

A 7-nm All-Digital Leakage-Current-Supply Circuit for Analog LDO Dropout Voltage Reduction

Keith A. Bowman^{ID}, Senior Member, IEEE, Samantak Gangopadhyay, Francois I. Atallah, Member, IEEE,
Hoan H. Nguyen^{ID}, Jihoon Jeong, Daniel J. Yingling, Anthony Polomik, Mahesh Harinath,
Nathaniel Reeves, Amer Cassier, Brad R. Appel, and Arijit Raychowdhury^{ID}, Senior Member, IEEE

Abstract—A 7-nm leakage-current-supply (LCS) circuit monitors leakage across process and temperature variations and controls PFET block-head switches (BHSs) to supply the slow-changing leakage current while a high-bandwidth analog low-dropout (LDO) voltage regulator supplies the fast-changing dynamic current. By decreasing the LDO maximum current demand (I_{MAX}), LCS test-chip measurements demonstrate a 70 mV (44%) reduction in the minimum dropout voltage ($V_{DO,MIN}$), resulting in a wider voltage range of LDO usage for power savings of 14%–22%.

Index Terms—Low-dropout (LDO) voltage, LDO headroom voltage, leakage monitor, leakage-current-starved ring oscillator (RO).

I. INTRODUCTION

Commercial system-on-chip (SoC) processors integrate a limited number of input voltage (V_{IN}) rails due to high costs. For this reason, many cores share the same V_{IN} as illustrated in Fig. 1. Low-dropout (LDO) voltage regulators [1]–[6] and dedicated phase-locked loops (PLLs) allow each core on a shared V_{IN} to employ a unique supply voltage (V_{DD}) and clock frequency (F_{CLK}). The core requiring the highest target V_{DD} and F_{CLK} sets the shared V_{IN} value as depicted by CPU-0 in Fig. 1. This core operates at V_{IN} via the PFET block-head switches (BHSs) between V_{IN} and V_{DD} while disabling the LDO. A core with a lower target V_{DD} and F_{CLK} always operates at the lower F_{CLK} to reduce power. To enable the LDO to lower power further, the target V_{DD} must satisfy the LDO minimum dropout voltage ($V_{DO,MIN}$). If the target $V_{DD} \leq V_{IN}-V_{DO,MIN}$, the software driver enables the LDO while setting the LDO reference voltage (V_{REF}) equal to the target V_{DD} to lower the core V_{DD} as described by CPU-1, CPU-2, and CPU-3 in Fig. 1. Otherwise, the core V_{DD} operates at V_{IN} via the PFET BHS while disabling the LDO.

In commercial SoCs, test time contributes a large portion of the overall cost to design, manufacture, and test a high-volume product. For this reason, SoC companies focus on reducing test time. Premium-tier SoC processor testing performs maximum F_{CLK} (F_{MAX}) tests at target V_{DD} values or minimum V_{DD} (V_{MIN}) tests at target F_{CLK} values for the CPU and DSP via the BHS as illustrated with the green curve in Fig. 2. To reduce costs, SoC processor tests traditionally do not include F_{MAX} or V_{MIN} tests in LDO mode. Rather, the test program verifies the LDO operation at the F_{MAX} or V_{MIN} conditions as measured with the BHS. As a result, high-volume commercial SoC products require the LDO $F_{MAX} \geq$ BHS F_{MAX} or LDO $V_{MIN} \leq$ BHS V_{MIN} to avoid high test costs. Otherwise,

Manuscript received August 15, 2019; revised September 20, 2019; accepted October 4, 2019. Date of publication October 14, 2019; date of current version December 20, 2019. This article was approved by Associate Editor Xin Zhang. (Corresponding author: Keith A. Bowman.)

K. A. Bowman, F. I. Atallah, H. H. Nguyen, J. Jeong, D. J. Yingling, A. Polomik, M. Harinath, and B. R. Appel are with Qualcomm Technologies, Inc., Raleigh, NC 27617 USA (e-mail: kbowman@qti.qualcomm.com).

