International Journal of Parallel and Distributed Systems and Networks, Vol. 4, No. 3, 2001

ENERGY AND PERFORMANCE CONSIDERATIONS FOR SMART DUST

L. Doherty,* B.A. Warneke,* B.E. Boser,* and K.S.J. Pister*

Abstract

We discuss the energy budget of nodes in a distributed wireless sensor network. Sensor nodes can currently compute for 1 pJ/instruction, communicate via RF at 100 nJ/bit, and sense the environment for 4 nJ/sample. Each class of expenditures is developed according to current empirical performance and physical limits. As technology progresses, cubic millimeter scale nodes will be able to survive on scavenged solar power. The limitations imposed by the energy budget are examined in the context of three network scenarios: building environment monitoring, earthquake sensing, and vehicle tracking. The scenarios are limited by either sensing or communication energy requirements, with computation being comparatively inexpensive. Algorithmic improvements to reduce communication and sensing will prolong the functionality of the sensor network.

Key Words

Smart dust, energy, MEMS, SOC

1. Introduction

Exponential technological improvements in sensing, computation, and communication have fomented the emergence of wireless sensor networks. Each node in a sensor network can sense the environment, make computations, and communicate its findings to other nodes in the network or a centralized observer, but all on a limited basis. Individual node capabilities are constrained by the energy available to perform tasks. Each sensor sample, computation, and bit communicated consumes valuable energy. The collective capability of a sensor network to monitor an environment with unprecedented spatial resolution is the aggregation of the limited capabilities of individual sensor nodes, and depends on the possibility of fabricating, deploying, and maintaining hundreds or thousands of nodes at low cost.

In an integrated solution to the autonomous sensor problem, size and cost are strongly correlated; prices of silicon devices are based on wafer area and are on the order of 5 cents/mm^2 [1]. The goal of the Smart Dust project [2] is to fabricate sensor nodes of cubic millimeter dimensions, thereby generating nodes costing pennies apiece. The components of Smart Dust are shown in Fig. 1, and a recently produced prototype is shown in Fig. 2. The MEMS-based (Micro-ElectroMechanical System) sensors and integrated circuitry for processing and RF communication can all be mass produced on the same silicon substrate. Although the integration of all three systems on a chip has not yet been demonstrated, we will look at the fundamental limits of sensor networks composed of such ideal components. Larger sensor nodes such as those in Fig. 3 have been demonstrated in networks [3, 4] with off-the-shelf components at the cubic inch scale, but these typically cost hundreds of dollars per unit to produce and require orders of magnitude higher power than dedicated silicon devices.

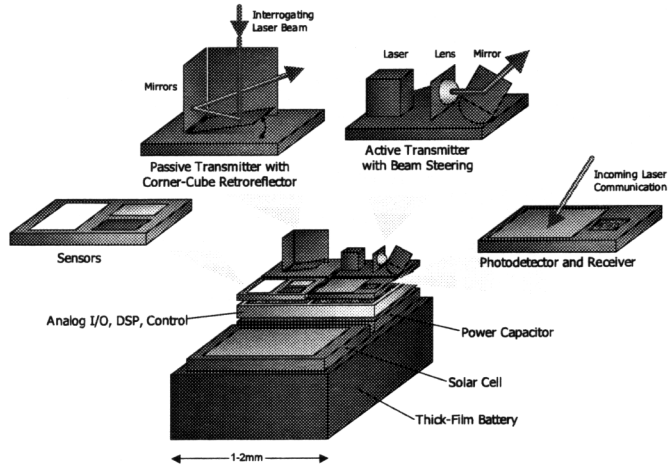

Figure 1. Smart Dust model. The node is broken up into energy sources, a digital processor, sensors, an analog-to-digital converter, and communication devices.

The parameters of greatest interest in most wireless sensor networks are sensor performance, power, and cost. In this article, we focus on the power aspect of sensor nodes, assuming that MEMS technology and integration will satisfy sensing and cost requirements. Two questions are posed. The first is, "How much energy is available to sensor nodes?" Once an apposite answer is found, we ask "What functionality can a sensor node achieve within this energy budget?" We consider these two questions from both the standpoint of current technological capabilities and the perspective of fundamental limits, for peer-to-peer sensor networks.

