Power Aware Wireless Microsensor Systems

Anantha Chandrakasan, Rex Min, Manish Bhardwaj, Seong-Hwan Cho, and Alice Wang
Massachusetts Institute of Technology, Cambridge, MA
{*anantha, rmin, manishb, chosta, aliwang*}*@mtl.mit.edu*

Abstract

Distributed networks of thousands of collaborating microsensors promise a maintenance-free, fault-tolerant platform for gathering rich, multi-dimensional observations of the environment. As a microsensor node must operate for years on a tiny energy source, innovative energy management techniques are needed. Widespread device deployment makes battery replacement infeasible, requiring energy to be scavenged from the environment—e.g., conversion of ambient vibrations to electric energy. Computation and communication must be optimized for very low duty cycles, making issues such as standby leakage and start-up overhead critical. All levels of the communication hierarchy, from the physical and link layer to routing protocols, must be tuned for energy efficiency. A total-system approach is required for reliable, self-powered microsensor networks that deliver maximal system lifetime in the most challenging environments.

1. Wireless Microsensor Networks

The idea of wireless microsensor networks has garnered a great deal of attention and interest. A distributed wireless microsensor network [1] consists of hundreds to several thousands of small sensor nodes scattered throughout an area of interest. Each node individually monitors its environment and collects data as directed by the user, while the network collaborates as a whole to deliver high-quality observations to a central base station. The fusion of observations from different perspectives offers a high-resolution, multi-dimensional picture of the environment that is not possible with fewer sensors. The sheer number of nodes naturally leads to the network's fault-tolerance and robustness to the loss of individual nodes, making maintenance unnecessary. Nodes can be deployed simply by scattering them about the region of interest or dropping them by air; the nodes can organize themselves into networks without user intervention.

These advantages, as well as the nodes' small size, make sensor networks ideal for any number of inhospitable or inaccessible locations where deployment is difficult, wires impractical, and maintenance impossible. For instance, microsensors deployed in hostile environments can monitor climate, classify moving vehicles, or provide an early warning of chemical or radiation hazards. A microsensor network distributed around the body, or perhaps within it, can offer a rich fusion of vital signs to medical professionals. Sensors within machines such as copiers and industrial robots, can detect and report emerging faults without the usual tangle of wires [2, 3, 4]. Microsensor networks promise to revolutionize how data is gathered.

A microsensor node integrates sensing, processing, and communication sub-systems. Several researchers have demonstrated operational nodes with low-power commercial, off-the-shelf (COTS) components. A representative example is depicted in Figure 1, which contains an on-board acoustic sensor and A/D, a StrongARM processor for data and protocol processing, power regulators for dynamic energy management, and a 2.4 GHz Bluetooth compatible radio. This node integrates these components onto stackable 55mm x 55mm boards as illustrated by Figure 2. Energy dissipation of this COTS-based microsensor node is reduced through a variety of techniques including fine-grain shutdown of inactive compo-

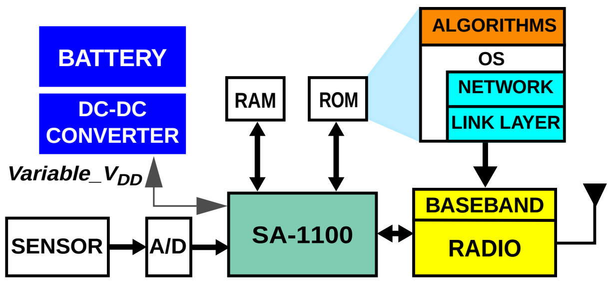

Figure 1: Architectural overview of the first generation MIT sensor node (The MIT μAMPS project) based on low-power off-the-shelf components. The node allows algorithms to gracefully scale its energy consumption by modifying hardware parameters.

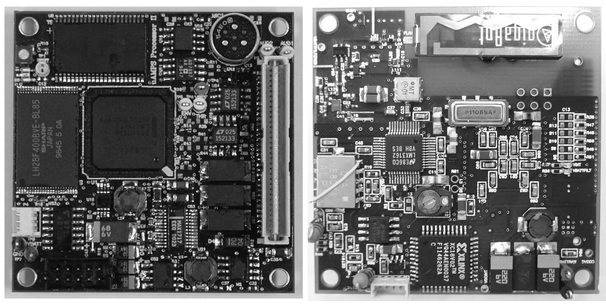

Figure 2: The μAMPS-1 sensor node: sensor and processor sub-system (left) and radio communication module (right).

nents through operating system control, dynamic voltage and frequency scaling of the processor core, approximate-processing algorithms, and adjustable radio transmission power based on required range. More than one order of magnitude variation in power dissipation is exhibited among the different power management states. In the peak operating condition, however, the node dissipates hundreds of milliwatts.

Dynamic voltage scaling (DVS) is a technique used for active power management in which the supply voltage and clock frequency of the processor are varied depending on the computational load [5]. The processor's supply voltage is reduced to the lowest possible level that meets the required performance constraints. DVS is highly applicable to processors within sensor nodes, as processor load can vary significantly based on the node's operational mode (*e.g.*, sensing/processing *vs.* data relay) and event statistics. Figure 3 depicts the measured energy consumed per operation for the StrongARM processor with respect to the processor clock frequency and supply voltage. The graph illustrates the advantage of voltage scaling at reduced processing loads. Note that at a fixed supply voltage, the leakage energy per operation increases as the allowed switching time per operation increases. The supply voltage is scheduled by the application-layer through an embedded operating system and is controlled at the physical-layer by an efficient variable voltage DC-DC converter.

