2.7 A 660pW Multi-Stage Temperature-Compensated Timer for Ultra-Low-Power Wireless Sensor Node Synchronization

Yoonmyung Lee, Bharan Giridhar, Zhiyoong Foo, Dennis Sylvester, David Blaauw

University of Michigan, Ann Arbor, MI

Recent work in ultra-low-power sensor platforms has enabled a number of new applications in medical, infrastructure, and environmental monitoring. Due to their limited energy storage volume, these sensors operate with long idle times and ultra-low standby power ranging from 10s of nW down to 100s of pW [1-2]. Since radio transmission is relatively expensive, even at the lowest reported power of 0.2mW [3], wireless communication between sensor nodes must be performed infrequently. Accurate measurement of the time interval between communication events (i.e. the synchronization cycle) is of great importance. Inaccuracy in the synchronization cycle time results in a longer period of uncertainty where sensor nodes are required to enable their radios to establish communication (Fig. 1), quickly making radio dominate the energy budget. Quartz crystal oscillators and CMOS harmonic oscillators exhibit very small sensitivity to supply voltage and temperature [4] but cannot be used in the target application space since they operate at very high frequencies and exhibit power consumption that is several orders of magnitude larger (>300nW) than the needed idle power. A gate-leakage based timer was proposed [5] that leveraged small gate leakage currents to achieve power consumption within the required budget (< 1nW). However, this timer incurs high RMS jitter (1400ppm) and temperature sensitivity (0.16%/ºC). A 150pW program-and-hold timer was proposed [6] to reduce temperature sensitivity but its drifting clock frequency limits its use for synchronization. The quality of a timer is not captured well by RMS jitter since it ignores the averaging of jitter over multiple timer clock periods in a single synchronization cycle. Instead, we propose the uncertainty in a single synchronization cycle of length T as new metric and use this synchronization uncertainty (SU) to evaluate different timer approaches. The timer period is a random variable X(n), with mean and sigma, μ and σ. Given a synchronization cycle time T, consisting of N timer periods, we define SU as the standard deviation of T as given by $\sqrt{T/\mu} \times \sigma$, assuming X(n) is Gaussian. Note that a smaller clock period increases N and results in more averaging and a lower SU with fixed jitter (σ/μ).

The timer in [5] has a high SU since it is triggered with a low gain Schmitt trigger and it has a long period (~10s). To combat this, we introduce: 1) a multi-stage structure with a high-gain triggering buffer, 2) boosted capacitance charging, 3) the use of zero threshold voltage transistor (ZVT) for faster gate leakage discharge and 4) closed-loop temperature compensation to reduce temperature sensitivity. The structure of the proposed multi-stage gate-leakage based timer and its waveforms are shown in Fig. 1. In a stage, a load capacitor (C_L) is charged with the combined gate leakage current of a ZVT and a PMOS transistor. As C_L is charged, the output driving the next stage is triggered by a buffer stage, which shows higher gain than a traditional Schmitt trigger previously used [5]. This places the next stage in a charging state while the current stage discharges. At any given time, only one stage is in a charging state while all others discharge. This allows *n-1* more discharging time than charging time in an *n* stage timer and increases the voltage swing on C_L (Q[n]). Longer discharge time lowers the slope at node Q[n] at the end of discharging state (from -238mV/s to 20mV/s for n from 3 to 10), which makes the initial capacitor node voltage for next following charging stage less sensitive to uncertainty. Each stage has low and high supply voltage domains and the use of two voltage domains allows us to boost the gate-leakage current with a higher supply voltage which steepens the charging transition on Q[n] by 5× and reduces uncertainty at the triggering point (Fig. 1).

We achieve temperature sensitivity compensation by exploiting the opposite temperature dependencies of gate leakage in ZVT and PMOS during the charging state as shown in Fig. 2. By using arrays of ZVT and PMOS transistors and selecting an appropriate combination of transistor sizes for charging, linear temperature dependency can be eliminated. This compensation scheme results in a residual second order dependency. To minimize the impact of this second order dependency, we propose an adaptive scheme in which, for each temperature range, a controller automatically selects a pre-stored transistor size configuration which minimizes the second order dependency (top left of Fig. 2 and Fig 5, left). The optimal configurations are determined and stored during post-silicon testing. Each time when the sensor node processor wakes up, it computes time by calculating the elapsed time using the stored period for proceeding configuration and the number of cycles during the last standby state. The transition between configurations occurs synchronously when the first stage starts a new charging state; this allows an exact period calculation and prevents noise injection during capacitor charging. Un-selected ZVTMOS transistors are driven to 400mV to minimize leakage by placing them in accumulation mode.

