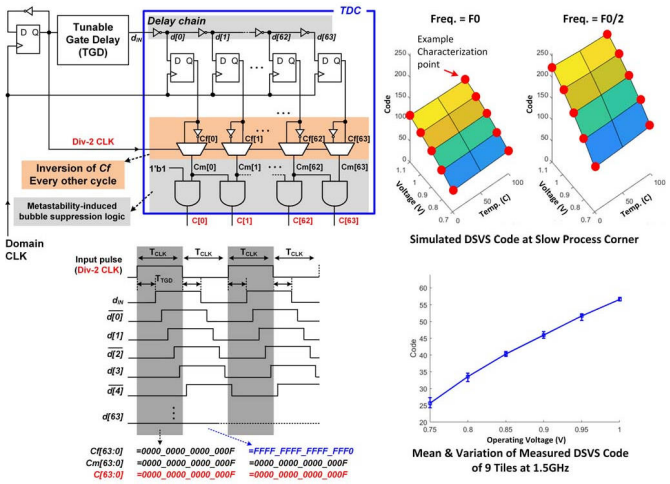

Fig. 6. Droop detector presented in [5]

2) Tunable Replica Circuits: When the stages in the delay chain replicate the critical path of a datapath, it is called a "tunable replica circuit" [6]. The delay measurement from a tunable replica circuit represents a direct measurement of the critical path delay that accounts for process, temperature, and voltage variations. This technique was used in [7] to achieve combined voltage and frequency control for droop detection and mitigation.

3) Mixed-Signal Detectors: The voltage-based droop detector presented in [8] (Fig. 7) avoids using an external voltage reference by generating a local reference. A 6-bit resistor-ladder digital-to-analog converter (DAC) scales down VDD by a programmable amount. The DAC output is then low-pass filtered and used as the reference for a clocked comparator. The clocked comparator compares VDD to its internal reference threshold; droops faster than the low-pass filter time constant trigger the comparator and signal a droop event. Unfortunately, this AMS technique is not easily realizable with digital flows.

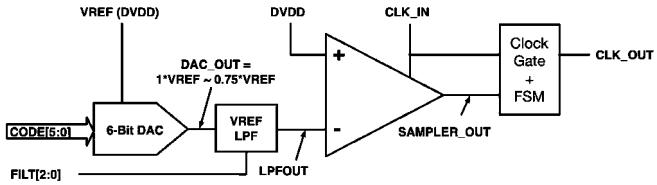

Fig. 7. Droop detector and clock gate presented in [8]

III. Comparison in ASAP7

In order to better understand the capabilities and shortcomings of very different types of droop detectors, we present a design methodology for two chosen droop detectors that we use to design detectors for two example specifications. Our designs are implemented in ASAP7, a freely-accessible predictive 7-nm FinFET PDK, with the goal of better understanding the challenges of designing droop detectors in an advanced digital process technology.

A. Design Examples

In real systems, droop detectors are designed with a mitigation strategy in mind, since droop detectors must detect droops in time for mitigation techniques to activate and prevent system reliability degradation. Since mitigation strategies are beyond the scope of this report, we focus on targeting single-cycle detection of the shortest and deepest "first droop" events. Our two design examples are:

1) Fast Design Example: Our fast design example represents a hypothetical high-performance desktop or server processor with a maximum boost frequency of 4 GHz. At 4 GHz, a droop detector has only 125 ps (half a period) to detect and signal a droop event.

2) Slow Design Example: Our slow design example represents a hypothetical low-power embedded or mobile processor with a core frequency of 800 MHz. At 800 MHz, a droop detector has 625 ps to detect and signal a droop event.

B. Design Metrics

At the time of writing, there does not appear to be a well-established figure of merit or basis for comparing different types of supply droop detectors. This report does not attempt to invent one. However, a droop detector can be treated as a sensor, and two ways to characterize a sensor are its *sensitivity* and *selectivity*. We will also consider secondary issues regarding *design ease*.

1) Sensitivity: The sensitivity of our droop detector can be measured by the minimum voltage droop that can be detected given our time constraints. A more sensitive droop detector will be able to measure a smaller droop in a fixed amount of time than a less sensitive one.

2) Selectivity: Everything is a sensor for something, but selective droop detectors should sense primarily voltage droops. To measure selectivity, we compare the sensitivity across different process and temperature corners.

3) Design Ease: Droop detectors that can be implemented entirely using existing digital synthesis and place-and-route tools are easy to design. Droop detectors that require non-standard blocks or external voltage routing may not fit nicely into existing digital flows.

C. NAND Divider Detector Design Methodology

The NAND divider detector has relatively few design variables, and they are relatively independent of the design specifications. We designed a single NAND divider detector for both the fast and slow design examples.