Despite the advances made, nodes designed with COTS components are too large and costly for truly dense microsensor networks. Nodes must be a cubic centimeter at most, and cost under $1, to be commercially viable for networks of thousands of nodes. Fortunately, advances in the integration of MEMS sensors, digital electronics, and radio circuits are reducing size and cost simultaneously. High-density sensor networks are capable of generating a tremendous amount of data—much of it redundant. Communication protocols will need to prevent individual nodes from saturating limited radio communication bandwidth.

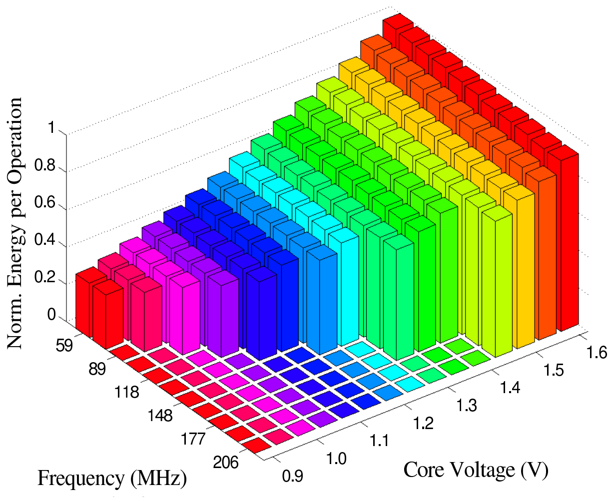

Figure 3: Dynamic Voltage Scaling on the SA-1100.

What have been overlooked thus far, however, are energy and lifetime considerations, quite possibly the greatest challenge to the design of efficient microsensor networks. Replacing the batteries of thousands of nodes in a hostile or inaccessible environment is simply not possible; nodes must be designed to operate without maintenance for years from a tiny on-board battery. When possible, energy must be scavenged from the environment to power the nodes. Improvements to energy efficiency require that COTS-based nodes give way to more integrated designs that explicitly consider the unique operational challenges of microsensor applications. Hardwired (or application-specific) processors must be used in favor of programmable solutions as they offer more than three orders of magnitude reduction in energy dissipation to implement a given function.

Environmental microsensors process data at low rates (bits/sec to kbits/sec) and with modest latencies during the data gathering process. Digital circuits should be designed for low-speed operation using deep voltage scaling and an emphasis on standby leakage reduction. Low data rates and high node density imply short packets and transmission distances. These characteristics amplify radio inefficiencies such as startup time, modulation circuitry, and media access, requiring special attention to radios for microsensors. The wireless network allows the microsensors to be highly collaborative entities, so we must look beyond node hardware to the protocols that drive inter-node communication. Even the data forwarding protocols must be designed with their impact on hardware energy consumption in mind. Using all of these techniques, we seek to reduce a node's energy consumption from hundreds of milliwatts to hundreds of *micro*-watts—a three order reduction—at which point it becomes possible to design *infinite*-lifetime nodes through environmental energy harvesting.

2. Ultra-Low-Energy Computing

The ability to perform low-energy computation is critical for extending the lifetime of the entire sensor network. The energy per operation, a key figure of merit in digital circuits, continually improves with process and supply voltage scaling. Energy is minimized through the use of highly dedicated computational fabrics and through careful conditioning of logic based on signal statistics. Use of extremely low-energy computational fabrics allows distribution of signal processing among the nodes in the network, instead of performing centralized computation. Specifically, data from multiple sensors can be cheaply aggregated among sensing nodes in the network, resulting in lower effective data communicated to the central basestation.

2.1 Energy-Aware Computing

Energy scalability is an important trend that involves the system adapting to time-varying operating conditions [6, 7, 8]. This is in contrast to current low-power approaches, which target the worst-case operating scenario. An energy-aware circuit monitors its available energy resources and dynamically adapts hardware parameters to meet latency and performance requirements. Hardware knobs that can be varied range from circuit parameters such as bit-precision and supply voltage to system parameters such as the numbers of operations performed (e.g., filter length).

For instance, an arithmetic circuit such as a multiplier is subject to diversity in operand width. Multiplier circuits are typically designed for a fixed operand size, such as 32 bits per input; calculating an 8-bit multiplication on a 32-bit multiplier results in unnecessary switching of the high-order bits. This excess switching would not have occurred if the 8-bit multiplication had been performed on an 8-bit multiplier.

As small operands can result in inefficient computation on larger multipliers, an architectural solution that improves energy awareness is the incorporation of additional, smaller multipliers of varying sizes, as shown in Figure 4. Incoming multiplications are routed to the smallest multiplier that can compute the correct result, reducing the energy overhead of unused bits. An ensemble of point systems, each of which is energy-efficient for a small range of input widths, takes the place of a single system whose energy consumption does not scale as gracefully with input diversity. The size and composition of the ensemble is an optimization problem that accounts for the probabilistic distribution of the inputs and the routing energy overhead. For an operand bitwidth distribution typical of a speech application, the ensemble of Figure 4 consumes 57% less energy than a monolithic multiplier [8].

