A Method for Mitigation of Droop Timing Errors Including a 500 MHz Droop Detector and Dual Mode Logic

Yizhak Shifman, *Graduate Student Member, IEEE*, Inbal Stanger, *Graduate Student Member, IEEE*, Netanel Shavit, *Graduate Student Member, IEEE*, Ramiro Taco, *Member, IEEE*, Alexander Fish, *Member, IEEE*, and Joseph Shor, *Senior Member, IEEE*

***Abstract*— A technique to mitigate timing errors induced by power supply droops is featured. We propose an inverter-based droop detector as well as dual mode logic (DML) to achieve a droop-resistant timing response. The droop detector is based on capacitor ratios and is thus less sensitive to process/voltage/temperature (PVT) and to random offset than the prior art. The DML can alter its power/performance ratio based on the droop level input it receives from the detector, such that the critical timings are preserved. A prototype instantiating a demo of the scheme was fabricated in a TSMC 65 nm process, incorporating a simultaneous three-level detector and a DML-based ripple carry adder (RCA). The droop detector consumes 62 μW, has a response time of 2 ns, and an accuracy of 0.9% of Vdd, making it one of the fastest, most accurate, and lowest power droop detectors in its class. The RCA can maintain timing for voltage droops up to 400 mV. A potential supply level reduction of up to 12% was demonstrated for the RCA, and a similar reduction could be achieved with larger-scale DML digital circuits as well.**

***Index Terms*— Adaptive circuits, droop detector, droop mitigation, dual mode logic (DML).**

I. Introduction

THE supply voltage (Vdd) level of systems-on-a-chip (SoCs) is a critical parameter that determines their power consumption and performance. Typically, decoupling capacitors of varying sizes are placed along the power delivery network (PDN) on the board, the package, and the die to stabilize Vdd across a wide frequency range. However, the combination of these capacitors with the parasitic inductance of the PDN may introduce resonant frequencies to the network. Typically, three such resonances are observed, due to the respective resonances of the board, the package, and the die, at frequencies ranging between 0.5 and 100 MHz [1], [2]. During current surges, these inductive resonances may result in Vdd droops at the corresponding frequencies.

Manuscript received February 11, 2021; revised May 10, 2021; accepted June 10, 2021. Date of publication July 5, 2021; date of current version January 28, 2022. This article was approved by Associate Editor Piero Malcovati. This work was supported by the Israel Ministry of Science. *(Yizhak Shifman and Inbal Stanger contributed equally to this work.) (Corresponding author: Yizhak Shifman.)*

The authors are with the Alexander Kofkin Faculty of Engineering, Bar-Ilan University, Ramat Gan 52900, Israel (e-mail: yizhak.shifman@biu.ac.il).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/JSSC.2021.3091586.

Digital Object Identifier 10.1109/JSSC.2021.3091586

Traditionally, these droops were treated by the addition of a guardband to the supply voltage [1], such that even in the presence of a droop, the Vdd level is still sufficient to enable the digital circuits to withstand the timing constraints of the system. This method, however, imposes significant additional power consumption [3], which scales with Vdd^2 and therefore with the square of the added guardband. Recently, more power-efficient solutions have been suggested to detect and mitigate such droops.

Droop detection techniques can roughly be divided into the two categories of digital and analog. Digital detectors typically translate droops to delay variations of logic gates. Some of these detectors utilize a delay line whose delay is dependent on the ac Vdd level [4]–[8], while others use a ring oscillator [9], [10] to modulate the Vdd droops on its output frequency. These solutions are fully digital and therefore simpler to design and implement, and provide a high-resolution indication of the supply level. The delay or the frequency of these detectors, however, is also highly dependent on other parameters, such as the dc Vdd level, temperature, and aging, and thus require extensive backend characterization and calibration across a matrix of these three dimensions. The accuracy of these solutions is highly dependent on this 3-D calibration as well as noise, which limits their utility in real-time applications.

Analog detectors [3], [11] utilize an analog circuit, such as a comparator, to determine when the Vdd crosses a predefined threshold, which indicates a droop. For example, Floyd *et al.* [11] utilize a precise voltage reference and four parallel comparators to generate a four-level detection signal. The analog detectors may only require calibration at one or two temperatures but do require a precision reference voltage.

To mitigate the detected droops, previously proposed methods include adaptive clocking [8], [11], instruction throttling [6], and charge injection [4], [5]. During adaptive clocking, the clock frequency is temporarily reduced to withstand the larger gate delays during a droop. Instruction throttling achieves a similar effect by delaying the execution of instructions, and charge injection temporarily increases the supply level, for example, by un-gating additional power gates. A disadvantage of adaptive clocking and instruction throttling is that they also temporarily reduce the performance of the device. If the droops occur very frequently, this performance loss

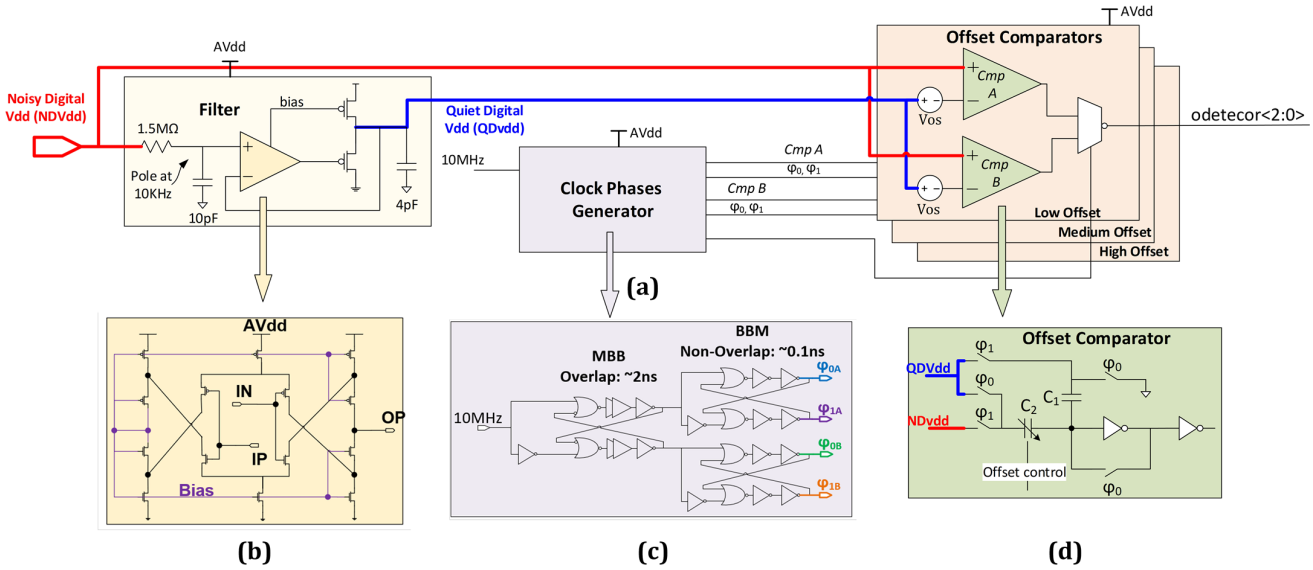

Fig. 1. (a) Block diagram of the droop detector top-level, comprising a parallel three droop levels detector, a filter, and a phase generator. (b) Transistor-level diagram of the VCDA instantiated in the filter. (c) Diagram of the clock phase generator, consisting of one MBB circuit and two BBM circuits. (d) Schematic of the detector's offset comparator.

could be significant. A disadvantage of charge injection is that it requires the power supply to be regulated on-die, and cannot support an externally generated supply, for example. Additionally, all of these solutions require the mitigation scheme to tune a circuit, such as the PLL or the regulator that typically serves multiple clients within the SoC. This may have undesired outcomes for some of these clients. It also renders the integration of the mitigation system more complicated, as multiple end-cases and cross-effects have to be considered.

In this article, a novel droop-induced timing errors mitigation system is suggested, which outperforms these previous works in a few important metrics. The proposed system comprises an inverter-based droop detector and dual mode logic (DML) [12]–[18]. Compared to the prior art, the detector is more accurate, consumes less power, more process/voltage/temperature (PVT) independent, and does not require 3-D characterization or precise analog references. To mitigate the timing errors caused by the droops, the critical paths are identified. Upon detection of droops, they alternate between the static and dynamic DML operating modes, which alter their delay, such that the timing constraints are maintained. Compared to the prior art, this solution does not degrade the performance of the system, does not tune a global system property, and does not require a specific power regulator, which makes it easier to implement. Additionally, as DML introduces an enhanced power/performance tradeoff relative to conventional CMOS, this solution also enables the integration of these DML advantages into the system. A TSMC 65 nm prototype was manufactured and measured, consisting of a droop detector and a DML ripple carry adder (RCA). The droop detector demonstrated a >500 MHz droop frequency detection coupled with a $\pm 0.9\%$ accuracy, and the RCA was shown to tolerate droops of >200 mV at Vdd $= 1.2$ V with no performance degradation.

The remainder of this article is organized as follows. In Section II, the droop detector architecture and its modes of operation are introduced. In Section III, the DML architecture and its operation to achieve droop mitigation are explained. Section IV provides the details of the manufactured proof-of-concept prototype. Section V presents the measurement results and Section VI discusses these results and compares them to the prior art. Finally, Section VII concludes the article.

