A 225 μm^2 Probe Single-Point Calibration Digital Temperature Sensor Using Body-Bias Adjustment in 28 nm FD-SOI CMOS

Martin Cochet, *Member, IEEE*, Ben Keller, *Member, IEEE*, Sylvain Clerc, *Member, IEEE*, Fady Abouzeid, Andreia Cathelin, *Senior Member, IEEE*, Jean-Luc Autran, *Senior Member, IEEE*, Philippe Roche, *Senior Member, IEEE*, and Borivoje Nikolić, *Fellow, IEEE*

***Abstract*—An embedded digital temperature sensor based on a single-ended probe is implemented in a 28 nm fully depleted silicon-on-insulator process. The nMOS-only ring-oscillator probe uses single-point calibration based on body-bias tuning of its well for process compensation. Nonlinearity compensation is implemented on-chip in custom digital logic, resulting in an area-efficient (225 μm^2 per probe, 11 482 μm^2 for the full system) sensor while achieving −1.4 °C/ +1.3 °C accuracy using 2.0 nJ/sample and maintaining functionality over a 0.62–1.2 V range, making it suitable for temperature monitoring in digital systems-on-chip.**

***Index Terms*—Body-bias calibration, CMOS sensor, digital curvature correction, temperature sensor, wide voltage range.**

I. Introduction

Compact and fast temperature sensors are necessary for thermal management of modern systems-on-chip (SoCs), in which numerous sensors are implemented on the same die to monitor temperature gradients [1]. Temperature sensing is often performed by comparing signals (voltages or currents) that are proportional to absolute temperature with signals that are complementary to absolute temperature (CTAT). The output is then obtained by comparing the frequency of two oscillators operating with different current bias [2], supply voltages [3], or transistor thresholds [4]. However, many modern processes do not offer thin-oxide devices with CTAT characteristics, requiring the use of large, high-voltage, thick-oxide devices for temperature measurements. As a result, alternate methods have been proposed, such as measuring the frequency F of a single oscillator [5], [6]. In these designs, the process- and voltage-dependency of dF/dT must be compensated. The method in [5] offers no supply rejection and requires averaging over a large number of samples, resulting in high (240 nJ/sample) power consumption. The method in [6] includes supply regulation, but requires a two-point calibration to compensate for process-dependent dF/dT curvature, which is impractical for low-cost SoCs. To compensate dF/dT nonlinearity, recently published sensors use either off-chip processing [6] or large (58 000 μm^2) microcontroller as a compute engine [3]. The ideal solution for thermal monitoring of large SoCs entails a compact sensor and embedded single-point calibration that achieves moderate temperature resolution requirements. We present a temperature sensor architecture that provides voltage scalability, a compact footprint, and single-point calibration, with 11 instances implemented across an SoC.

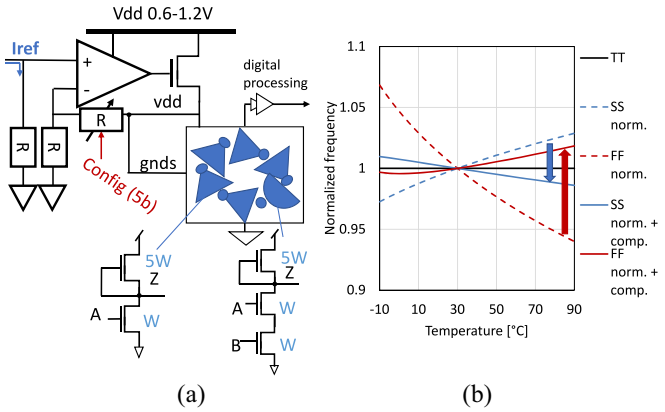

Fig. 1. (a) Schematics of the temperature sensing probe and (b) simulated process variation reduction by vdd/gnds compensation.

II. System Description

The system includes a current reference divider, 11 oscillator probes, and a digital processing unit. Fig. 1 illustrates the operating principle of the proposed probe. To reject supply noise, the probe includes a custom-designed LDO. Single-point calibration based on frequency normalization only to compensate for process variation [5] is not sufficient due to the inherent dF/dT process-corner dependency of the oscillator. Simulations find that the normalized oscillator frequency changes by up to $-6.8\%/+6.0\%$ across corners, which would lead to −0.9 °C / +1.9 °C error.

