A 12.4μm^2 133.4μW 4.56mV/^{o}C Resolution Digital On-Chip Thermal Sensing Circuit in 45nm CMOS Utilizing Sub-Threshold Operation

Basab Datta and Wayne Burleson
Electrical and Computer Engineering Department
University of Massachusetts, Amherst, MA USA
E-mail: {bdatta, burleson}@ecs.umass.edu

Abstract
Dynamic thermal management (DTM) schemes rely on physical sensors to provide them with feedback to ensure an accurate and closed-loop throttling mechanism. Power-density trends in current-generation, high-performance processors motivate the need for multiple low-area, low-power and high-sensitivity temperature monitoring circuits. To this end, we have proposed and implemented in 45nm CMOS, a novel, digital, on-chip thermal sensing circuit with a high thermal sensitivity of 4.56mV/ºC, low power dissipation of 133.4μW and a compact layout occupying 12.4μ^2**.** The design is robust towards process and supply induced noise incurring a 1-σ accuracy loss of <=1.24ºC for a +/-10% random variation in the physical and environmental parameters.

Keywords
Thermal sensor, sub-threshold, low-power, high-resolution

1. Introduction
Aggressive device scaling has been the primary driving factor behind the uninhibited increase in processor and memory speeds over the past decade. While the increased transistor density has enabled system designers to meet the block-level functionality demands of modern processors it has also caused a tremendous increase in the power dissipated per unit area across the die. In order to meet performance demands, operating voltages have not scaled at the same rate as other technology parameters; any improvement in performance now comes at the cost of increased power-density. Thus, a combination of ideal device-scaling and sub-optimal electrical scaling trends coupled with the extensive usage of power-hungry micro-architecture techniques has pushed the power-envelope with every new technology generation. A significant fraction of this power is converted to heat resulting in elevated die temperatures which have a critical impact on all key circuit metrics: lifetime, reliability, speed, power and cost [1]. Thermal gradients are formed as a result of the spatially and temporally non-uniform power-density distributions that arise due to non-uniform on-die activity [2].

Higher operating temperatures increase the device and interconnect-delays which may lead to timing failures [3]. All circuit reliability issues (e.g. electro-migration, time-dependent dielectric breakdown and negative bias temperature instability) have strong temperature dependence and hence, elevated temperatures degrade the mean-time-to-failure of circuits [4]. Different thermal coefficient of expansion of the different on-chip materials causes mechanical stress which can eventually lead to cracks in the chip/package interface [3]. Higher die temperatures increase the leakage-power dissipated which could potentially lead to thermal runaway [5]. Many core architectures with localized power-consumption could aggravate thermal issues and severely limit the overall system performance. [6].

Dynamic Thermal Management (DTM) techniques are used to control on-chip temperature surges and optimize performance [7]. The prominent DTM techniques include clock-gating, dynamic frequency and voltage-scaling (DVFS) and thread-migration [8]. Most DTM schemes are reliant on accurate on-chip temperature sensors to provide them with the necessary feedback. Hotspots tend to shift as a function of the workload necessitating the usage of multiple, low-power, digital thermal sensors to meet the high-level die temperature control requirements. This trend of multiple monitoring circuits is evident in recent processors such as the CELL, Itanium and Power5 processors [9] as well as in multi-core systems like in AMD Opteron [7].

In the recent past, significant research effort has been expended towards developing architectural techniques that tackle: (i) Inaccurate sensor-placement due to unpredictable workloads and (ii) Sensor noise caused by within-die process-variations and variations in operating parameters [1][2][7][8][10]. However, most such model-based runtime thermal prediction schemes still require accurate real-time temperature information [11] to maintain advantages of closed-loop thermal monitoring. The accuracy of the temperature measurements directly affects the performance of the thermal management unit [12]. It has been suggested that better sensor-measurement accuracy translates into power-savings and better system performance (1ºC accuracy equivalent to 2W of CPU power) [13].

In this paper, we present a novel thermal sensor which operates in sub-threshold mode in the sampling mode and hence, has a high resolution of 4.56mV/ºC and a comparatively low active power consumption of 133.4μW in 45nm CMOS-SOI process. Our sensor takes advantage of the increased temperature sensitivity of the MOS device in the sub-threshold region of operation to allow for a higher-than-nominal thermal sensitivity (~1%/ºC). We have analyzed the robustness of our design towards various on-chip noise-sources and implemented the complete sensing system in 45nm CMOS-SOI technology. The rest of the paper is organized as follows. In section 2 we discuss the

related works in the area of on-chip thermal-sensing. In section 3, we discuss the relative advantages of thermal sensing in sub-threshold. In section 4, we present our sub-threshold thermal-sensor and discuss its operation mechanism along with the digitization procedure. In section 5, we analyze the robustness of the proposed design to process-variations and supply-noise. In section 5, we present the implementation details of a test-setup for the proposed sensor. Finally, in section 6 we present our conclusions.