* Berkeley Sensor and Actuator Center, 497 Cory Hall, Berkeley, CA 94720 USA; e-mail {ldoherty,warneke,boser,pister}@eecs.berkeley.edu

(paper no. 204-0390)

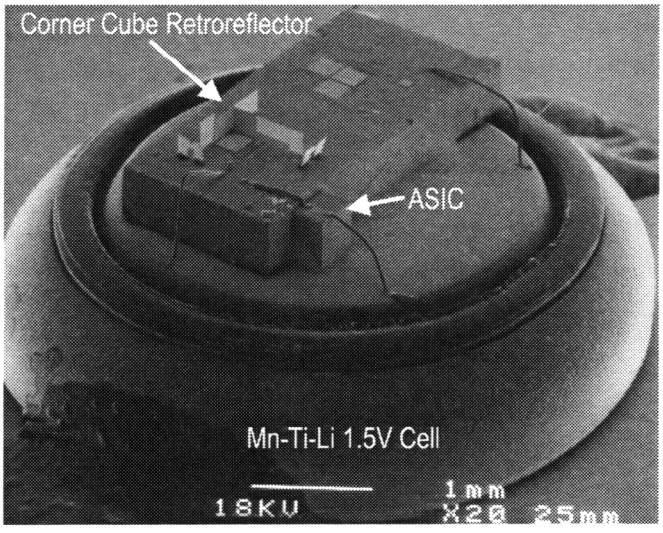

Figure 2. Current Smart Dust prototype—a $63mm^3$ autonomous bidirectional communication node. The bulk of the volume is a rechargeable hearing aid battery that powers the device. The circuitry (ASIC) constitutes only a small portion of the total size. Communication from this node is accomplished passively with an actuated micromirror.

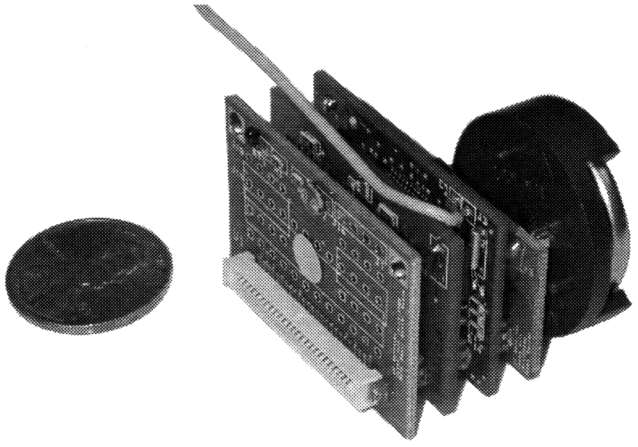

Figure 3. Photograph of a Rene RF sensor node. The lithium coin cell on the right powers the node, which is comprised of a control board and two sensor boards. The yellow wire extending from the control board is the antenna. Photo courtesy Robert Szewcyzk.

The article is organized as follows. In Section 2, we examine the current capabilities of nonrenewable and renewable energy sources for sensor nodes. In Section 3, we detail the energy requirements of pertinent digital operations in an optimized low-power processor. In Section 4, the current and theoretical energy limits of analog sensing and analog-to-digital conversion (ADC) are explored. In Section 5, the energy cost of transmitting data through radio frequency (RF) and optical methodologies is considered. In Section 6, we apply the estimated costs to three disparate sensor network scenarios to determine the feasibility of accomplishing useful tasks within the energy budget. In Section 7, we summarize the results.

2. Energy Sources

The most likely and simplest type of power storage for wireless sensor nodes is the lithium energy cell. The latest generations are rechargeable and capable of energy densities around 300 W-Hr/kg, or 2,000 J/cc. Other probable candidates include thin-film vanadium oxide and molybdenum oxide [5] that are fabricated using spin-casting sol-gel techniques. Not only do these materials achieve energy densities comparable to lithium cells, but their fabrication technique lends itself well to integration into silicon fabrication processes. For higher instantaneous current requirements, the Ultracapacitor [6] provides a new means of capacitive energy storage with a density of nearly 10 J/cc. This is less than 1% of the energy density of lithium cells, but has the capability of delivering the energy in seconds due to significantly less source impedance than the chemical storage cells.

"Solar" radiation is the most available source for energy scavenging. This photoelectric energy can come not only from the sun, but also from indoor lighting. Full sunlight gives around $1mW/mm^2$ (approximately $1J/day/mm^2$), and bright indoor illumination is roughly $1\mu W/mm^2$. Conversion efficiency is around 30% for the best cells. For applications where duty-cycling is acceptable, solar cells or other power scavenging sources can be used to trickle-charge a capacitor or energy cell, after which the stored energy can be used at much higher power rates than the charging pace.

Vibration has been proposed as a scavengeable energy source [7]. Indeed, vibration spectra of office windows, copy machines, and industrial motors reveal that there is useable energy here—typically on the order of 10 μW/g of mass of the converter. The mass of a cubic millimeter of silicon is about 2 mg, rendering this energy source substantially less valuable than solar power at the sizes of interest.