Rather than using a two-point calibration [6], this letter compensates the frequency change detected at calibration by simultaneously adjusting the supply and body bias voltage of the oscillator, which results in equalizing dF/dT. This approach takes advantage of the body-biasing flexibility of the fully depleted silicon-on-insulator (FD-SOI) process [7]. To avoid compensating independently for the supply (vdd), the pMOS p-well voltage (vdds), and the nMOS n-well voltage (gnds) with three different regulators, the nMOS-only oscillator is designed using a wide diode-connected nMOS pull-up. This way, gnds can directly be tied to the supply vdd and the same supply rejection LDO is reused for process compensation. By tuning vdd and gnds (using optimal values of 0.68 V in the SS corner and 0.45 V in the FF corner), the slope change across corners is reduced to $-1.8\%/+1.4\%$ in simulation, corresponding to an error of -0.4 °C / +0.6 °C. To perform this tuning, the body-bias and supply are trimmed manually by

Manuscript received November 6, 2017; revised December 15, 2017; accepted January 12, 2018. Date of publication January 24, 2018; date of current version February 21, 2018. This paper was approved by Associate Editor Ana Rusu. *(Corresponding author: Martin Cochet.)*

M. Cochet was with STMicroelectronics, 38920 Crolles, France, also with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with CNRS, IM2NP (UMR 7334), Aix-Marseille University, 13397 Marseille, France. He is now with IBM Research, Yorktown Heights, NY 10598 USA (e-mail: martin.cochet@ibm.com).

B. Keller and B. Nikolić are with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA.

S. Clerc, F. Abouzeid, A. Cathelin, and P. Roche are with STMicroelectronics, 38920 Crolles, France.

J.-L. Autran is with CNRS, IM2NP (UMR 7334), Aix-Marseille University, 13397 Marseille, France.

Digital Object Identifier 10.1109/LSSC.2018.2797427

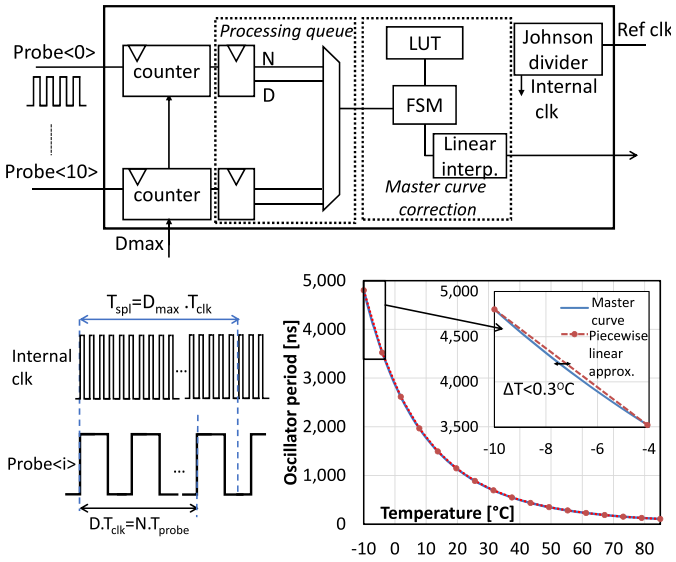

Fig. 2. Block diagram of the digital processing logic, a waveform showing the counting operation, and the piece-wise linear approximation of the master curve.