2. Related Work

On-chip thermal sensors can be classified based on their output type: current, voltage, delay, frequency and leakage-decay period. They can be also classified based on the constituent circuit used to perform thermal sensing. This includes differential amplifiers [14], 4-T SRAM cells [15], ring-oscillators [16], CMOS lateral/substrate bipolar transistors [17], delay cells with time-to-digital conversion [18] and diodes [19]. Thermal diodes have remained the de facto industrial standard for thermal sensing providing a reasonably high sensitivity of -1.6mV/ºC at the cost of a high area/power overhead [19]. Several frequency output thermal sensors have been proposed in literature making use of the temperature dependent parameters of MOS transistor/ microelectronic structures present on an IC [20]. These however occupy a large silicon area, have a considerably high power-requirement and are significantly affected by process-variations rendering their usage as sensors unreliable in the sub-90nm regime. Leakage based thermal sensors have been proposed in [21] and [22]. Leakage is a highly process-dependent parameter and at higher temperatures the sensor-response curves become asymptotic making it difficult to discern consecutive measurements and compromising the accuracy of the sensor. Time-to-Digital conversion based thermal sensors [18][23] have a limited measurement range due to non-linearity of the output at the higher temperatures. Recent advances in ADC design and development of precision band-gap reference voltage generators have allowed PNP substrate bipolar transistor based thermal sensors (integrated with σ-Δ converters) to have an accuracy of <=0.5ºC over a wide temperature range of -50ºC to 120ºC [24].

Among digital techniques, recent work done attempts to i) minimize area (sensor, pad and ADC area), ii) improve immunity to supply variations iii) minimize calibration error due to process-variations and iv) minimize usage of analog I/O. However, the underlying circuits employed in these more 'robust' techniques remain the same as ones discussed previously. In [22], the temperature dependence of transistor-leakage current is utilized followed by current-to-time conversion. The sensor (parallel combination of PMOS transistors) although small in size will be very susceptible to process-variations. In [23], a delay line-based approach is taken followed by time-to-digital conversion. Although more robust to process due to elimination of common-mode noise, this technique will have a large area/power requirement due to a large number of ancillary circuits required including delay-elements, charge-pump, multiplexers and a phase-detector. In [25], a frequency to digital converter based sensor has been proposed. The resolution of this sensor degrades drastically (~3ºC) in the high temperature zone (range of concern).

3. Thermal Sensing in Sub-Threshold

The sub-V_{th} I_{on} has direct temperature dependence as opposed to the threshold-voltage and electron-mobility induced temperature dependence of the super-V_{th} I_{on}. The well-established equation used to model sub-V_{th} ON current is given as [26]:

$$I_{sub} = \mu_0 C_{ox} \frac{W}{L}(n-1)\frac{k^2T^2}{q^2} e^{\frac{(V_{GS}-V_{th})q}{n(kT)}} \left[1 - e^{-V_{DS}q/kT}\right] \quad (1)$$

where 'n' is the subthreshold slope factor ($1 + C_d/C_{ox}$) and (kT/q) is the thermal-voltage. Unlike super-threshold, where typically due to the high gate-overdrive the mobility effect dominates and hence the transistor drain current decreases with temperature, in sub-V_{th}, the I_{ON} increases almost exponentially with temperature. Thus, in sub-threshold, circuit speed increases with temperature. Using a 16-stage invX2 inverter-chain as a representative circuit and the 45nm PTM model [27] we found that the delay sensitivity to temperature is 0.177%/ºC in super-threshold while it is 1.86%/ºC in sub-threshold (~10X increase in sensitivity).

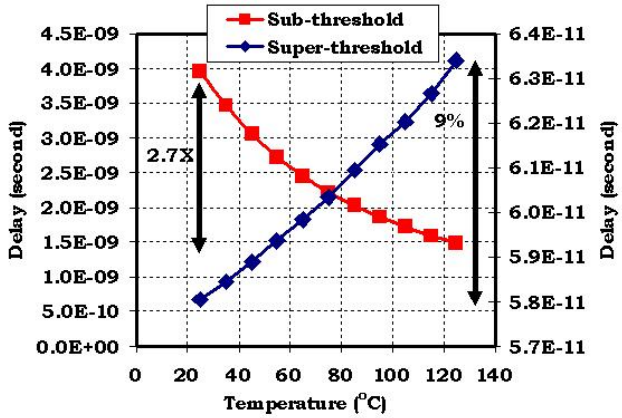

Fig 1 Thermal impact on delay of 16-stage inverter chain

This result indicates that sub-V_{th} circuits are excellent candidates for thermal sensing from the sensitivity perspective. An additional advantage is that since active power dissipation has a quadratic dependence on the supply-voltage, the power consumed by a sub-threshold thermal-sensing circuit will be significantly lesser than its super-threshold counterpart. However, the fact that the sub-V_{th} current has an exponential dependence on threshold voltage makes these circuits much more vulnerable to process-variations. In 45nm technology, process-induced variability was found to be 2.5-3X in sub-threshold as compared to that observed in super-threshold. Also, operating traditional thermal sensing circuits like ring-oscillators or delay-lines directly in sub-threshold mode by scaling down the supply voltage will place routing constraints on the power-distribution network (dual power-delivery lines will be required for sub-threshold sensors in super-threshold blocks). In our proposed design, we attempt to achieve the following objectives: 1) low area/power overhead 2) high temperature sensitivity 3) high tolerance to process/power-supply noise and 4) operation at nominal VDD to enable

seamless embedding of a large number of sensors in different functional-blocks.