Even more exotic energy sources have been demonstrated or proposed, including utilizing the excess heat from micro rocket engine [8] combustion, micro radioactive energy sources, and harvesting ATP for *in vivo* applications.

3. Energy Expenditure in Digital Circuitry

Digital integrated circuit design techniques allow computations to be performed at a relatively low level of energy consumption in comparison to other system operations. It is thus often beneficial to expend extra computational energy to reduce the number of operations required in other parts of the system in order to diminish the overall system energy consumption. To make high-level decisions about what operations can or should be performed on a sensor node and what algorithms should be utilized, we look at the energy consumption of each instruction that the processor can perform (the reciprocal of MIPS/W). This expenditure depends on the architecture of the processor and the consumption of particular blocks within the processor,

Table 1
Leakage Power for Selected Blocks in a 0.25μm CMOS Process

	$Vdd = 0.5V$		$Vdd = 1.0V$	
	$\|V_{bs}\| = 0.0V$	$\|V_{bs}\| = 0.5V$	$\|V_{bs}\| = 0.0V$	$\|V_{bs}\| = 0.5V$
8-bit ripple-carry adder	98pW	5.5pW	260pW	18pW
Microprocessor core	6nW	0.34nW	16nW	1.1nW
1024x8 SRAM	53nW	4.3nW	160nW	15nW

which is further dependent on device level consumption.

3.1 Energy Consumption of Devices

Fundamentally, energy is dissipated by two mechanisms: static leakage currents and switching activity. Subthreshold leakage current flows through any transistor with a potential difference between its source and drain, even when the gate is turned off. For a given device, the current depends on the source/drain potential and the source/channel potential. To minimize this leakage, we thus want to minimize the supply voltage to minimize the source/drain potential of devices that must be active, and power-down devices whenever they are not needed. For a 0.25μm CMOS process with $V_{DD} = V_{DS} = 1.0V$, the leakage current for a minimum-sized n-channel MOSFET is 14.4pA, and for $V_{DD} = V_{DS} = 0.5V$ it is 9.9pA. An 8-bit adder example with 244 transistors has leakage power for various supply voltages, as given in Table 1, which also shows that increasing the reverse bias between the source and the channel (set by the well potential) will decrease the subthreshold leakage current exponentially.

The reverse-biased junction leakage currents consume attoamps per MOSFET source/drain junction - negligible compared to the subthreshold leakage. However, as these currents do vary exponentially with temperature, they can become a more significant source of leakage current when the device is operated at high temperatures.

Static power consumption is particularly important in ultra-low power design when a large digital system is considered. For example, a small microprocessor for a sensor node may contain some 15,000 transistors in the core and another 50,000 in a 1kb SRAM. Table 1 lists the appropriate estimated static consumption.

Switching energy consumption in non-adiabatic digital blocks is attributed to two sources: dynamic and short circuit. Dynamic energy is consumed through the charging of node capacitances within a circuit and is given by $E_{01} = C_L V_{DD}^2$ for a transition from ground to V_{DD} with a load capacitance C_L. Dynamic energy consumption is reduced by using the lowest supply voltage that meets the performance requirements and minimizing the node capacitances and number of transitions through layout, circuit, logic, and architecture techniques. Short-circuit energy consumption occurs when the power supply rails are temporarily connected through a low impedance path by transistors undergoing a transition. Energy loss through this mechanism is similarly reduced by minimizing the number of transitions, along with properly shaping the transition waveforms, sizing the transistors, and choosing appropriate logic styles. As an example of the magnitude of dynamic energy consumption, a minimum-sized inverter in the 0.25μm CMOS process driving a load of four identical inverters consumes 0.49fJ for $V_{DD} = 1.0V$ and 0.16fJ for $V_{DD} = 0.5V$. Although this is not necessarily indicative of the average energy per transition within the processor core, we will assume that it is typical behaviour in the following sections.

3.2 Energy Consumption of Particular Blocks

The processor control unit is implemented with random logic to minimize the amount of overhead as compared to the various programmable logic methods of implementing the control unit. To model the energy consumption of the random logic, the NAND gate will be used as an example. From simulation data, the typical energy consumption for a NAND gate within the control would be 10fJ with a load of four gates. Assuming that on average five control signals need to be asserted, and an average logic depth of two, approximately eight logic gates will be activated each cycle, yielding an energy consumption of about 100fJ to decode an instruction and drive the datapath.