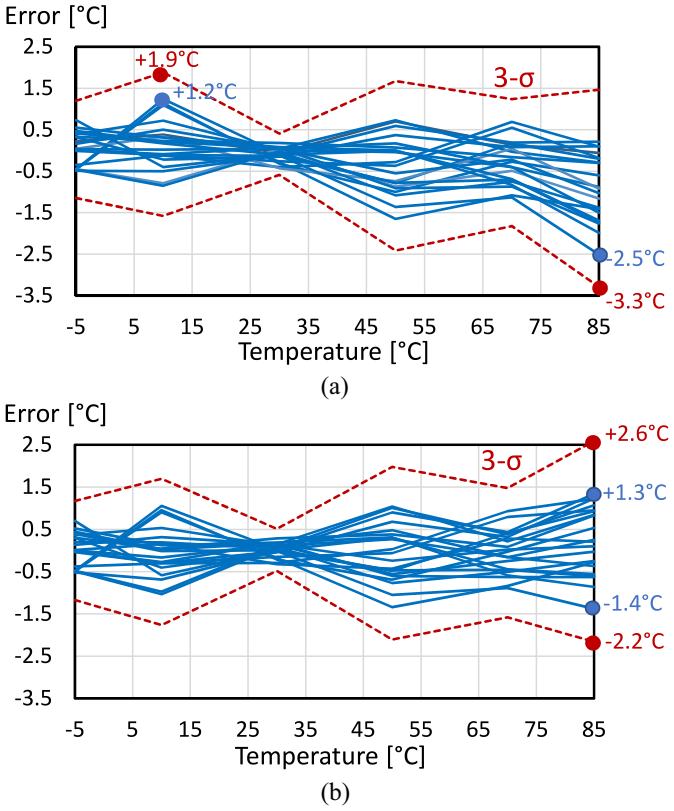

Fig. 3. Measured temperature accuracy from 21 probes. (a) Accuracy after calibration. (b) Accuracy after calibration plus systematic error correction.

adjusting the value of the LDO feedback resistor (in 32 steps) during calibration until the oscillator frequency matches the target value. The oscillator period is dominated by the current of the pull-up nMOS in subthreshold operation, resulting in an exponential dependency of the delay on temperature and device threshold [6].

As the pull-up transistor results in a low slew, the steeper falling edges are used to count the oscillator period, reducing jitter. Transistor length was optimized to balance oscillation frequency with device noise, a tradeoff between sampling speed and repeatability. A device length of 40 nm was found to be optimal. The LDO is designed to minimize area and power overhead by taking advantage of the very low current drawn by the oscillator (below 1 μA). Rather than using a full two-stage operational amplifier, a simple complementary self-biased differential amplifier was used [8]. The stability of the LDO is guaranteed by a load capacitor placed at the output of the operational amplifier to provide the desired phase margin. To reduce area overhead, it is implemented as a metal-on-metal capacitor, placed over the probe.

As most complex SoCs require a reference clock of known frequency and current references for their various building blocks, the temperature sensor can reuse these references for the digital logic and the LDO, respectively. In this implementation, the on-chip reference generation is composed of a pair of current mirrors to generate 11 1 μA references for each LDO, from a 48 μA reference. The use of a current rather than a voltage reference guarantees matching across distant sensors even in the presence of ground level mismatch across the die.

Fig. 2 presents the digital logic used for data processing. The digital logic includes one frequency counter per sensor to measure the frequency for each sensor simultaneously. As the oscillator frequency dependence on temperature is exponential, counting a fixed number of oscillator cycles would lead to a sampling time increase of 25× at 85 °C compared to −5 °C. On the other hand, counting only the number of probe cycles in a fixed number of reference cycles would lead to large quantization error at low temperatures. The frequency counter is therefore implemented as a finite state machine which outputs both the number of probe cycles N and the duration of those cycles D within a fixed number of reference cycles $D_{\max}$. In this way, the probe frequency is computed as D/N, with a quantization error of only one reference cycle ($1/D_{\max}$, which is 0.1% as implemented) and a fixed sampling time.

As the frequency-to-temperature conversion is nonlinear, a second processing stage is implemented to correct for the curvature error. Instead of using an analytical calculation [3], [6], the linearization includes a coarse 4-bit look-up table (LUT) that provides the four most significant bits (MSBs) and a linear fit for the four least significant bits (LSBs). This piece-wise linear approximation allows for a very efficient implementation, while guaranteeing an error under 0.3 °C for the read-out value. The 8-bit output corresponds to an LSB resolution of 0.35 °C, which is below the measured repeatability and accuracy. As the binary LUT search with linear fitting procedure only requires five cycles, all 11 probes can be processed sequentially without impacting sampling time. The digital logic also includes a Johnson counter, providing a division factor of the input clock by 2 to 16, allowing for a flexible reference clock frequency. The nominal 125 MHz reference frequency was used for measurement.

III. Measurement Results

The temperature measurement system was implemented in 28 nm ultrathin body and buried oxide FD-SOI technology. Fig. 3 presents the measured temperature accuracy in the −5 °C to 85 °C range at 0.9 V supply for 21 sensors across seven dice. First, the LUT registers are programmed based on the simulated relative frequency master curve, corrected by a 2-point calibration measurement of a single sensor. Then, each sensor is measured at one point (30 °C) and calibrated by choosing the LDO code leading to the lowest error. This results in a min/max error of −2.5 °C / +1.2 °C, or −3.3 °C/+1.9 °C for a worst case of plus or minus three standard deviations around the mean (3-σ). A systematic correction can be implemented to subtract the average measured error across the chips for each temperature. The