1) Input high-pass network: The first step is to design the input high-pass network. The size of the input AC-coupling capacitor *C* and the width of the NAND divider *W(INV)* affect the characteristics of the input high-pass filter that couple voltage droops into the high-pass node. The location of the pole should be set low enough to capture the lowest droop frequency of interest. For a fixed *W(INV)*, larger *C* improves the minimum detectable droop but increases the layout area of the detector.

In our design example, we select a standard two-input NAND gate and an AC-coupling capacitor with C = 300 fF. A two-terminal capacitor is required for AC coupling, ruling out the bypass capacitor standard cells. ASAP7 also does not provide MIM capacitors, so the AC-coupling capacitor must be implemented as a MOM capacitor. A 4-layer MOM capacitor can achieve roughly 2 fF/um sq. [9], so a 300 fF MOM capacitor requires 13x13 um of layout area.

2) Inverter chain: An odd number of inverters form an inverting amplifier chain from the high-pass node to the input of the latch. The number of stages *N(INV)* of the amplifier chain can be optimized for minimum delay. Our design makes the somewhat arbitrary choice of using *N(INV)* = 3.

3) Output latch: ASAP7 provides only one asynchronous latch, so this completes our design example.

What is achievable with this design? We performed transient simulations on the pre-extraction droop detector using Virtuoso (R) Spectre(R). After resetting the latch, the simulation applies droops of varying magnitudes to the detector input. It measures the delay between the start of the droop event and when the detector output goes high. Fig. 8 plots this droop detector delay against droop magnitude across the industrial temperature range. The curves end on the left when the droops are too small to trigger the detector; droops less than 20 mV in magnitude are not detectable under any condition. This detector is most sensitive at low temperatures, and its sensitivity varies by about 20 mV across the industrial temperature range. For higher magnitude droops, temperature variations only have the effect of increasing delay. For droops greater than 100 mV, the delay variation is almost negligible, clocking in at under 10 ps. Fig. 9 plots the detector delay across process corners. Process corners do not substantially change the minimum detectable droop, only the delay.

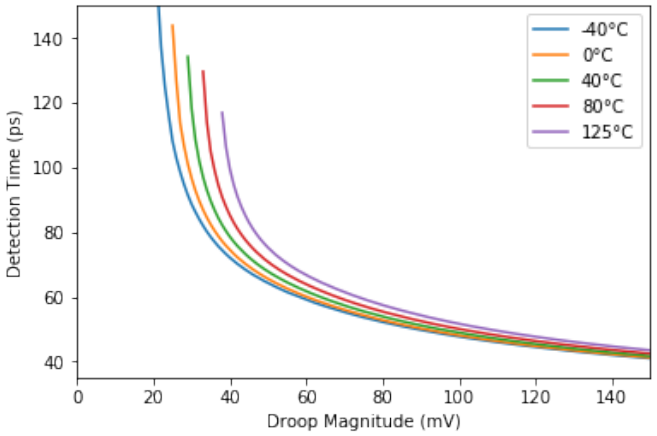

Fig. 8. NAND divider detector design example delay vs. droop magnitude simulated across the industrial temperature range

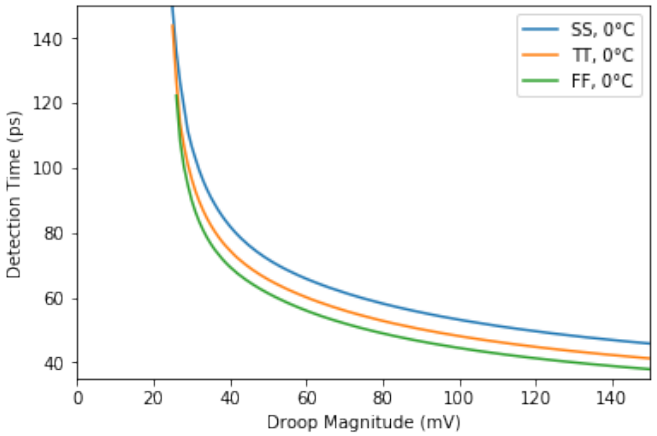

Fig. 9. NAND divider detector design example delay vs. droop magnitude across simulated process corners

The performance of this detector for both design examples is summarized in Section IV. For both design examples, the worse-case minimum detectable droop is 38 mV, which represents just a little over 5% for a 0.7 V nominal supply. Even at worse-case, this detector responds in under 100 ps at the minimum detectable droop, easily meeting both design requirements.

D. Ring Oscillator Detector Design Methodology

The fully-digital ring oscillator droop detector has even fewer design variables, but they do depend on the design specifications.