The ensemble concept is applicable to a broad class of digital circuits such as discrete-time filters and memory arrays. A processor using DVS can also be considered an ensemble of point systems of fixed voltage, with a "virtual ensemble" synthesized by the voltage scaling circuitry.

2.2 Optimal Supply and Threshold Scaling

As discussed above, dynamic voltage supply scaling efficiently reduces power dissipation in response to variations in computational activity. Further energy reductions are possible by coordinating a second parameter, the device threshold voltage V_{th}, with the supply voltage. Device thresholds can be adjusted through substrate biasing in triple-well CMOS technology. Just as V_{dd} scaling exploits the trade-off between propagation delay and switching energy, V_{th} scaling can exploit a trade-off between propagation delay and sub-threshold leakage power. For every clock frequency, there exists an optimal pair of parameters (V_{th}, V_{dd}) that minimizes energy.

Figure 5 illustrates the impact of varying the power supply and threshold voltage on the performance and energy dissipation of submicron circuits. Diagonal lines indicate the (V_{th}, V_{dd}) values that will support operation of a 16-bit adder for widely varying levels of performance between 10 kHz and 100 MHz. Circular contours represent the increasing amount of energy consumed for each (V_{th}, V_{dd}) around a minimum value at (460 mV, 300 mV). The minimum-energy operating points for each clock frequency, then, occur where each frequency plot is tangent to the (lowest) energy contour. These optimal points are joined by a dotted line that represents the optimal (V_{th}, V_{dd}) for each clock frequency.

The optimal (V_{th}, V_{dd}) selections are most intuitive at the extremes of performance. The highest performance can only be achieved with a low V_{th} and high V_{dd}, at the expense of high switching and leakage energies. In the kilohertz regime where high circuit latencies are tolerable, not only is V_{dd} lowered for reduced switching energy,

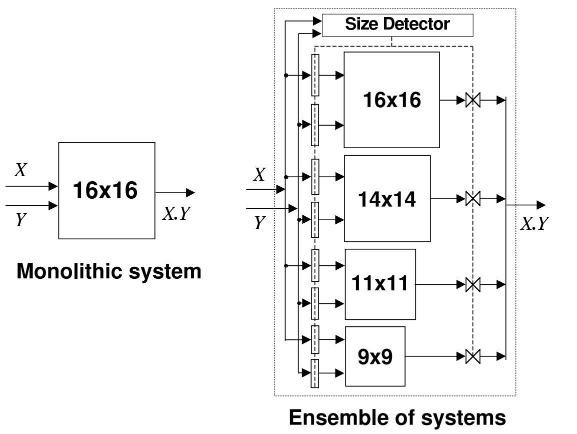

Figure 4: An ensemble of point multipliers routes incoming operands to the smallest multiplier capable of accepting the operands, resulting in reduced switched capacitance compared to a monolithic circuit.

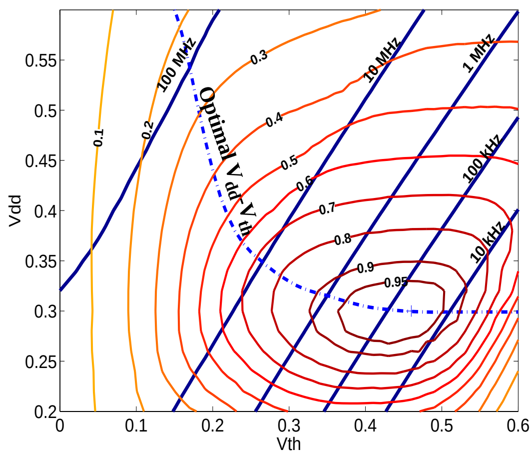

Figure 5: Average energy dissipation and performance for a 16-bit adder. The dotted line indicates the adder's optimal (V_{th}, V_{dd}) operating points for varying clock frequencies.

but V_{th} is simultaneously raised to suppress leakage. In fact, the supply voltage is scaled *below* the threshold voltage; load capacitances are switched by subthreshold leakage currents. As leakage currents are orders of magnitude lower than drain currents in the strong inversion regime, both performance and active power dissipation are substantially reduced. Given the low performance demands of microsensor nodes, operation in the subthreshold operating regime is an exciting possibility.

While a single adder has been considered here, this analysis can be expanded to find optimal operating points for larger digital circuits. A truly energy-scalable microsensor node would feature variable clock frequency, supply voltage, and threshold voltage, with dynamic power management used to vary the clock frequency and select (V_{th}, V_{dd}) for ultra low power operation. Practical approaches for adaptive supply and threshold voltage are presented in [9]. Both open loop (e.g., lookup tables) and closed loop (critical path replica circuits) approaches can adjust these parameters in an optimal fashion.