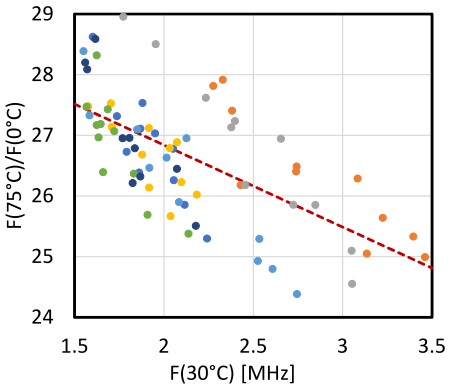

Chip	r
1	-0.87
2	-0.90
3	-0.96
4	-0.76
5	-0.95
6	-0.84
7	-0.92
All chips	-0.67

Fig. 4. Measured correlation between the measured frequency $F(30\ °C)$ before calibration and the normalized slope dF/dT, computed as $F(75\ °C)/F(0\ °C)$.

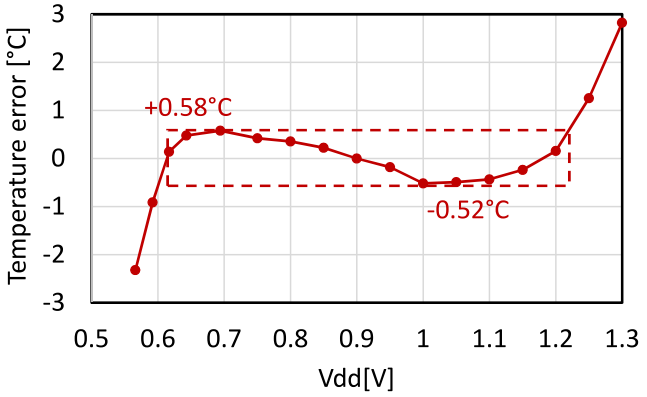

Fig. 5. Measured operation across supply voltages with temperature error caused by the oscillator's frequency change.

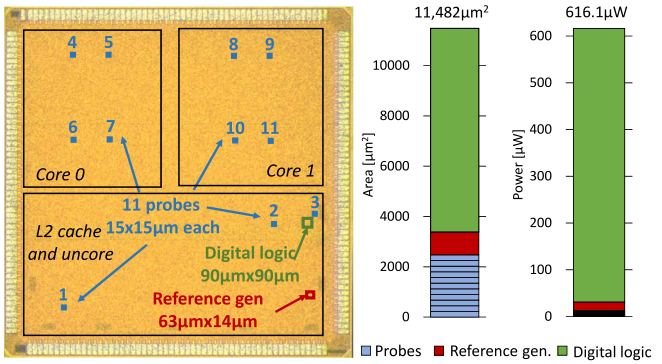

Fig. 6. Annotated chip micrograph, system area, and power breakdown.

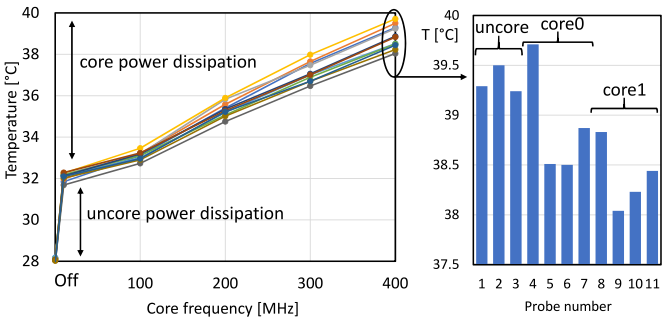

Fig. 7. Runtime *in situ* temperature measurements of a dual-core RISC-V processor.

programmable LUT values enable this to be integrated for subsequent measurements. When applying this correction, the error is reduced to -1.4 °C / $+1.3$ °C min/max and -2.2 °C / $+2.6$ °C (3-σ).

Because the circuit relies on fixed frequency and current references, it can suffer inaccuracy due to their temperature dependence. However, due to the strong superlinear temperature dependency of the oscillation period, the sensitivity to an error in the measured frequency (due, for example, to jitter/drift in the sampling clock T_{ref}) is low (0.21 °C/% at 25 °C, worst case of 0.36 °C/% at 85 °C). In comparison, the sensitivity of other single-ended CMOS sensors, which are based on a linear frequency/temperature dependency, are ten to twenty times higher (2.4 °C/% for [2] and 8.3 °C/% for [5]). The sensitivity to a change in the reference current is also low: I_{ref} is 0.32 °C/%, i.e., 0.66 °C/µA referred to the predivided reference. Additionally, systematic errors caused by T_{ref} or I_{ref} sensitivity to temperature can be compensated by incorporating them into the master curve, even if they are nonlinear.