A test chip was designed and fabricated in 0.13μm CMOS with the proposed multi-stage gate-leakage timer (MGT). Measured results in Fig. 3 show that as the number of stages increases, duty cycle (the ratio of charging time to timer period) decreases, increasing voltage swing and reducing jitter by up to 8.1×. Boosting of the charging gate-leakage current (Fig. 2) leads to higher jitter, particularly for low stage counts. However, the key SU metric for an interval of 1 hr is reduced by 3× due to the shorter clock period, which enhances statistical averaging. Together with multi-staging and boosted charging, SU is reduced by 3.6×. With small stage counts (<5) power consumption increases. This is due to the higher average node voltage of Q[n] resulting in higher leakage current for the triggering buffer (Fig. 4). With high stage counts (>7), power increases due to static leakage of added stages. A proposed MGT with 9 stages was tested for 24 hours allowing us to compute the SU for a large number of synchronization intervals. We also tested a baseline 3-stage MGT without boosted charging or ZVT transistors. The SU distribution had expected value of 196ms for 1 hour synchronization intervals. It also shows that the proposed timer reduced the expected SU by 4.1× compared to the baseline. Since the period of the timer is not truly Gaussian, the measured SU was larger than the theoretical calculation based on jitter. Power supply sensitivity was 0.42%/mV from 650mV to 750mV for low supply and was 0.49%/mV from 1.15V to 1.25V for high supply. This necessitates the voltage regulation using an ultra-low power voltage reference such as the one proposed in [7].

The period of the temperature compensated MGT for -20-60ºC with selected configurations is shown in Fig. 5. A five configuration scheme and its temperature range is shown as an example (Fig. 5, left). For each configuration, period deviation as a function of temperature is shown and worst period deviation was 0.28% (Fig. 5, top right). With a single configuration, the maximum deviation in period over -20-60ºC was 3% while the use of 10 configurations reduced this to 0.25%, giving an effective temperature sensitivity of 31ppm/ºC. Measured results from a closed loop, temperature compensated MGT is shown in Fig. 6 when the temperature oscillates between 20 and 30ºC. The closed loop temperature compensation reduces SU by 4.8×. A second test chip where the proposed closed loop temperature compensation was implemented on-die was also tested and Fig 6, top left, shows how the configurations track with temperature.

Acknowledgement:
The authors thank MuSyC for partial funding support and Samsung Scholarship Foundation.

References:

[1] B. Warneke, et al., "An Ultra-Low Energy Microcontroller for Smart Dust Wireless Sensor Networks," *ISSCC Dig. Tech. Papers*, pp. 316-317, Feb. 2004.
[2] G. Chen, et al., "Millimeter-Scale Nearly Perpetual Sensor System with Stacked Battery and Solar Cells," *ISSCC Dig. Tech. Papers*, pp. 288-289, Feb. 2010.
[3] M. Crepaldi, et al., "An Ultra-Low-Power Interference-Robust IR-UWB Transceiver Chipset Using Self-Synchronizing OOK Modulation," *ISSCC Dig. Tech. Papers*, pp. 226-227, Feb. 2010.
[4] M. McCorquodale, et al., "A 0.5-to-480MHz Self-Referenced CMOS Clock Generator with 90ppm Total Frequency Error and Spread-Spectrum Capability," *ISSCC Dig. Tech. Papers*, pp. 350-351, Feb. 2008.
[5] Y. Lin, et al., "A sub-pW timer using gate leakage for ultra low-power sub-Hz monitoring systems," *IEEE Custom Integrated Circuits Conf.*, pp. 397-400, Sep. 2007.
[6] Y. Lin, et al., "A 150pW Program-and-Hold Timer for Ultra-Low-Power Sensor Platforms," *ISSCC Dig. Tech. Papers*, pp. 326-327, Feb. 2009.
[7] M. Seok, et al., "A 0.5V 2.2pW 2-Transistor Voltage Reference," *IEEE Custom Integrated Circuits Conf.*, pp. 577-580, Sep. 2009.

