

Design Considerations for Ultra-Low Energy Wireless Microsensor Nodes

Benton H. Calhoun, *Student Member, IEEE*, Denis C. Daly, *Student Member, IEEE*,
Naveen Verma, *Student Member, IEEE*, Daniel F. Finchelstein,
David D. Wentzloff, *Student Member, IEEE*, Alice Wang, *Member, IEEE*,
Seong-Hwan Cho, *Member, IEEE*, and Anantha P. Chandrakasan, *Fellow, IEEE*

Abstract—This tutorial paper examines architectural and circuit design techniques for a microsensor node operating at power levels low enough to enable the use of an energy harvesting source. These requirements place demands on all levels of the design. We propose an architecture for achieving the required ultra-low energy operation and discuss the circuit techniques necessary to implement the system. Dedicated hardware implementations improve the efficiency for specific functionality, and modular partitioning permits fine-grained optimization and power-gating. We describe modeling and operating at the minimum energy point in the subthreshold region for digital circuits. We also examine approaches for improving the energy efficiency of analog components like the transmitter and the ADC. A microsensor node using the techniques we describe can function in an energy-harvesting scenario.

Index Terms—Integrated circuits, energy-aware systems, low-power design, wireless sensor networks.

1 INTRODUCTION

A wireless microsensor network consists of tens to thousands of distributed nodes that sense and process data and relay it to the end-user. Applications for wireless sensor networks range from military target tracking to industrial monitoring and home environmental control. The distributed nature of microsensor networks places an energy constraint on the sensor nodes. Typically, this constraint is imposed by the capacity of the node's battery. For this reason, most microsensor networks duty cycle, or shutdown unused components whenever possible. In this paper, duty cycling refers generically to alternating between an active mode and a low-power sleep mode. Although duty cycling helps to extend sensor network lifetimes, it does not remove the energy constraint placed by the battery. For some applications, a limited lifetime is sufficient and battery power is the logical choice. A 1cm³ Lithium battery can continuously supply 10 μ W of power for five years [1]. This tutorial focuses on applications demanding higher peak power or longer lifetime in an environment where changing batteries is impractical or impossible, therefore requiring a renewable energy source.

Research into energy scavenging suggests that micro-sensors can utilize energy harvested from the environment. Energy harvesting schemes convert ambient energy into electrical energy, which is stored and utilized by the node. The most familiar sources of ambient energy include solar power, thermal gradients, radio-frequency (RF), and mechanical vibration. Table 1 gives a comparison of some energy harvesting technologies. Power per area is reported because the thickness of these devices is typically dominated by the other two dimensions. The power available from these sources is highly dependent on the node's environment at any given time. However, these examples show that it is reasonable to expect 10s of microwatts of power to be harvested from ambient energy. Barring significant advances in energy scavenging technology, the high instantaneous power consumption of an active wireless transceiver (milliwatts for Mbps) requires micro-sensors to retain local energy storage. Coupling energy-harvesting techniques with some form of energy storage can theoretically extend microsensor node lifetimes indefinitely.

Using a rechargeable energy reserve with energy-harvesting implies several constraints for improving node efficiency. First, the standby *power* of the node must be less than the average power supplied by the energy-harvesting mechanism. If this is not the case, then energy-harvesting cannot recharge the battery and the nodes will expire. Second, the node should use as little energy as possible during active operation. Minimizing energy per operation allows decreased energy storage capacity (size, weight, cost) and/or a higher duty cycle (better performance). Third, the node should transition gracefully to and from standby mode with very little time or energy overhead, increasing the efficiency of duty cycling for extremely short periods of time in the active mode. Last, a microsensor node should be *power-aware*. A power-aware node scales performance of individual blocks

- B.H. Calhoun, D.C. Daly, N. Verma, D.F. Finchelstein, D.D. Wentzloff, and A.P. Chandrakasan are with the Massachusetts Institute of Technology, 50 Vassar St., 38-107, Cambridge, MA 02139. E-mail: {bcalhoun, ddaly, nverma, dfinchel, ddw, anantha}@mitl.mit.edu.
- A. Wang is with Texas Instruments, 12500 TI Boulevard, MS 8723, Dallas, TX 75243. E-mail: aliwang@ti.com.
- S.-H. Cho is with the Department of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology (KAIST), 373-1 Guseong-dong, Yuseong-gu, Daejeon 305-701, Korea. E-mail: chosta@ee.kaist.ac.kr.

Manuscript received 14 Mar. 2004; revised 30 Sept. 2004; accepted 8 Oct. 2004; published online 15 Apr. 2005.

For information on obtaining reprints of this article, please send e-mail to: tc@computer.org, and reference IEEECS Log Number TCSI-0090-0304.

TABLE 1
Example Power Densities of Energy Harvesting Mechanisms

Technology	Power Density [$\mu\text{W}/\text{cm}^2$]
Vibration - electromagnetic [2]	4.0
Vibration - piezoelectric [1]	500
Vibration - electrostatic [3]	3.8
Thermoelectric (5°C difference) [4]	60
Solar - direct sunlight [5]	3700
Solar - indoor [5]	3.2

gracefully to accommodate power requirements, resulting in appropriately varying energy consumption. Time-varying fluctuations in the environment make this feature necessary for functionality in an energy-harvesting system. Additionally, power awareness contributes to meeting the second constraint of reducing active energy per operation.

The key challenge of next generation nodes is meeting these requirements through aggressive optimization in all layers of design. First generation microsensor nodes were built using mostly commercial parts that were limited in their power-aware capabilities [6], [7]. Section 2 describes a new architecture for a general wireless microsensor node and compares it to the off-the-shelf approach. Section 3 presents approaches for meeting the design constraints in digital circuits. Section 4 evaluates the radio transmitter for microsensor nodes. Section 5 describes the sensor front end including the ADC, and Section 6 contains a summary and conclusions.

2 MICROSENSOR NODE ARCHITECTURE

Over the last decade, several academic and industrial research groups have been actively designing wireless microsensor nodes. The $\mu\text{AMPS-1}$ sensor node, a representative node example, provides a hardware platform for distributed microsensor networks using commercial, off-the-shelf (COTS) components. The sensor node processor uses dynamic voltage scaling (DVS) to minimize energy consumption for a given performance requirement. The radio transmit power adjusts to one of six levels, depending on the physical location of the target nodes. Power consumption of the node varies from 3.5mW in the deepest sleep state up to almost 2W (1.1W of which goes into the transmitter power amplifier) with the processor running at the fastest clock rate and the radio transmitting at the highest power level. Fig. 1 shows the instantaneous power consumption of a $\mu\text{AMPS-1}$ node as it collects data samples from the microphone, performs a line-of-bearing (LOB) calculation on the collected data, and relays the results of this calculation to other nearby nodes. Using generic components makes the power too high for the constraints we have described, so a customized architecture is necessary.

The energy savings of a custom approach come from modularizing the sensor node by considering common tasks for sensor network applications. Key tasks which can be implemented in hardware include the fast Fourier transform (FFT), finite impulse response (FIR) filters,




Fig. 1. Power-scaling measurements on the $\mu\text{AMPS-1}$ node (courtesy of N. Ickes).

encryption, source coding, channel coding/decoding, and interfaces for the radio and sensor. In order to achieve energy efficiency throughout the entire system, the hardware modules can use independent voltage supplies and operate at different clock frequencies. The drawbacks of this architecture are the increase in system complexity and area, the need for additional data transfers between the DSP and specialized modules, and the difficulty of interoperability across different voltage and clock islands.

Fig. 2 shows our proposed architecture for an energy-efficient sensor node. The digital architecture contains a simple DSP that executes arbitrary programs. The DSP communicates with the specialized modules through a shared bus and the DMA schedules the transfer of data between modules and the bus. Data memory is accessible by both the specialized modules and the DSP.

Dynamic voltage scaling (DVS) can be used to trade energy for computational latency for each module. A module's supply voltage should be set to the lowest possible value that satisfies its speed requirements. However, there is a supply voltage below which computations become less energy efficient due to leakage currents [8]. When no computation is taking place, the supply voltage should be shut off from the CMOS logic to reduce leakage power. The analog modules

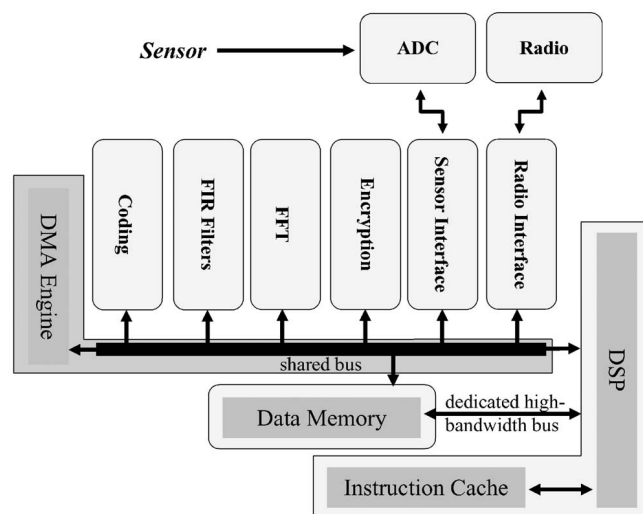


Fig. 2. Proposed architecture of an energy-efficient sensor node.

require the same dynamic performance controls as the digital modules. Both the sensor and radio must have an “always-on,” low-power standby mode that allows for basic threshold detection of a wake-up signal. For instance, an audio sensor might operate in a low-power mode until sound of a certain magnitude is detected.