S. Gangopadhyay and A. Raychowdhury are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA.

N. Reeves and A. Cassier are with Qualcomm Technologies, Inc., San Diego, CA 92121 USA.

Digital Object Identifier 10.1109/LSSC.2019.2947239

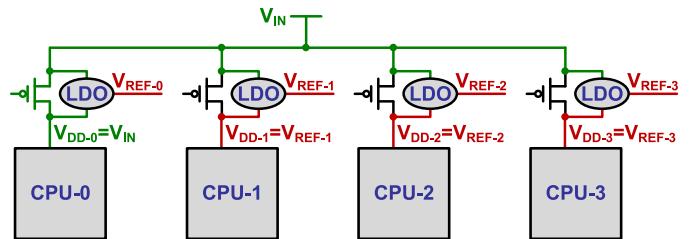


Fig. 1. Example of SoC processor cores sharing a single V_{IN} . CPU-0 requires the highest target V_{DD} and F_{CLK} , thus setting the V_{IN} value and connecting to V_{IN} via the PFET BHSs with the LDO disabled. CPU-1, CPU-2, and CPU-3 have lower target V_{DD} and F_{CLK} values with the target $V_{DD} \leq V_{IN}-V_{DO,MIN}$, allowing LDO usage with V_{REF} set to the target V_{DD} .

the software drivers permanently disable the LDO. Since the LDO operates with much less decoupling capacitance to suppress V_{DD} noise as compared to the BHS, the processor cores need a high-bandwidth LDO with a fast-transient response. Fig. 2 contains two hypothetical curves in red corresponding to separate LDOs with fast and slow transient responses. In this example, the LDO with a slow-transient response provides an LDO $F_{MAX} \leq$ BHS F_{MAX} , resulting in the software driver disabling the LDO operation. The LDO with a fast-transient response in Fig. 2 satisfies the constraint of LDO $F_{MAX} \geq$ BHS F_{MAX} , allowing LDO usage when the target $V_{DD} \leq V_{IN}-V_{DO,MIN}$. For these reasons, high-volume premium-tier SoC CPU and DSP cores prefer high-bandwidth analog LDOs for fast-transient response [1], [2].

A key challenge with analog LDOs in SoC cores is the large $V_{DO,MIN}$ (e.g., 150–200 mV [2], [3]) to supply the core maximum current demand (I_{MAX}) at worst-case dynamic and leakage conditions. The large $V_{DO,MIN}$ limits LDO usage. An approach for reducing $V_{DO,MIN}$ is increasing the power-PFET width. For the analog LDO design in this letter [6], however, the bandwidth decreases as the power-PFET width increases as described in Fig. 3. For this design, the maximum V_{DD} -droop specifications limit the power-PFET width. Alternatively, the stability of an output-pole-dominant analog LDO constrains the power-PFET width. Although all-digital LDOs reduce $V_{DO,MIN}$, these designs suffer from low gain and high output ripple, thus degrading core performance. Hybrid LDOs [1] employ digital and analog loops to tradeoff the strengths and weaknesses of traditional digital and analog designs. The challenge with hybrid LDOs is managing the complex load sharing between the analog and digital loops while maintaining high bandwidth and stability.

This letter describes an all-digital leakage-current-supply (LCS) circuit in a 7-nm [7] test chip to only supply the leakage current. Although leakage exponentially varies across process and temperature variations, the current change is slow, allowing the LCS to compensate for the leakage while avoiding load sharing with the analog LDO. At worst-case process and temperature conditions, leakage significantly contributes to I_{MAX} . LCS reduces the LDO I_{MAX} , and consequently $V_{DO,MIN}$, resulting in higher LDO usage for core power savings [5], [6].

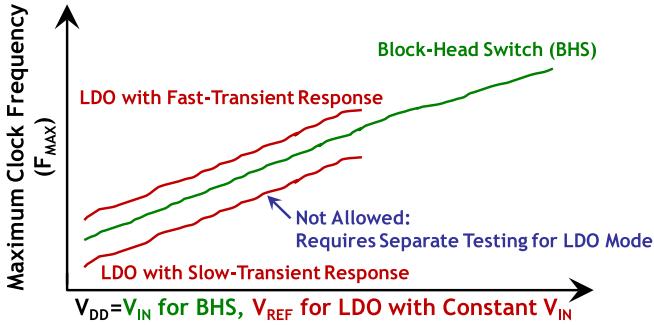


Fig. 2. Illustration of the maximum F_{CLK} (F_{MAX}) versus V_{IN} for BHS and V_{REF} for two separate LDOs with a constant V_{IN} to highlight the LDO testing implications in commercial high-volume SoC processors.