II. Droop Detection by Offset Comparators

A block diagram of the proposed droop detector is presented in Fig. 1(a). The detector employs: 1) a supply filter and buffer; 2) three pairs of parallel-clocked offset comparators; and 3) a clock phase generator. The detector's input is the Noisy Digital Vdd (NDVdd) and it outputs a 3 bit thermometer code indicating the droop size, where a "000" indicates a 'no droop' state and a "111" indicates a droop that exceeds the highest droop threshold.

A. Supply Filter

The supply filter provides the nominal Quiet Digital Vdd (QDVdd) value as a reference for the offset comparators. It comprises a low-pass RC filter with a 10 KHz pole ($R = 1.5$ MΩ, $C = 10$ pF), a decade below the lowest anticipated resonance frequency [1], [2], and a unity-gain buffer which drives its output to the comparators and charges their capacitors. The unity gain buffer is a very-wide-common-mode-range differential amplifier [VCDA, Fig. 1(a)] [19] followed by a PMOS source follower, which was required due to the VCDA's high output impedance. The VCDA was selected for its bias-less architecture, its wide common mode range, and its low power consumption. The source follower drives the current required for charging the capacitors at the inputs to the offset comparators. A PMOS follower was selected due to the required high output voltages. A 4.5 pF capacitor was instantiated at the output of the buffer to filter the switching noises of the comparators and avoid errors due to the limited

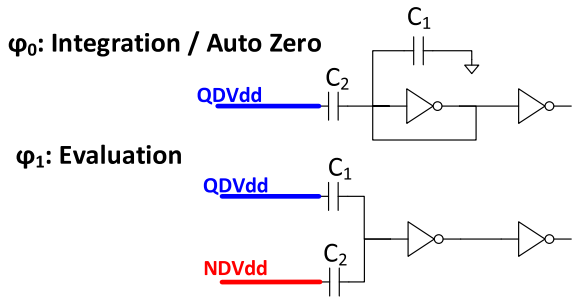

Fig. 2. Simplified comparator diagram for each of its operation phases, integration, and evaluation.

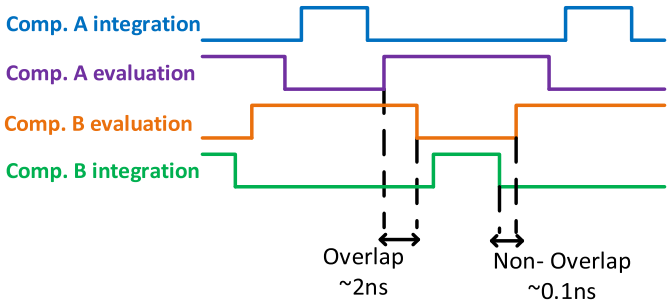

Fig. 3. Timing diagram of the four phases at the output of the clock phase generator.

slew rate of the buffer. The bandwidth of the buffer with the 4.5 pF capacitor at its output ranges (in simulations) between 150 KHz and 10 MHz across PVT.

B. Offset Comparator

For the offset comparator [Fig. 1(d)], we propose a new inverter-based comparator architecture. This comparator is auto-zeroed for offset suppression and utilizes the C_1/C_2 capacitance ratio to achieve the required offsets. During its integration phase (Fig. 2), φ_0, it is auto-zeroed, with C_2 charged to QDVdd minus the inverter trip point V_{trip}, and C_1 charged to V_{trip}. In its evaluation phase, φ_1, C_1 is connected to QDVdd and C_2 is connected to NDVdd, such that the charge is redistributed between the capacitors. By charge conservation, the inverter trips if the droop size, i.e., QDVdd–NDVdd, equals QDVdd$\cdot$ C_1/C_2. A higher droop size yields a "0" at the comparator output. In other words, the comparator offset is given by

$$V_{\text{os}} = \text{QDVdd} \cdot \frac{C_1}{C_2}. \tag{1}$$

Note that the desired outcome of (1) is that V_{OS} scales with QDVdd, such that the offsets are a percentage of NDVdd. The thresholds are tuned by switching the capacitance legs of C_2. A second inverter is instantiated after the auto-zeroed inverter to provide additional gain and amplify the signal to the logic levels. Both the inverters are identically sized to about X4 of the minimum, and their layout is matched to suppress mismatch.

This offset comparator architecture is less PVT-dependent than the prior art since: 1) the inverter structure provides a high gain across operation conditions and 2) C_1 and C_2 are MOM capacitors, which have a negligible leakage, a low offset, and a virtually PVT-independent ratio. In addition, its random offset is nearly eliminated by auto-zeroing. While the second inverter is not auto-zeroed, its design and layout match the first, and it is assumed that due to that and to the high gain provided by the first inverter, the error introduced by its trip-point mismatch is negligible.

For each of the three parallel thresholds, two comparators are instantiated in parallel such that when one evaluates (φ_1), the other integrates (φ_0), and the output of the evaluating comparator is selected. In the prototype discussed in Section IV, three comparator pairs are used, but a user could implement a different number as a function of the system's requirements.

C. Phase Generator

The clock phase generator, in Fig. 1(c), utilizes a 10 MHz clock and instantiates a make-before-break (MBB) circuit followed by two break-before-make (BBM) circuits to produce the four clock phases, $\varphi_{0\text{A}}$, $\varphi_{1\text{A}}$, $\varphi_{0\text{B}}$ and $\varphi_{1\text{B}}$, required for the two comparators. As depicted in Fig. 3, the two evaluation phases, $\varphi_{1\text{A}}$ and $\varphi_{1\text{B}}$, overlap by about 2 ns to mask transition effects between the integration and evaluation phases and prevent a dead-zone in the output. The evaluation and integration phases are separated by about 0.1 ns to prevent discharge of the capacitors upon transition. In addition to the four phases, this block also generates a synchronized selector for the output multiplexors of the comparator pairs, such that the output of the evaluating comparator is driven to the detector's output. Note that if leakage of the floating node at the input to the comparator is of concern, the clock frequency could be increased with no performance penalty besides additional power consumption.

The entire detector is supplied by a quiet Analog Vdd (AVdd), which is generally present in most systems as the IO or PLL supplies [20]. Such requirement is common to almost every droop detector which is not delay-based [3], [9]–[11]. It is also possible to add a simple voltage regulator to the circuit as was done, for example, in [21] where a similar auto-zeroed inverter was used in a temperature sensor.

III. Droop Mitigation By DML

DML is a logic family which can switch between a low-power static mode, and a high-performance dynamic mode [12]–[18]. At the gate level, a generic DML gate consists of a conventional static CMOS gate with an additional clocked transistor for dynamic operation [Fig. 4(a) and (b)]. To enable an interface between static and dynamic domains, DML gates may also be implemented with a footer [Fig. 4(c)] or a header [Fig. 4(d)] clocked transistor [12]. In the first phase of the dynamic mode, the output is pre-charged (or pre-discharged) by the clocked transistor and in the evaluation phase it is either unchanged or discharged (charged). The network parallel to the clock transistor is usually sized for minimum capacitance, while the sizing of the complementary network is optimized

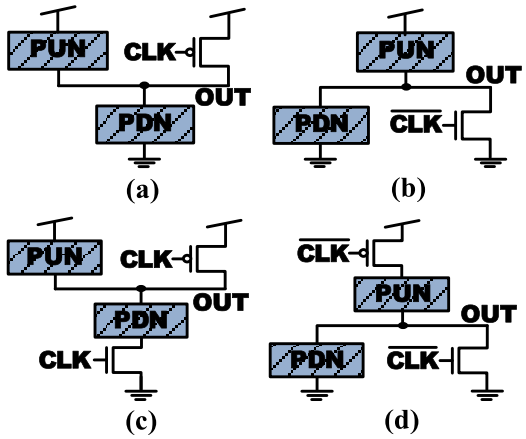

Fig. 4. Topologies of DML. (a) Type-A unfooted. (b) Type-B unheaded. (c) Type-A footed. (d) Type-B headed DML gates.

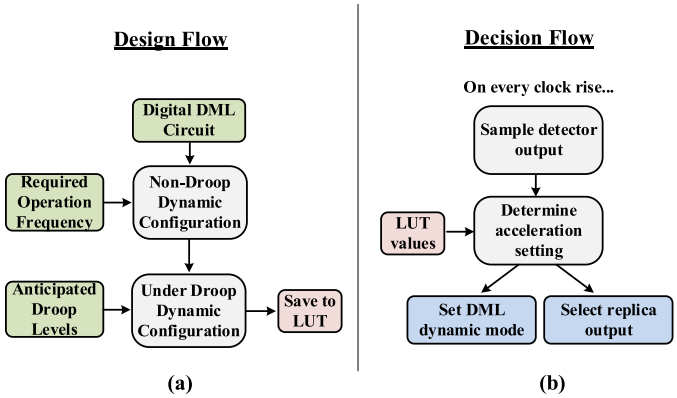

Fig. 5. Flowcharts describing (a) design of the DML droop mitigation and (b) acceleration decision algorithm, executed in every clock cycle.

for speed. Compared to a conventional CMOS, this sizing method reduces the input capacitance of the gate, such that higher speed in the dynamic mode and lower energy consumption in static mode are achieved. The DML design can be mixed, such that only the most critical paths are placed in the dynamic mode, while the rest of the gates save power in the static mode [13]. The paths can also be segmented, so that parts of the design can be gradually accelerated as needed from the static to dynamic modes to transition from low power to higher performance states. A critical path replica should be utilized during this segmentation to obtain optimal precharge/evaluation phases [4], [18]. The replica is entirely in the static mode and its inputs are configured to trigger the critical path, such that a signal toggling is observed at the output of the appropriate segment according to the DML setting. One example of a DML implementation is the 28 nm FD-SOI multiplier-accumulator (MAC) presented in [13]. Comparing to CMOS, this circuit provided an energy reduction of 35% at the expense of 34% frequency reduction in the static mode, and a 25% area reduction. In the mixed DML mode, it outperformed the CMOS speed by 46% and reduced its energy per operation by 9%. Additionally, in this circuit, only 4% more area is required for a critical path replica.