2.3 Idle-mode Leakage Control

Microsensors typically spend most of their time in a standby mode, waiting for significant events to occur. Hence, powered components dissipate leakage energy over long periods of time. One approach to reducing idle mode energy dissipation is simply to shut off all unused electronics during idle mode. However, any energy savings from shutdown can be negated by the potentially large latencies and energy overheads required to power up the node from its off state. Idle mode energy is therefore best addressed at its source, the leakage currents flowing through idle circuits.

Multiple-Threshold CMOS (MTCMOS), for instance, reduces idle mode leakage by employing high-V_{th} transistors to gate the power supplies to the logic blocks which are designed with low-V_{th} transistors [10]. Designing sequential MTCMOS circuits is challenging since state is lost during sleep mode while the power supplies are floating. MTCMOS designs are prone to “sneak” (unexpected) leakage paths through low-V_{th} gates. Leakage feedback flip-flops utilize leakage to hold state while avoiding sneak leakage paths [11]. Figure 6 shows a static leakage feedback flip-flop with the low-threshold devices highlighted. This circuit achieves performance close to a traditional low-V_{th} flip-flop while retaining the low leakage of a high-V_{th} flip-flop. Future digital systems must exploit multiple and variable threshold devices for leakage control.

3. Low Duty Cycle Radio Communication

Microsensors’ long idle periods and low data rates imply node-to-node communication with a low duty-cycle and brief transmissions. The communication subsystem for wireless microsensors must therefore be optimized for these conditions. For short range transmission at GHz carrier frequencies, the power consumption of communication is dominated by the radio components (frequency synthesizer, mixers, etc.) rather than the actual transmit power radiated into the air. To conserve power, it is therefore essential that radio electronics be turned off during idle periods. Unfortunately, GHz-band frequency synthesizers require significant time and energy overhead to transition from the sleep state to the active state. For short packet sizes, the transient energy consumed during start-up can be significantly higher than the energy required by the electronics during the actual transmission.

The effect of the start-up transient on energy efficiency is shown in Figure 7. The energy required to transmit a bit with 0 dBm output power is plotted *versus* packet size. The upper plot represents a COTS radio while the lower plot represents an “ideal” radio, a lower bound that consists of only the power radiated (at 100% efficiency). For the shorter packet sizes characteristic of microsensor networks, the difference between the two curves increases dramatically due to the fixed start-up cost of the radio. While the constant offset from the lower bound can be reduced with low-power transmitter techniques, the inefficiency introduced by short packet sizes can only be improved by reducing start-up time.

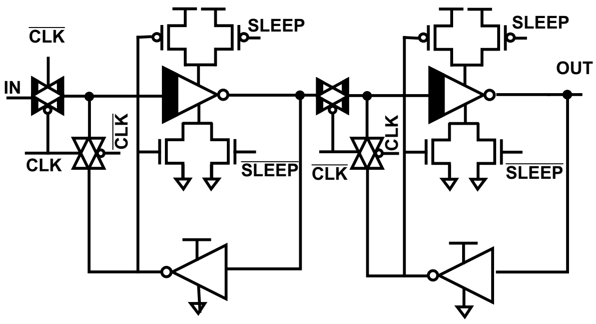

Figure 6: Leakage Feedback Flip-flop [11].

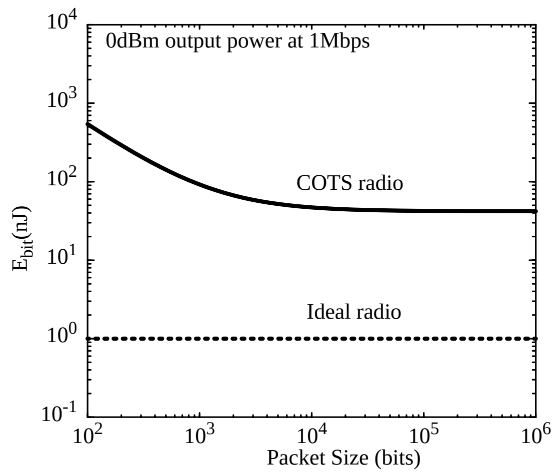

Figure 7: Effect of start-up time on transmitter’s energy.

3.1 Fast Start-up Low Power Transmitter

The start-up time of the transmitter is dominated by the frequency synthesizer due to the time required to stabilize its PLL. A popular approach to reduce the settling time is the use of a variable loop bandwidth [12]. The PLL is started with a wide loop bandwidth and is transitioned to a narrower loop bandwidth as the loop approaches lock. As this method requires simple overhead circuitry, it is attractive for low power PLL applications.

The on-time of the transmitter must be reduced to lower the energy utilized per bit. One promising architecture for continuous phase-modulated signals is an indirect modulation method that uses Σ–Δ in a fractional-N synthesizer [13]. This architecture eliminates the need for mixers or DACs in the heterodyne scheme. Another compact architecture for continuous phase modulation is closed loop, direct VCO modulation. This architecture requires a low gain varactor on the VCO and supports simple BFSK modulation.

The modulator shown in Figure 8 employs closed loop, direct VCO modulation to achieve a high data rate, variable loop bandwidth for a swift transient response, and a Σ–Δ for reduced power consumption in the divider with fine resolution in channel selection. The result of the variable loop bandwidth technique on start-up time is shown in Figure 9, where the VCO control voltage is plotted for both fixed and variable loop bandwidth. Variable loop bandwidth reduces the start-up time by a factor of four.