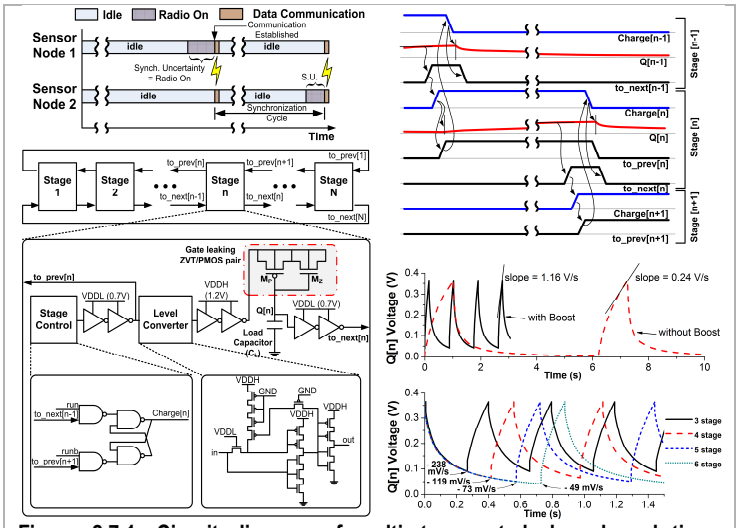

Figure 2.7.1: Circuit diagram of multi-stage gate-leakage-based timer (bottom left) and simulated waveform showing effect of multi-stage and boosted charging (bottom right).

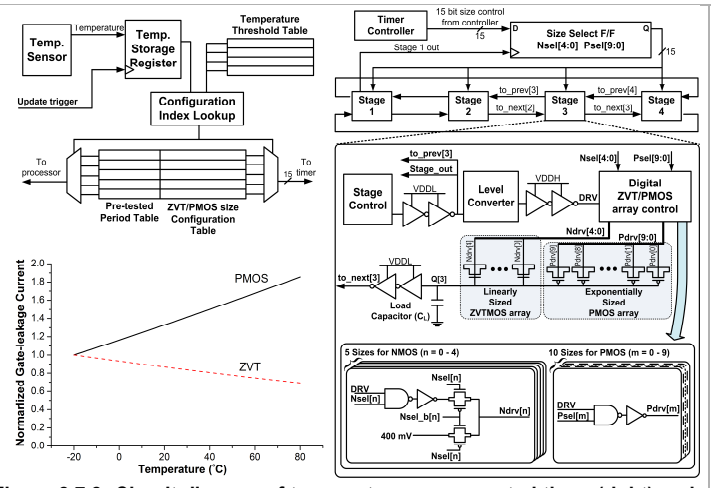

Figure 2.7.2: Circuit diagram of temperature compensated timer (right) and its controller (top left). Compensation is realized by exploiting opposite temperature dependency of ZVTMOS and PMOS gate leakage current as shown in simulated data (bottom left).

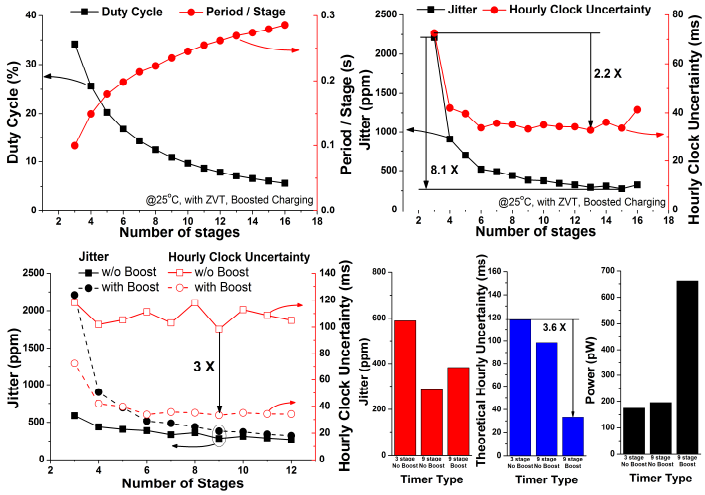

Figure 2.7.3: Measured results shows that, with larger number of stages, duty cycle decreases (top left) and jitter is reduced (top right). With boosted charging, uncertainty is reduced by 3× (bottom left).