1) Ring oscillator stages: The first design variable is the number of stages in the ring oscillator, *N*. In a ring oscillator, *N* must be an odd number greater than one. Increasing *N* decreases the frequency of oscillation, and as we will show later, decreases the sensitivity of the droop detector. The frequency of a ring oscillator is highly dependent on technology parameters and is best found through simulation.

We chose an initial $N = 3$ ring oscillator and simulated its pre-extraction transient response at various supply voltages with a small timestep and an appropriate initial condition to ensure reliable onset of oscillation. The period of oscillation is measured after settling and plotted against the supply voltage across the industrial temperature range in Fig. 10. In the TT corner, the period of oscillation nearly doubles from 15 ps at nominal supply to 30 ps at a 200 mV droop. At the same time, period changes by around 15% across the industrial temperature range. This means that temperature changes across the operating range change the detector sensitivity by roughly 30 mV in the TT corner. This, however, is not as problematic as the period of oscillation across process corners shown in Fig. 11. At nominal supply voltage, period varies by as much as 25% from 14 ps to 18 ps due to process variations. The situation is worse at higher droops, varying as much as 40% for a 200 mV droop. While increasing *N* could decrease the transistor-to-transistor variation in the ring oscillator, it would not improve chip-to-chip variations and would increase the period of oscillation.

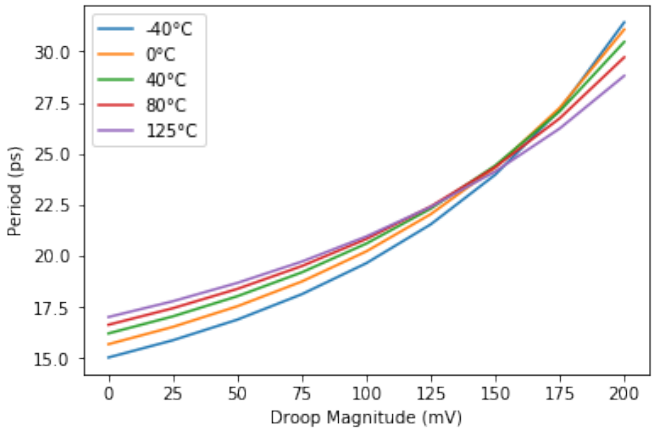

Fig. 10. Simulated pre-extraction ASAP7 ring oscillator period vs supply voltage across industrial temperature range

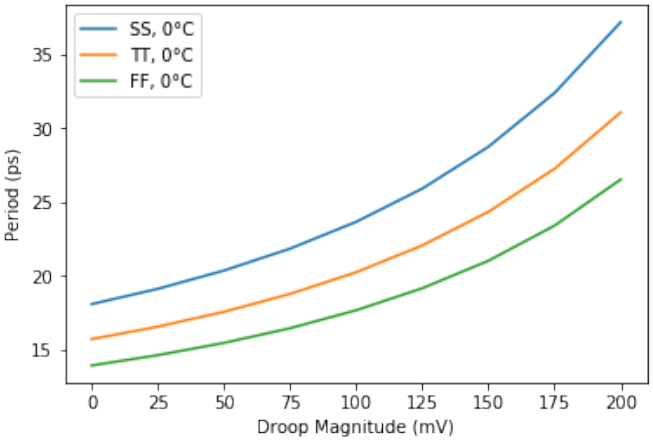

Fig. 11. Simulated pre-extraction ASAP7 ring oscillator period vs supply voltage across process corners

2) Counter threshold: The counter threshold *M* sets a lower bound on the ring oscillator frequency before which the detector signals a droop event. Leaving some room for margin, the maximum period of the $N = 3$ ring oscillator across temperature and process variations is $T_{nom} = 20$ ps. At nominal supply voltage, $M = T_{detection}/T_{nom}$, where $T_{detection}$ represents the half-cycle detection time. This corresponds to roughly $M = 6$ counts during a 125 ps half-cycle in the fast design example and $M = 31$ counts during a 625 ps half-cycle in the slow design example.

How does this detector perform? In the ring oscillator detector, *M* acts like a time amplifier. If the ring oscillator period changes by ΔT, the change is only detectable after *M* cycles if $M\Delta T > T_{nom}$. Rearranging, the minimum resolvable time difference after *M* cycles is:

$$\Delta T_{min} = T_{nom}/M$$

Substituting for *M*, we find that:

$$\Delta T_{min} = T_{nom}^2/T_{detection}$$

Lower ΔT_{min} represents a more sensitive droop detector. From this equation, it is evident that decreasing T_{nom} by reducing *N* or increasing $T_{detection}$ improves detection sensitivity. Of course, both T_{nom} and the relationship between ΔT and the droop magnitude have non-linear temperature and process dependence. For example, in the fast design example, ΔT_{min} is about 2.84 ps. From Fig. 10, this period difference corresponds to 70 mV in the worse-case temperature corner. However, from Fig. 11, the period can vary by as much as 5 ps solely due to process variations; the sensitivity is limited by process variations. This brings the worse-case sensitivity of the fast design example to 106 mV. And, even though sensitivity improves for the slow design example where *M* is greater, the minimum detectable droop is still in the process-variation-limited regime. The performance of this detector for both design examples is also summarized in Section IV.