In this work, the above DML advantages are utilized to mitigate the timing errors resulting from supply voltage droops. The core idea is that the dynamic mode under a lower, or drooped, supply level, is at least as fast as the static mode under a higher, or undrooped, supply level. This enables the mitigation of droops by switching selected critical-path gates to the dynamic mode during a droop, such that timing constraints are maintained, and the circuit output is uncorrupted.

To enable effective droop mitigation, only the required gates are switched to the dynamic mode during the droop. This is done by identifying the critical timing paths and only accelerating the gates in these paths. Gradual acceleration is enabled by adopting the DML segmentation above. The identified paths are segmented, such that during a small droop, only some of the gates are accelerated by switching to the dynamic mode, but during a larger droop, more gates are accelerated. These critical paths could be identified using standard timing tools.

A flowchart of the overall proposed design process is presented in Fig. 5(a). First, a digital block is designed with DML gates and the requisite operation frequency is selected. Then, the nominal operation (no droop) is configured. The critical timing paths are identified and some of the DML gates in these paths are switched to the dynamic mode to enable operation at the required frequency. Finally, operation in the droop condition is determined. For a given nominal supply level (DVdd) and droop size, the critical paths are identified, and the level of acceleration is determined, such that during a droop, more gates in the critical paths are in the dynamic mode than in the nominal condition. A lookup table (LUT) is populated with the required acceleration levels for each DVdd and droop size. A flowchart of the correction algorithm is presented in Fig. 5(b). On every rising edge of the system clock, the thermometric 3 bit output of the detector is sampled. This input is considered together with the LUT data, to determine a cycle-specific acceleration mode. This DML acceleration mode is driven to two blocks: first, to the main DML circuit where it controls the gates' states, and second, to the replica where it selects the correct DML clock.

In [13], a MAC was demonstrated which could be accelerated in the dynamic mode to 46% faster than CMOS. If a droop-aware approach had been implemented, part of this acceleration could be used to match the performance of a CMOS MAC, while the rest could have compensated for droops.

IV. Proof-of-Concept Prototype

A system-level block diagram of the prototype is shown in Fig. 6(a). Droops are generated by an on-die droop inducer, which is a design-for-test circuit. NDVdd feeds a DML 128 bit RCA and is driven to a droop detector. The generated droops are sensed by the detector, which drives its output, "odetector," to the DML control block, which also receives the LUT data, "lut_data." This block determines the level of acceleration of the DML gates in the RCA, according to the droop size and the corresponding LUT value. The acceleration setting, "DML_mode," is driven to the replica, where it selects the replica-generated clock signal, and to the RCA device under test (DUT) itself, where it controls the states of the individual

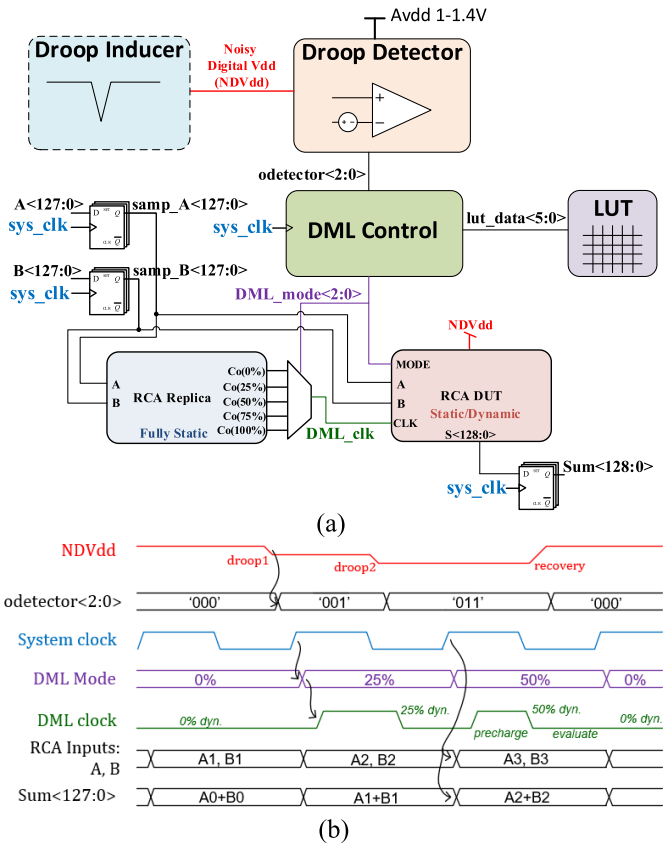

Fig. 6. (a) Block diagram of the droop mitigation scheme. (b) Schematic timing diagram depicting droop detection and mitigation.

TABLE I
DROOP DETECTOR CAPACITORS AND OFFSETS

Pair / Trim*		Low Offset		Medium Offset		High Offset	
	C_1	C_2	Offset**	C_2	Offset**	C_2	Offset**
#0	5.2fF	79fF	6.6%	41fF	12.6%	31fF	16.7%
#1		99fF	5.3%	51fF	10.1%	36fF	14.3%
#2		119fF	4.4%	62fF	8.4%	42fF	12.5%
#3		138fF	3.8%	72fF	7.2%	47fF	11.1%
#4		198fF	2.6%	--	--	--	--

*The offset of each comparator pair is set independently
** % of DVdd

DML gates. A⟨127:0⟩ and B⟨127:0⟩ are the inputs to the RCA, driven by the test environment. They are sampled by the system clock on each rising edge and driven to both the replica and the RCA DUT. The output of the DUT is also sampled on the rising edge, and the data is driven to the test environment to verify its correctness.

The parallel three-level droop detector, as in Fig. 1(a), detects the induced droops and outputs a 3 bit thermometer code, indicating the droop size. The thresholds of the three detectors are configurable, and set according to the value of C_2, as in (1). Table I depicts the capacitor sizes and the thresholds for each comparator pair. The thresholds are the design targets, calculated by substituting the capacitor values in (1). The trim

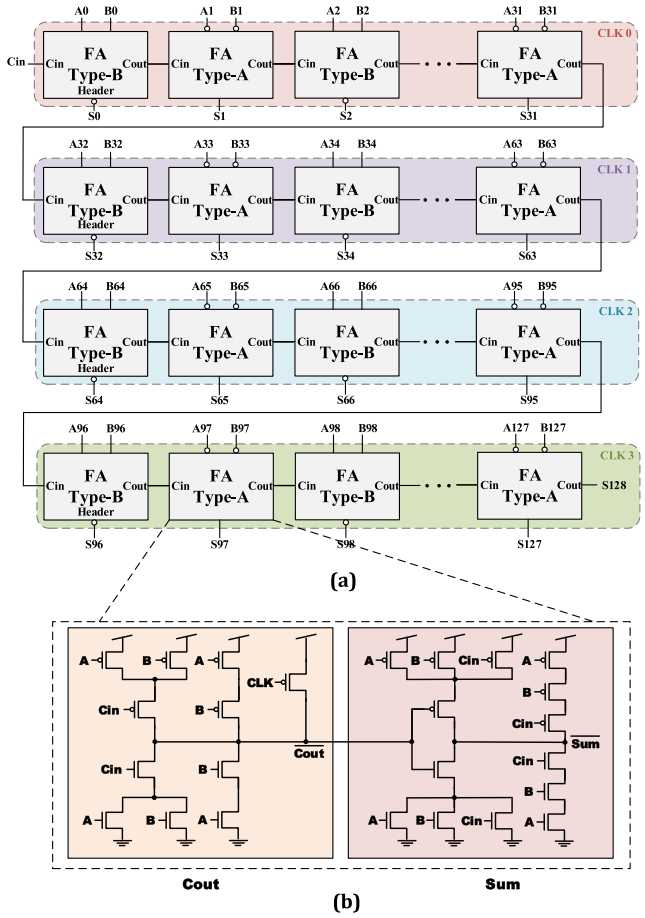

Fig. 7. (a) DML 128 bit RCA block diagram. The RCA comprises 128 DML FAs which are segmented into four segments, to enable acceleration in steps of 25%. (b) Transistor-level block diagram of a Type A DML single-bit full adder.

value of each pair is independent of the other pairs. Overall, the thresholds of the detector lie between 2.6% and 16.7% of the QDVdd.