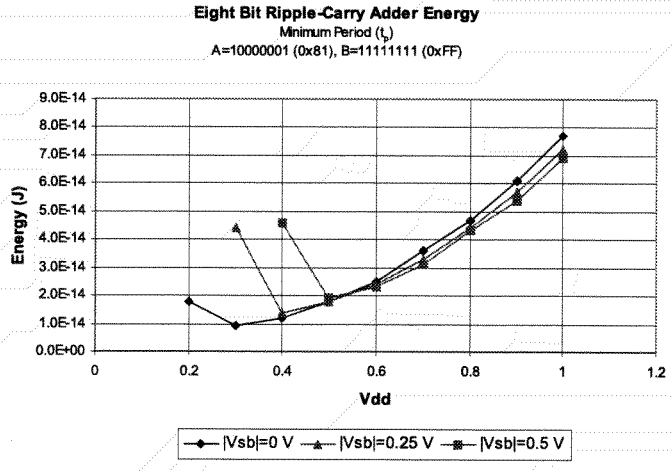

Figure 4. Worst-case energy consumption of an eight-bit ripple-carry adder integrated over the minimum transition period at each operation point.

As another example, consider the block required to add two numbers. A ripple-carry adder architecture minimizes

Table 2
Register Energy Consumption in a 0.25μm CMOS Process

	Stored Value	Number of Gate Loads	$V_{dd} = 0.5V$	$V_{dd} = 1.0V$
D flip-flop	1	0	2.0fJ	7.5fJ
	0	0	2.8fJ	13fJ
	1	4	2.2fJ	8.5fJ
	0	4	3.2fJ	15fJ
	1	8	2.3fJ	9.5fJ
	0	8	3.5fJ	17fJ
8-bit register	—	0	7.6fJ	41fJ
	—	4	10.8fJ	47fJ
	—	8	11.6fJ	53fJ

the number of transitions. For a 0.25μm process, energy requirements are shown in Fig. 4 for the worst-case addition of 0x81 and 0xFF (requiring propagation of a carry the length of the adder chain). Assuming that the power supply to the circuit can be shut off entirely, an optimum point for V_{DD} exists as illustrated. Simulating with an equal integration time for each supply voltage level reveals the impact of leakage current on energy consumption independent of frequency, showing whether it is better to operate quickly and then wait quietly, or just operate at the minimum speed possible. If the circuit cannot be shut off, the latter case is preferable.

Another common block in the processor is the register, comprising a bank of D flip-flops. A single D flip-flop in the 0.25μm CMOS process consumes an average of 10.25fJ for $V_{DD} = 1.0V$ each time the state changes. However, the register flip-flops are typically loaded by several tri-state buffers that selectively drive the buses. A D flip-flop loaded by four such buffers consumes an average of 11.75fJ. For an 8-bit register with an average of four bits changing states, this results in an energy consumption of 47fJ for $V_{DD} = 1.0V$. Table 2 summarizes more such results.

Traditionally, the buses in a microprocessor cause significant power consumption because of the large capacitive load they place on the components. However, in the small processor paradigm, the bus loading is not too substantial. Furthermore, by adding more buses to the architecture, the load on each bus can be reduced as well as the number of states in the control. For a 300 μm datapath, the wire loading will be 2.5 fF with a single gate loading of 2 fF; the load can easily dominate the energy required to drive the bus. Fig. 5 shows the energy consumed by a low-to-high transition on a single wire driven by a minimum-sized tri-state inverter at $V_{DD} = 1.0V$ loaded by various numbers of inverters representing the gate loading on the bus; each register adds the same load as one inverter. Each additional load inverter consumes an extra 0.9fJ. Across an 8-bit bus one could expect that on average only half the lines would make a transition, but through proper design of the datapath, instruction set, and programs there can be a large amount of correlation on the bus lines that would reduce the number of transitions.

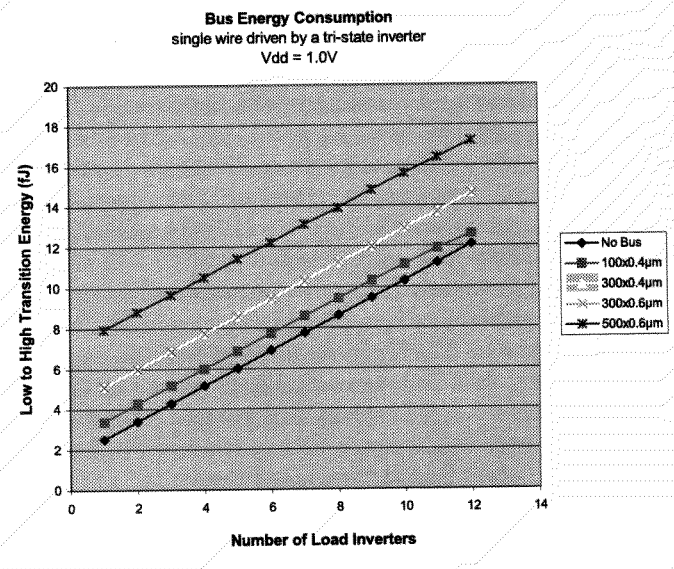

Figure 5. Energy consumption of a single wire of a bus making a low-to-high transition for various lengths and loads. The wire is driven by a minimum-size tri-state inverter, and the energy consumption of the load inverters is not included in the values.