IV. Conclusion

Table I summarizes the design metric comparison between the NAND divider and the ring oscillator detectors designed

TABLE I
SUMMARY OF DESIGN METRIC COMPARISON

Design Metric	NAND Divider	Ring Oscillator
Min. Detectable Droop (mV)	38	106
Temp. Sensitivity (uV/K)	110	360 / 450
Proc. Sensitivity (mV)	2	7.4 / 1.9
Fully Synthesizable	No	Yes

in Section III for our two design examples. The minimum detectable droop is a measurement of sensitivity. It is given in mV at the overall worse-case temperature and process corner. This represents the minimum droop that the sensor can detect reliably across all temperature and process variations. We measure selectivity with two different metrics, one considering temperature variations, and the other considering process variations. Temperature selectivity is measured by taking the difference between the high and low temperature minimum detectable droop in the TT corner and dividing through by the temperature range. The resulting figure is reported in uV/K, where lower numbers are better. Process selectivity is measured in a similar way, but considering only the temperature T = 0C and taking the difference between the minimum detectable droop at the FF and SS corners. The resulting figure is reported in mV, where lower numbers are better. Finally, design ease is approximated by whether or not the design is fully synthesizeable. Where two figures are given, the first and second number represent the design metric for the fast and slow design examples, respectively. The same NAND divider detector was used to satisfy both design requirements. Only the counter threshold M was changed between the two ring oscillator detector designs.

From the design metrics considered in this study, it is clear that the NAND divider droop detector performs better than the ring oscillator droop detector in ASAP7, both in terms of its detection speed and selectivity for both example designs. As noted in Section III, the design metrics chosen here substantially penalize the ring oscillator detector for its process variability. The upshot is, though a fully-digital ring oscillator droop detector may be attractive for some applications, the NAND divider detector outperforms the ring oscillator detector by nearly 3x in each of the metrics discussed here for fast single-cycle droop detection in ASAP7.

REFERENCES

[1] J. Crossley et al., "BAG: A designer-oriented integrated framework for the development of AMS circuit generators," 2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), San Jose, CA, 2013, pp. 74-81.

[2] K. L. Wong, T. Rahal-Arabi, M. Ma and G. Taylor, "Enhancing microprocessor immunity to power supply noise with clock-data compensation," in IEEE Journal of Solid-State Circuits, vol. 41, no. 4, pp. 749-758, April 2006.

[3] J. Oh et. al, "A 480mA Output-Capacitor-Free Synthesizable Digital LDO Using CMP-Triggered Oscillator and Droop Detector with 99.99% Current Efficiency, 1.3ns Response Time, and 9.8A/mm 2 Current Density," ISSCC 2020.

[4] Y. D. Kim et. al, "A 7nm High-Performance and Energy-Efficient Mobile Application Processor with Tri-Cluster CPUs and a Sparsity-Aware NPU," ISSCC 2020.

[5] S. Bang et. al, "A Fully Synthesizable Distributed and Scalable All-Digital LDO in 10nm CMOS," ISSCC 2020.

[6] J. Tschanz, K. Bowman, S. Walstra, M. Agostinelli, T. Karnik and V. De, "Tunable replica circuits and adaptive voltage-frequency techniques for dynamic voltage, temperature, and aging variation tolerance," 2009 Symposium on VLSI Circuits, Kyoto, Japan, 2009, pp. 112-113.

[7] X. Sun et. al, "A Combined All-Digital PLL-Buck Slack Regulation System with Autonomous CCM/DCM Transition Control and 82% Average Voltage-Margin Reduction in a 0.6-to-1.0V Cortex-M0 Processor," ISSCC 2018.

[8] H. Mair et. al, "A 7nm FinFET 2.5GHz/2.0GHz Dual-Gear Octa-Core CPU Subsystem with Power/Performance Enhancements for a Fully Integrated 5G Smartphone SoC," ISSCC 2020.

[9] P. Chiu, M. Ker, "Metal-layer capacitors in the 65 nm CMOS process and the applicationfor low-leakage power-rail ESD clamp circuit," Journal of Microelectronics Reliability, 2014.