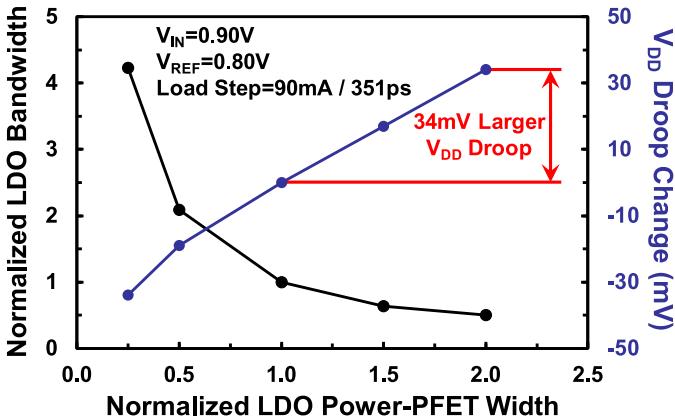


Fig. 3. Simulated analog LDO bandwidth and V_{DD} droop change versus power-PFET width.

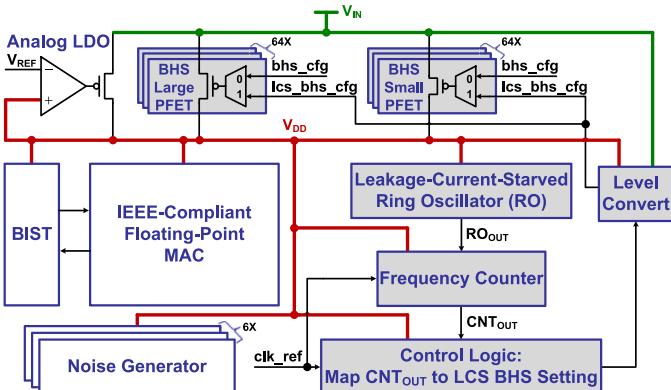


Fig. 4. Test-chip block diagram of the LCS circuit integrated with a BHS and an analog LDO to power a CPU IEEE-compliant floating-point MAC, BIST, and noise generators.

II. TEST-CHIP OVERVIEW

As illustrated in Fig. 4, the test chip features the LCS integrated with a BHS and analog LDO between V_{IN} and V_{DD} to power a CPU IEEE-compliant floating-point multiply accumulate (MAC) unit, an on-die built-in self-test (BIST), and on-die programmable noise generators. The LDO is a symmetrical operational transconductance amplifier followed by a power PFET. With post-layout extraction and no external capacitor, the phase margin is 92° at the unity-gain bandwidth [6]. The MAC performs single- and double-precision IEEE floating-point multiply, fused multiply add, and register load/store instructions to represent core functionality. An on-die PLL generates the MAC F_{CLK} . The BIST executes functional tests and verifies MAC outputs for F_{MAX} testing. Noise generators

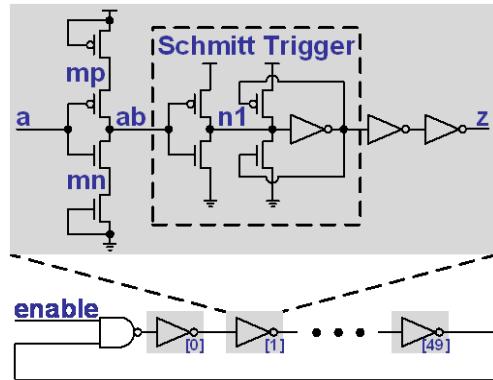


Fig. 5. Leakage-current-starved RO schematic.

allow the injection of V_{DD} droops as observed during the normal operation of a high-performance CPU. The test chip is manufactured in a 7-nm FinFET CMOS technology [7]. Test-chip measurements are performed with an Advantest 93 K production tester.