Another key component of the sensor node processor is self-incrementing/decrementing registers, which are used in the program counter and the timers that control the polling and power cycling of the other components of the node, such as the sensors and communication units. At $V_{DD} = 1.0V$ a 10-bit self-incrementing register uses an average of approximately 46fJ per update and increment cycle.

A final major element of the sensor node is the static RAM used to store the program and sensor data. One particularly notable low-power 2kx16 SRAM [9] achieved

Table 3
Energy Consumption of Various Processor Instructions

Instruction	Basic Processor (Calculated)	Basic Processor (Simulated)	ARM8 [3]	Description
MOVI	0.047	0.38	4.3	Move immediate value into a register
MOV	0.047	0.59	5.1	Move from register to register
ADD	0.068	0.79	4.6	Add two registers and store in a third
LW	0.053	0.44	4.6	Load from memory location specified by one register into a second register
SW	0.047	0.46	4.6	Store from a register to the memory location specified by a second register
MUL			31.9	Multiply two registers and store in a third
B - fail	0.041	0.51	4.8	Conditional branch not taken
B - pass	0.047	0.60	4.8	Conditional branch taken
JMP	0.047	0.43	4.8	Unconditional jump

Note. Energy consumption is based on (a)calculations from the component consumption values, (b) the simulated consumption of a full processor using a standard cell library, and (c) simulations of an ARM8 core. Energy consumption of the SRAM is not included. Values are given in pJ per datapath bit.

an energy consumption of 9pJ per access at 1V in a 0.25μm dual-V_t CMOS process. This result was one to two orders of magnitude below most other low-power SRAMs in the literature [10]. A low-power processor design would follow the Harvard architecture and utilize two separate memories - one for the program and one for data. By separating the memory accesses, the design can take advantage of correlation on the address buses, as the program is executed in a fairly sequential manner and sensor data are streamed linearly to memory. Temporal correlation reduces the number of transitions on the buses and the decoder logic of the memory.

3.3 Energy Consumption of Particular Instructions

The energy consumption of the various high-level instructions is the sum of the consumption of the various blocks of the sensor node processor. Each execution cycle begins with an instruction fetch that consists of driving the program counter onto the program memory address bus, reading the instruction from memory, driving the program memory data bus, and loading the instruction register. The instruction is then decoded and the particular control lines are driven to execute the instruction. For an arithmetic logic unit (ALU) operation, the two source registers would then drive the two ALU busses and the ALU registers would latch the data. This set of registers prevents the ALU from calculating when it should not and thereby wasting energy. The ALU then performs the calculation and drives the result bus, which is loaded into the destination register. Meanwhile, the program counter has been incremented to be ready for the next cycle.

The second column of Table 3 lists the results of adding up the energy consumption estimated from the previous sections of each step for typical instructions. The base cost of every instruction for the instruction fetch and decode is 36fJ/bit. Although this shows the nearly ideal case for the technology used in this article, the third column shows simulated values for a simple datapath and control that follow many of the design techniques outlined above. However, it was implemented with a standard cell library that is not optimized for low power. Switching this same design to a library optimized for low energy consumption would reduce the energy consumption by a factor of 10 based on simulations of a set of 28 flip-flops and a few logic gates. For comparison, the energy consumption of the same instructions simulated on an ARM8 [11] core is presented in the fourth column. The original ARM8 simulations were performed at 3.3V, so to provide a better comparison the results have been scaled to 1V assuming pure capacitive energy consumption. Furthermore, because the ARM8 is a 32-bit processor, the results for all columns are presented per bit slice of the datapath. The values shown in Table 3 do not include the actual SRAM energy consumption, as this is relatively independent of the processor architecture yet is significantly dependent on the size of the block. For the SRAM mentioned above, an extra 0.6pJ/bit would be added to each instruction, plus a second 0.6pJ/bit for the

memory access instructions. As a final reference point, the eight-bit CoolRisc 81 core [12] achieves an average energy consumption of 2.8 pJ/instruction/bit on a 1.5V supply and has capabilities appropriate for a sensor node.