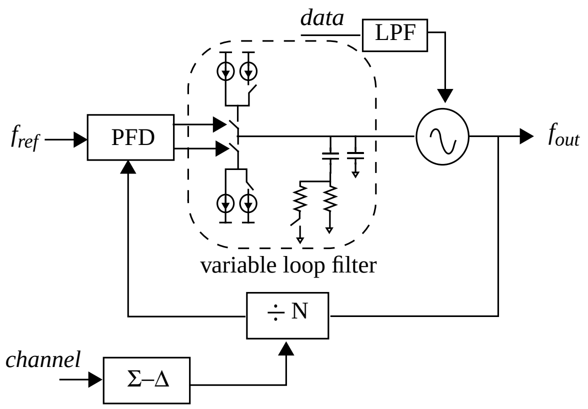

Figure 8: Low power, fast start-up transmitter.

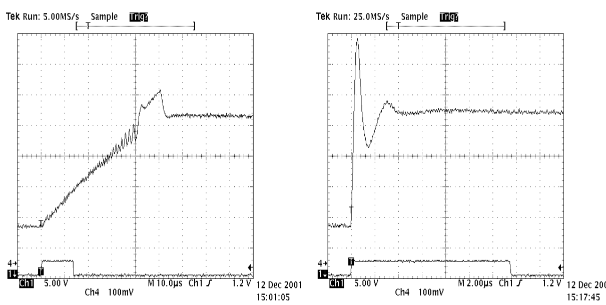

(a) Fixed loop (b) Variable loop

Figure 9: Start-up times of variable and fixed loop modulator.

3.2 Low Power Modulation

The radio's start-up time impacts the choice of an energy-efficient modulation technique for the microsensor communication subsystem. The energy efficiency of communication is traditionally improved by sending multiple bits per transmitted symbol (i.e., an *M*-ary modulation technique). Architectures that implement binary and *M*-ary modulation are compared in Figure 10. The binary modulation architecture assumes a continuous phase modulation scheme that can be implemented using direct or indirect modulation of the VCO, which eliminates the need for the mixer and DACs seen in the *M*-ary scheme.

While *M*-ary modulation reduces the transmit time of the radio by sending multiple bits per symbol, it may not necessarily achieve lower energy consumption for microsensors due to the increased complexity of the modulation circuitry. The effect of start-up time is shown in Figure 11, which plots the ratio of energy consumption between the two modulation schemes. In both schemes, the packet size is 100 bits and a data rate of 1 Mbps is assumed. *M*-ary modulation achieves lower energy than binary modulation only if the start-up time is small compared to the transmit time, reiterating the need for microsensor radios to be designed for fast start-up.

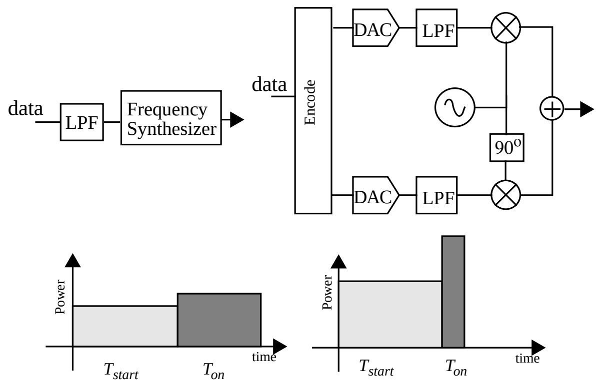

(a) Binary modulation (b) *M*-ary modulation

Figure 10: Comparison of binary and *M*-ary modulation.

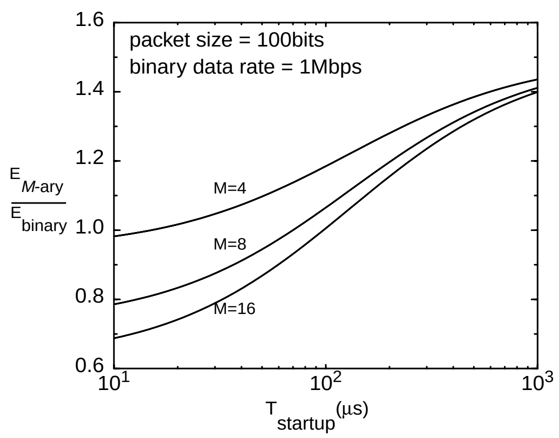

Figure 11: Impact of startup time on the efficiency of *M*-ary modulation.

4. Energy-Efficient Communication

The energy consumption of node-to-node communication depends not only on the processing and radio hardware, but also the communication protocols that drive this hardware. It is essential to consider how protocols and software impact hardware energy consumption.

4.1 Energy of Multihop Communication

The energy of on-chip communication is approximately linear with distance, for the capacitance of one-dimensional interconnect scales linearly with distance. The energy required for inter-node communication, however, scales with distance as d^2 to d^4. Since the path loss of radio transmission scales with distance in a greater-than-linear fashion, communication energy can be reduced by dividing a long transmission into several shorter ones. Intermediate nodes between a data source and destination can serve as relays that receive and rebroadcast data. This concept, known as *multihop* communication, is analogous to the use of buffers over a long, on-chip interconnect.

