8.2 Batteryless Sub-nW Cortex-M0+ Processor with Dynamic Leakage-Suppression Logic

Wootaek Lim, Inhee Lee, Dennis Sylvester, David Blaauw

University of Michigan, Ann Arbor, MI

Recent low-voltage design techniques have enabled dramatic improvements in miniaturization and lifetime of wireless sensor nodes [1-3]. These systems typically use a secondary battery to provide energy when the sensor is awake and operating; the battery is then recharged from a harvesting source when the sensor is asleep. In these systems, the key requirement is to minimize *energy per operation* of the sensor. This extends the number of operations on one battery charge and/or reduces the time to recharge the battery between awake cycles. This requirement has driven significant advances in energy efficiency [1-2] and standby power consumption [3].

Batteries suffer from limited endurance (e.g., 5k discharge cycles [4] limiting lifetime to 3.5 months with a 30 min wakeup period) and scalability challenges in the sub-5 mm range due to sealing requirements [5]. This paper therefore focuses on a battery-less sensor system that operates directly from the energy harvesting source. In these systems, power is consumed as it is obtained, and hence the key requirement is to limit the *maximum power draw*, thereby reducing the size of the required harvesting source. While significant advances have been made in low power systems [6], the minimum power draw per logic gate remains in the 1 - 30pW range, resulting in 10s of nW consumed by a microcontroller. This in turn requires a relatively large harvesting source, limiting the ability to scale a sensor system to true miniature sizes (e.g., an 4mm² solar cell @240 lux is needed to produce 30nW [7]). Note that reducing supply voltage further in these systems is ineffective since they become leakage power dominated. Robustness concerns also often limit voltage scalability.

This paper proposes a new logic implementation, referred to as dynamic leakage-suppression logic (DLSL) that consumes 10fW active power per gate, making two orders of magnitude improvement over recently published work. Power is reduced through a super-cut-off feedback mechanism, and minimum power is achieved at 350 – 550mV supply voltage. This supply voltage range eliminates the need for ultra-low voltage operation, which increases robustness. It also allows the circuits to be directly connected to various harvesting sources without DC-DC conversion. DLS logic is used to implement a Cortex M0+ processor that consumes 295pW, which is the lowest reported to date for a microcontroller. We show full functionality across –5 to 65°C and demonstrate autonomous operation when powered by a 0.09mm² solar cell in room lighting (240lux).

Fig. 1 shows a DLS inverter and its steady-state node voltages at VDD = 0.4V. The output voltage of the gate is fed back to the bottom PMOS and top NMOS, placing all leaking transistors in a super-cutoff state. When IN = 0, the leakage current is contributed by the pull-down logic M_{NB} and M_{PB} . Since the gate of M_{PB} is connected to a high OUT voltage, node *n2* settles to roughly half VDD, placing both M_{N2} and M_{P2} into super-cutoff. The same dual super-cutoff effect occurs with M_{NT} and M_{PT} when IN = VDD.

During dynamic operation of DLS logic (Fig. 1), the output node transitions using the leakage currents of the top and bottom transistors, which are in initially in super-cutoff (M_{NT} for rising and M_{PB} for falling output transitions). As IN transitions from 0V to VDD, M_{NB} switches from super-cutoff to weak-inversion and starts to equalize the voltage of *n2* and OUT. This has two effects: 1) M_{PB} switches from super-cutoff to a traditional cutoff bias point. 2) As M_{NB} pulls *n2* up, OUT is also being discharged, as is *n1* to some degree. This causes M_{NT} and M_{PT} to become super-cutoff, sharply reducing the leakage from VDD to OUT. At the same time, the leakage through M_{PB} , which is no longer super-cutoff, continues to pull OUT low, further suppressing the leakage of M_{NT} and M_{PT} and accelerating the overall discharge of OUT. Due to this super-cutoff feedback effect, DLS logic naturally has different rising and falling switch points, resulting in hysteresis and a 1.45× increase in static noise margin over a standard CMOS inverter (Fig 2).

At VDD=0.4V, DLS logic obtains $320 \times$ lower leakage than a standard lowleakage inverter with 2-stacked transistors (Fig. 2). There are two key reasons for this improvement: 1) In DLS logic, two series-connected devices are in super-cutoff whereas in the 2-stacked topology there is only one. 2) In DLS logic, the complementary nature of the PMOS and NMOS super-cutoff transistors causes the intermediate node (*n*1 or *n*2) to settle to half VDD, providing a very strong super-cutoff effect that increases with higher VDD. In comparison, in stacks of two NMOS or PMOS transistors the intermediate node settles to ~20mV at VDD =0.4V, resulting in a much weaker super-cutoff effect that is independent of VDD to first order. The strong super-cut-off effect also increases the I_{ON}/I_{OFF} ratio compared to a stacked topology, improving static robustness (Fig. 2).

Since DLS logic operates using leakage currents, it is sensitive to threshold voltage shifts due to process variations. Fig. 3 shows how transistor sizing for optimal robustness is performed for a DLS 2-input NAND gate. In simulation, the transistor threshold voltages are skewed by an increasing value δ until the gate fails to transition. Fig. 3 shows the value of δ (normalized to σ_{VTH} for the target CMOS process) found for different sizing combinations of M_{NT} and M_{PB} . As mentioned, the critical point in the transition occurs when transistors M_{NT} and M_{PB} flip between cutoff and super-cutoff, while the internal transistors controlled by A and B are in weak-inversion. Hence, the robustness of the gate principally depends on the sizing of M_{NT} and M_{PB} and all internal transistors can be set to minimum size to reduce input loading. Increasing W_{PB} and decreasing W_{NT} improves the OUT falling transition while degrading its rising transition. The $(\delta_{HL} \times \delta_{LH})^{0.5}$ plot seeks to co-optimize robustness in both transitions and the selected design point exceeds $\delta = 4\sigma_{VTH}$ to ensure robust switching.