4. Proposed Sub-V_{th} Thermal Sensor

In order to take advantages of the increased thermal sensitivity in sub-V_{th} while at the same time operate using the normal supply-voltage we propose a novel design shown in Fig.2. Transistor P1 is diode-connected and behaves like a current-source (always in saturation). Transistors N1 and N2 form a current-mirror and since all the transistors are sized identically (same pull-up to pull-down ratio as well), current into node 'V_A' mirrors the current-flow across P1. When the 'switch' (P3) is closed 'V_{OUT}' follows 'V_A' at one threshold-drop below V_{DD} (for a V_{DD} of 1.0V in 45nm, its ~0.4V). The series transistor stack of N4-N7 forms the temperature sensitive element of this design.

During tracking mode, the signal track/hold (T/H) is held 'low' and the node-cap of OUT retains its charge. Transistors N3-N7 are all diode-connected but since V_A is below the threshold-voltage of the NMOS devices, N4 (driven by V_A) and the remaining stack operate at the edge of linear/sub-threshold mode. As long as the switch is closed, voltage at node 'OUT' follows the voltage at node 'A'. Since the output node is actively driven, the NMOS stack only causes a dc voltage-drop of the output node during tracking mode. For sampling the sensor, we make the T/H signal 'high'. With the 'switch' (P3) open, the loop becomes open i.e. the output node is no longer actively driven and hence, the voltage at 'OUT' quickly drops below the NMOS threshold voltage forcing the entire device-stack into the sub-threshold region. N4 continues to be driven by V_A but since its drain-voltage is below its threshold-voltage the entire stack is forced to operate in sub-threshold. Current through the NMOS stack can now be described by the sub-V_{th} current equation (1) which has an exponential dependence on temperature (increases exponentially with temperature rise).

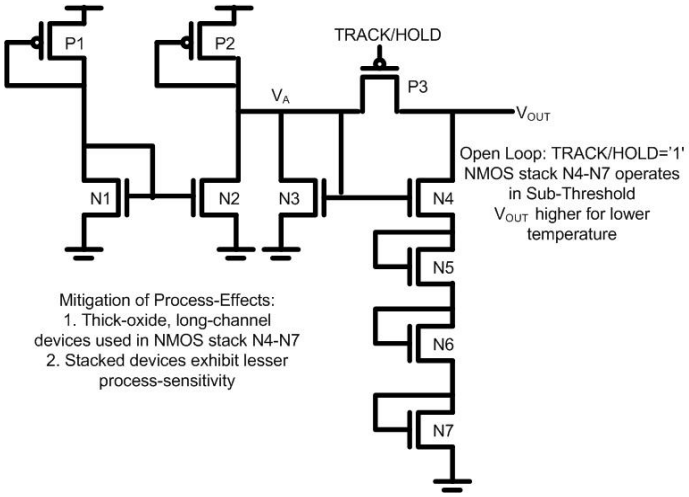

Fig.2 Schematic of Sub-V_{th} Thermal Sensing Circuit

The voltage level acquired by the floating node (V_{OUT}) at the end of the sampling interval has a strong dependence on the temperature of the sensing circuit. It decreases linearly with rise in temperature due to the temperature-dependent sub-threshold current flowing through the NMOS stack. Thus, the main idea of this sensor is that a tracked voltage is modulated by a temperature-dependent sub-V_{th} current during sampling. To mitigate the inherently high process-sensitivity of sub-V_{th} circuits we used: 1) stacked devices which have been shown to have much lesser process-induced current-variability and 2) thick-oxide, long-channel devices in the temperature-sensitive stack to further minimize process-effects.

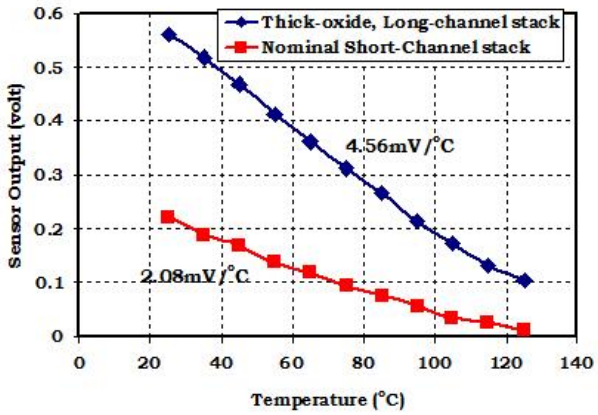

Fig.3 Sensitivity of Sub-V_{th} Thermal Sensor (V_{OUT})

In 45nm PTM technology [27], the nominal and long-channel design variants of the NMOS stack showed a high-sensitivity of 2.08mV/ºC and 4.56mV/ºC respectively (Fig.3). Since we use the nominal supply-voltage during tracking mode and there is no supply to output connection during the sampling mode, for a low sampling frequency the design exhibits the same robustness towards supply-noise as any super-threshold design. The sensor requires only 10 transistors and hence is comparable in size to contemporary 'small' digital designs. It has a relatively low active power-consumption of 133.4µW for the sizing assumed (P1-P3 = 855nm, N1-N3= 551nm, N4-N7=312nm).