The logic DUT is a 128 bit DML RCA. The RCA is segmented into four independent DML clock domains, such that five different DML clock modes (fully static, 25% dynamic, 50% dynamic, 75% dynamic, and fully dynamic) can be enabled, to achieve gradual acceleration. The 128 bit DML RCA, depicted in Fig. 7(a), is composed of 128 DML Type-A and Type-B interleaved full adders (FAs) propagating through a carry signal. The critical path of the RCA is obtained when a carry signal propagates through all the adders, thus, the Mirror FA [22] implementation was used. A generic Mirror FA is composed of two parts, the Sum, for summation, and the C_{out}, for the carry out calculation, both of them inverted. A transistor-level implementation of a single mirror DML FA in Type-A is presented in Fig. 7(b). The DML FA is based on the architecture of a CMOS mirror FA, with an additional clock transistor in parallel to the pull-up (Type-A) or pull-down (Type-B) networks of the C_{out} part, to boost the critical path. This facilitates a precharge (for Type-A) or pre-discharge (for Type B) of the carry signal during the clock pre-charge phase, such that during the evaluation phase the carry signal is changed, if needed, according to the inputs.

Since this FA implementation delivers the inverted value of the carry, for an RCA structure it reduces the critical path delay by eliminating one inverter for each FA. However, the use of the mirror adders requires inverting the inputs for even FAs and inverting the outputs for odd FAs in the RCA implementation.

To allow for proper DML operation in the partially dynamic modes, footed or headed DML gates [Fig. 4(c) and (d)] are used at the interface of each DML clock domain. Additionally, the partially dynamic modes are designed such that the static gates precede the dynamic gates, so that the propagation through the static gates occurs during the pre-charge of the dynamic gates. Note that in this demo, the entire logic DUT is a critical path and the phase-generating replica is a copy of the DUT in the static mode, and therefore doubles its area. In most systems, such as the MAC in [13], the replica would only include critical paths, which are a small percentage of the total system, and its area overhead should be considerably smaller. For example, in [13] the area of the gates in the critical paths is only 4% of the design's area. The DML RCA lies between registers, so that the result needs to be obtained within a single tunable clock cycle.

The 3 bit thermometer code is driven from the detector to a DML control circuit. The DML control utilizes this code, together with the input from a LUT, to determine the required DML clock mode, such that the RCA is accelerated based on the droop detector's output to preserve the circuit's frequency. Because the droop detector output is not synchronized, the DML control is sampled every clock cycle, such that the acceleration level of the DML remains constant throughout the cycle. The correction is thus accomplished within a resolution of one clock cycle. The LUT is preloaded with the acceleration values for different droop sizes and DVdd values, based on worst-case delay measurements across temperature.

The timing diagram in Fig. 6(b) depicts an operation example of the prototype. It assumes that according to the LUT, once odetector⟨0⟩ is "1," the DML is accelerated to 25% and once odetector⟨1⟩ is "1," the acceleration is 50%. Upon a rising edge of sys_clock, the odetector is sampled and DML_mode is set, which determines DML_clock. Note how the precharge phase of DML_clock, in green, changes due to the dynamic settings: for 25% dynamic it is longer, since the precharge phase represents the delay of 3/4 of the gates (which are in a static mode), while for 50% it is shorter as it represents the delay of only 1/2 of the gates. For 0% dynamic, a precharge phase is not required as no gates are in dynamic mode. Eventually, this operation results in a correct summation at Sum⟨128:0⟩.

A droop inducer was integrated to enable prototype testing. The inducer is capable of generating square droops, with controllable duration and size. For clarity, the duration of the droop is denoted as its width, and its size is also denoted as its depth. Fig. 8 presents a block diagram of the inducer, which comprises an analog multiplexer that switches between two external supplies and is controlled by one of two controllable-width pulse generators. The supply with the higher level, VddH, represents the nominal NDVdd and the supply with the lower level, VddL, represents the drooped NDVdd. During an induced droop, NDVdd switches from

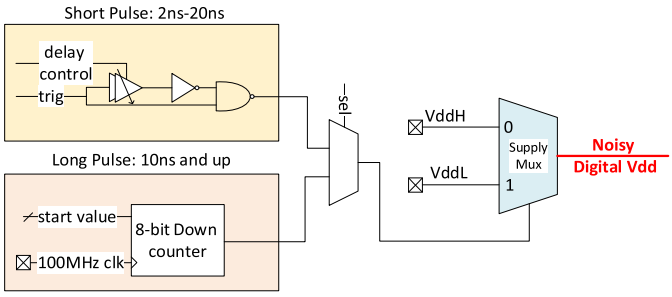

Fig. 8. Droop inducer, comprising an analog mux which switches between two external Vdds. The mux selector is a variable length pulse, starting at 2 ns.

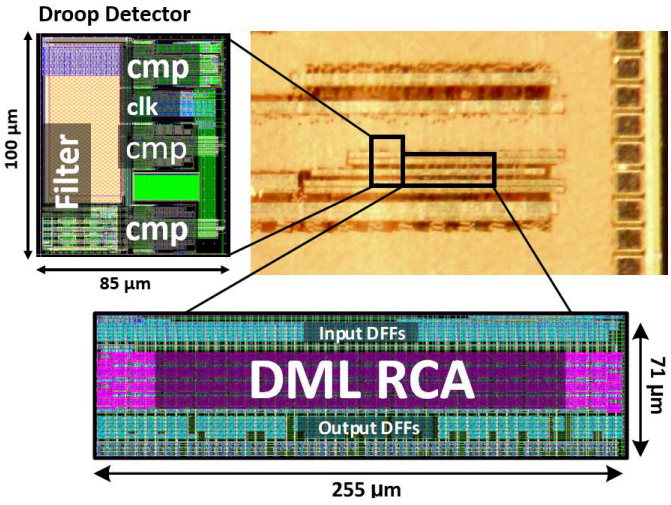

Fig. 9. Die micrograph and layout screenshots of the droop mitigation prototype, fabricated in TSMC 65 nm process node.

VddH to VddL and back to the VddH. The mux selection signal is a variable-width pulse generated by one of two mechanisms. The first is based on a variable delay which can apply pulses from 2 to 20 ns in steps of 2 ns, and the second is a down counter, which counts the cycles of a 100 MHz clock and can impose droops of 10 ns or longer in steps of 10 ns. This structure makes it possible to control both the width of the droop, by changing the selection pulse width and the depth of the droop, by tuning the analog inputs of the mux. The fast transition to droops, about 100 ps (simulated) enables accurate testing of the delay of the detector. The mux structure enables thorough characterization of the detector and the mitigation system across a wide and controlled range of droops.

V. Measurement Results

A demonstration of the droop mitigation system was designed and manufactured in TSMC 65 nm. A die micrograph and the layout of the prototype are shown in Fig. 9. Several chips were measured across a temperature range of −10 °C to 90 °C, an AVdd range of 1–1.4 V, and a DVdd range of 0.6 V to AVdd-0.1V. Since the filtered NDVdd is buffered by the unity gain buffer which is fed by AVdd (Fig. 1), a 100 mV difference must be maintained between AVdd and DVdd to prevent an offset in the buffer output.

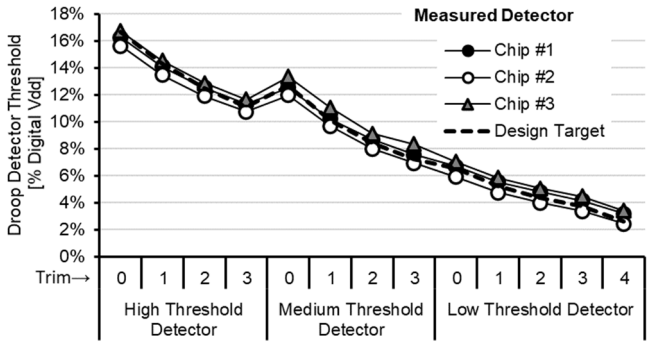

Fig. 10. Measured droop detector thresholds (trip points) across three chips relative to the design target.

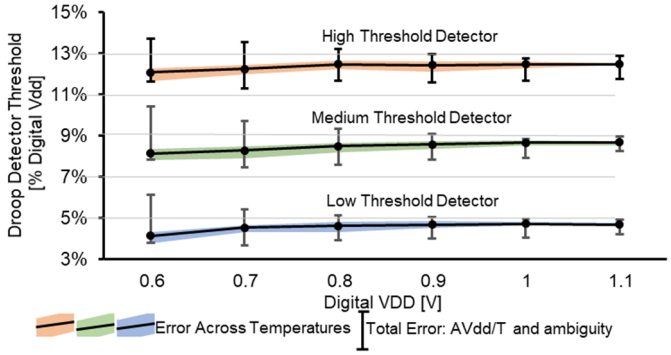

Fig. 11. Measured droop detector thresholds (trip points) including the errors from all error sources: A/DVdd, temperature, noise, and mismatch.

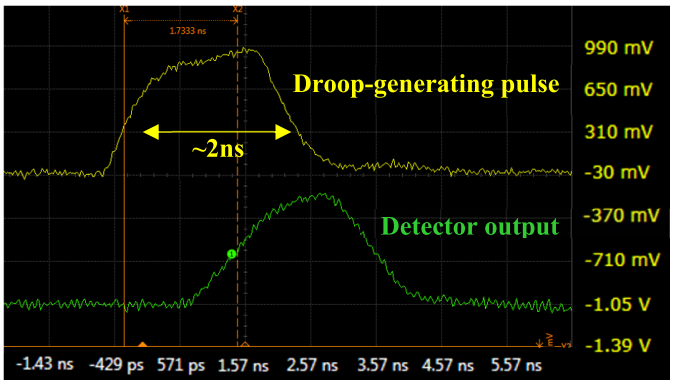

Fig. 12. Measured oscilloscope waveforms depicting the detection of a 2 ns droop. Droop-inducing pulse (top). Detector (comparator) output (bottom).