Signal-processing algorithms implemented directly in hardware almost always consume less energy than a software implementation; it is preferable for major functions needed for certain scenarios to be realized in specialized hardware attached to the processor. One such block is a finite impulse response (FIR) filter. Using distributed arithmetic, efficient FIR filters have been built that allow power to be scaled with the precision of the filtering. Amirtharajah [13] showed such a filter that consumed 300nW, 240nW, and 230nW with 8, 4, and 2 bit input samples, respectively, at 1.2k samples/s with a 1.1V supply. Another key signal-processing capability in sensor systems are fast Fourier transforms (FFT). Baas [14] has built an FFT that can perform a 1,024-point FFT for 3.1μJ with a 1.1V supply.

4. Power Dissipation of Sensor Interfaces

Fig. 6 shows the block diagram of a sensor for smart dust. It consists of a transducer, optional amplifier, and analog-to-digital converter (ADC). The output is sent to the smart dust microcontroller for processing, storage, or transmission. We predict power dissipation and energy consumption of the amplifier and ADC and apply these to relevant scenarios.

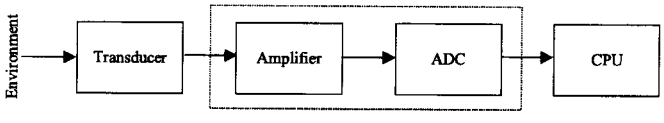

Figure 6. Sensor interface diagram. The transducer converts a physical signal from the environment to an analog electrical signal. Energy is consumed both in amplifying this electrical signal and in the analog-to-digital conversion (ADC) as enclosed by the dashed line. The digital signal can then be interpreted by the microcontroller.

4.1 Amplification of Sampled Signals

The transducer translates the physical input into an electrical signal. The amplitude of this signal depends strongly on the type of sensor and measurement. For example, typical temperature transducers have sensitivities around 1mV/°C. Many applications require resolutions around 1°C, and hence the amplifier and ADC must be capable of resolving $\Delta V = 1$ mV. Other types of sensors produce smaller signals. For example, typical acceleration transducers have sensitivities around 1mV/G ($G = 9.8m/s^2$, the earth's gravitational acceleration). For 1mG resolution the amplifier must now resolve a $\Delta V = 1\mu V$ input. As ADCs usually require larger inputs, this signal must first be amplified.

As in digital systems, the power dissipation of analog circuits is a function of operating speed and resolution. It is also proportional to clock frequency except at very high speeds. However, the power-resolution trade-off is more complicated. Unlike digital systems, where resolution is proportional to the bit width of buses and computational units, analog circuits carry all information in a single pair of wires. Because the peak signal amplitude is constrained by the supply voltage or transducer design, higher resolution requires resolving ever-smaller changes in voltage or current. For example, to resolve a 1V signal to 16 bits, changes of less than $1V/2^{16} = 15\mu V$ must be resolved. These small signals are easily corrupted by errors such as coupling from other circuits, for example, through the supply or adjacent wires. Although these effects can be minimized with careful design, another source of error—thermal noise, also called Johnson noise or Brownian motion—is fundamental to all circuit operation [15]. In circuits that are limited by thermal noise the power-resolution trade-off is exponential: increasing the resolution by 1 bit, and thus decreasing the noise voltage by one-half, requires a fourfold increase of operating currents, quadrupling power dissipation. This is remarkably different from digital systems, where adding a single bit to an already wide data path typically results in only a fractional power penalty. It is also not a function of technology: device scaling does not reduce the power dissipation of circuits that are limited by thermal noise. For example, if the resolution of an analog-to-digital converter dissipating 100mW is increased from 15 to 16 bits while keeping the supply voltage, operating speed, and all other parameters constant, the power dissipation increases to 400mW.

In practice, the increase would be less significant, as not all elements of the converter are actually dealing with the full accuracy signal. For example, the power increase of biases and digital support circuits would be marginal, and just like in a digital adder, only a small amount of additional circuitry may be required to convert the extra bit. If only half the power in the original converter is dissipated in systems that are limited by thermal noise, the power "only" increases from 100mW to 250mW plus the added power in the rest of the system. The latter is difficult to estimate because the relationship between power dissipation and resolution or accuracy of circuits where thermal noise is not the dominant error source is a strong function of architecture and technology, and general statements cannot be made.

Just as in digital systems, the power dissipation of analog circuits can be reduced significantly with power-down modes. However, additional complications arise. Thermal noise is a broadband signal, being essentially constant up to high frequencies. If the analog-to-digital converter samples this noise, the entire noise power is added to the signal. Because many sensors are interrogated at low rates, thermal noise can be reduced significantly by limiting its bandwidth before sampling. Unfortunately, when such a filter is turned on it takes a finite time to settle to the correct value that is proportional to the accuracy in bits.