III. LEAKAGE-CURRENT-SUPPLY DESIGN

The LCS design in Fig. 4 contains a leakage-current-starved ring oscillator (RO) to generate an RO frequency (F_{RO}) to track leakage, a frequency counter to measure F_{RO} , control logic to map the frequency counter output (CNT_{OUT}) to a BHS configuration (lcs_bhs_cfg) to supply the leakage, and a BHS. The digital RO in Fig. 5 maps leakage current to an easily measurable F_{RO} while avoiding the circuit complexities, additional bias or reference voltages, and analog-to-digital converters in prior analog leakage sensors [8], [9]. The RO includes a NAND gate and 50 inversion-delay stages with each stage consisting of a leakage-current-starved inverter followed by a Schmitt trigger (ST) to provide a critical hysteresis for increasing the RO-delay dependency on leakage. Fig. 6 highlights the importance of the ST with simulations of the RO with and without the ST. As the input (a) of the current-starved inverter transitions from 0 V to V_{DD} , the output (ab) experiences a charge-sharing effect with node mn that quickly changes the ab voltage from 0.8 V to 0.67 V, resulting in a rapid 130 mV drop. The leakage from the bottom NFET slowly completes the ab voltage transition from 0.67 V to 0 V. Without the ST circuit in Fig. 6(b), the delay dependency on leakage only occurs when ab changes from 0.67 V to 0.40 V to transition the next inverter stage. The ST creates a hysteresis to require a larger ab voltage change beyond 0.4 V to transition $n1$. With the ST circuit in Fig. 6(a), the leakage discharges ab from 0.67 V to 0.32 V to transition $n1$, thus increasing the required ab voltage change by 80 mV (30%) as compared to the RO without the ST. Although ab experiences ~ 50 mV increase from the coupling capacitance as $n1$ transitions high, this does not affect the $n1$ transition due to the ST hysteresis. A similar behavior happens when the input transitions from V_{DD} to 0 V. As a result, the bottom NFET and top PFET leakage currents of the current-starved inverter dominate the stage switching delay to accurately map leakage to F_{RO} . The inverters after the ST circuit in Fig. 5 improve the input slope to the next RO stage.

Another concern for the RO without the ST is the significant slowdown in the ab voltage transition near $0.5V_{DD}$ [e.g., 0.4 V in Fig. 6(b)] due to the coupling capacitance as $n1$ rises. For this reason, the $n1$ transition is highly sensitive to voltage noise, especially, at lower temperatures where the voltage change is slower due to less leakage. At these conditions with voltage noise, $n1$ may transition up, then down, and eventually up again from a single input (a) transition. This multiple- $n1$ -transition effect may lead to an inaccurate F_{RO} . The inclusion of the ST avoids this issue to enhance the RO stability.

As described in Fig. 7, the frequency counter contains a reference counter with a reference clock (clk_{ref}) at 20 MHz and an RO counter with an RO clock (clk_{ro}) at F_{RO} . RO_{OUT} in

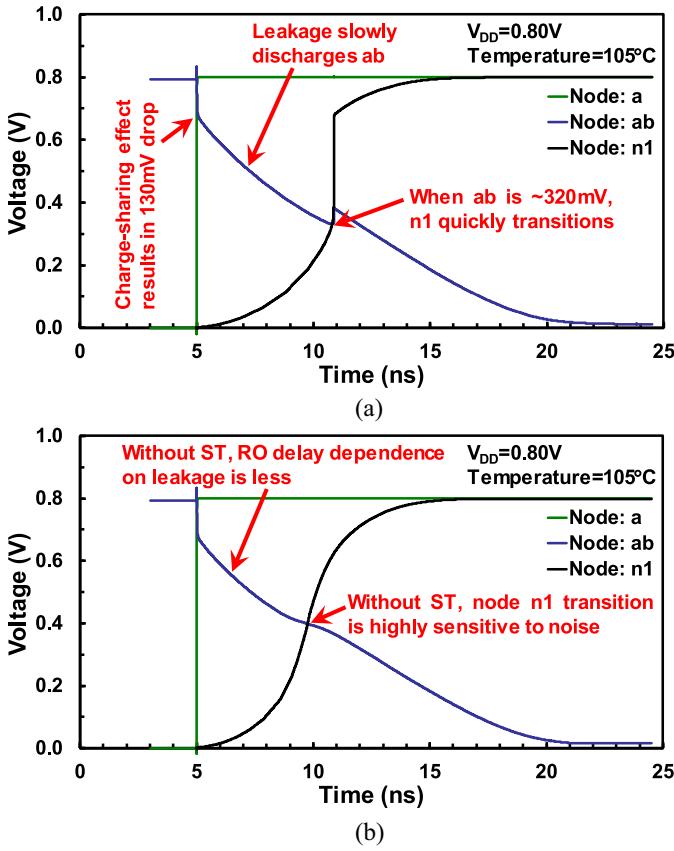


Fig. 6. Simulated node voltages versus time for the leakage-current-starved RO (a) with the ST and (b) without the ST.