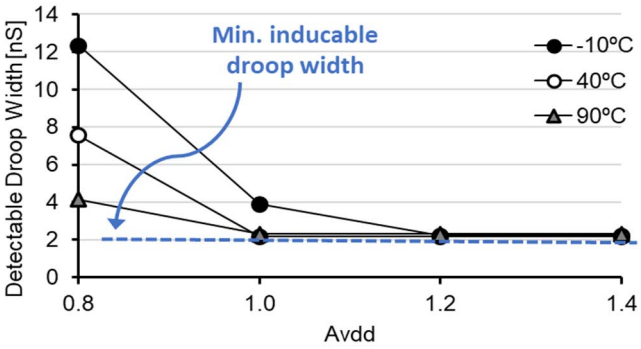

Fig. 13. Measured minimum detectable droop width across AVdd and temperature. For AVdd $\geq$ 1 V, the entire required droop frequency range is covered.

A. Stand-Alone Droop Detector

Fig. 10 presents the design target, given by the capacitance ratio calculated by (1) and in Table I, relative to the measured droop detector thresholds across the three chips, measured at a nominal condition (40 °C, AVdd = 1.2 V and DVdd = 1 V). It shows that the chip-to-chip and chip-to-target distributions are tight. The maximal chip-to-chip difference is 1.4% of DVdd, and the maximal chip-to-target difference is 1.1% of DVdd. Since there are multiple thresholds per comparator, this difference can be calibrated at a single temperature.

The sources of inaccuracy of the detector include charge injection from the switches to the capacitors, thermal noise, Vdd/temperature variations, and gain mismatch. For 0.8 V $\leq$ DVdd $\leq$ 1.1 V, these errors accumulate up to ~0.9% (7 mV) ambiguity of the comparator trip-point, i.e., its threshold. If the DVdd range is extended down to 0.6 V, the error increases up to ~2.3% (14 mV), in part due to offsets in the buffer and in part since when the DVdd is lower, a given error in mV translates to a larger percentage error. The measured total error of the detector is presented in Fig. 11, for step #2 of each of the three sub-detectors in one of the measured chips. For each DVdd value (x-axis), the error across all the measured range of AVdds and temperatures is presented. The colored bands around the curves represent the temperature-related error only, which translates to an error of up to 0.43% in the worst case, for DVdd = 0.6 V. The total error is given in Fig. 11 in the error bars. The largest error contributor is the decreased comparator gain for AVdd = 1 V. In addition, simulations of skewed Si, which was not available for us to measure, demonstrate a low error of up to 2.2% across PVT. In the worst-case PVT, the leakage-induced degradation at the input to the comparator was about 1.5 mV during the 10 MHz clock period. Monte Carlo simulations show an error with a standard deviation of 0.35% of DVdd. Note that relative to the design target, no favorable error direction, either positive error or negative, could be observed in the measurements (Fig. 11). This indicates that the contribution of deterministic offsets, such as charge injection or parasitic capacitances, is minor.

In order to measure the speed of the detector, the delay time and the minimum detectable droop width were tested. As mentioned, the minimum droop width that the droop inducer was capable of producing was 2 ns. This can characterize the most important droops which are generally less than 100 MHz [1], [2]. In Fig. 12, oscilloscope waveforms of detection of a ~2 ns droop are presented. The top waveform is the pulse that generates the droop, i.e., the "select" signal of the analog supply multiplexer in the droop inducer (Fig. 8). The bottom waveform is the output of the droop detector. Note that due to a limited probing functionality, the transient waveform of NDVdd could not be captured. The minimum detectable droop widths across AVdd and temperature are charted in Fig. 13. For AVdd $\geq$ 1 V, the detector was capable of detecting the entire required droop frequency range. The delay times of the detector are depicted in Fig. 14

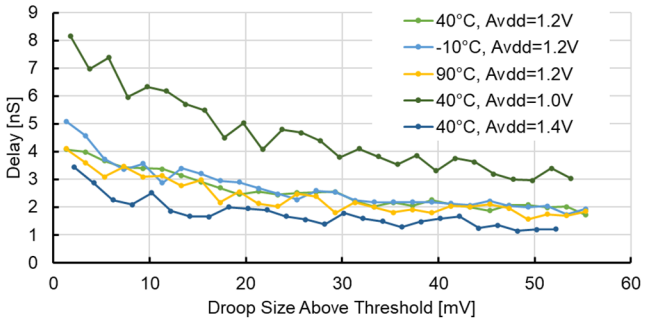

Fig. 14. Measured droop detector delay for varying droop sizes.

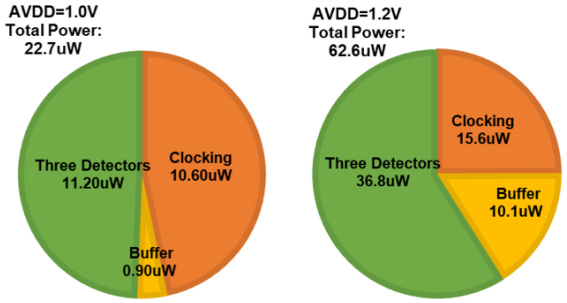

Fig. 15. Measured droop detector power break-up for AVdd = 1 and 1.2 V.

across AVdd and temperature and for varying droop depths. Cross-PVT simulations for droop size of 20 mV above the threshold show delay that varies between 1.1 and 2.7 ns for AVdd = 1.3 V, where simulations of typical material show a delay of 1.7 ns, very close to the measured delay at that condition, about 2 ns. The measured power consumption of the detector for typical material at 40 °C ranged from 22.7 μW for AVdd = 1 to 62.6 μW for AVdd = 1.2 V, as in Fig. 15. Simulated power consumption of skewed Si for the entire detector (three comparators, filter, and clocking) across voltage and temperature, ranged from 16.1 μW (SS, −10 °C, 1 V) to 116 μW (FF, 110 °C, 1.2 V). There is a clear tradeoff between power and speed demonstrated here, where for AVdd = 1.4 V the delay was less than 2 nS when the droop depth was more than 15 mV above the threshold, whereas for AVdd = 1.2 V the delay increased to about 2–3 nS and for AVdd = 1 the delay was 5 nS (Fig. 14). This is a result of the improved inverter–amplifier bandwidth.

B. DML RCA

To characterize the RCA, the adder was tested across the applicable DVdd and temperature range. For each DVdd, temperature, and DML acceleration setting, the delay time of the RCA was measured. This was done by gradually decreasing the clock cycle time until an erroneous response was captured in the output FFs. The results for T = 40 °C are depicted in Fig. 16. For each DVdd between 0.6 and 1.2 V, the delay decreased by roughly 10× between the 0% and 100% dynamic mode. It should be noted that a target delay could be obtained for different power supplies by an application of different DML clock modes. This demonstrates

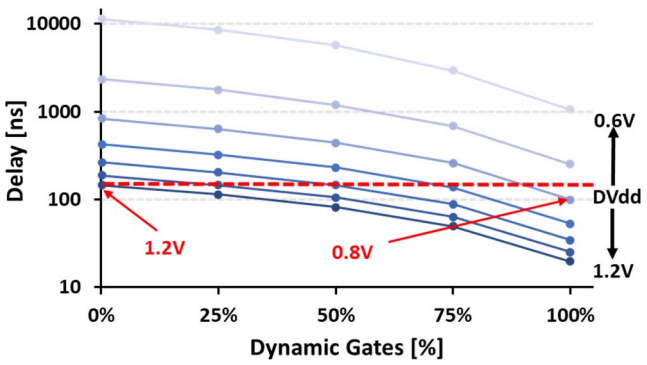

Fig. 16. Measured RCA delay time for varying DVdd and acceleration settings.

TABLE II
PROTOTYPE LUT

Droop Size [% DVdd]	Detector Output	DVdd [V]					
		1.2	1.1	1	0.9	0.8	0.75
0%-4%	000	0	0	0	0	0	0
4%-8%	001	50	25	25	50	50	50
8%-12.5%	011	50	50	50	75	75	100
12.5%-18%	111	75	50	75	100	100	100
		Required DML Acceleration [%]					

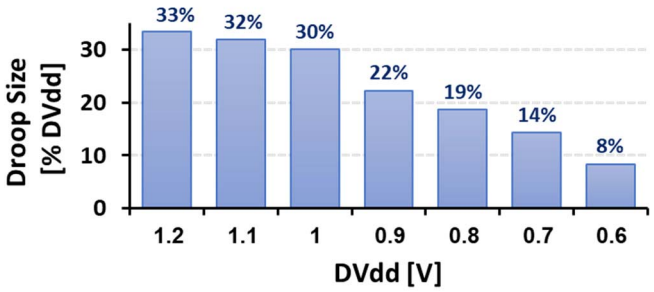

Fig. 17. Measured maximal correctable droop size from the DML RCA perspective. While these values exceed the detector's threshold by up to X2, the robustness of the DML as a droop mitigator is demonstrated.

the potential of DML gates in mitigating droops, as discussed in Section III. To derive the LUT, for each DVdd, the reference frequency was set to be the frequency in the fully static mode. The cross-temperature measurements were considered, and the worst-case required acceleration was utilized for each DVdd and anticipated droop size. The LUT of the required DML accelerations for different DVdd's and anticipated droops is presented in Table II. Note that the non-monotonicity is due to the different trim levels of the three comparator pairs.