2.1 Memory Hierarchy for Specialized Processing Modules

The proposed architecture requires overhead energy to move data along the bus between the memory and the processing modules. This energy can be reduced by dividing the memory hierarchy into a large store (the main memory) and smaller memories inside the modules. Traditional caches apply this principle to reduce access time, but we examine the impact on access energy. Generally, using a single memory means each memory access in the algorithm will be directed to a larger SRAM and will occur over a larger, more heavily loaded bus. This alternative also forces all of the processing modules to run at the same clock speed, which prevents the use of fine-grained DVS at the module level. Embedding a local memory in the module alleviates these issues, but has the overhead of transferring the inputs and outputs of the algorithm over the main bus between the large and small SRAMs. Having an additional local module memory will also add to the total standby current consumed by the node, although its power supply can be shut off if preserving the cache contents is not a requirement.

To quantify this overhead, we compared a single standalone 8k-word memory to an 8k-word memory coupled with a local memory consisting of either 1,024 or 128 words. A commercial memory compiler generated all of the memories for the experiment. The results of the experiment show that a local 1,024-word memory reduces total energy only when it performs more than 6,700 memory accesses. For fewer accesses, the energy overhead of transferring data into the local memory and then back to the main memory dominates the savings from accessing a smaller memory. The experiment showed that 700 memory accesses were necessary for a local 128-word cache to reduce the total energy. The FFT processor of [8] does $N \cdot \log(N)$ memory accesses during an N -bit FFT. Thus, a local data cache saves energy for the FFT processor because it accesses the memory more times than the break even scenarios we found in the above experiment.

In conclusion, specialized modules can reduce their total memory energy by using a local data cache for two reasons. First, the hardwired algorithms will access a much smaller SRAM, leading to a decrease in energy per access. Second, this cache memory can act as a synchronizer between the main memory and the specialized module, allowing independently scaled voltage and frequency for different modules in order to reduce overall circuit energy.

3 DIGITAL CIRCUIT TECHNIQUES

Digital circuit design for the microsensor space must focus primarily on the energy and power constraints we have presented, rather than solely on maximizing performance. The unpredictable environment of microsensor networks,

coupled with less stringent performance requirements, allows a trade-off of speed for reduced energy at both the architecture and circuit levels.

3.1 Energy-Aware Architectures

At the architectural level, designing for energy awareness can allow a sensor node to minimize energy consumption in the variable environment of a microsensor network. Energy-aware design is in contrast to low-power design, which targets the worst-case scenario and may not be globally optimal for systems with varying conditions. The energy-awareness of a system can be increased by adding additional hardware to cover functionality over many scenarios of interest and to tune the hardware such that the system is energy-efficient over a range of scenarios. Energy efficiency reduces the average energy per operation under varying performance requirements and, thus, relaxes the energy storage requirement for the microsensor node. This section provides an example of an energy-aware implementation of the widely used Fast Fourier Transform (FFT) algorithm.

The FFT calculates the frequency content of time-domain data. It is used in frequency domain beamforming, source tracking, harmonic line association, and classification. To achieve energy-awareness, the FFT implementation includes tunable structures, such as memory size and variable bit precision, to handle a variety of scenarios efficiently. This example design implements a real-valued FFT (RVFFT) that scales between 128-point and 1,024-point FFT lengths and operates at both 8 and 16-bit precision. The Baugh Wooley (BW) multiplier design provides an example of an energy-scalable bit precision datapath in Fig. 3. The RVFFT uses four BW multipliers to perform complex multiplication. When 16-bit multiplication is needed, the entire multiplier is used. However, if 8-bit multiplication is needed, only the MSB quadrant of adders is required. In this case, the 8-bit inputs feed directly to the MSB quadrant and the LSB inputs are gated to eliminate switching in the unused adders.

Variable FFT length is another hook designed into the FFT processor. The control logic to the FFT scales the number of butterflies with FFT length. The processor stops early to save energy for smaller FFT lengths. Also, the dedicated memory is designed to scale the memory size with FFT lengths. For example, 128-point FFT processing only requires a 128W×32b memory. Therefore, using a nonscalable memory designed for the 1,024-point FFT dissipates additional energy overhead for 128-point processing. A scalable memory that uses the correct memory size for 128-point processing and the entire memory for 1,024-point processing is shown in Fig. 4.

The energy scalable FFT architecture was simulated in a 0.18 μ m CMOS process at 1.5-V operation, and the simulated energy dissipated is shown in Table 2. The simulation results show a definite advantage for an energy-scalable architecture over a nonscalable architecture. The scalable architecture is more energy-efficient for all but the high quality point (1,024 point, 16-bit). At the high quality point, the scalable design sees a disadvantage due to the overhead logic. However, the scalable implementation uses 2.7 times lower energy at the low quality point (128 point, 8-bit). The scalable FFT processor was fabricated in a standard 0.18 μ m




Fig. 3. Eight-bit and 16-bit scalable Baugh Wooley multiplier architecture. The 8-bit multiplier is reused for the 16-bit multiplication, thereby adding scalability without a large area penalty.

CMOS process and standard ASIC flow to demonstrate these energy-scalable architectural techniques. At 1.5-V operation, when compared to a StrongARM SA-1100 implementation, the FFT processor shows over a 350X measured energy reduction. This result is evidence of the significant energy savings that can be achieved by using dedicated hardware modules.

Energy-scalable architectures are designed with many hooks that allow the processor to gracefully scale energy with quality and to achieve global energy-efficiency. These techniques enabled variable bit-precision and variable FFT lengths in an FFT processor and increased the energy-awareness of the system with minimal area and energy overhead.

3.2 Subthreshold Operation

When minimizing energy is the primary system requirement, the subthreshold region gives the minimum energy solution [9], [10] for most circuits. Subthreshold circuits use a supply voltage, V_{DD} , that is less than the threshold voltage, V_T , of the transistors. In this regime, subthreshold leakage currents charge and discharge load capacitances,

limiting performance but giving significant energy savings over nominal V_{DD} operation. Fig. 5 gives an example of subthreshold operation for a $0.18\mu\text{m}$ CMOS technology. The left-hand plot shows the measured frequency of a ring oscillator versus V_{DD} . Once V_{DD} drops into the subthreshold region, the on-current of the transistors becomes exponential with voltage and the I_{on}/I_{off} ratio reduces quickly. This causes the delay to increase exponentially. The right-hand plot shows an oscilloscope plot of an FIR filter operating at 150mV and 3.2kHz.

Fig. 6 gives measurements from a subthreshold FFT processor that shows how minimum energy operation does not necessarily occur at minimum voltage operation. The $0.18\mu\text{m}$ CMOS chip implements a 1,024-point, 16-bit FFT [8]. A new subthreshold design methodology using a modified standard logic cell library, custom multiplier, and memory generators was employed to implement the processor without additional process steps or body-biasing. The processor operates down to 180mV, where it runs at 164Hz and 90nW. The figure shows the minimum energy point for the 16b, 1,024-pt FFT processor at 350mV, where it dissipates 155nJ/FFT at a clock frequency of 10kHz. As V_{DD} decreases, the switching energy reduces quadratically. But, propagation delay increases exponentially in the subthreshold region, allowing leakage current to integrate longer for each operation. The resulting increase in leakage energy causes the minimum energy point. For 8-bit operation, the minimum energy point moves to higher V_{DD} . Since the minimum energy point clearly changes for different scenarios, we present a model for finding minimum energy operation in the subthreshold region.

3.2.1 Subthreshold Energy Modeling

In order to develop a model for subthreshold operation of arbitrary circuits, we first examine the subthreshold propagation delay of a characteristic inverter:

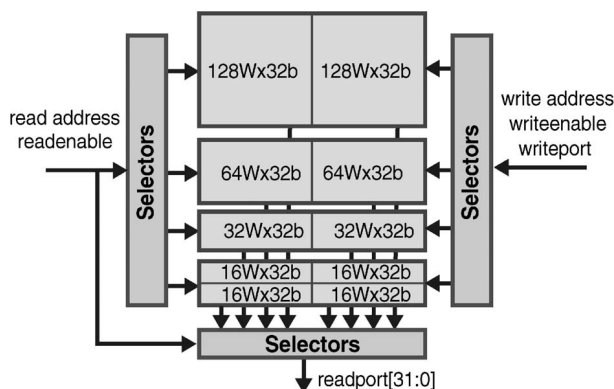


Fig. 4. Scalable FFT memory that enables variable memory size.