Figure 12 illustrates multihop communication to a base station across a distance d using h hops. The power consumed by this communication $P(h, d)$ is

$$P(h,d) = (h-1)P_{rxElec} + h\left[P_{txRad1}\left(\frac{d}{h}\right)^r + P_{txElec}\right] \quad (1)$$

where P_{rxElec} and P_{txElec} represent the power required by the receive and transmit electronics, P_{txRad1} is the radiated power required for a successful one-meter transmission, and r is the path loss exponent. Since the last hop is always received by an energy-unconstrained base station, there are h transmitting and h-1 receiving nodes. Figure 13 evaluates equation (1) over varying total transmission distances d and one to four hops, using representative power consumption parameters for a COTS-based sensor node.

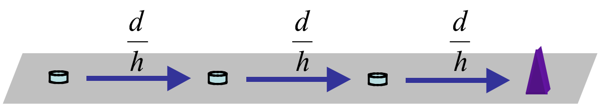

Figure 12: Multihop routing.

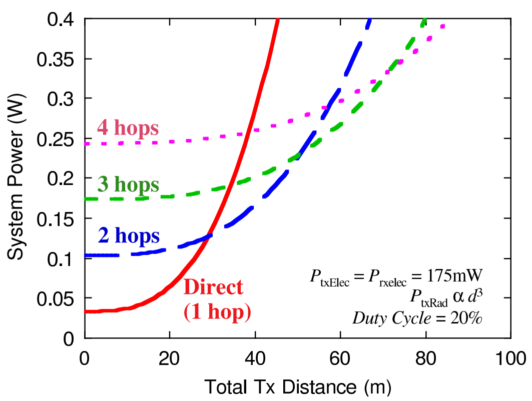

Figure 13: As the distance between a source node and base station increases, the number of intermediate hops required for maximal energy-efficiency increases.

The introduction of relay nodes is clearly a balancing act between reduced transmission energy and increased receive energy. Hops that are too short lead to excessive receive energy. Hops that are too long lead to excessive path loss. In between these extremes is an optimum transmission distance called the *characteristic distance* d_{char} [14]. The characteristic distance depends only on the energy consumption of the hardware and the path loss coefficient; d_{char} alone determines the optimal number of hops. For typical COTS-based sensor nodes, d_{char} is about 20 meters.

The existence of a characteristic distance has two practical implications for microsensor networks. First, it is often impractical to ensure that all nodes are spaced exactly d_{char} apart. Nodes may dropped by air, or their deployment constrained by terrain or physical obstacles. The deployed nodes may be placed as in Figure 14, a line of nodes and a base station separated a distance of either d or $2d$, with $d < d_{char} < 2d$. As the figure illustrates, there are three possible multi-hop policies from the farthest node to the base station. Considering that none of the inter-node distances is exactly equal to d_{char}, what is the minimum-energy policy?

The optimal solution turns out to be a rotation of roles over time. The final numerical result depends heavily on the node energy models that quantify the trade-off between the path loss of transmission and the power dissipation of the radio electronics. For the energy models used by [14], the optimal policy dictates that communication occur through each one of the one-hop routes 24.5% of the time, and through the two-hop route 51% of the time. This rotation of policies effectively *dithers* the transmission distance so that it approaches d_{char} when the actual nodes are not d_{char} apart.

The second practical implication of a fairly large d_{char} is that there are a large class of applications for which the

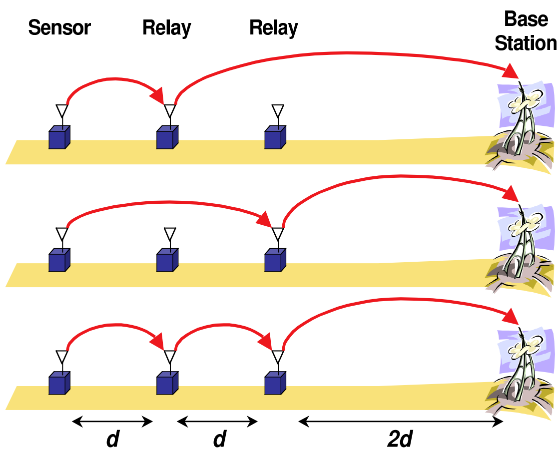

Figure 14: Three multihop policies are possible when two intermediate nodes stand between the transmitting sensor and base station.

entire network diameter will be less than d_{char}. For these applications, the best communication policy is not to employ multihop at all; direct transmission from each node to the base station is the most energy-efficient communication scheme. For today's radio hardware, the typical d_{char} of 20 meters exceeds the size of many interior spaces. Hence, until advances in low-power receiver technology lead to a reduction in d_{char}, most indoor microsensor networks will not save energy using a multihop routing protocol.

4.2 Communication API

Communication protocols, such as multihop routing, must take advantage of a microsensor node's energy scalability and awareness. Energy aware communication is achieved by allowing well-defined performance metrics for communication to be traded gracefully for energy savings in the hardware.