Fig. 4 drives clk_{ro} . The measurement starts when the reference trigger signal ($\text{ref}_{\text{start}}$) transitions high to load a programmable value ($\text{ref}_{\text{cnt}}_{\text{preload}}$) into the reference counter. On the subsequent clk_{ref} cycles, the reference counter decrements by one from $\text{ref}_{\text{cnt}}_{\text{preload}}$ to zero. When the synchronized RO trigger signal (ro_{start}) transitions high, the RO counter initially resets to zero and then increments by one per clk_{ro} cycle. After the reference counter reaches zero, the synchronized ro_{start} transitions low to disable the RO counter, which holds the RO counter output ($\text{ro}_{\text{cnt}}_{\text{q}}$) constant to drive CNT_{OUT} in Fig. 4. The frequency counter measures F_{RO} over a target delay (e.g., 1 ms). Slow F_{RO} values require a sufficiently long target delay for an accurate F_{RO} measurement. The $\text{ref}_{\text{cnt}}_{\text{preload}}$ equals the product of the target delay and clk_{ref} (e.g., $\text{ref}_{\text{cnt}}_{\text{preload}} = 1 \text{ ms} \times 20 \text{ MHz} = 20000$). Thus, F_{RO} is the final $\text{ro}_{\text{cnt}}_{\text{q}}$ divided by the target delay (e.g., $F_{\text{RO}} = \text{ro}_{\text{cnt}}_{\text{q}}/1 \text{ ms}$). Once the F_{RO} measurement finishes, the frequency counter asserts an interface signal to trigger the control logic to map CNT_{OUT} to a target $\text{lcs}_{\text{bhs}}_{\text{cfg}}$ to supply the leakage.

Operating at 20 MHz, the control logic performs the $\text{CNT}_{\text{OUT-lcs}_{\text{bhs}}_{\text{cfg}}}$ mapping with look-up tables based on post-silicon characterization of F_{RO} versus leakage and BHS settings versus leakage as described in Section IV. The LCS reconfigures the conventional BHS for power gating and LDO enablement to reduce the power-PFET area overhead by separating the BHS into 64 large-width power PFETs and 64 small-width power PFETs with independent control. The 64 small-width PFETs provide the equivalent drive current of one large-width PFET. The mapping of CNT_{OUT} to $\text{lcs}_{\text{bhs}}_{\text{cfg}}$ occurs after each F_{RO} measurement (e.g., 1 ms). Otherwise, the control logic is clock gated.

IV. TEST-CHIP MEASUREMENTS

The silicon measurements from the 7-nm test chip in Fig. 8 demonstrate the leakage-current-starved F_{RO} closely tracking leakage

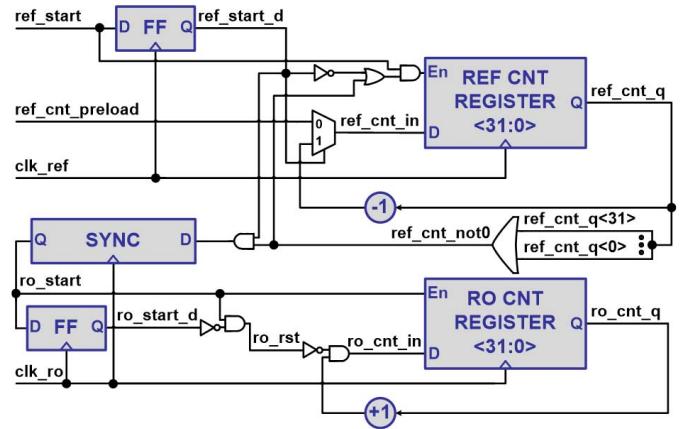


Fig. 7. Frequency counter schematic.