TABLE 2
Comparing a Nonscalable RVFFT to the Scalable RVFFT Method

		Non-scalable		Scalable	
Energy/FFT	FFT Length	8-bit	16-bit	8-bit	16-bit
	1024-point	1320nJ	1448nJ	575nJ	1491nJ
	512-point	607nJ	750nJ	240nJ	629nJ
	256-point	269nJ	334nJ	103nJ	269nJ
	128-point	118nJ	147nJ	44nJ	116nJ

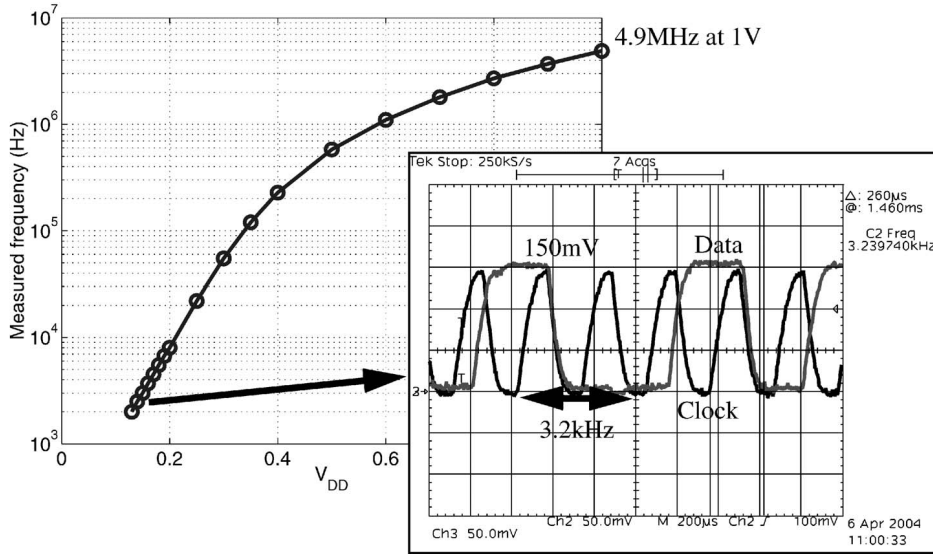


Fig. 5. Measured frequency versus V_{DD} and 150mV operation for 0.18 μm test chip.

$$t_d = \frac{KC_g V_{DD}}{I_{o,g} e^{-\frac{V_{GS}-V_{T,g}}{nV_{th}}}} \quad (1)$$

The switched capacitance of the inverter is C_g , K is a fitting parameter, and n is the subthreshold slope factor. Since the current in (1) accounts for transitions through both NMOS and PMOS devices, the terms $I_{o,g}$ and $V_{T,g}$ are fitted

parameters that do not correspond exactly with the MOSFET parameters of the same name. Operational frequency is simply:

$$f = \frac{1}{t_d L_{DP}}, \quad (2)$$

where L_{DP} is the depth of the critical path in characteristic inverter delays. Dynamic energy (E_{DYN}), leakage energy (E_L), and total energy (E_T) per cycle are expressed in (3), (4), and (5), assuming rail-to-rail swing ($V_{GS} = V_{DD}$ for "on" current).

$$E_{DYN} = C_{eff} V_{DD}^2, \quad (3)$$

$$E_L = W_{eff} I_{o,g} e^{-\frac{V_{T,g}}{nV_{th}}} V_{DD} t_d L_{DP} = W_{eff} K C_g L_{DP} V_{DD}^2 e^{-\frac{V_{DD}}{nV_{th}}}, \quad (4)$$

$$E_T = E_{DYN} + E_L = V_{DD}^2 \left(C_{eff} + W_{eff} K C_g L_{DP} e^{-\frac{V_{DD}}{nV_{th}}} \right). \quad (5)$$

Equations (3), (4), and (5) extend the expressions for current and delay of an inverter to arbitrary larger circuits, sacrificing accuracy for simplicity since the fitted parameters cannot account for all of the details of every circuit. Thus, C_{eff} is the average total switched capacitance of the entire circuit, including the average activity factor over all of its nodes. Likewise, W_{eff} estimates the average total width, normalized to the characteristic inverter, that

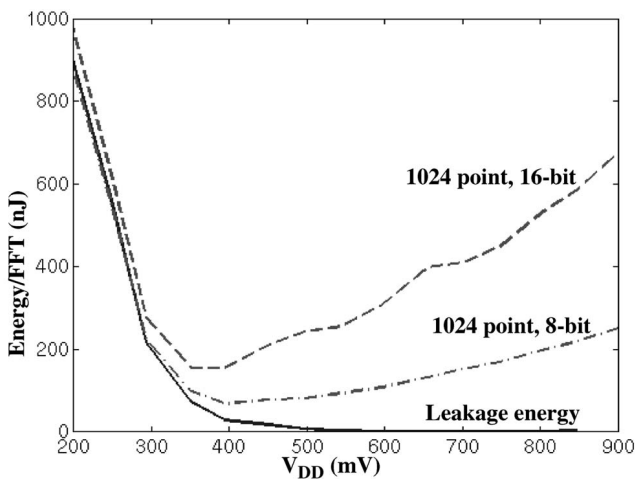


Fig. 6. Measured energy per operation for the FFT processor versus V_{DD} .

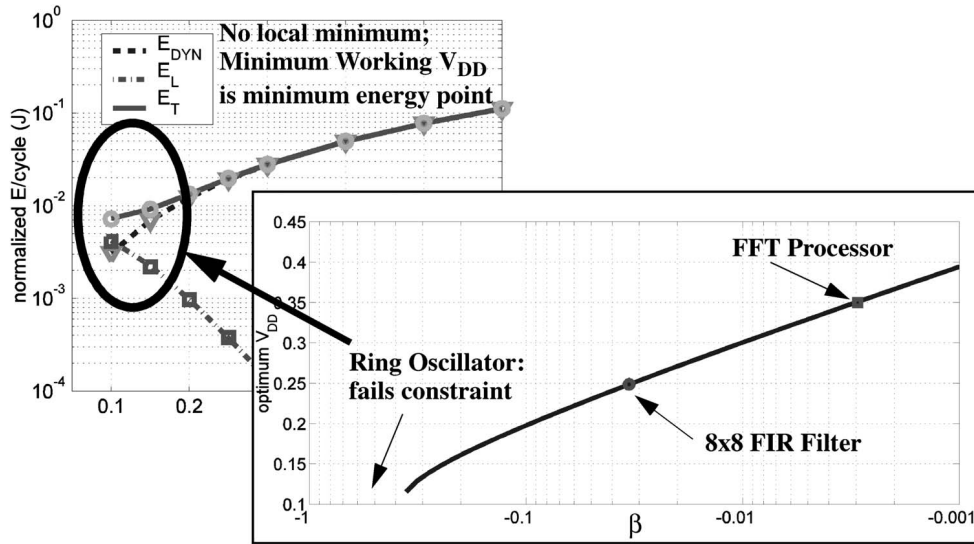


Fig. 7. V_{DD} optimum calculated with (7). β for ring oscillator ($L_{DP} = 21$) fails constraint. β for an 8×8 parallel FIR filter and for the FFT processor in [11] are also shown.

contributes to leakage current. Treating this parameter as a constant ignores the state dependence of leakage. Solving this set of equations provides a good estimate of the optimum for the average case and shows how the optimum point depends on the major parameters. Differentiating (5) and equating to 0 allows us to solve for V_{DDopt} :

$$\begin{aligned} \frac{\partial E_T}{\partial V_{DD}} &= 2C_{eff}V_{DD} + 2KC_gL_DW_{eff}V_{DD}e^{-\frac{V_{DD}}{nV_{th}}} \\ &+ \frac{-KC_gL_DW_{eff}V_{DD}^2e^{-\frac{V_{DD}}{nV_{th}}}}{nV_{th}} = 0. \end{aligned} \quad (6)$$

The following equation gives an analytical solution for V_{DDopt} :

$$V_{DDopt} = nV_{th}(2 - \text{lambert}W(\beta)), \quad (7)$$

$$\beta = \frac{-2C_{eff}}{W_{eff}KC_gL_{DP}}e^2 > -e^{-1}. \quad (8)$$

Equations (7) and (2) give the optimum supply voltage and frequency for subthreshold circuits consuming the minimum energy for a given V_T [11]. Equation (7) shows that the optimum V_{DD} value is independent of V_T . Instead, it is set by the relative significance of dynamic and leakage energy components as expressed in (8). The Lambert W Function, $W = \text{lambert}W(x)$, gives the solution to the equation $We^W = x$, just as $W = \ln x$ is the solution to $e^W = x$, and (8) gives the constraint for the solution. For a short ring oscillator with activity factor of one, C_{eff} and W_{eff} equal one and L_{DP} is small, so β does not meet the constraint in (8). Mathematically, this means that the derivative of E_T never equals zero, as shown in the left-hand plot in Fig. 7. The markers on this plot represent simulation data and the lines show the model. Physically, the leakage component for the high activity ring oscillator remains insignificant compared to dynamic energy over all supply voltages. The true optimum V_{DD} in this case is the lowest voltage for which the circuit functions. Fig. 7 also shows the β values for the FFT processor in [8] and for an

8-bit, 8-tap FIR filter. Circuits with relatively more leakage energy than dynamic energy have less negative β and, thus, higher optimum V_{DD} .