The performance of communication can be quantified by three parameters: range, reliability, and latency. Range represents the distance to the recipient, reliability indicates the likelihood that the transmitted data is properly received, and latency measures the time required for the end-to-end communication. Applications can facilitate energy conservation by relaxing any of these parameters, allowing the communication hardware to trade performance for energy savings. Transmission range, for instance, can be reduced with a variable-power transmit amplifier. Reliability can be adjusted with variable-strength forward error correction (FEC). Finally, DVS and clock frequency scaling can adjust the latency of digital computation (e.g., required for FEC).

The remaining task is to set hardware "knobs" such as supply voltage, clock frequency and amplifier power such that the performance parameters requested by communication software are satisfied with minimal energy expenditure. Relating latency, reliability, and range to actual hardware energy consumption is a challenging task. Many parameters interact: range and reliability are closely linked, for instance, since a radio transmission becomes less reliably received as it travels farther from its sender. FEC strength impacts the energy consumption of both processor and radio: a stronger code not only consumes more digital processing resources, but also potentially increases the number of transmitted bits.

While communication software and hardware each expose the necessary parameters for graceful energy scalability, the two levels essentially speak very different languages. Communication software requests performance in terms of meters and bit error rates, not supply voltages and power levels. Something must bridge the gap.

The solution is a layer of power-aware "middleware" between the communication hardware and software. As illustrated in Figure 15, the middleware layer exposes an application programming interface (API) to communication software that allows the specification of constraints on latency, reliability, range, and total energy. The middleware translates these software constraints into the minimum-energy hardware policies that satisfy them.

Figure 16 presents a realization of range and reliability scaling by the middleware based on the energy consumption parameters of the COTS node of Figure 2. Given specifications of transmission distance and tolerable bit error rate from the application, the middleware selects the least-energy FEC scheme and transmission power level supported by the hardware. *R* and *K* represent the rate and constraint length of a convolutional code; a lower *R* and higher *K* afford increasing resilience to errors. The modeled hardware supports two transmission power levels, +0 dBm and +20 dBm. Five convolutional codes and two power levels are considered, yielding ten possible operational policies. As increasing range or reliability are demanded, the middleware responds by either switching to a more robust coding scheme or increasing transmitter power.

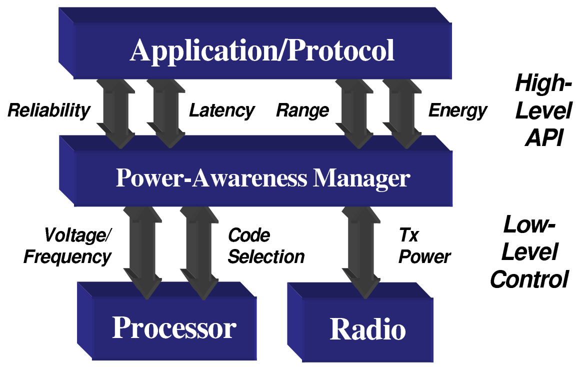

Figure 15: Power-aware middleware translates communication performance requirements into optimal hardware settings.

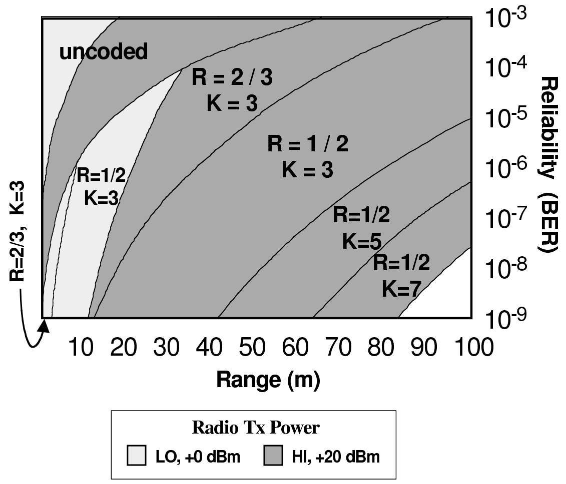

Figure 16: Example middleware operational policy for a COTS microsensor node. As an application's range and reliability demands change, the middleware chooses the least-energy convolutional code (rate *R*, constraint length *K*) and transmit power that meet the demands.

5. Energy Scavenging

As the power dissipation of entire sensor systems is reduced to hundreds of microwatts, it becomes possible to operate nodes with energy derived from the environment. Various schemes have been proposed to eliminate the need for batteries in a portable digital system by converting ambient energy in the environment into electrical form. The harvested electrical energy can be stored and utilized by the node's electronic circuits. The most familiar sources of ambient energy include solar power, thermal gradients, RF, and mechanical vibration [15].

Advances in MEMS technology have enabled the construction of a self-powered system in which a MEMS devices acts a power source for a digital load [16]. The MEMS device is a variable capacitor that converts mechanical vibration into electrical energy. The capacitor plates are charged and then moved apart by vibration, resulting in the conversion of mechanical energy into electrical energy. The device consists of three basic parts: a floating mass, a folded spring, and two sets of interdigitated combs. A plan view of the MEMS generator and micrograph showing the device detail are shown in Figure 17. With appropriate regulation circuitry, this device delivers 10 μW of power.