The Cortex M0+ processor is synthesized with a DLS standard-cell library containing an inverter, 2-input NAND and NOR, and a D flip-flop with asynchronous reset. The standard cells are verified with 500k Monte-Carlo simulation and Fig. 4 presents the key timing, leakage, and robustness results. A 32-bit RISC ARM Cortex M0+ processor is implemented using the DLS library, including an on-chip clock generator, address decoder and128B latch-based (for low VDD robustness) instruction and data memories [8]. All logic is implemented using a standard-cell approach with fully automatic place & route. The latch-based memory has DLS logic-based read-in/out path and a negative-edge write scheme to ensure timing violation-free write operation. The measured waveforms in Fig 4 show the execution of a simple toggle program on the M0+ processor.

As expected, leakage power is dominant at the operating frequency (~15Hz) of this system (Fig. 5). The lowest functional operating voltage is 0.16V. As the supply voltage increases from 0.16V (with a fixed frequency), the power consumption decreases exponentially due to the stronger super-cutoff effect. The minimum power consumption is 295pW, occurring at 0.55V. At higher VDD, DIBL effects and p-n junction diode leakage cancel out the increased super-cutoff effect. Dynamic power increases quadratically with VDD as expected, reaching 13.5% percent of the total power at VDD = 1.15V.

The system is fully functional across -5 to 65° C. Due to the high temperature sensitivity of subthreshold current, total power increases exponentially from 50pW at -5° C to 4.4nW at 65° C. Due to the relatively constant operating frequency, energy per operation follows the power consumption curve with a minimum of 44.7pJ/instruction at the minimum power point, VDD = 0.55V. Fig. 5 gives the power distribution across 28 different dies, with sigma/mean of 6.35%. To highlight its extremely low power consumption, the core was operated when powered directly by a 0.09mm² bulk Silicon solar cell. Powered only by this solar cell, the processor operates at 12Hz at 0.32V, consuming 970pW during program execution. The minimum light intensity required for the solar cell to successfully power the processor is 240 lux, which is equivalent to dim indoor light. Fig. 6 compares the proposed low power processor to prior work in ultralow power digital systems, showing an 80× improvement in active power per gate. Fig. 7 includes the test chip die photo.

References

[1] A. Wang and A. Chandrakasan, "A 180mV FFT processor using subthreshold circuit techniques," ISSCC, 2004.

[2] S. Hanson, et al., "A Low-Voltage Processor for Sensing Applications With Picowatt Standby Mode," JSSC, Apr. 2009.

[3] Y. Lee, et al., "A Modular 1mm³ Die-Stacked Sensing Platform with Optical Communication and Multi-Modal Energy Harvesting," ISSCC, 2012.

[4] Cymbet Corp., "EnerChip™ Bare Die Batteries Data Sheet," 2014.

[5] R. Hahn, et al., "Development of near hermetic silicon/glass cavities for packaging of integrated lithium micro batteries," MEMS/MOEMS, Apr. 2009.

[6] N. Lotze and Y. Manoli, "A 62mV 0.13µm CMOS standard-cell-based design technique using Schmitt-trigger logic," ISSCC, 2011.

[7] W. Jung., et al., "A 3nW fully integrated energy harvester based on selfoscillating switched-capacitor DC-DC converter," ISSCC, 2014.

[8] D. Jeon, et al., "A Super-Pipelined Energy Efficient Subthreshold 240 MS/s FFT Core in 65 nm CMOS" IEEE JSSC, vol. 47, Jan. 2012.











Figure 8.2.5: Measurement results of battery-less sub-nW Cortex M0+ processor.



Figure 8.2.2: Comparison of DLS inverter with standard CMOS inverter.



Figure 8.2.4: System block diagram and measured waveforms of the system running a toggle program.

	This work	[1]	[2]	[3]	[6]
Technology	0.18µm	0.18µm	0.13µm	0.18µm	0.13µm
Architecture	ARM Cortex M0+	16-bit FFT processor	8-bit processor	ARM Cortex M0	8-bit multiplier
Operating Voltage	160mV ~ 1.15V	180mV ~ 1.8V	450mV ~ 900mV	600mV	62mV ~ 1.2V
Operating Frequency	2Hz ~ 15Hz	164Hz ~ 5.5MHz	40kHz ~ 3MHz	160kHz ~ 330kHz	5kHz ~ 650kHz
Minimum Operating Power	295pW @550mV, 2Hz	90nW @180mV,164Hz	100nW @450mV, 40kHz	3.3μW @600mV, 1kHz [†]	17.9nW @62mV, 5.2kHz
Gate Count	40,800 ^{††}	156,750 ^{††}	23,120	35,240	775
Active Power Per Gate	7.23fW @550mV	574fW @180mV	4.32pW @450mV	93.6pW @600mV	2.31pW @62mV
Minimum Energy Per Operation	44.7pJ/inst @550mV	151nJ/FFT @350mV	2.8pJ/inst @350mV	17.2pJ/inst @260mV	1.31pJ/inst @260mV
Area	2.04mm ²	5.46mm ²	0.90mm ²	1.7mm ²	0.04mm ²

[†]Operating frequency at minimum reported power of the system.

 $^{\rm th}{\rm Gate}$ Count is obtained by dividing total number of transistors by 4 for CMOS and 6 for DSL, based on the number of transistors of 2-input NAND gate.

Figure 8.2.6: Performance summary and comparison.



Figure 8.2.7: Die photograph of 180µm test chip.