3.2.2 Subthreshold Optimization

As a case study, we examine the parallel 8-bit, 8-tap FIR filter introduced in Fig. 7. The simulations of the filter use netlists extracted from synthesized layout in a $0.18\mu\text{m}$ CMOS process. The synthesis flow incorporates a standard cell library that was modified to enable operation down to 100mV at the typical corner. Calibrating the model requires three parameters. First, C_{eff} was determined by measuring average supply current for an extended NanoSim simulation and solving $C_{eff} = I_{avg}/(fV_{DD})$. Simulating the total delay on the circuit's critical path provides the logic depth, L_{DP} , relative to the inverter delay. Last, W_{eff} is determined by simulating the circuit's leakage current in steady state and normalizing to the characteristic inverter. Since W_{eff} is a function of circuit state, averaging the circuit leakage current for simulations over many states improves the total leakage estimate. The analytical solution for optimum V_{DD} matches the simulated value with less than 0.1 percent error. However, this solution can change based on operating parameters. Any increase in leakage energy relative to active energy pushes the optimum V_{DD} and frequency higher. Likewise, any decrease in E_{LEAK} or increase in E_{DYN} will lower the optimum V_{DD} .

Designing subthreshold circuits has several limitations. First, subthreshold circuits show high sensitivity to process variations because of the exponential dependence of current on V_T . Fine-grained threshold voltage control can counteract the effects of process variation using adaptive back biasing or leakage-controlled feedback circuits. Second, the minimum energy point is sensitive to variables such as activity factor, temperature, and duty cycle [11]. This sensitivity makes careful analysis of a circuit's operating environment important for selecting the optimum V_{DD} and it suggests the benefits of closed loop control of V_{DD} . The

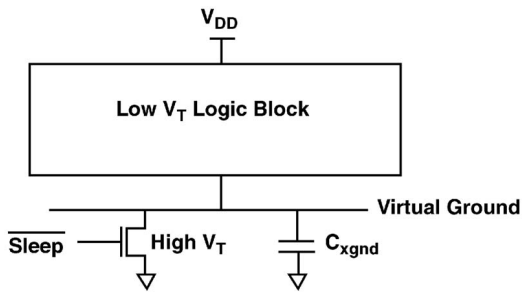


Fig. 8. MTCMOS power-gating circuits for standby power reduction.

low V_{DD} also makes the circuits more prone to soft errors, although redundancy techniques can compensate for soft errors in critical circuits (e.g., [12]).

Although some design challenges remain for subthreshold circuits, they provide significant reductions in energy, making them ideal for microsensor nodes. The FFT and FIR circuits we have discussed show energy savings at the minimum energy point of approximately 20X and 36X, respectively, relative to operation at a V_{DD} of 1.8V.

3.3 Standby Power Reduction

In the energy-aware FFT architecture described earlier, signals are gated to improve energy efficiency. This technique reduces active power dissipation, but leakage power is not affected. As nanometer CMOS processes are leveraged to improve performance and energy-efficiency, leakage mitigation becomes an increasingly important design consideration. Deep submicron processes have increased subthreshold leakage, gate leakage, gate-induced drain leakage, and reverse biased diode leakage [13]. The literature contains many techniques for standby power reduction. Two promising approaches for microsensor nodes are multithreshold CMOS (MTCMOS) and standby voltage scaling.

Fig. 8 shows how MTCMOS circuits reduce standby leakage power by severing a circuit from the power rails with high V_T sleep devices [14]. Sizing the sleep transistor has received a lot of attention since oversizing limits the leakage savings while undersizing restricts performance [15], [16]. Likewise, designing sequential MTCMOS circuits takes special care to reduce leakage during sleep without losing state [17], [18]. Most MTCMOS designs use large sleep devices at the block level, but local sleep devices allow circuit partitioning into local sleep regions. Any unused circuit regions can enter sleep mode while surrounding circuits remain active. This approach only provides savings if all leakage currents are prevented during sleep mode. A careful design methodology can prevent subtle leakage paths from occurring at the MTCMOS interface to active circuits. A fabricated 0.13- μm , dual V_T CMOS test chip shows a low power FPGA architecture with over 8X measured standby current reduction [19]. The local sleep regions reduce active chip leakage by up to 2.2X for some configurations. The test chip uses sequential elements that allow power gating without the loss of data.

A second promising approach to standby power reduction is standby voltage scaling. This involves lowering the voltage supply to the circuits in standby to reduce power.

The voltage component of power decreases linearly and the current also decreases due to DIBL. Since the DIBL effect is more pronounced in nanometer technologies, the potential savings from this approach will improve with scaling. Clearly, the voltage cannot be reduced beyond the point where critical data in the memories and sequential logic is lost. The limit to voltage scaling was theorized to be 3-4 times the thermal voltage for early technologies [20] because the magnitude of the gain of the logic gates drops below one, eliminating bistable operating points. Fig. 9 shows the simulated data retention capability of a flip-flop in a 90nm CMOS technology. For each DC sweep, the input to the flip-flop is the worst-case value that encourages node N1 to pull away from its correct voltage. The simulations show that the flip-flop will lose its data below about 150mV. Measurements of the flip-flop from a 90nm test chip show that the flip-flop retains its state to 110mV. One approach to guard against losing data in sequential elements during standby mode is to use canary flip-flops. These flip-flops are sized to lose their data at higher supply voltages than the core flip-flops. Monitoring the failure pattern of these flip-flops gives an estimate of how close the core is to failing [21] and allows scaling closer to the failure point than an open-loop approach. Fig. 10 shows measured standby power savings for a 32-bit Kogge-Stone adder block on a 90nm CMOS test chip. Clearly, a significant reduction in power is possible for scaling anywhere near the failure point. Standby power scaling can be used in conjunction with MTCMOS circuits that use sequential elements that hold their state even during power gating, producing power savings of orders of magnitude.

Since memories have low activity factor and must retain state during long idle periods, it is important that they have low standby power. Standby voltage scaling can be applied to SRAMs. In [22], the $3\text{-}\sigma$ data retention voltage of a 0.13 μm CMOS SRAM was measured to be 250mV. Other techniques for designing low standby power memories are available. The lowest reported standby current per cell is 16.7fA in [23]. Based on these references, it is conceivable to predict that a 4MB SRAM using low leakage techniques (16.7fA/cell, assuming no DIBL for worst-case) combined with voltage scaling (to 250mV) can sustain standby power of only 0.13 μW . Combining standby voltage scaling with other leakage reduction techniques can bring the standby power of digital circuits well within the stringent requirements for energy-harvesting microsensors.

4 RADIO SUBSYSTEM

A microsensor radio shares the same key design constraints as the other circuit blocks, including low standby power consumption, fast switching into and out of standby, and energy efficient operation when active. However, since radios operate at significantly higher frequencies than the rest of the microsensor node and consume milliwatts of power when on, they have their own specific constraints and limitations. We present the design of an energy efficient wireless transmitter as an example of how to optimize a microsensor radio.