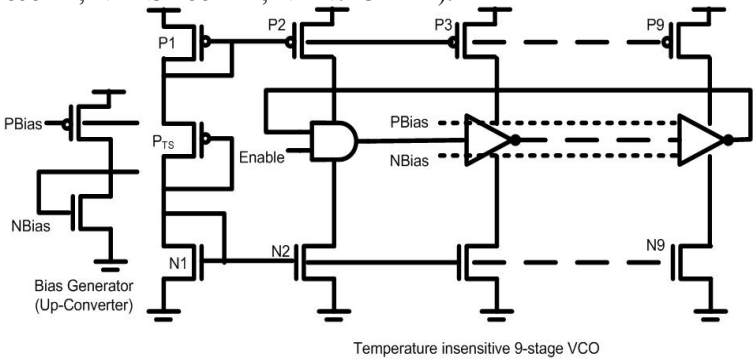

Fig.4 Digitization using temperature-insensitive current-starved 9-stage VCO (V_{OUT} = PBias)

For digitization we adopted the VCO-based approach taken in the Itanium Family of processors [17] to save on area/power. Typically current starved VCOs are extremely sensitive to temperature variations and hence can potentially add noise to the temperature-reading during digitization. Therefore, we used the temperature-insensitive design shown in Fig.4 wherein the diode-connected transistors P_1, P_{TS} and N_1 provide temperature compensation to the remaining circuit. Transistors P_1 and N_1 were sized in such a manner that the drain current of P_{TS} remained invariant to temperature and hence, the current mirrored into the current-

starved branches was also temperature insensitive. This particular VCO design offered a low temperature sensitivity of 0.0071%/ºC as opposed to a sensitivity of 0.79%/ºC observed in a regular current-starved VCO.

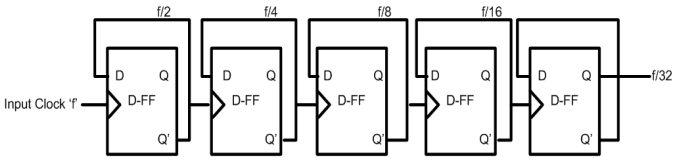

Figure 5 Divide-by-32 Ripple Divider

The sensor output served as the bias for the VCO (P-devices specifically) while bias for the N-devices was generated using an up-converter. We used a gated VCO wherein the 'track/hold' signal was used as the 'enable' signal and we performed digitization only when sampling the sensor to minimize power-consumption and reduce noise disturbances to other circuit blocks. The VCO frequency was converted to a digital value by using a frequency-divider (divide-by-32) (Fig.5) followed by a 10-bit counter. A simple ripple-divider comprising of five D-flip-flops was used to perform the frequency-division. Frequency-division minimizes the size of the subsequent counter needed and also minimizes chances of reflection/interference during external probing.

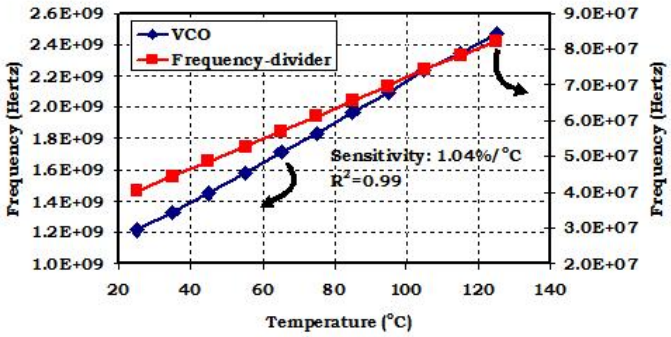

Fig.6 Digitization using VCO and frequency-divider

As is evident from Fig.4, bias for the VCO becomes stronger as temperature increases (PBias decreases and NBias increases) and hence, the oscillation frequency increases linearly with temperature (Fig.6). We were able to achieve a relatively high post-digitization temperature sensitivity of 1.04%/ºC. The thermal sensing circuit was found to have a low active-power consumption of 133.4 µW and the VCO and divide-by-32 frequency-divider had an active power consumption of 253.2µW and 318.6µW respectively in 45nm technology.

5. Analysis of Process Variation and Supply Noise Tolerance

Process variations are the variations in the environmental and physical design parameters which cause the circuit-performance to deviate from its nominal value. Environmental variables are the dynamic operating parameters such as temperature and supply-voltage. Intra-die physical variations include variations in device electrical parameters and variation in interconnect-dimensions. Device variations are MOS parameter fluctuations which include effective channel-length (L_{eff}), threshold-voltage (V_{th}), gate-oxide thickness (T_{ox}) and drain/source parasitic resistance (R_{dsw}). The photo-lithographic and etching process causes variations in interconnect-geometry like line-width and spacing. To study the effects of intra-die process-variation on the response of proposed sensor, we modeled effective channel-length and threshold voltage variations using a Gaussian distribution with a maximum variation (3-σ) of +/-10% and performed 50 Monte-Carlo simulations. HSPICE™ (Synopsys) was the main simulation platform used in our analyses. 45nm technology files were obtained from the Predictive Technology Model website [27].