According to these measurements, the maximal correctable droop was calculated for each DVdd level, as presented in Fig. 17. Droops as high as 33% could be corrected for DVdd = 1.2 V. While these droop sizes may exceed realistic anticipated droops as well as the droop detector's highest threshold, they demonstrate the robustness of the DML droop mitigation concept.

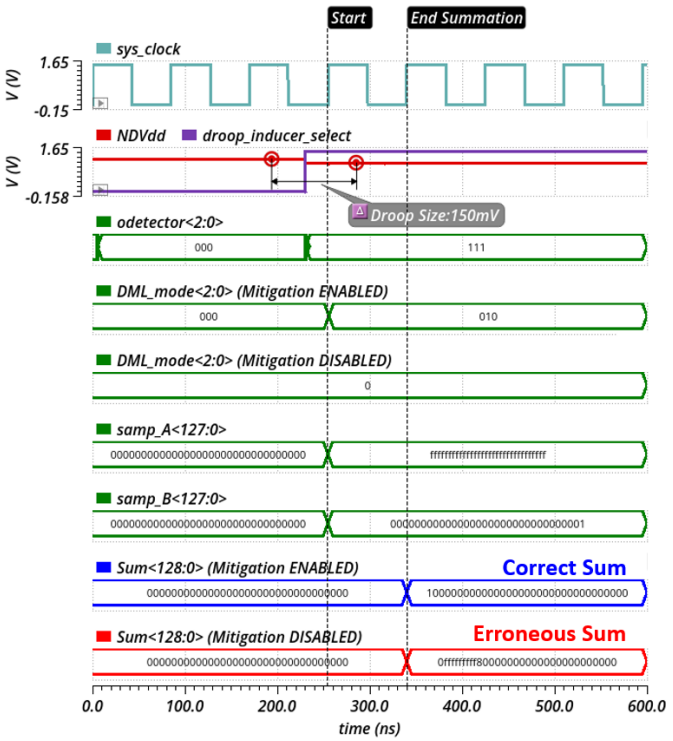

Fig. 18. Simulated waveforms of the RCA and the mitigation system (Fig. 6) when the mitigation is enabled and the correct sum is obtained, and when it is disabled and incorrect sum is sampled.

C. Droop Mitigation System

In this prototype, it is assumed that the required operating frequency is the fastest frequency that the fully static mode can support. This selection provides the lowest energy consumption for each additional operation. A simulated example of the operation of the mitigation system is provided by the waveforms presented in Fig. 18, which is a simulation of the system in Fig. 6. In this example, a 150 mV droop was injected and two scenarios were simulated. In the first, the mitigation system was enabled and in the second it was disabled. The clock period was set to 85 ns, the delay of the RCA when no droop is present, such that if the mitigation was disabled, the RCA's delay exceeds the clock period and an incorrect sum was sampled. Observe how when the mitigation was enabled, DML_mode was switched from "000" to "010" while when it was disabled, it remained "000." The inputs to the RCA are set to trigger the critical path, where "A⟨127:0⟩" is "all 1" and "B⟨127:0⟩" is "00…001." Fig. 18 presents the sampled version of A and B, after the flip-flops. When the mitigation is enabled, the correct summation result, "100…00" was sampled on "sum⟨128:0⟩" (blue waveform), while when it is disabled, an incorrect result is sampled (red waveform).

To measure the overall scheme performance, different droop depths and LUT values were applied, and the fraction of dynamic gates required to compensate for each given droop was found across the range of correctable droops. The results are presented in Fig. 19. Droop sizes larger than the maximal detectable droop size are also correctable, at the expense of larger power consumption during lower size droops. For

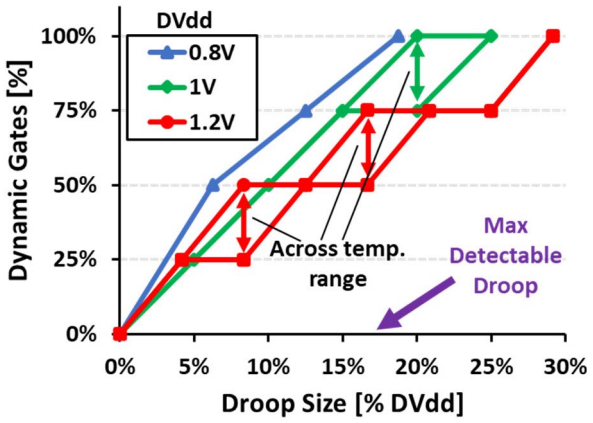

Fig. 19. Measured required LUT values of acceleration for % droop across DVdd and temperature to stand given droop sizes.

example, assume that the threshold of the detector is set to its maximum, 16.7% of the DVdd, and the required acceleration to withstand this droop size is 75% dynamic, as is the case for DVdd = 1.2 V. At this acceleration, droops of up to 25% DVdd could be corrected. A user, however, may wish to correct droops of 30% as well, and could set the LUT to accelerate to 100% dynamic when the 16.7% threshold is crossed, and his target is thus achieved at a tradeoff of excessive acceleration and power for droops of 16.7%–25% of DVdd.

To calculate the achievable power reduction of the DML RCA, the operation frequency was assumed to be the maximal frequency for DVdd = 1 V in the fully static mode. The benchmark for power reduction is the RCA power consumption when no mitigation is available. In this case, the minimal required DVdd of 1 V is increased by the maximal anticipated droop size, as well as 25 mV additional guardband. For example, if droops of up to 100 mV are anticipated, the required DVdd is 1.125 V. When a droop mitigation system is available, the minimal DVdd is not increased by the maximal droop size, but a larger guardband is required due to the minimum detectable droop size, so the DVdd equals 1 V + 50 mV guardband. The power savings are the differences between the power consumptions for the higher and lower DVdd's for the anticipated droops. The results are depicted in Fig. 20. If the anticipated droops are 15% of the DVdd, a 35% power reduction can be achieved.

Fig. 20 also depicts the potential power reduction if the RCA is operated in 25% or 50% dynamic mode, to enable faster operation frequency. Note that negligible power increases occur in the dynamic mode because of the lower voltage during the droops (e.g., an additional power of only 0.7% was measured as a result of the transition of 50% of the gates to the dynamic mode for NDVdd = 1.05 V and Droop = 150 mV).

VI. Discussion

A performance summary and a comparison to the recent prior art are presented in Table III. The DML mitigation scheme does not degrade the digital logic's performance, as do adaptive clocking and instruction throttling, and the DVdd and power reduction enabled by this scheme are the largest

TABLE III

PERFORMANCE SUMMARY AND COMPARISON WITH PRIOR-ART. (A) DROOP MITIGATION SYSTEM. (B) DROOP DETECTOR

	This work	[5] JSSC'20	[6] ISSCC'18	[11] ISSCC'17	[4] JSSC'17	[23] JSSC'16	[3] JSSC'04
A							
Process	65nm	10nm	14nm	14nm	22nm	16nm	90nm
Mitigation Scheme	Dual Mode Logic	Charge Injection	Instruction Throttling	Adaptive Clocking	Dynamic Power Gating	Adaptive Clocking	No Mitigation Presented
Digital Vdd Reduction	12% (DVdd=1.0V, Droop=166mV)	6.7%	--	--	11%	--	--
Power Reduction	25% (DVdd=1.0V, Droop=100mV) [1]	7%-14.5%	8%-14%	--	10%	5%(0.9V) – 13%(0.6V)	--
Performance Degradation During Droop	Not Degraded	Not Degraded	Degraded	Degraded	Not Degraded	Degraded	--
B							
Architecture	Offset Comparators	Delay Measurement	Delay Measurement	Flash A/D	Delay Measurement	Delay Measurement	Calibrated Comparator
Power [µW]	**22 (AVdd=1.0V) 62 (AVdd=1.2V)**	1480	--	--	--	2500	--
Area [um^2]	8,500	--	12,000	12,875	--	2590	--
Area [F^2] [2]	**2M**	--	61M	65M	--	10M	--
Detector Vdd Range [V]	1.0-1.4 (AVdd)	0.6-1	--	1.5	0.7	0.6-0.9	0.7-1.3
Digital Vdd Range [V]	0.8-1.2	0.6-1	--	--	0.7	0.6-0.9	1
Detectable Droop Magnitude [% Dig. Vdd]	2.6%-16.6%	0%-30%	--	--	--	10%	27%
Response Time	3.5ns (AVdd=1.0V) 2ns (AVdd=1.2V)	2ns	--	2.3ns	--	2.4ns	--
Threshold Temperature Dependence [% / 100°C]	**±0.4% (Dig. Vdd ≥ 0.8V)**	Not Reported [3]	Not Reported [3]	Not Reported	Not Reported [3]	Not Reported [3]	Not Reported
Accuracy	±0.9% (Dig. Vdd ≥ 0.8V)	±2.9%	±1%	±1% (Simulated)	--	--	--

1 – For the measured prototype. In a full DML system, this performance could be maintained if all the critical paths are identified and adjusted during droops.
2 – F is the minimum feature size of the process. The size in F^2 is given for fair comparison, as the compared works were manufactured in different processes.
3 – No specific data were stated in the paper. Delay-based circuits are known to have a substantial temperature dependence.