Following on the above arguments, the energy per sample of a thermal noise limited circuit is:

$$E = \frac{NV_{DD}}{(\Delta V)^2}\frac{1}{F_{Amp}} \quad (1)$$

In this expression, N is the required accuracy in bits, V_{DD} the supply voltage, ΔV the resolution of the system, and F a factor that captures the peculiarities of a particular design. The value of F_{Amp} depends on the particular amplifier design and operating temperature, but for low power solutions it is about $10^{19}/V \cdot J$ at room temperature.

The resolution requirements depend strongly on the type of sensor and the application. Accelerometers and pressure sensors, for example, often produce signals in the μV range or less that must be resolved. Thermometers generate much larger outputs, and mV resolution usually suffices. The energy per sample for $V_{DD} = 1V$ and $N = 10$ bits is approximately 1μJ for $\Delta V = 1\mu V$ and 1pJ for $\Delta V = 1mV$. In the latter case the actual energy consumption would probably be significantly higher, dominated by considerations other than noise, such as minimum transistor size or the overhead for power cycling.

The resolution ΔV depends, of course, on the application requirements, but is also a strong function of the transducer. For example, the sensitivity of typical temperature sensors is 1 to 10 mV/°C. As for many applications a one degree resolution is adequate, amplifier power is likely negligible and 1nJ/sample or lower energy consumption should be achievable. Commercial low-power temperature sensors such as the LM20 [16] have one to two orders-of-magnitude higher energy consumption. Much of the overall power dissipation comes from the buffer driving off-chip circuits and from the need to address a wide market segment (e.g., a -55 to +130°C input range), constraints that may not apply to a particular application and highlight the importance of custom hardware for very low power systems.

Other sensors have much lower sensitivity. For example, typical MEMS accelerometer designs have sensitivities in the 0.1 to 10μV/mG range. Hence, for milli-G resolution, μV signals must be resolved, resulting in an energy per sample of the amplifier alone of around 1μJ. The ADXL202 two-axis accelerometer consumes about 10μJ per sample and axis for 5mG resolution [17].

4.2 ADC Energy Consumption

ADCs follow the same energy–performance relationship when dominated by thermal noise. However, the smallest signal in a converter is usually not set by the sensor output level, but by the supply and the number of bits. Assuming a somewhat optimistic full V_{DD} input range, the energy per sample is now:

$$E = \frac{2^{2N}}{V_{DD}}\frac{1}{F_{ADC}} \quad (2)$$

Unlike digital systems, analog circuit power dissipation increases when the supply voltage is reduced if the system is limited by thermal noise. Of course, the energy per sample cannot be made arbitrarily small by simply increasing V_{DD}—then the circuit is no longer limited by thermal noise and the equation is invalid. As has been pointed out already, the energy is a very strong function of the converter resolution: adding a single bit to a thermal noise limited converter quadruples the energy consumption.

Recently published power-efficient converters achieve $F_{ADC} = 5 \times 10^{13}/V \cdot J$. This translates into 4nJ for a 10-bit resolution conversion and 17μJ for a 16-bit one. For resolutions below 10 bits the expression gives optimistic results, as these converters are not limited by thermal noise. Most currently available converters operate from a 5V supply, with some 10-bit parts using 3V, but lower supply voltages likely result in higher power dissipation, particularly for the higher resolution parts. A more detailed study of analog-to-digital converters can be found in [18].

Overall energy consumption is therefore around 4nJ for a 10-bit A/D conversion for sensors with output in the mV range, such as, for example, thermometers or humidity sensors, and on the order of a μJ or more for sensors that produce outputs in the μV range, including pressure sensors, microphones, and accelerometers. Gyroscopes produce yet smaller outputs and have even higher power dissipation.

The empirical results mentioned in this section neglect the energy required to power-up the amplifier and ADC prior to taking a sensor sample. Most devices are designed to sample continuously; the power-up does not factor into the energy consumption when amortized over a large number of samples. In the sensor network paradigm, however, we may wish to sample on demand without paying large overhead. We are currently designing an analog interface to a 1mm^2 MEMS accelerometer with 8-bit resolution. Including power-up and power-down, the device will require 300pJ/sample independent of the sampling frequency below 10kHz. Furthermore, each bit of the digital signal can be determined seriatim: if a lower resolution sample is all that is required, 37.5pJ can be expended for each bit of precision starting from the MSB.

Because the power dissipation of analog thermal noise limited systems is a weak function of feature size but inversely proportional to supply voltage, technology scaling will result in increased overall power dissipation in high-resolution analog interfaces.