Table 1 Device Parameters and +/-3σ variations in 45nm

Parameter	Nominal + 3σ deviation
L_{eff} (nm)	17.5 +/- 10%
V_{th} (V)	(0.466,-0.418) +/- 10%

First, we varied the number of devices in the temperature-sensitive stack and determined the resolution of the resultant design and the variability induced by simultaneous random variations in L_{eff} and V_{th}. Fig. 7 suggests that a stack of size 4 and upwards yields the best results in terms of variability (~ 74% decrease from a stack of size 2). Resolution improves albeit marginally with increase in the size of the stack (~23%). Increasing the size of the stack beyond 4 yielded very little improvement from the variability perspective and hence we kept it fixed at 4 for all subsequent analysis.

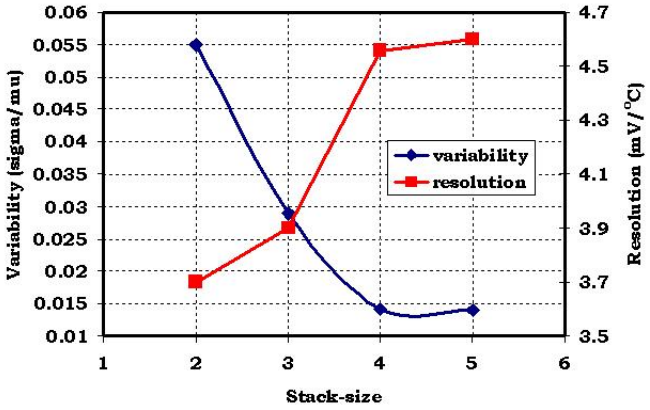

Fig.7 Process induced variability in sensor response for various stack-size and corresponding resolution

Increasing device-size is a popular design optimization technique to improve robustness towards process-noise. For our design, we increased the device-size of the temperature-sensitive NMOS stack from its nominal value (315nm) to ~2.4X (750nm). For each Monte Carlo simulation we determined the deviation from the nominal response caused by the process-variation and then calculated the resultant loss in accuracy assuming 1ºC loss for every 4.56mV deviation.

Fig.8 shows the histogram plot of accuracy-loss for various device-sizing. It is evident that we can significantly reduce the accuracy loss due to process-effects by selectively upsizing the temperature sensitive devices. For device-sizing of 750nm, a 1-σ deviation in response will result in an error of 1.24ºC. Upsizing the device-stack does not affect the post-digitization sensor resolution (in the range of 0.9-1% even for high temperatures) and has a negligible active power-overhead (only 2.5%) as indicated by Fig.9. Thus,

using large, thick-oxide, long-channel devices in a stacked configuration for the temperature sensitive element in the proposed design provides strong tolerance against random variations in device parameters.

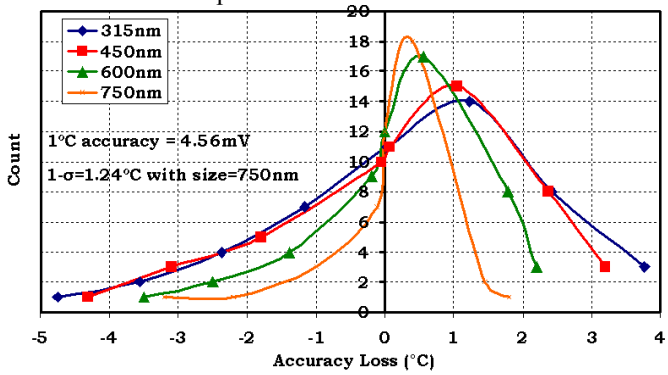

Fig.8 Accuracy loss due to random process variations

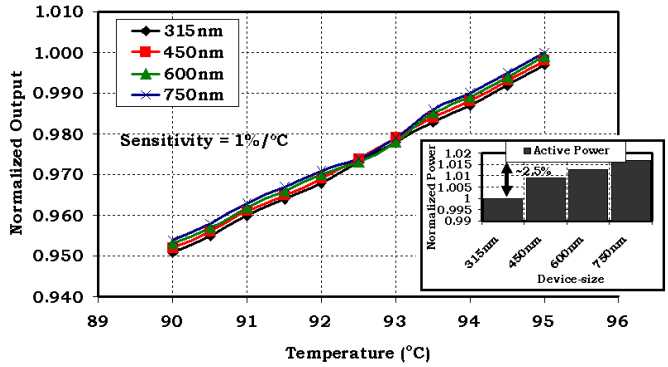

Fig.9 Impact of upsizing on post-digitization sensitivity and dynamic power (normalized output for each sizing)

As CMOS feature sizes shrink and device integration density increases, on-chip power supply droop caused by large dI/dt becomes an increasing concern. To study the impact of voltage-supply on the sensor-response, we systematically varied the supply by 10% on either side of the nominal value. We calculated the measurement error as the error incurred when calibrating against an ideal design without any supply deviation. The worst-case measurement error was found to be ~3.6% (Fig.10) which for the pre-digitized sensor output can result in an accuracy loss of 4.42°C.