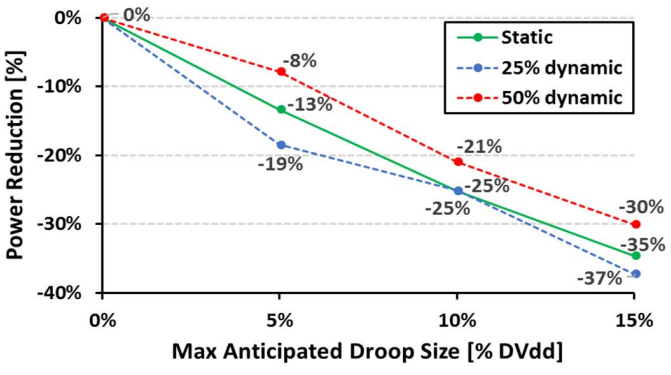

Fig. 20. Measured power savings of DML RCA versus the maximal anticipated droop size for various non-droop operation modes: static, 25% dynamic, and 50% dynamic.

presented in the literature. The proposed scheme exhibits the lowest detector area, the lowest detector power, and the highest detector accuracy. The delay time is comparable to the prior art, despite the fact that this work is in older technology. Prior detectors that rely on delay measurements require complex 3-D characterization across the dc voltage, temperature, and aging conditions. This droop detector is insensitive to these factors and only requires a calibration point at a single temperature to compensate for a random mismatch. In addition, the detector does not require a precision external reference as do other analog works, except for a quiet analog supply which generally exists in any system.

Although the measured demo represents a simplified droop mitigation system, the concept can be adapted to a complete system. Note, first of all, that the droop detector was not simplified, and a user could adjust the number of comparators and the thresholds according to needs.

In a full system, implementation aspects of the replica should be considered. First, in this demo, the replica doubled the power consumption of the RCA in the static mode since the entire RCA is a critical path. However, in a fully functional

system, the replica will only be a small percentage of the gates and will operate in the static mode alone. For example, the MAC presented in [13] had only 4% of the gates in its critical path. Therefore, it is assumed that critical path replicas would add roughly 3%–5% to the power and area of the primary DML circuit.

Additionally, the replica in this demo provided the same delay of the original path as it was an exact copy of the RCA. In more complex circuits, special care has to be taken during the replica design. An insufficient precharge phase may result in an increased overall path delay and an incorrect result, while a too long precharge phase results in an increased delay and reduced performance. The replica should comprise the same logic gates as the original path with similar loading gates, and interconnects as similar as possible for maximal matching across PVT. To confirm an error-free operation, it is recommended to slightly increase the replica's delay such that random mismatch may not result in an insufficient precharge phase. An analysis of replicas designed for the three most critical paths of [13] showed a delay difference of up to 2% between the replicas and the paths (without interconnects) and demonstrated the applicability of this method.

Clocking area and power are also overheads of DML. The clock switch increases the output capacitance of the gates and a distribution network must be implemented for the clock signals. However, as was previously shown [13], the benefits of the sizing reduction of the DML gates in terms of area and power are greater than the price paid for the clock distribution.

In multi-core, large SoCs, droops are not only temporal but also spatial. For example, one core could experience a droop while others not. To address this, multiple mitigation systems may have to be instantiated in a single SoC, where each includes a detector, DML gates and replicas, and DML control. This case would be similar for most prior-art solutions, where multiple detectors would have to be instantiated. In this respect, the proposed detector is advantageous over the prior art, since it is the smallest, lowest power and requires a relatively simple calibration compared to most of the prior art (especially delay-based detectors). From the mitigation standpoint, some of the other solutions may be globally implemented at the circuit level. For example, adaptive clocking may be able to adjust a single clock source, if it is applied to the entire chip. The proposed mitigation may have to be instantiated separately in each sensitive area. While this is more complex to design at the circuit level, it may be simpler to integrate at the global SoC level. When a global property, such as a clock frequency, is changed, multiple implications and end-cases have to be considered, whereas the integration of the DML is seamless from that perspective.

VII. Conclusion

Supply droops may result in a major increase in the power consumption of modern SoCs, since a voltage guardband is required to maintain timing. Recently, droop detection and mitigation schemes have been proposed which make it possible to decrease this guardband and significantly reduce the power consumption.

In this article, a droop detection and error mitigation scheme was proposed that outperforms previous schemes. This scheme comprises a droop detector and a mitigation system. The core of the detector is a novel inverter-based, auto-zeroed offset comparator, whose offset is determined by the ratio of MOM capacitors and scales to the DVdd. The detector filters the noisy Vdd to generate a clean reference, such that an external accurate analog reference is not required. This offset comparator architecture, to the best of our knowledge, was never introduced before in droop detectors. Its advantages over prior comparator-based droop detectors are the following: 1) it compares the noisy Vdd to its nominal value, which is easy to achieve using a filter. It does not require any external accurate reference; 2) it relies on the capacitors' ratio, which makes it less sensitive to PVT; and 3) it is compact, mismatch resistant, and low power. An additional advantage over digital, delay-based droop detectors is that it does not require a complex 3-D calibration (across temperature, supply voltage, and aging).

The mitigation system utilizes DML gates that can switch either to static mode (slow) or to dynamic mode (fast) on the fly. A digital design based on DML gates can maintain a given performance target during droop events. During droop events, the critical timing paths can adopt a proper configuration in which segments of the critical timing paths switch from the static to dynamic mode, to maintain delay time and avoid timing errors. This mitigation approach directly modifies the logic gates instead of modifying a global SoC property, making it easier to integrate. In addition, it does not degrade the performance of the system, even temporarily.

A prototype manufactured in TSMC 65 nm demonstrated the applicability of the proposed scheme. This prototype incorporated a detector capable of detecting three tunable droop thresholds simultaneously, with a DML RCA. Accurate and fast detection was shown, with a worst-case accuracy of $\pm 0.9\%$ of DVdd and a delay time of 2 ns, better or similar to the prior art. The power and area of the detector, 62 μW and 2 M F^2, respectively, were substantially lower than the prior art. The manufactured mitigation scheme, comprising the detector and the RCA, was shown to mitigate errors caused by droops as high as 16% of the DVdd, with a potential DVdd reduction of 12%. The mitigation concept demonstrated by the RCA may be adopted to larger scale digital circuits and still provide similar power improvement by following the same design principles.

References

[1] E. A. Burton *et al.*, "FIVR—Fully integrated voltage regulators on 4th generation Intel Core SoCs," in *Proc. IEEE Appl. Power Electron. Conf. Expo. (APEC)*, Fort Worth, TX, Mar. 2014, pp. 432–439.

[2] P. N. Whatmough, S. Das, Z. Hadjilambrou, and D. M. Bull, "Power integrity analysis of a 28 nm dual-core ARM cortex-A57 cluster using an all-digital power delivery monitor," *IEEE J. Solid-State Circuits*, vol. 52, no. 6, pp. 1643–1654, Jun. 2017.

[3] A. Muhtaroglu, G. Taylor, and T. Rahal-Arabi, "On-die droop detector for analog sensing of power supply noise," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 651–660, Apr. 2004.

[4] M. Cho *et al.*, "Postsilicon voltage guard-band reduction in a 22 nm graphics execution core using adaptive voltage scaling and dynamic power gating," *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 50–63, Jan. 2017.

[5] S. Bang *et al.*, "An all-digital, VMAX-compliant, stable, and scalable distributed charge injection scheme in 10-nm CMOS for fast and local mitigation of voltage droop," *IEEE J. Solid-State Circuits*, vol. 55, no. 7, pp. 1898–1908, Jul. 2020.

[6] C. Vezyrtzis *et al.*, "Droop mitigation using critical-path sensors and an on-chip distributed power supply estimation engine in the z14 enterprise processor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2018, pp. 300–302.

[7] A. Drake *et al.*, "A distributed critical-path timing monitor for a 65 nm high-performance microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2007, pp. 398–399.

[8] K. Bowman *et al.*, "A 16 nm auto-calibrating dynamically adaptive clock distribution for maximizing supply-voltage-droop tolerance across a wide operating range," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2015, pp. 1–3.

[9] E. Alon, V. Stojanovic, and M. A. Horowitz, "Circuits and techniques for high-resolution measurement of on-chip power supply noise," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 820–828, Apr. 2005.

[10] A. Roy and B. H. Calhoun, "Exploring circuit robustness to power supply variation in low-voltage latch and register-based digital systems," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Montreal, QC, Canada, May 2016, pp. 273–276.

[11] M. S. Floyd *et al.*, "26.5 adaptive clocking in the POWER9 processor for voltage droop protection," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2017, pp. 444–445.

[12] I. Levi and A. Fish, "Dual mode logic—Design for energy efficiency and high performance," *IEEE Access*, vol. 1, pp. 258–265, 2013.