Measured data presented in Fig. 4 show the measured correlation ($|r| = 0.67$) between the oscillation frequency at the calibration point and the normalized slope dF/dT. In accordance with the simulated results of Fig. 1(b), fast chips at 30 °C have a lower relative frequency temperature dependency, so they benefit from forward biasing compensation, which increases the temperature dependency. Repeatability of measurements for each sensor was also measured: by averaging five samples, the circuit achieves 27.8 Ks/s while guaranteeing a repeatability between measurements of 0.76 °C rms.

Fig. 5 shows the voltage rejection of the probe at 25 °C, which can operate between 0.62 V and 1.2 V with less than $+2.7\%/-2.4\%$ frequency change. This translates to $+0.58$ °C / -0.52 °C error or 1.9 °CV^{-1}. The LDO and current mirrors were designed to reject noise below the sampling frequency (27.8 kHz), while higher-frequency noise is averaged out over the integration time.

Fig. 6 presents the power and area breakdown of the system along with the test chip micrograph. The eleven sensors are spread throughout the SoC, with four placed each processor core and three in the uncore that includes a shared cache. The small probe size of 225 μm^2 makes it easy to integrate many sensors throughout the die without impacting the physical design of the full system. As both the current reference and processing logic scale linearly with the number of probes, the total system area is 1044 μm^2/probe. The total power consumption of the system is 616.1 µW, or 56 µW, i.e., 2.0 nJ/S over the 11 probes when accounting for 5-sample averaging.

Fig. 7 shows the temperature readout of the 11 probes while one core runs a matrix multiplication benchmark at different frequencies. The uncore and core temperatures are captured across the 11 sensors and measure a detectable 10 °C–12 °C temperature change caused by the increased power dissipation, as well as a temperature gradient across the cores and uncore of up to 1.7 °C. Measurements presented in Figs. 3–5 were performed with the processor clock gated to avoid self-heating artifacts.

Fig. 8 and Table I compare the proposed work performance with other recent oscillator-based state-of-the-art digital temperature sensors. As in previous publications [2]–[6], [9], the area comparison is performed on absolute values, rather than normalized with process, because the sensor performance is dependent on process mismatch and variability, which does not scale linearly with recent technology nodes. The novel vdd/gnds calibration method we propose offers the

TABLE I
COMPARISON OF THE PROPOSED ARCHITECTURE WITH STATE-OF-THE-ART OSCILLATOR-BASED TEMPERATURE SENSORS

	This work		ISSCC 2009 [5]	JSSC 2016 [4]	ISSCC 2017 [6]	TCAS 2013 [2]	CICC 2015 [3]
Method	Single CMOS		Single CMOS	Diff CMOS	Single CMOS	Diff CMOS	Diff CMOS
Technology	28 nm FD-SOI		130 nm	65 nm	180 nm	65 nm	40 nm
Calibration	1pt (+ 2pt on 1 die)		1pt	2pt	2pt	1 pt	2pt + batch
Supply	0.62 V-1.2 V		1.2 V	0.85 V-1.05 V	1.2 V	1 V	1 V, 0.5 V
Sensitivity [C/V]	1.9		NA	34	0.13	NA	NA
External references	LDO Iref Clock ref.		Clock ref.	none	Clock ref.	LDO ref. Clock ref.	none
Temp range	−5 °C-85 °C		0 °C-100 °C	0 °C-100 °C	−20 °C-100 °C	0 °C-110 °C	−40 °C-100 °C
Systematic correction	No	Yes	No	Yes	Yes (quadratic)	No	No
Accuracy [°C]	-2.5/+1.2	-1.4/+1.3	-1.8/+2.3	-0.9/+0.9	-0.22/+0.19	-1.5/+1.5	-0.95/+0.97
Resolution [°C rms] (a)	0.76		NA	0.3	0.07	0.94	0.12
Conversion rate	27.8 kS/s (b)		5 kS/s	45 kS/s	0.125 kS/s	496 kS/s	50 kS/s
Power	56 µW (0.9 V) (b)		1200 µW	154 µW (c)	0.075 µW (c)	500 µW	241 µW
Energy	2.0 nJ/S (b)		240 nJ/S	3.4 nJ/S (c)	0.6 nJ/S (c)	1.0 nJ/S	4 nJ/S
Area [μm^2]	1.0k (b)		120k	8.2k	8.9k (c)	8.2k	58k