To study the effect of random supply noise on sensor performance, we modeled it as a Gaussian distribution with a 3-σ variation of +/-10% and performed 30 Monte-Carlo simulations to calculate the resultant loss in accuracy at different temperatures of measurement (1°C=4.56mV). The 1-σ accuracy loss was found to be 1.2°C (Fig.11) which is small enough to prevent the false-triggering of DVFS schemes. In the proposed sensor, the supply-voltage primarily affects the initial charge stored in the parasitic capacitance of the output node which is then discharged during sampling. There is little impact of voltage droop during actual polling of the sensor. The results shown in Fig.10 and Fig.11 are for a sampling frequency of 100MHz. Since the time-scale of on-chip thermal events is quite slow i.e. of the order of 10ms [22], the sensor response time needs to be 1ms or lesser in order to effectively respond to sudden temperature rises. For such a slow sampling rate (order of few kHz), most components of high and mid-frequency droops will be eliminated which will allow the output node to track a stable voltage and hence the worst-case accuracy loss due to supply-droop will be even lesser than the observed values.

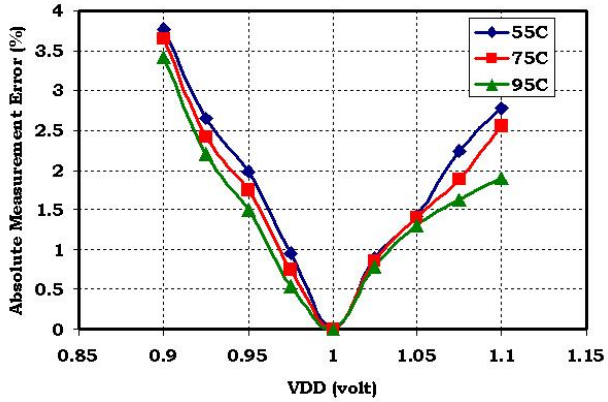

Fig.10 Measurement error in sensor-response due to systematic power-supply deviation

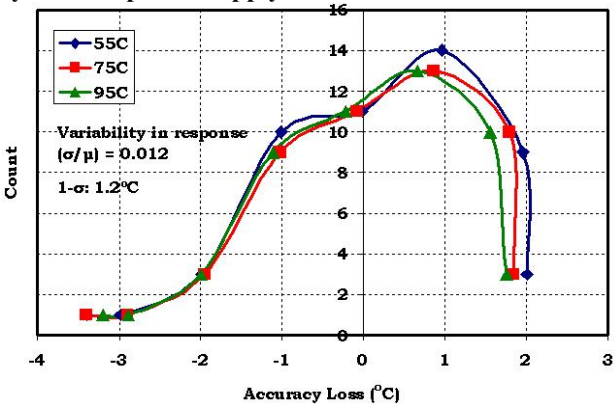

Fig.11 Accuracy loss due to power-supply noise

6. Implementation in 45nm CMOS-SOI: Measurement and Calibration

The proposed sensor was implemented using a 45nm CMOS-SOI design-kit provided by MOSIS. Fig.12 gives a top-level view of the complete sensing system we used in our test setup. Our test-setup comprised of 4 thermal sensors which could be enabled sequentially by applying the appropriate sensor selection bits.

A 2:4 decoder and 4:1 multiplexer were used to enable and sample a particular sensor. We used separate blocks for digitization of each sensor output to minimize noise-susceptibility of the mixed-signal components. We also eliminated the counter block since off-chip measurements could easily be taken using an oscilloscope after applying appropriate buffering. A divide-by-32 frequency divider was used in order to limit the sensor-output within the MHz range. This minimizes risk of reflection observed in the GHz frequency range with meter-long probes. For an actual on-chip implementation in a microprocessor, the frequency-divider will reduce the size of the subsequent counter and also the size of registers used to store the digitized output. A cascade of buffers was used to drive the output pad. Diode-connected shunting devices were used to sink the charge accumulated during the plasma-etching process. A VDD of 1.0V was used and the sensors were sampled at 100MHz.

Post-layout simulations were run to validate the design and obtain area/power results (Table 2).

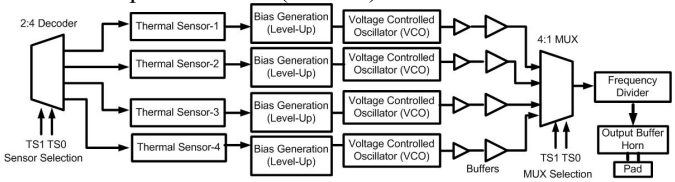

Fig.12 Test setup for proposed thermal sensors

Our proposed design has a low area and power overhead of 12.4μ^2 and 133.4μW respectively. The entire sensing system is relatively compact at ~2200μ^2. To sample a particular sensor we apply appropriate selection bits to the decoder whose output activates the sensor as well as the corresponding VCO. The same selection bits are applied to the multiplexer as well to ensure that noise from other sensors do not corrupt the output channel and only the desired sensor-output is channeled through the I/O buffers.