[13] R. Taco, I. Levi, M. Lanuzza, and A. Fish, "An 88-FJ/40-MHz [0.4 v]-0.61-pJ/1-GHz [0.9 V] dual-mode logic 8 × 8 bit multiplier accumulator with a self-adjustment mechanism in 28-nm FD-SOI," *IEEE J. Solid- State Circuits*, vol. 54, no. 2, pp. 560–568, Feb. 2019.

[14] N. Shavit, I. Stanger, R. Taco, M. Lanuzza, and A. Fish, "A 0.8-V, 1.54-pJ/940-MHz dual-mode logic-based 16×16-b booth multiplier in 16-nm FinFET," *IEEE Solid-State Circuits Lett.*, vol. 3, pp. 314–317, 2020.

[15] I. Levi, A. Belenky, and A. Fish, "Logical effort for CMOS-based dual mode logic gates," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 22, no. 5, pp. 1042–1053, May 2014.

[16] I. Levi, A. Kaizerman, and A. Fish, "Low voltage dual mode logic: Model analysis and parameter extraction," *Microelectron. J.*, vol. 44, no. 6, pp. 553–560, Jun. 2013.

[17] N. Shavit, I. Stanger, R. Taco, and A. Fish, "Process variation-aware datapath employing dual mode logic," in *Proc. IEEE SOI-3D-Subthreshold Microelectron. Technol. Unified Conf. (S3S)*, Burlingame, CA, USA, Oct. 2018, pp. 1–3.

[18] I. Stanger, N. Shavit, R. Taco, M. Lanuzza, and A. Fish, "Silicon evaluation of multimode dual mode logic for PVT-aware datapaths," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 67, no. 9, pp. 1639–1643, Sep. 2020.

[19] M. Bazes, "Two novel fully complementary self-biased CMOS differential amplifiers," *IEEE J. Solid-State Circuits*, vol. 26, no. 2, pp. 165–168, Feb. 1991.

[20] M. Yuffe *et al.*, "A fully integrated multi-CPU, processor graphics, and memory controller 32-nm processor," *IEEE J. Solid-State Circuits*, vol. 47, no. 1, pp. 194–205, Jan. 2012.

[21] T. Oshita, J. Shor, D. E. Duarte, A. Kornfeld, and D. Zilberman, "Compact BJT-based thermal sensor for processor applications in a 14 nm tri-gate CMOS process," *IEEE J. Solid-State Circuits*, vol. 50, no. 3, pp. 799–807, Mar. 2015.

[22] A. T. Tran and B. M. Baas, "Design of an energy-efficient 32-bit adder operating at subthreshold voltages in 45-nm CMOS," in *Proc. Int. Conf. Commun. Electron.*, Aug. 2010, pp. 87–91.

[23] K. A. Bowman *et al.*, "A 16 nm all-digital auto-calibrating adaptive clock distribution for supply voltage droop tolerance across a wide operating range," *IEEE J. Solid-State Circuits*, vol. 51, no. 1, pp. 8–17, Jan. 2016.

Yizhak Shifman (Graduate Student Member, IEEE) received the B.Sc. degree (*magna cum laude*) in computer engineering and the MBA degree (*magna cum laude*) from The Hebrew University of Jerusalem, Jerusalem, Israel, in 2009 and 2010, respectively, and the M.Sc. degree (excellence) in electrical engineering from Bar-Ilan University, Ramat Gan, Israel, in 2019.

Since 2007, he has been with Intel Corporation, Jerusalem, designing a wide range of clocking and high-speed circuits for Ethernet physical link interfaces (PHYs). Since 2019, he has been conducting his doctoral research on ultralow-power analog circuits for security, power management, and biosensors at Bar-Ilan University. He has authored nine journal and conference papers and holds five issued and pending patents. His current research interests are in the areas of clocking circuits, time-interleaved analog-to-digital converters (ADCs), power management, physical unclonable functions, and biosensing.

Inbal Stanger (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from Bar-Ilan University, Ramat Gan, Israel, in 2018 and 2020, respectively, where she is currently pursuing the Ph.D. degree.

Her current research interests include the design of low-power and high-frequency digital circuits under extreme PVT variations and cryogenic temperatures.

Netanel Shavit (Graduate Student Member, IEEE) received the B.Sc. (*summa cum laude*) and M.Sc. degrees in electrical engineering from Bar-Ilan University, Ramat Gan, Israel, in 2017 and 2019, respectively, where he is currently pursuing the Ph.D. degree.

His current research interests are low-power and high-frequency logic families and digital circuit design.

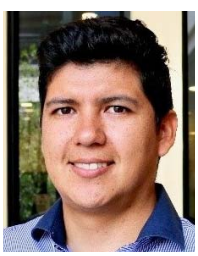

Ramiro Taco (Member, IEEE) received the M.S. degree in electrical engineering from the University San Francisco de Quito, Quito, Ecuador, in 2012, and the Ph.D. degree in electrical engineering from the University of Calabria, Rende, Italy, in 2017.

In 2017, he joined Emerging Nanoscaled Integrated Circuits and Systems Labs, Bar-Ilan University, Ramat Gan, Israel, as a Post-Doctoral Fellow. In 2020, he joined the Faculty of Engineering, University of San Francisco of Quito, Quito, as a Faculty Member and the Director of the Institute of Micro and Nanoelectronics (IMNE). His research interests include high-speed and energy-efficient mixed-signal VLSI chips and energy-efficient memories.

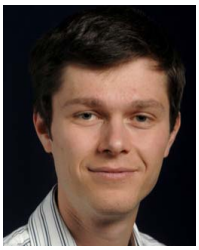

Alexander Fish (Member, IEEE) received the B.Sc. degree in electrical engineering from Technion, Israel Institute of Technology, Haifa, Israel, in 1999, and the M.Sc. and Ph.D. (*summa cum laude*) degrees from Ben-Gurion University (BGU), Beersheba, Israel, in 2002 and 2006, respectively.

He was a Post-Doctoral Fellow with the ATIPS Laboratory, University of Calgary, Calgary, AB, Canada, from 2006 to 2008. In 2008, he joined BGU, as a Faculty Member with the Electrical and Computer Engineering Department. There he founded the Low Power Circuits and Systems (LPC&S) Laboratory, specializing in low-power circuits and systems. In July 2011, he was appointed as a Head of the VLSI Systems Center at BGU. In October 2012, he joined the Faculty of Engineering, Bar-Ilan University, Ramat Gan, Israel, as an Associate Professor and the Head of the nanoelectronics track. In March 2015, he founded the Emerging Nanoscaled Integrated Circuits and Systems (ENICS) Labs. He is currently a Full Professor and the Heads of the EnICS Impact Center. He has authored over 200 scientific papers in journals and conferences. He also submitted more than 30 patent applications of which 21 were granted. He has authored three book chapters and two books as an Editor. He was a coauthor of papers that won the Best Paper Finalist awards at IEEE ISCAS and ICECS conferences. His research interests include power reduction methodologies for high-speed digital and mixed-signal VLSI chips, energy-efficient SRAM, and eDRAM memory arrays, CMOS image sensors and biomedical circuits, systems and applications, cryogenic design for quantum computing, and hardware security.

Dr. Fish is a member of the VLSI Systems and Applications and Biomedical Systems Technical Committees of IEEE Circuits and Systems Society. He founded and served as an Editor-in-Chief for the MDPI *Journal of Low Power Electronics and Applications* (JLPEA) from 2012 to 2018 and was an Associate Editor for the IEEE SENSORS JOURNAL. He is currently an Associate Editor for IEEE ACCESS, *Microelectronics Journal* (Elsevier), and *Integration*, the VLSI journal (Elsevier). He is on the Technical Program Committee of the European Solid-State Circuits Conference. He also served as a Chair of different tracks of various IEEE conferences. He was a Co-Organizer of many special sessions at IEEE conferences, including IEEE ISCAS, IEEE Sensors, and IEEE conferences.

Joseph Shor (Senior Member, IEEE) received the B.A. degree in physics from the Queens College, New York, NY, USA, in 1986, and the Ph.D. and M.S. degrees in electrical engineering from Columbia University, NY, USA, in 1993 and 1988, respectively.

From 2004 to 2015, he was at Intel Israel, as a Principal Engineer, and the Head of the Analog Team at Intel Yakum, Yakum, Israel. From 1999 to 2004, he worked at Saifun Semiconductor, Netanya, Israel, as a Staff Engineer where he established the analog activities for Flash and EEPROM NROM memories. From 1994 to 1999, he was a Senior Analog Designer at the DSP Division, Motorola Semiconductor, Herzlia, Israel. From 1988 to 1994, he was a Senior Research Scientist at Kulite Semiconductor, Leonia, NJ, USA, where he developed processes and devices for silicon carbide and diamond microsensors. He is currently an Associate Professor of electrical engineering with Bar Ilan University, Ramat Gan, Israel. He has authored more than 70 papers in refereed journals and conference proceedings in the areas of analog circuit design and device physics. He holds over 50 issued patents and several pending patents. His present interests include analog circuits, switching and linear voltage regulators, sensors, PLLs and IO circuits, microprocessors, and security.

Dr. Shor was a member of the ISSCC Technical Program Committee (TPC) from 2014 to 2018, has been a member of the TPC of ESSCIRC since 2018, and is an Associate Editor for the IEEE SENSORS JOURNAL. He has been a Guest Associate Editor of JSSC and SSCL.