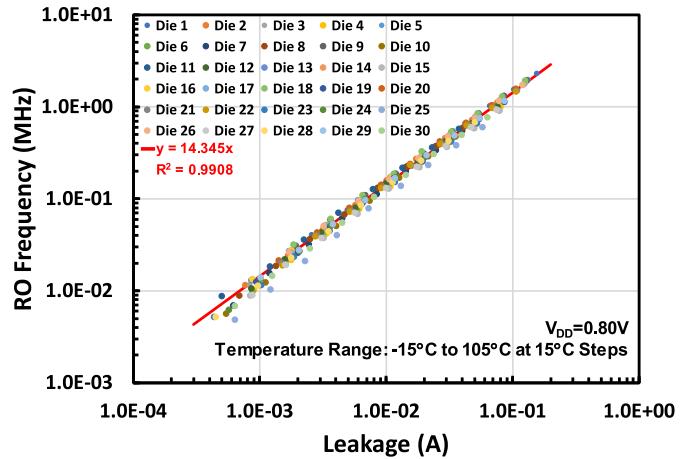


Fig. 8. Measured LCS RO frequency (F_{RO}) versus leakage for 30 dies at 0.8 V and temperature ranging from -15 °C to 105 °C at 15 °C steps.

V_{DD} (V)	0.85	0.80	0.75	0.70	0.65	0.60	0.55
Slope = RO Frequency / Leakage (MHz / A)	13.237	14.345	15.545	16.872	18.342	20.000	21.913
R^2	0.99	0.99	0.99	0.99	0.99	0.99	0.99

V_{IN} (V) : V_{DD} (V)	0.90 : 0.80	0.90 : 0.75	0.80 : 0.70	0.80 : 0.65	0.70 : 0.60	0.70 : 0.55
Slope = BHS Settings / Leakage (# BHS Large PFETs / A)	153.44	100.87	176.85	117.88	218.01	146.97
R^2	0.99	0.99	0.99	0.99	0.99	0.99

Fig. 9. Measured slope and R^2 results from F_{RO} versus leakage across V_{DD} and LCS BHS settings to supply the leakage for V_{IN} and V_{DD} combinations.

changes across the process and temperature variations of 30 dies from -15 °C to 105 °C at 15 °C steps for a total of 270 data points (i.e., 30 dies at 9 temperatures) at 0.8 V. Across the 30 dies, the leakage power contribution to the total power ranges from 2% to 10% at 30 °C and from 26% to 52% at 105 °C. The number of large and small BHS power PFETs to supply the leakage is measured across V_{IN} and V_{DD} combinations by using an analog voltage comparator and configuring the LCS control logic to operate as a low-bandwidth digital LDO to adjust $\text{lcs}_{\text{bhs}}_{\text{cfg}}$ accordingly. From Fig. 9, linear-regression models derived from measurements of F_{RO} versus leakage across V_{DD} and BHS settings to supply the leakage across V_{IN} and V_{DD} combinations for the 30 dies from -15 °C to 105 °C consistently result in an R^2 of 0.99. This post-silicon characterization enables one set of look-up tables for every part to map F_{RO} to $\text{lcs}_{\text{bhs}}_{\text{cfg}}$ in the control logic to avoid the expensive test time of per part calibration.

The worst-case process, temperature, and activity conditions define the $V_{DD,\text{MIN}}$ applied to software drivers for every part in commercial SoC processors. At these conditions, leakage is a major contributor to

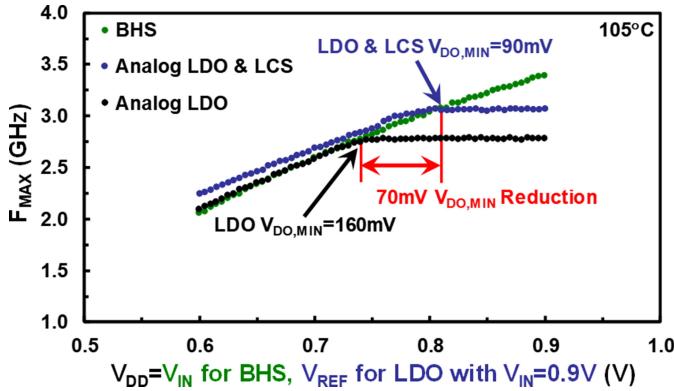


Fig. 10. Measured F_{MAX} for the IEEE-compliant floating-point MAC versus V_{IN} for BHS and V_{REF} for LDO with $V_{IN} = 0.9$ V at 105 °C to demonstrate the LCS minimum dropout voltage ($V_{DO,MIN}$) reduction.