Table 2 Area/Power overhead @ 100MHz sampling frequency in 45nm CMOS-SOI

Component	Area (μ^2)	Power (μW)
2:4 Decoder	60.5	15.7
Thermal Sensor	12.4	133.4
Bias Generator	3.6	3.6
VCO	65	253.2
4:1 Multiplexer	71.2	112.5
Frequency Divider	384.3	318.6
I/O Buffers	340.5	100.4
Total	2200	897.6

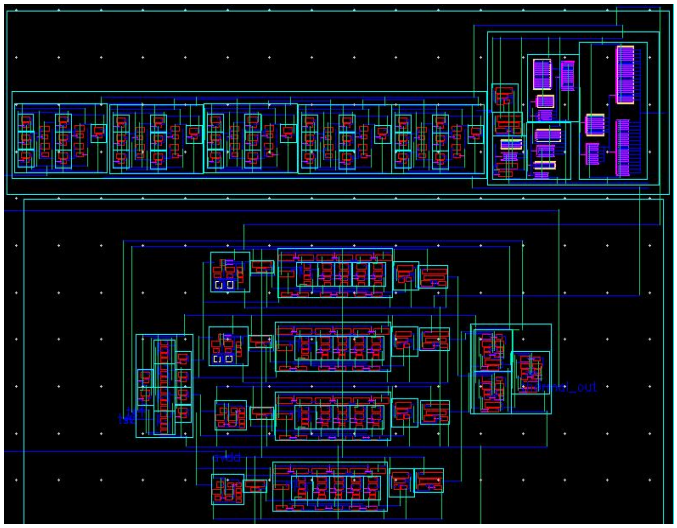

Fig.13 Layout of test setup in 45nm CMOS-SOI

Digitization using an ADC requires a large-area and also suffers from high sensitivity to different circuit parameters. In our VCO-based digitization approach, a linear mapping between sensor-output frequency and temperature is maintained even in the high temperature zone (typically affected by loss of resolution). Post-digitization the sensor-output has a high resolution of 8410kHz/ºC which ensures a discernible deviation between close temperature points. Our VCO design is such that the shift in its gain and frequency range caused by temperature drift is nominal. A combination of slow sensor-sampling (sampling interval of the order of μs) and temporal averaging of the sensor-response could be used to mitigate/average-out effects of dynamic variations on sensor operation. The multiple advantages of proposed sensor include – small area, high-sensitivity, low-power and high PV-tolerance. This is beneficial towards thermal research at the architectural level since now a PV-tolerant, 'low-overhead' measurement entity is available to make informed decisions for optimized DTM schemes.

Table 3 compares the proposed design with other contemporary thermal sensing circuits. The comparison metrics suggest that our design performs favorably by presenting an excellent tradeoff between resolution Vs area/power-overhead.

Table 3 Comparison with state-of-the-art designs

Sensor	Power	Area (mm^2)	Resolution	Process
[18]	1.5 μW @ 5 samples/s	0.09	0.6ºC	0.35μm
[14]	50 μW	0.013	<1ºC	0.18μm
[15]	221 μW	0.0017	1ºC	0.18 μm
[23]	High	0.12	0.66ºC	0.13 μm
[22]	Low	0.01225	<1ºC	65nm
[16]	2.5mW	2.3	N/A	0.7 μm
[28]	0.49mW	0.175	0.16ºC	0.35 μm
This Work	133.4μW	12.4e-6	0.1ºC	45nm

Proper calibration is necessary to ensure accurate temperature measurement in the presence of varying environmental and physical parameters. In our test-setup we adopted the 2-point calibration approach typically used with thermal diodes. We measure the output frequency (F_1 and F_2) at 2 temperature points (T_1 and T_2) and calculate the calibration constants as shown below.

$$F_{OUT} = mT + b \quad (2)$$

$$m = \frac{F_2 - F_1}{T_2 - T_1} \qquad b = F_{OUT} - mT \quad (3)$$

$$T_{read} = \frac{F_{measured} - b}{m} \quad (4)$$

Expectedly, the calibration constants will shift as a result of temperature and aging effects on the VCO. Hence, for usage in actual microprocessors, a more prudent approach will be to follow the calibration procedure used in [19] wherein a temperature-compensated, precision band-gap voltage reference is used to cycle the VCO through a fixed set of voltages and a frequency/output-count to voltage mapping table is obtained. Henceforth, any frequency/output-count reading will lie between two particular points in the calibration-table constructed previously and the corresponding temperature can be obtained by interpolation (enabling higher resolution) (Fig.14).