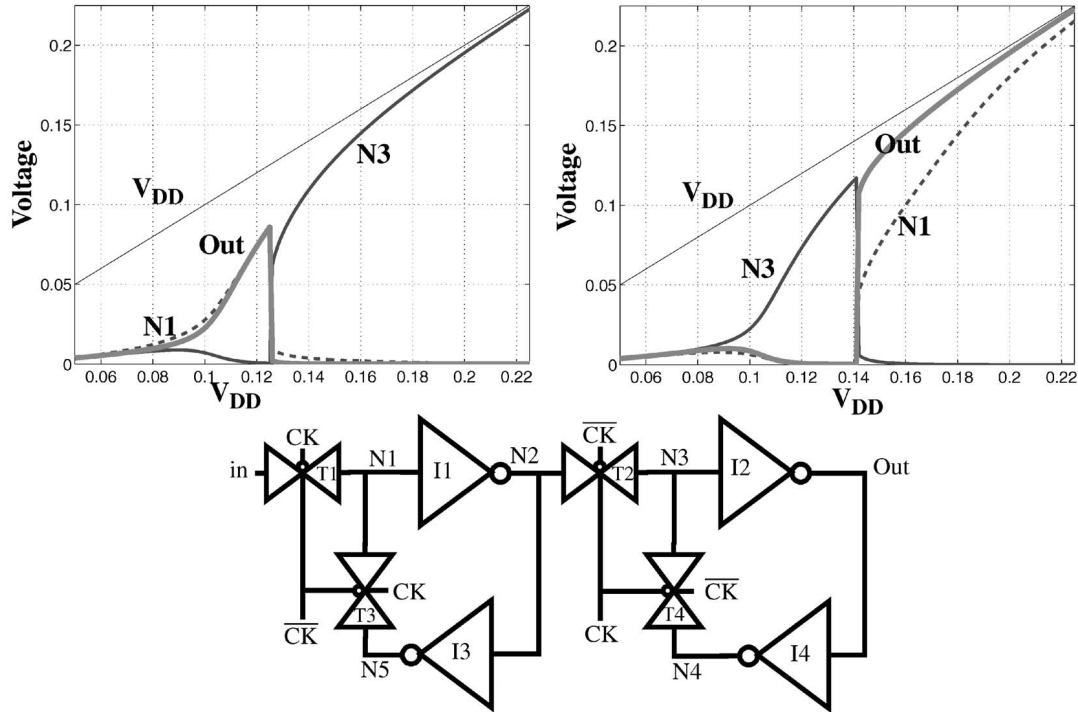


Fig. 9. Simulated data retention for standby power reduction.

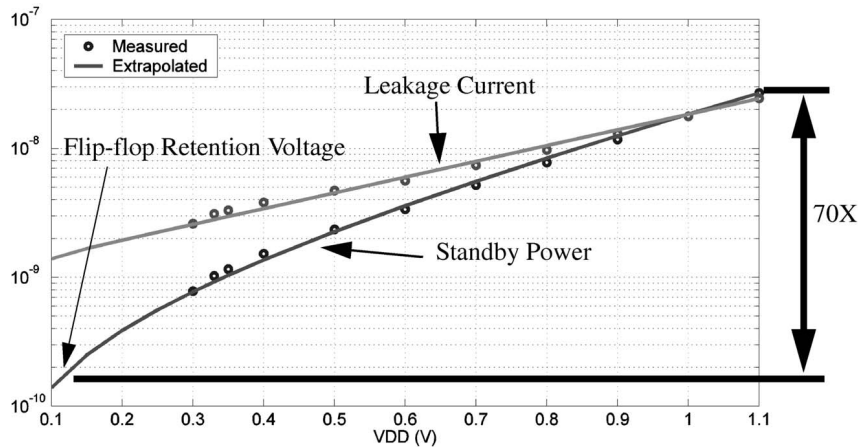


Fig. 10. Measured power savings on a 90nm test chip.

4.1 Design Considerations for an Energy Efficient Transmitter

An energy model for data transmission supplies the basis for determining a suitable transmitter architecture for sensor networks and for examining the trade-offs associated with duty cycling. The following equation expresses the energy consumed to transmit one packet:

$$E_{tx} = \frac{(P_{tx} + P_{out})L}{R} + P_{tx}T_{start}. \quad (9)$$

P_{tx} is the power consumption of the modulator, P_{out} is the power consumption of the output power amplifier, L is the size of the transmitted packet (bits), R is the data rate (bps), and T_{start} is the start-up time, which is the time it takes for the transmitter to wake up from the sleep state to the active state. For short range transmission at GHz

frequencies, the modulator components (frequency synthesizers, mixers, etc.), rather than the power amplifier, dominate power consumption ($P_{tx} \gg P_{out}$). Hence, for short packet sizes ($T_{start} > L/R$), the start-up energy, $P_{tx}T_{start}$, significantly increases the overall transmission energy. Fig. 11 illustrates the effect of start-up time on energy efficiency by plotting the energy to transmit a bit ($E_{bit} = E_{tx}/L$) versus packet size. The solid line presents a COTS radio, while the dotted line represents an ideal radio, a lower bound that consists of only the power amplifier at 100 percent efficiency ($E_{bit} = P_{out}/R$). Lowering the modulator power consumption (P_{tx}), increasing the data rate (R), and increasing the efficiency of the power amplifier all reduce the energy per bit for large packet sizes. However, the inefficiency introduced for short packet sizes can only be improved by reducing the start-up time. Therefore, implementing an energy efficient




Fig. 11. Impact of start-up time on transmitter’s energy consumption (0 dBm output power at 1 Mb/sec).

transmitter for a microsensor implies designing a *high data rate, low power, and fast start-up* transmitter.

4.2 Implementation of an Energy Efficient Transmitter

Energy efficient transmitters for short range communication primarily use continuous phase modulation schemes to reduce power consumption by removing the need for analog mixers and digital-to-analog converters (DACs). This approach is in contrast to traditional homodyne or heterodyne transmitters that mix a baseband or low-frequency signal with a local oscillator to generate the required RF waveform.

A continuous phase modulation architecture typically involves modulating a phase-locked loop’s (PLL) voltage-controlled oscillator (VCO) through direct modulation or modulating the PLL’s frequency divider ratio through indirect modulation. One drawback to indirect modulation is that high data rates cannot be easily achieved due to the low-pass filtering caused by the PLL. To overcome this obstacle, direct VCO modulation modulates data at the input of the VCO. These techniques allow a transmitter to consist only of digital circuitry, a PLL, a filter, and a power amplifier, thereby minimizing the power consumption of the modulator (P_{tx}).

The frequency synthesizer dominates the start-up time of a transmitter due to the inherent feedback loop in PLLs.

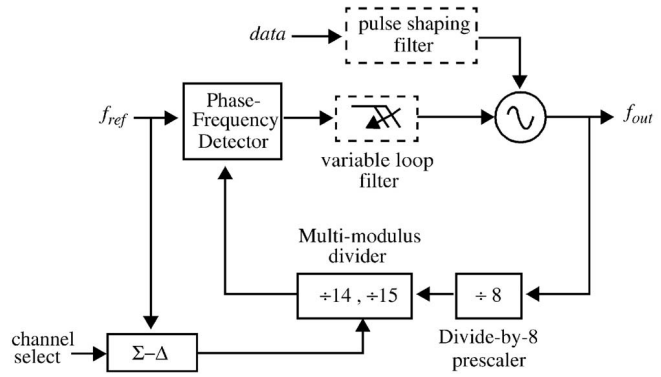


Fig. 12. Simplified block diagram of the energy efficient transmitter.

Minimizing the start-up time is critical for sensor nodes because of the typically short packets. A popular approach to increase the locking speed is the variable loop bandwidth method [24], where the PLL starts with a wide loop bandwidth and switches to a smaller loop bandwidth as the loop approaches lock. This method requires minimal overhead circuitry and, hence, is attractive for low-power PLL applications.

A proof of concept chip for variable loop bandwidth and closed loop direct VCO modulation has been implemented using 0.25 μm CMOS [25]. Fig. 12 shows a simplified block diagram of the proposed transmitter. The frequency synthesizer is a fourth-order PLL with a third-order $\Sigma\Delta$ modulator for fractional channel selection of the reference frequency. High data rate FSK modulation is achieved by directly modulating the VCO in a closed loop. The multi-stage variable loop bandwidth technique provides a fast start-up time. A high reference frequency (f_{ref}) maximizes the initial loop bandwidth in the variable loop bandwidth scheme. The prototype chip achieves a 20 μs start-up time with a data rate of 2.5 Mb/sec while consuming 22 mW at a carrier frequency of 6.5 GHz. The effect of the variable loop bandwidth technique on the start-up time appears in Fig. 13, where control voltage is plotted for a fixed loop bandwidth




Fig. 13. Measured start-up times of (a) fixed loop and (b) variable loop frequency synthesizers.

and variable loop bandwidth. The variable loop bandwidth method reduces the start-up time by about a factor of 4.

The transmitter described above has a startup energy of 440 nJ and the approximate transmission energy per bit is 10.8 nJ. The transmitter takes 1,520nJ to transmit a 100-bit packet. If these numbers are translated to average continuous power consumption assuming a 0dBm power amp at 20 percent efficiency, then the transmitter consumes 15 μ W for an average sensor data rate of 1kbps (i.e., 10 packets/sec, 1 packet = 100bits). These values compare favorably with other low power transmitters, especially for short packet sizes where startup energy dominates the total energy consumption [26], [27], [28].

4.3 Low Energy Communication Techniques

The energy required for radio communication scales with distance as d^n , where d is the distance and n is the path loss exponent, which typically ranges between 2 and 4. Communication energy may be reduced by dividing a long-distance transmission into several shorter ones. Intermediate nodes between a data source and destination serve as relays that receive and rebroadcast data. This concept, known as multihop communication [29], is analogous to using buffers over a long, on-chip interconnect.