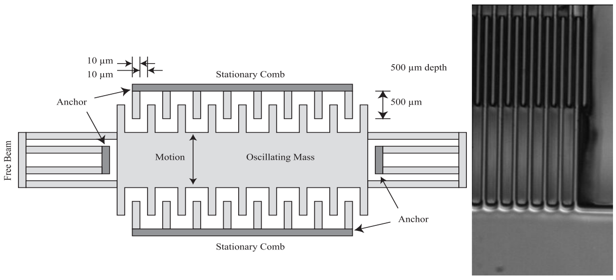

Figure 17: MEMS generator plan view and device detail [16].

6. Conclusion

Wireless microsensor networks utilize hundreds to thousands of tiny, inexpensive, and densely placed nodes to achieve unprecedented sensing resolution and fault-tolerance. The nodes' limited energy capacity and extended lifetime requirements demand that all aspects of a microsensor node be designed in a power aware fashion. Both digital and analog node circuitry must reflect the distinctive operational characteristics of microsensors, notably long idle times and low performance demands. Power aware design cannot end with hardware; software and communication protocols for microsensors must actively contribute to energy savings through energy-efficient operational policies and sacrifices in performance. As with all emerging wireless applications, the future of microsensor networks is reliant on designs that successfully meld unique operational demands with innovative circuit design for maximal energy efficiency.

Acknowledgments

The authors thank Nathan Ickes, Fred Lee, and Piyada Phanaphat for their contribution to the μAMPS-1 node. The authors acknowledge the contributions of James Kao, Stephen Kosonocky, Jose Mur-Miranda, and Jeff Lang to this paper. This research is sponsored by the DARPA Power Aware Computing/Communication Program and the Air Force Research Laboratory, Air Force Materiel Command, USAF, under agreement number F30602-00-2-0551; by the ARL Collaborative Technology Alliance, through BAE Systems, Inc. subcontract RK7854; by ABB Corporation; and by Hewlett-Packard under the HP/MIT Alliance. R. Min is supported by an NDSEG Fellowship, A. Wang by a Lucent GRPW Fellowship, and M. Bhardwaj by an IBM Fellowship.

References

[1] G. Asada *et al.*, "Wireless Integrated Network Sensors: Low Power Systems on a Chip," *IEEE ESSCIRC '98*, pp. 9-16, September 1998.

[2] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next Century Challenges: Scalable Coordination in Sensor Networks," *ACM MobiCom '99*, pp. 263-270, August 1999.

[3] J. Kahn, R. Katz, and K. Pister, "Next Century Challenges: Mobile Networking for Smart Dust," *ACM MobiCom '99*, pp. 271-278, August 1999.

[4] J. Rabaey *et al.*, "PicoRadios for Wireless Sensor Networks: The Next Challenge in Ultra-Low-Power Design," *IEEE ISSCC*, pp. 200-201, July 2002.

[5] V. Gutnik, A. Chandrakasan, "Embedded Power Supply for Low-Power DSP," *IEEE Trans. on VLSI Systems*, pp. 425-435, December 1997.

[6] A. Sinha, A. Wang and A. Chandrakasan, "Energy Scalable System Design," *IEEE Trans. on VLSI Systems*, pp. 135-145, April 2002.

[7] E. Shih *et al.*, "Physical Layer Driven Protocol and Algorithm Design for Energy Efficient Wireless Sensor Networks," *ACM MobiCom* 2001, pp. 272-287, July 2001.

[8] M. Bhardwaj, R. Min, A. P. Chandrakasan, "Quantifying and Enhancing Power-Awareness of VLSI Systems," *IEEE Trans. on VLSI Systems*, pp. 757 -772, Dec 2001.

[9] M. Miyazaki, J. Kao and A. Chandrakasan, "A 175mV Multiply-Accumulate Unit using an Adaptive Supply Voltage and Body Bias (ASB) Architecture," *IEEE ISSCC 2002*, pp. 58-59, February 2002.

[10] S. Mutoh *et al.*, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS", *IEEE Journal of Solid-State Circuits*, pp. 847 - 854, August 1995.

[11] J. Kao and A. Chandrakasan, "MTCMOS Sequential Circuits", *IEEE ESSCIRC 2001*, pp. 332-335, September 2001.

[12] F.M. Gardner, *Phaselock techniques*, Wiley, NY, 1979.

[13] T. Riley, M. Copeland, and T. Kwasniewski, "Σ–Δ modulation in fractional-N frequency synthesis," *IEEE JSSC*, pp. 553-559, May 1993.

[14] M. Bhardwaj, T. Garnett, A. Chandrakasan, "Upper Bounds on the Lifetime of Sensor Networks," *IEEE International Conference on Communications*, pp. 785 -790, 2001.

[15] R. Amirtharajah and A. P. Chandrakasan, "Self-powered Signal Processing Using Vibration-based Power Generation," *IEEE JSSC*, pp. 687-695, May 1998.

[16] S. Meninger, J.O. Mur-Miranda, *et al.*, "Vibration-to-Electric Energy Conversion," *IEEE Trans. on VLSI Systems*, pp. 64-76, February 2001.