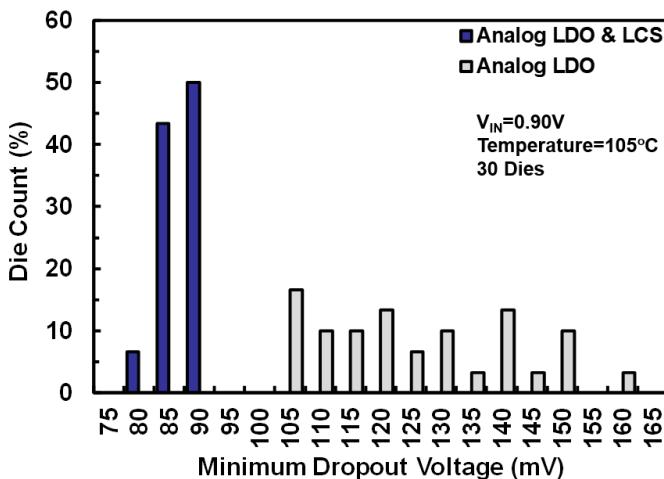


Fig. 11. Measured $V_{DO,MIN}$ distribution for 30 dies.

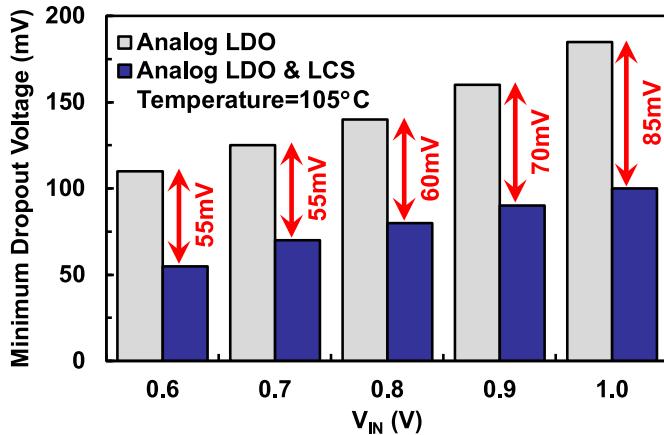


Fig. 12. Measured $V_{DO,MIN}$ versus V_{IN} .

I_{MAX} . In Fig. 10, MAC F_{MAX} measurements versus V_{IN} for BHS and V_{REF} for LDO with $V_{IN} = 0.9$ V at 105 °C reveal the LCS $V_{DO,MIN}$ benefits. Since the LDO cannot supply I_{MAX} until satisfying $V_{DO,MIN}$, the LDO F_{MAX} remains constant and below the BHS F_{MAX} until $V_{REF} = 0.9$ V- $V_{DO,MIN}$. LCS reduces $V_{DO,MIN}$ by 70 mV (44%). At low V_{REF} , LCS assists the LDO transient response for a higher F_{MAX} . From Fig. 11, LCS decreases the $V_{DO,MIN}$ variation from 55 mV to 10 mV (5.5×) across 30 dies. In Fig. 12, measurements demonstrate 43%–50% lower $V_{DO,MIN}$ for LCS across V_{IN} .

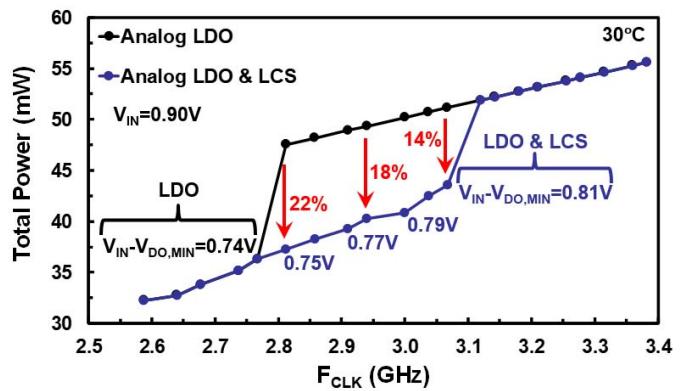


Fig. 13. Measured total power versus F_{CLK} with V_{IN} at 0.9 V, where the target V_{DD} changes with F_{CLK} to represent the V_{DD} - F_{CLK} values across DVFS states. V_{DD} remains at 0.9 V unless the target $V_{DD} \leq V_{IN}-V_{DO,MIN}$ to allow LDO operation at the target V_{DD} .