In order to analyze how power trades off for multihop routing, the network is assumed to be capable of transmitting a distance d using an arbitrary number of hops h such that all individual hop distances equal d/h . The power consumed by multihop may then be modeled as

$$P(h, d) = h \left[\alpha + \beta \left(\frac{d}{h} \right)^n \right], \quad (10)$$

where α is the sum of the distance-independent components of communication power, such as the receiver, bias currents, and startup time, and β is the sum of the distance-dependent terms, such as the power amplifier and path loss.

Introducing relay nodes creates a balancing act between reduced $\beta(d/h)^n$ and increased α . Hops that are too short lead to excessive distance-independent overhead. Hops that are too long lead to excessive path loss. Between these extremes is an optimum transmission distance, called the characteristic distance, d_{char} [30]. This distance determines the optimal number of hops and depends only on the energy consumption of the hardware and the path loss coefficient.

An example radio is presented to demonstrate the trade-off between multihop routing and direct transmission. Performance characteristics were taken from the datasheet for a Zigbee compliant transceiver with an internal power amplifier (PA) capable of transmitting a maximum RF power of 0dBm. This power level sets an upper bound on the distance that the transceiver can transmit for a fixed channel model. To illustrate the effects of multihop routing, we assume that a hypothetical external PA can be added to the transceiver, with a gain of 20dB and fixed efficiency of 20 percent. This increases the maximum total RF power of the radio to +20dBm, extending the upper bound on transmission distance.

Fig. 14 is a plot of the total DC power, including multihop power overhead, required to transmit over a set

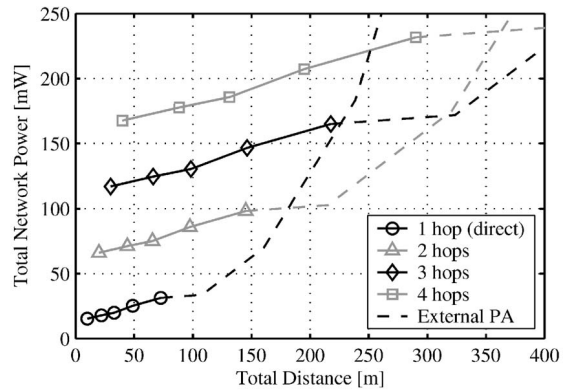


Fig. 14. Multihop curves showing crossover points.

total distance in one to four hops using the example radio described above. The transmit distance for each hop equals the total distance divided by the number of hops. Each h -hop curve has two sections, distinguished by the solid and dashed lines. The solid lines represent the power and achievable distance for the transceiver alone. The dashed lines represent the same for the transceiver with the external PA. The PA provides an increase in distance, but at the cost of exponentially increasing power. The actual transmit distances depend on a number of parameters. Here, a path loss exponent of 2.9 and receiver sensitivity of -94dBm are assumed.

The energy-optimal number of hops is indicated by the lowest curve in Fig. 14 for a given distance. The point at which the h -hop and $h+1$ -hop curves intersect is the characteristic distance described above. Beyond this point, the additional distance provided by higher power does not sufficiently compensate for the additional energy overhead. Before this point, the distance-independent power consumption in the transceiver negates the benefits of reducing the power of the PA and introducing additional hops.

A protocol-level technique such as multihop has disadvantages. The increased protocol complexity of multihop may result in additional energy consumption. Since each data packet is received and rebroadcast by relay nodes, the end-to-end delay increases and data throughput decreases. Moreover, reliance on multihop requires that relay nodes be in range of every transmitter, placing a minimum-density constraint on node placement. For applications with stringent latency and bandwidth requirements, or low or variable node density, multihop may not be an option. In addition, notice, in Fig. 14, that the portions of the curves pertaining to the transceiver alone with no external PA (solid lines) do not intersect. This implies that, over the range of RF power levels available from the transceiver, direct transmission is always more energy efficient than multihop routing. This result is common to most short-range radios in use today [30].

To minimize the need for multihop routing, next-generation power amplifiers should be designed for high efficiency over a wide range of operating points, which would increase the range of distances for which direct transmission is energy-efficient. A series of power amps highly optimized for a small range of outputs is one




Fig. 15. Block diagram of the sensor front-end.

potential solution, reflecting an “ensemble of point systems” that enable continuous, wide-ranging scalability [31].

5 ADC AND SENSOR SUBSYSTEM

In sensor nodes, where a low-power DSP performs application-level processing, a front-end analog-to-digital conversion system acquires data from the physical sensor. Since the ADC requirements are tightly coupled to a generally unpredictable environment, the ability to dynamically compromise features and performance in favor of power reduction is a valuable characteristic. In the limit, the ADC subsystem may act only as a threshold detector. This requires downstream data processing units to tolerate the compromises and to provide feedback to the ADC subsystem regarding the desired operating mode. Factors affecting that decision feedback might include characteristics of the sensing environment or the availability of harvested energy. This section examines a number of dimensions along which scaling could have a significant effect on overall power for the sensor front-end and ADC.

The design of a low power ADC subsystem requires consideration of the entire front-end, not just the ADC. Fig. 15 shows a very simple ADC subsystem. The components shown include a sensor, a low-noise preamplifier, an antialiasing filter, an ADC, and a DSP. Here, the DSP may be used for the application of ADC linearity calibration coefficients [32], offset/gain error cancellation, or digital decimation filtering.

To demonstrate the tight coupling between the performance modes of each component, consider the implementation of two 8-bit 10kS/sec ADC subsystems. In the first case, the data is oversampled ($F_S = 40\text{kS/s}$) by the ADC to ease filtering requirements (passive first order), while, in the second case, the ADC operates at near Nyquist rate (15kS/s) enabled by a high-order antialiasing filter (active eighth order). In this analysis, the dynamics of the sensor, preamplifier, and filter are all considered. The power figures used for the components are derived from currently available low power parts.¹ Table 3 shows that, for the given performance point, the oversampling subsystem consumes only about 54 percent the power of the near Nyquist one. Of course, this analysis is highly subject to the considered performance point and the available parts. Nonetheless, it is apparent that system optimizations depend on the trade-offs available at the component level.

In the case of low-event sensor nodes, optimizations in three critical states have been identified. These include

1. As used here, the power consumption of the physical sensor would be prohibitive to a self-powered node. Innovations in sensor technology or customized raw transducers are required.

TABLE 3
Comparison of Oversampling and Near Nyquist Sensing Front-Ends

Component	Oversampling Front-End Power [μW]	Near Nyquist Front-End Power [μW]
Sensor	40	40
Preamp	20	20
Filter	0	80
ADC	20	7.5
DSP	0.3	0
Total	80.3	148

minimization of standby power, minimization of energy overhead during activation (i.e., standby-to-operating state transition), and minimization of operating energy. Standby currents can typically be reduced to leakage levels in ADCs. Activation energy can also be quite low, particularly in certain ADC architectures. For instance, successive approximation register (SAR) converters require just one active component, the comparator, which can be turned on quickly [33] with only the CV^2 energy overhead of charging load capacitors. Reduced operating energy, which is equivalent to reduced operating power at a given sampling rate, can be achieved through two approaches. The first relies on improved circuit techniques and innovations leading to more efficient ADCs. The second relies on trading performance and features for power savings.

A figure of merit for ADCs has been suggested that considers power, sampling rate, and effective number of bits [34]:

$$FOM = \frac{P}{F_S \times 2^{ENOB}}. \quad (11)$$

Here, we see that power, P , is related directly to sampling rate, F_S . This is expected since power in both digital and analog circuits scales linearly with speed. The exponential power relationship associated with effective number of bits (ENOB) is a more empirical result and is addressed below.

Fig. 16 shows the energy consumed during a single conversion by two theoretical SAR ADCs. Here, 8-bit and 10-bit converters with approximately the same FOM and sampling rate (100kS/sec) are considered. In the 10-bit case, the conversion cycle is divided into 11 phases: one to sample the input and 10 to resolve each bit. Similarly, in the 8-bit case, the conversion cycle is divided into nine phases. As a result, the digital circuitry in the 10-bit converter must perform more operations and, therefore, consume more energy by a linear factor of 11/9. However, the analog circuitry, including the capacitor array DAC and comparator (which is composed of linear preamplifiers and a regenerative latch), exponentially increase in power consumption by approximately a factor of 4. As a result, the energy required to resolve bits in the 10-bit ADC is greater than the energy required to resolve bits in the 8-bit ADC.²

2. This factor does not consider an increase in unit capacitances required in order to improve matching characteristics of fabricated elements.