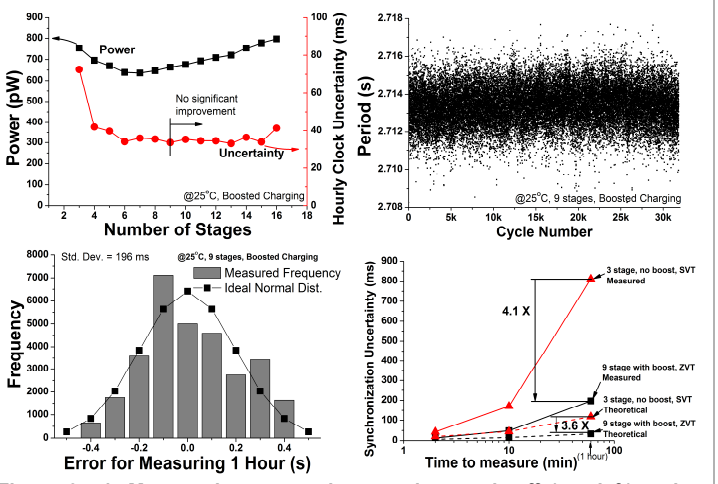

Figure 2.7.4: Measured power and uncertainty trade-off (top left) and a continuous 24-hour measurement (top right). Synchronization uncertainty (SU) distribution for 1 hour measurement (bottom left) and uncertainty reduction was 4.1× (bottom right).

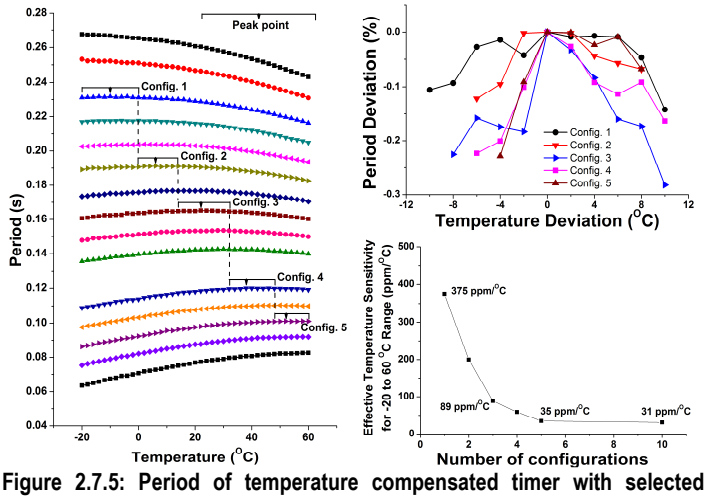

Figure 2.7.5: Period of temperature compensated timer with selected configurations (left) and period vs. temperature deviation for selected configurations (top right). Maximum period variation for -20°C to 60°C range decreases with use of more configurations (bottom right).

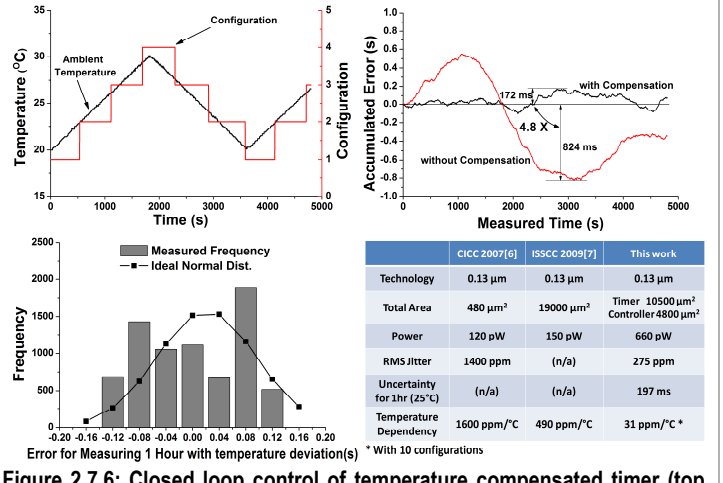

	CICC 2007[6]	ISSCC 2009[7]	This work
Technology	0.13 µm	0.13 µm	0.13 µm
Total Area	480 µm²	19000 µm²	Timer 10500 µm² Controller 4800 µm²
Power	120 pW	150 pW	660 pW
RMS Jitter	1400 ppm	(n/a)	275 ppm
Uncertainty for 1hr (25°C)	(n/a)	(n/a)	197 ms
Temperature Dependency	1600 ppm/°C	490 ppm/°C	31 ppm/°C *

* With 10 configurations

Figure 2.7.6: Closed loop control of temperature compensated timer (top left) and accumulated time measurement error with given temperature profile (top right). Synchronization uncertainty distribution (bottom left) and comparison with other works (bottom right).