Measurements in Fig. 13 of the total power across target V_{DD} - F_{CLK} values highlight the LCS power savings across dynamic voltage-frequency scaling (DVFS) states. Since commercial mobile SoC days-of-use (DoU) power is evaluated at 30 °C, the tester applies the same temperature to provide a realistic measurement of the LCS improvements. Although the power is measured at 30 °C, the worst-case temperature of 105 °C dictates the $V_{DO,MIN}$ across all temperatures. Starting at the highest F_{CLK} and moving toward lower values, the MAC operates at the target F_{CLK} while V_{DD} remains at 0.9 V unless the target $V_{DD} \leq V_{IN}-V_{DO,MIN}$. Once the target V_{DD} is low enough to satisfy $V_{DO,MIN}$, the system allows LDO operation at the target V_{DD} for an additional V_{DD} power reduction. This behavior is consistent with the software drivers in commercial mobile SoC processors. The LCS enables LDO operation at 810 mV and below, resulting in a 70-mV wider dynamic V_{DD} range of LDO usage for 14%–22% power savings at the same F_{CLK} .

REFERENCES

- [1] M. Saint-Laurent *et al.*, “A 28 nm DSP powered by an on-chip LDO for high-performance and energy-efficient mobile applications,” *IEEE J. Solid-State Circuits*, vol. 50, no. 1, pp. 81–91, Jan. 2015.
- [2] W.-C. Chen, T.-C. Huang, C.-C. Chiu, C.-W. Chang, and K.-C. Hsu, “94% power-recycle and near-zero driving-dead-zone N-type low-dropout regulator with 20mV undershoot at short-period load transient of flash memory in smart phone,” in *ISSCC Dig. Tech. Papers*, San Francisco, CA, USA, 2018, pp. 436–438.
- [3] Y. Lu, W.-H. Ki, and C. P. Yue, “A 0.65ns-response-time 3.01ps FOM fully-integrated low-dropout regulator with full-spectrum power-supply-rejection for wideband communication systems,” in *ISSCC Dig. Tech. Papers*, San Francisco, CA, USA, 2014, pp. 306–307.
- [4] Y.-J. Lee *et al.*, “A 200mA digital low-drop-out regulator with coarse-fine dual loop in mobile application processors,” in *ISSCC Dig. Tech. Papers*, San Francisco, CA, USA, 2016, pp. 150–151.
- [5] S. Gangopadhyay *et al.*, “Digitally-assisted leakage current supply circuit for reducing the analog LDO minimum dropout voltage,” in *Proc. CICC*, Austin, TX, USA, 2017.
- [6] K. Bowman *et al.*, “A 7nm leakage-current-supply circuit for LDO dropout voltage reduction,” in *Symp. VLSI Circuits Dig. Tech. Papers*, Kyoto, Japan, 2019, pp. 126–127.
- [7] S.-Y. Wu *et al.*, “A 7nm CMOS platform technology featuring 4th generation FinFET Transistors with a 0.027um² high density 6-T SRAM cell for mobile SoC applications,” in *IEDM Tech. Dig.*, San Francisco, CA, USA, 2016, pp. 43–46.
- [8] T. Kuroda *et al.*, “A 0.9-V, 150-MHz, 10-mW, 4 mm², 2-D discrete cosine transform core processor with variable threshold-voltage (VT) scheme,” *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1770–1779, Nov. 1996.
- [9] C. H. Kim, K. Roy, S. Hsu, R. Krishnamurthy, and S. Borkar, “A process variation compensating technique with an on-die leakage current sensor for nanometer scale dynamic circuits,” *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 14, no. 6, pp. 646–649, Jun. 2006.