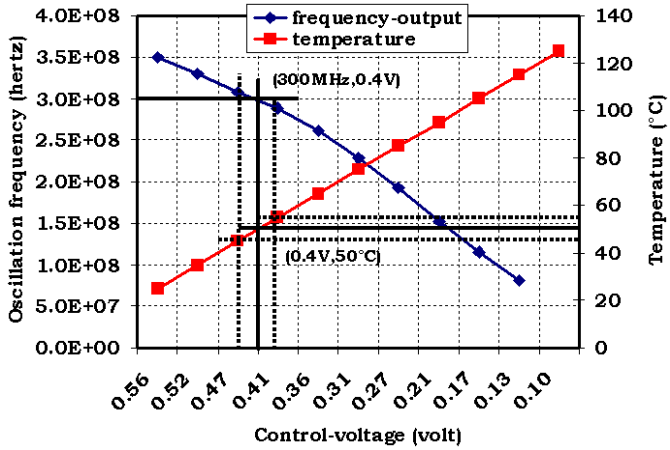

Fig.14 Temperature determination through interpolation between successive entries stored in calibration table to perform (freq→V_{out}) and (V_{out}→Temp) mapping

In figure 14 (built using post-layout simulations), the frequency reading of 300MHz was found to lie between the entries (320 MHz, 0.47V) and (280 MHz, 0.41V) and hence, mapped to 0.4V. The control-voltage obtained was then used to map to the sensor temperature of 50ºC lying between entries (0.47V, 40ºC) and (0.41V, 55ºC). To account for environment/aging induced variance in sensor response, re-calibration must be done periodically and the calibration table must be built afresh.

7. Conclusions

A digital CMOS thermal sensor targeted for on-die temperature measurement with high tolerance to process and voltage noise was implemented in 45nm CMOS-SOI technology. Due to utilization of sub-threshold mode of operation, the design offers a high thermal sensitivity of 4.56mV/ºC and has a low active power consumption of 133.4 µW. The use of predominantly digital and mixed-signal components ensures a compact layout measuring 2200μ^2 for the entire prototype sensing system. The design exhibits good tolerance towards process and power-supply induced noise, typically incurring a 1-σ accuracy loss of 1.24ºC and 1.2ºC respectively for a sampling frequency of 100MHz.

8. References

[1] A. Nowroz et al, "Thermal monitoring of real processors: Techniques for sensor allocation and full characterization", DAC, 2010

[2] R. Cochran et al, "Spectral techniques for high resolution thermal characterization with limited sensor data", DAC, 09

[3] D. Brooks et al, "Power, thermal and reliability monitoring in nanometer-scale microprocessors", Micro, 07

[4] M. Pedram et al, "Thermal modeling, analysis and management in VLSI circuits", Proceedings of IEEE, 06

[5] S.Lin et al, "Opportunities and implications for power and thermal management", IEEE Trans. On Electron Devices, 08

[6] W.Huan et al, "Many core design from a thermal perspective", DAC, 08

[7] Y.Zhang et al, "Accurate temperature estimation using noisy thermal sensors", DAC, 09

[8] R.Cochran et al, "Consistent runtime thermal prediction and control through workload phase detection", DAC, 2010

[9] M.Sasaki et al, "-1/0.8ºC Error, accurate temperature sensor using 90nm 1V CMOS for on-line thermal monitoring of VLSI circuits", IEEE Trans. Semiconductor Manufacturing, 2008

[10] S.Sharifi et al, "Accurate temperature estimation for efficient thermal management", ISQED, 08

[11] A.Coskun et al, "Temperature aware task scheduling in MPSoCs", DATE, 07

[12] K.Skadron et al, "Temperature aware microarchitecture", ISCA, 03

[13] E.Rotem et al, "Temperature measurement in the Intel Core Duo Processor", Intel White Paper, 06

[14] A. Syal et al, "Sensing temperature in CMOS circuits for thermal testing", VTS, 04

[15] S Kaxiras et al, "4-T Decay Sensors: A new class of small, fast, robust and low-power, temperature/leakage sensors", ISLPED, 04

[16] C.Kyung-Kim et al, "CMOS temperature sensor with ring oscillator for mobile DRAM self-refresh control", Microelectronics Journal, 07

[17] A. Bakker, "High accuracy CMOS smart temperature sensors" Kluwer Academic Publishers

[18] P. Chen et al, "A time-to-digital converter based CMOS smart temperature sensor", IEEE JSSC, 2005

[19] R.McGowen et al, "Power and temperature control on a 90nm Itanium family processor", IEEE JSSC, 2006

[20] V.Szekely et al, "CMOS sensors for on-line thermal monitoring of VLSI circuits", IEEE TVLSI, 1997

[21] P.Ituero et al, "Leakage based on-chip thermal sensor for CMOS technology", ISCAS, 2007

[22] E.Saneyoshi et al, "A 1.1V 35µm x 35µm thermal sensor with supply voltage sensitivity of 2ºC/10%-supply for thermal management in the SX-9 supercomputer", IEEE Symposium on VLSI Circuits, 2008

[23] K.Woo et al, "Dual-DLL-Based CMOS all-digital temperature sensor for microprocessor thermal monitoring", ISSCC, 2009

[24] M.Pertijs et al, "A CMOS temperature sensor with a 3σ inaccuracy of +-0.5ºC from -50ºC to 120ºC", ISSCC, 2003

[25] H.Lee et al, "A 0.0018 mm^2 frequency to digital converter based CMOS smart temperature sensor", Analog Integrated Circuits & Signal Processing, Springer, 2009

[26] B.Datta et al, "Temperature effects on energy optimization in sub-threshold circuit design", ISQED, 2009

[27] Predictive Technology Model, ASU

[28] K.Makinwa et al, "A CMOS temperature to frequency converter with an inaccuracy of less than +/-0.5ºC (3σ) from -40ºC to 105ºC", IEEE JSSC, 2006