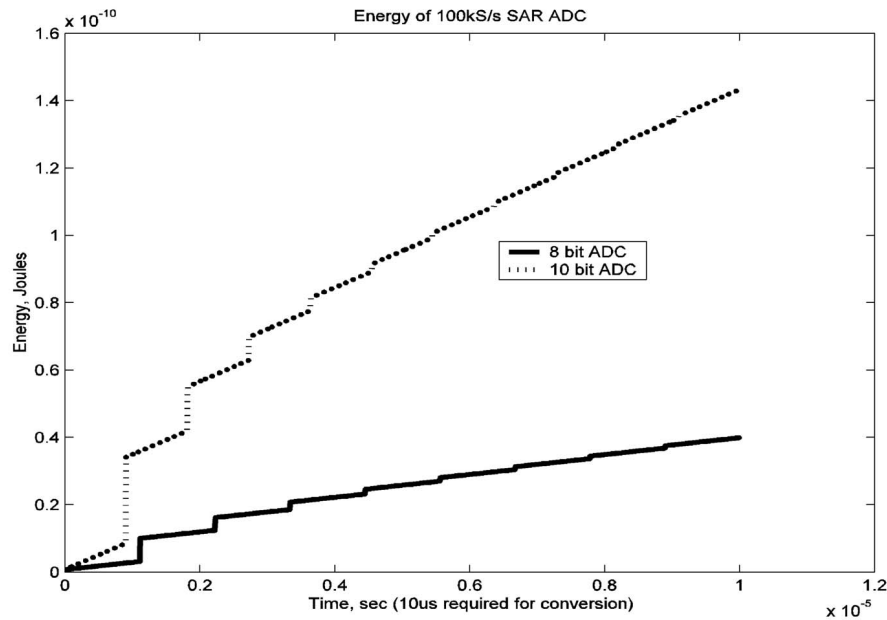


Fig. 16. Energy per conversion for two SAR ADCs.

As mentioned, scaling the ADC energy dynamically in response to reduced performance requirements lowers the total energy consumed by microsensor nodes. The FOM suggests that there are two main knobs for making a power trade-off: sampling rate and resolution. The ability to reduce standby energy and activation overhead to negligible levels implies that (dynamic) linear power reduction can be easily achieved with reduced sampling rate. Specifically, all blocks of the ADC (including clocking circuitry) can be turned off between data conversions.³

So far as resolution scaling is concerned, it is more difficult to achieve the exponential power savings suggested by the FOM. A critical limitation stems from the fact that, in digital circuitry, power scales linearly with resolution; for instance, a 16-bit adder is composed of twice as many full-adder cells as an 8-bit adder. Consequently, for the digital circuitry, higher resolution conversion is inherently more efficient with respect to the FOM.

Reducing the precision of analog circuits to recover exponential power savings is theoretically possible, but fraught with challenges. For instance, the precision of an amplifier is fundamentally set by its thermal noise floor. Since this is ultimately a function of load capacitances, which limit the noise spectral power, scaling precision requires variable capacitors of high quality (i.e., with small associated parasitics). Such devices are not readily available in low-cost digital processes. Of course, an alternative is to switch capacitor banks or entire amplifiers into and out of the signal path. Here, we face the limitations associated with analog switches that, in modern low-power design, require low voltage control signals. Nonetheless, highly scalable ADCs have been demonstrated and the associated design challenges are described in [35].

3. Some minimal circuitry to manage standby-to-active state transitions, such as a counter, might impose additional overhead.

6 CONCLUSION

This paper describes the challenges facing wireless microsensor design and presents a general microsensor node architecture. The challenge for next generation nodes is to further reduce energy consumption by optimizing energy-awareness over all levels of design. Subthreshold operation, power gating, and standby voltage scaling enable digital circuits to meet the low active energy and standby power requirements of microsensor nodes. Reducing startup time improves the energy efficiency of a transmitter for short packets and multihop routing reduces energy for long-distance communication. Since the ADC subsystem might be the front-end of a reactive sensor node, it is important to seek alternatives to full sleep modes. We analyzed the dimensions along which ADC performance might be compromised in order to recover power savings. Applying all of these techniques to a microsensor node makes energy-harvesting operation a possibility for microsensor networks.

ACKNOWLEDGMENTS

This effort sponsored by the US Defense Advanced Research Projects Agency (DARPA) Power Aware Computing/Communication Program and US Air Force Research Laboratory under agreement numbers F30602-00-2-0551 and F33615-02-2-4005 and through a subcontract with MIT Lincoln Labs, by Hewlett-Packard under the HP/MIT Alliance, and by Texas Instruments. Denis Daly and Daniel Finchelstein are supported by NSERC Postgraduate Scholarships.

REFERENCES

- [1] S. Roundy, P. Wright, and J. Rabaey, "A Study of Low Level Vibrations as a Power Source for Wireless Sensor Nodes," *Computer Comm.*, vol. 26, no. 11, pp. 1131-1144, July 2003.
- [2] H. Kulah and K. Najafi, "An Electromagnetic Micro Power Generator For Low-Frequency Environmental Vibrations," *Proc. 17th IEEE Int'l Conf. Micro Electro Mechanical Systems (MEMS)*, pp. 237-240, Jan. 2004.

- [3] S. Meninger et al., "Vibration-to-Electric Energy Conversion," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 9, no. 1, pp. 64-76, Feb. 2001.
- [4] H. Bottner et al., "New Thermoelectric Components Using Microsystem Technologies," *J. Micro Electro Mechanical Systems*, vol. 13, no. 3, pp. 414-420, June 2004.
- [5] "Panasonic Solar Cell Technical Handbook '98/99," Aug. 1998.
- [6] J.M. Kahn, R.H. Katz, and K.S. J. Pister, "Next Century Challenges: Mobile Networking for 'Smart Dust,'" *Proc. Mobicom 1999*, pp. 271-278, 1999.
- [7] J. Rabaey et al., "PicoRadio Supports Ad Hoc Ultra-Low Power Wireless Networking," *Computer*, pp. 42-48, July 2000.
- [8] A. Wang and A.P. Chandrakasan, "A 180 mV FFT Processor Using Subthreshold Circuit Techniques," *Proc. Int'l Solid-States Circuits Conf. (ISSCC)*, pp. 292-293, Feb. 2004.
- [9] A. Wang, A. Chandrakasan, and S. Kosonocky, "Optimal Supply and Threshold Scaling for Subthreshold CMOS Circuits," *Proc. Symp. VLSI*, pp. 5-9, 2002.
- [10] J. Burr and A. Peterson, "Ultra Low Power CMOS Technology," *Proc. Third NASA Symp. VLSI Design*, pp. 4.2.1-4.2.13, 1991.
- [11] B.H. Calhoun and A. Chandrakasan, "Characterizing and Modeling Minimum Energy Operation for Subthreshold Circuits," *Proc. Int'l Symp. Low Power Electronics and Design (ISLPED)*, Aug. 2004.
- [12] K. Osada et al., "SRAM Immunity to Cosmic-Ray-Induced Multierrors Based on Analysis of an Induced Parasitic Bipolar Effect," *IEEE J. Solid State Circuits*, vol. 39, no. 5, pp. 827-833, May 2004.
- [13] S. Borkar, "Design Challenges of Technology Scaling," *IEEE Micro*, vol. 19, no. 4, pp. 23-29, July/Aug. 1999.
- [14] S. Mutoh et al., "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," *IEEE J. Solid State Circuits*, vol. 30, no. 8, Aug. 1995.
- [15] J. Kao, S. Narendra, and A. Chandrakasan, "MTCMOS Hierarchical Sizing Based on Mutual Exclusive Discharge Patterns," *Proc. Design Automation Conf. (DAC)*, 1998.
- [16] M. Anis, S. Areibi, and M. Elmasry, "Design and Optimization of Multithreshold CMOS (MTCMOS) Circuits," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 22, no. 10, Oct. 2003.
- [17] S. Shigematsu, S. Mutoh, Y. Matsuya, Y. Tanabe, and J. Yamada, "A 1-V High-Speed MTCMOS Circuit Scheme for Power-Down Application Circuits," *IEEE J. Solid State Circuits*, vol. 32, no. 6, pp. 861-869, June 1997.
- [18] J. Kao and A. Chandrakasan, "MTCMOS Sequential Circuits," *Proc. European Solid State Circuit Conf. (ESSCIRC)*, 2001.
- [19] B. Calhoun, F. Honore, and A. Chandrakasan, "Design Methodology for Fine-Grained Leakage Control in MTCMOS," *Proc. IEEE Int'l Symp. Low Power Electronics and Design (ISLPED)*, 2003.
- [20] R. Swanson and J.D. Meindl, "Ion-Implanted Complimentary MOS Transistors in Low-Voltage Circuits," *IEEE J. Solid State Circuits*, vol. 7, no. 2, pp. 146-153, Apr. 1972.
- [21] B.H. Calhoun and A. Chandrakasan, "Standby Power Reduction Using Dynamic Voltage Scaling and Canary Flip-Flop Structures," *IEEE J. Solid State Circuits*, vol. 39, no. 9, Sept. 2004.
- [22] H. Qin, Y. Cao, D. Markovic, A. Vladimirescu, and J. Rabaey, "SRAM Leakage Suppression by Minimizing Standby Supply Voltage," *Proc. Int'l Symp. Quality Electronic Design (ISQED)*, 2004.
- [23] K. Osada, Y. Saitoh, E. Ibe, and K. Ishibashi, "16.7-fA/Cell Tunnel-Leakage-Suppressed 16-Mb SRAM for Handling Cosmic-Ray-Induced Multierrors," *IEEE J. Solid State Circuits*, vol. 38, no. 11, Nov. 2003.
- [24] F.M. Gardner, *Phase Lock Techniques*. New York: Wiley, 1979.
- [25] S.H. Cho and A.P. Chandrakasan, "6.5GHz CMOS FSK Modulator for Wireless Sensor Applications," *IEEE Symp. VLSI Circuits*, pp. 182-185, 2002.
- [26] M. Perrott, T. Tewksbury, and C.G. Sodini, "A 27 mW CMOS Fractional-N Synthesizer Using Digital Compensation for 2.5Mb/s GFSK Modulation," *IEEE J. Solid State Circuits*, vol. 32, no. 12, pp. 2048-2060, Dec. 1997.
- [27] D.A. Hitko, T.L. Tewksbury, and C.G. Sodini, "A 1 V, 5 mW, 1.8 GHz, Balanced Voltage-Controlled Oscillator with an Integrated Resonator," *Proc. IEEE Intl Symp. Low Power Electronic Design (ISLPED)*, pp. 46-51, Aug. 1997.
- [28] N. Filiol et al., "A 22 mW Bluetooth RF Transceiver with Direct RF Modulation and On-Chip IF Filtering," *Int'l Solid State Circuits Conf. Digest of Technical Papers*, pp. 202-203, 2001.