[a]The resolution corresponds to the worst of the quantization step and of the rms repeatability.
[b]Normalized per probe. The aggregated power, conversion rate and area over the 11 probes are 616 µW, 305.8 kS/s and 11.482 μm^2 respectively
[c]Does not include linear correction. Excluding digital processing logic, this work uses 0.1 nJ/S per probe and 307 μm^2 per probe.

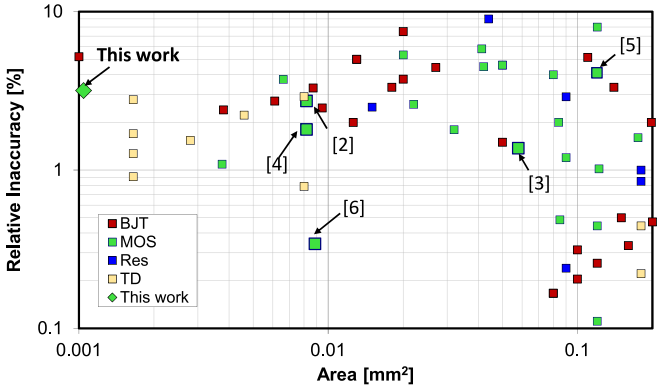

Fig. 8. Comparison of this letter with published temperature sensors [9].

most compact area while requiring only single-point per-die calibration. Compared to [2], which uses the same external references as this letter, the wide operating voltage range (0.62–1.2 V, compared to a fixed 1 V), improved repeatability, and lower area offer a benefit for integrated on-chip sensing.

IV. CONCLUSION

A compact temperature sensor with wide operating voltage range has been proposed. It leverages body-biasing and supply tuning calibration to compensate process variation, and has been demonstrated as part of a dual-core processor SoC. The sensor accuracy of −1.4 °C / +1.3 °C combined with a 0.62–1.2 V operating range and single-point calibration make this sensor a strong candidate for low-overhead temperature monitoring of voltage-scalable low-power SoCs.

REFERENCES

[1] D. E. Duarte, G. Geannopoulos, U. Mughal, K. L. Wong, and G. Taylor, "Temperature sensor design in a high volume manufacturing 65nm CMOS digital process," in *Proc. Custom Integr. Circuits Conf.*, San Jose, CA, USA, 2007, pp. 221–224.

[2] S. Hwang, J. Koo, K. Kim, H. Lee, and C. Kim, "A 0.008 mm^2 500 μW 469 kS/s frequency-to-digital converter based CMOS temperature sensor with process variation compensation," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 9, pp. 2241–2248, Sep. 2013.

[3] M. Saligane *et al.*, "All-digital SoC thermal sensor using on-chip high order temperature curvature correction," in *Proc. IEEE Custom Integr. Circuits Conf.*, San Jose, CA, USA, Sep. 2015, pp. 1–4.

[4] T. Anand, K. A. A. Makinwa, and P. K. Hanumolu, "A VCO based highly digital temperature sensor with 0.034 °C/mV supply sensitivity," *IEEE J. Solid-State Circuits*, vol. 51, no. 11, pp. 2651–2663, Nov. 2016.

[5] K. Woo *et al.*, "Dual-DLL-based CMOS all-digital temperature sensor for microprocessor thermal monitoring," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2009, pp. 68–69.

[6] K. Yang *et al.*, "A 0.6nJ $-0.22/+0.19$°C inaccuracy temperature sensor using exponential subthreshold oscillation dependence," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, San Francisco, CA, USA, Feb. 2017, pp. 160–161.

[7] A. Cathelin, "RF/analog and mixed-signal design techniques in FD-SOI technology," in *Proc. IEEE Custom Integr. Circuits Conf.*, Austin, TX, USA, Apr. 2017, pp. 1–53.

[8] 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.

[9] K. Makinwa, *Smart Temperature Sensor Performance Survey*, Delft Univ. Technol., Delft, The Netherlands, Feb. 2017. [Online]. Available: http://ei.ewi.tudelft.nl/docs/TSensor_survey.xls