- [29] L. Kleinrock, "On Giant Stepping in Packet Radio Networks," Packet Radio Temporary Note #5 PRT 136, Univ. of California Los Angeles, Mar. 1975.
- [30] M. Bhardwaj, T. Garnett, and A. Chandrakasan, "Upper Bounds on the Lifetime of Sensor Networks," *Proc. IEEE Int'l Conf. Comm.*, pp. 785-790, 2001.
- [31] M. Bhardwaj, R. Min, and A. Chandrakasan, "Quantifying and Enhancing Power-Awareness of VLSI Systems," *IEEE Trans. VLSI Systems*, pp. 757-772, Dec. 2001.
- [32] H.-S. Lee, D.A. Hodges, and P.R. Gray, "A Self-Calibrating 15 Bit CMOS A/D Converter," *IEEE J. Solid State Circuits*, vol. 19, no. 6, pp. 813-819, Dec. 1984.
- [33] G. Promitzer, "12-Bit Low-Power Fully Differential Switched Capacitor Noncalibrating Successive Approximation ADC with 1MS/s," *IEEE J. Solid State Circuits*, vol. 36, no. 7, pp. 1138-1143, July 2001.
- [34] R.H. Walden, "Analog-to-Digital Converter Survey and Analysis," *IEEE J. Selected Areas in Comm.*, vol. 17, no. 4, pp. 539-550, Apr. 1999.
- [35] K. Gulati and H.-S. Lee, "A Low-Power Reconfigurable Analog-to-Digital Converter," *IEEE J. Solid State Circuits*, vol. 36, no. 12, pp. 1900-1911, Dec. 2001.



threshold circuit operation. He is a student member of the IEEE.



He is a student member of the IEEE.



Networks, Ottawa, Canada, on high-performance internet switching applications. From April 2001 to September 2001, he worked as an ASIC designer at AcceLight Networks, Ottawa, Canada, on optical switching products. Finally, from January 2002 to September 2002, he worked as a high-speed circuit designer at NetLogic Microsystems, Mountain View, California, on advanced interface technology. He is a student member of the IEEE.


Benton H. Calhoun received the BS degree in electrical engineering with a concentration in computer science from the University of Virginia, Charlottesville, in 2000. He received the MS degree from the Massachusetts Institute of Technology, Cambridge, in 2002. He is currently pursuing the PhD degree at the Massachusetts Institute of Technology. His research interests include leakage reduction, sensor networks, energy-efficient circuits, and sub-

Denis C. Daly (S'02) received the BAsC degree in engineering science from the University of Toronto in 2003. He is currently working toward the SM and PhD degrees at the Massachusetts Institute of Technology. His research interests include low-power wireless transceivers, sensor networks, and energy-efficient circuits. From May 2003 to August 2003, he worked on high-speed signaling systems at Intel Laboratories, Hillsboro, Oregon


Naveen Verma received the BAsC degree in electrical engineering and computer engineering from University of British Columbia in 2003. He is currently working toward the Master's degree in electrical engineering at Massachusetts Institute of Technology, Cambridge. His research focuses on the design of ultra-low power mixed signal circuits and, in particular, analog to digital converters. From April 2000 to January 2001, he worked as an ASIC designer at Nortel




Daniel F. Finchelstein graduated from the University of Waterloo in 2003 with the BSc degree in computer engineering. He is now pursuing the SM and PhD degrees in electrical engineering at the Massachusetts Institute of Technology. His current research interests include low-energy circuits/architectures and wireless sensor networks.



David D. Wentzloff received the BSE degree from the University of Michigan, Ann Arbor, in December 2000. He received the SM degree from the Massachusetts Institute of Technology (MIT), Cambridge, in September, 2002, where his focus was in power electronics and electric machines. He is currently working toward the PhD degree. His research interests are in RF front-end circuit design for ultra-wideband communication. He is a student member of the IEEE



Alice Wang received the SB, MEng, and PhD degrees in electrical engineering and computer science from the Massachusetts Institute of Technology in 1997, 1998, and 2004, respectively. Since January 2004, she has been working at Texas Instruments, Dallas, Texas, where she is a senior IC designer for the Chip Technology Center in the Wireless Division. Her research interests include ultra-low voltage CMOS circuits and systems and power management techniques for deep submicron technologies. She is serving on the microprocessor forum committee for ISSCC 2005 and on the technical program committee for the 2004 International Symposium on Low Power Electronics and Design (ISLPED). She is a member of the IEEE.



Seong-Hwan Cho (S'94-M'04) received the BS degree in electrical engineering from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, in 1995, and the SM and PhD degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1997 and 2002, respectively. During the summer of 1999, he was with the IBM T.J. Watson Research Center, Yorktown Heights, New York.

In 2002, he joined Engim, Inc., where he was involved in data converters, phased-locked loops (PLL), and voltage-controlled oscillator (VCO) design. Since the fall of 2004, he has been with KAIST as an assistant professor in the Department of Electrical Engineering. His research interests include mixed-signal and analog circuits for low-power communication systems. He is a member of the IEEE.



Anantha P. Chandrakasan received the BS, MS, and PhD degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1989, 1990, and 1994, respectively. Since September 1994, he has been with the Massachusetts Institute of Technology, Cambridge, where he is currently a professor of electrical engineering and computer science. He has received several awards, including the 1993 IEEE Communications Society's Best Tutorial Paper Award, the IEEE Electron Devices Society's 1997 Paul Rappaport Award for the Best Paper in an EDS publication during 1997, the 1999 Design Automation Conference Design Contest Award, and the 2004 DAC/ISSCC Student Design Contest Award. His research interests include low-power digital integrated circuit design, wireless microsensors, ultra-wideband radios, and emerging technologies. He is a coauthor of *Low Power Digital CMOS Design* (Kluwer Academic, 1995) and *Digital Integrated Circuits* (Pearson Prentice-Hall, 2003, second edition). He is also a coeditor of *Low Power CMOS Design* (IEEE Press, 1998) and *Design of High-Performance Microprocessor Circuits* (IEEE Press, 2000). He served as a technical program cochair for the 1997 International Symposium on Low Power Electronics and Design (ISLPED), VLSI Design '98, and the 1998 IEEE Workshop on Signal Processing Systems. He was the Signal Processing Subcommittee chair for ISSCC 1999-2001, the Program vice-chair for ISSCC 2002, the program chair for ISSCC 2003, and the Technology Directions Subcommittee chair for ISSCC 2004. He was an associate editor for the *IEEE Journal of Solid-State Circuits* from 1998 to 2001. He serves on the SCS AdCom and is the meetings committee chair. He is the Technology Directions Chair for ISSCC 2005. He is a fellow of the IEEE.

He was the Signal Processing Subcommittee chair for ISSCC 1999-2001, the Program vice-chair for ISSCC 2002, the program chair for ISSCC 2003, and the Technology Directions Subcommittee chair for ISSCC 2004. He was an associate editor for the *IEEE Journal of Solid-State Circuits* from 1998 to 2001. He serves on the SCS AdCom and is the meetings committee chair. He is the Technology Directions Chair for ISSCC 2005. He is a fellow of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.