5. Communication

We are interested in computing the energy required to transmit a single bit of information between sensor nodes. Communication among sensor nodes can be accomplished by microscale components for both RF and optical media. The former is a higher energy solution requiring control of a small number of radio channels, whereas the latter requires less energy but intricate peer location control. We detail the current capabilities of each methodology, examine the theoretical limits, and estimate communication costs for the sensor network paradigm. RF communication could be considered as sensing and actuating with electromagnetic signals in the environment, but the operation is specific enough to justify a separate treatment.

5.1 Low-Power RF Communication

Current research seeks to develop new low-power radio solutions. For example, the LWIM project has a 1 mW radio goal in the 902-928 MHz band [4]. The picoRadio project [19] proposes power-dissipation levels below 100 μW, resulting in communication at 0.1 nJ/bit. As technology improves, the cost of radio transmission and reception will decrease, but will always be bounded by fundamental limits that still constrain operation in RF sensor networks.

It is difficult to generalize power consumption by communication systems, as many variables influence the performance of these systems. However, the fundamental limits are again related to thermal noise. For a receiver with a noise bandwidth B (roughly the symbol rate), the thermal noise power from the antenna is kTB where k is Boltzmann's constant and T is the temperature in Kelvin. The quality of the electronics in the receiver determines the ratio of actual noise performance to this theoretical limit, and is represented by the noise figure of the receiver, N_f. The strength of the received radio signal must exceed the noise by a factor governed by the downstream signal processing of the signal, and is given by SNR_{min}. Collecting all the factors, to successfully interpret an incoming radio message, the signal power received by the antenna must be greater than:

$$P_{r,min} = kTB \cdot N_f \cdot SNR_{min} \tag{3}$$

With clear line-of-sight, the power lost between the transmitter and the receiver is proportional to the square of the distance in wavelengths. In typical operation, however, reflections off the ground, buildings, trees, and so on, cause attenuation to be proportional to the fourth power of distance [20]. A 1GHz signal has a 30 cm wavelength; transmitting a distance of 300 m results in an attenuation of roughly 12 orders of magnitude. Traveling 3 km gives an attenuation of 16 orders of magnitude, taking a 1 W transmitted signal down to 0.1pW.

As alluded to above, the receive power P_r as a function of the transmit power P_t is given by:

$$P_r = \frac{P_t G_{ant}}{16\pi^2 (d/\lambda)^n} \tag{4}$$

where d/λ is the distance in wavelengths, n is the exponent of path loss ($n = 2$ in free space and $n = 2 - 7$ at ground level with an average of 4), and G_{ant} is the antenna gain. The antenna gain depends on the size scale. At 1 GHz, a quarter-wavelength antenna measures 7.5 cm long: suitable for cubic inch sensor nodes, but not for cubic millimeter scales. Unless specified otherwise, the examples below use carrier signals near 1 GHz.

As a concrete example, consider the GSM cellular telephone standard [21]. The noise bandwidth is roughly 200kHz (53 dB) for a 115kbps link. The receiver has about 8 times (9 dB) more noise than the thermal limit, and the downstream electronics need a signal-to-noise ratio of about 10 to achieve an adequately low bit error rate. In decibels relative to 1 mW (dBm), $kT = -174$ dBm at room temperature, and the receiver thus requires a sensitivity of: -174 + 53 + 9 + 10 = -102 dBm, less than 0.1pW. This received power minimum, however, requires the transmitting cellular phone to transmit Watts of RF power due to path loss. The GSM receiver uses 200mW; the transmitter uses 4W. At the 115 kbps data rate, GSM costs 2 μJ to receive a bit and 40μJ to send one. As illustrated in Fig. 7, there is an optimal number of hops for GSM-type transmission over given distance.

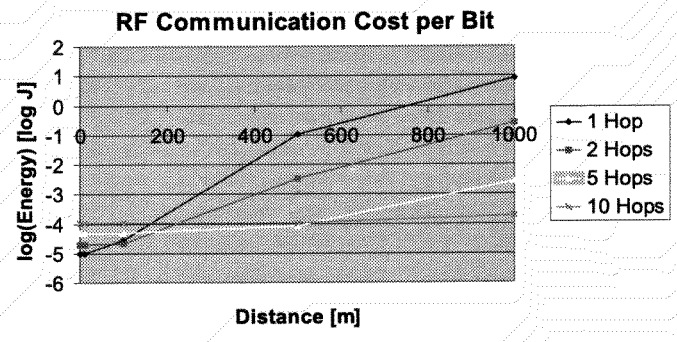

Figure 7. Energy requirement to transmit a single bit of data over various distances under the GSM specifications. Each hop contributes one transmission event with a power sufficient to provide 0.1pW at the total distance divided by the number of hops, and assumes that five nodes will receive the data with a cost of 2μJ each.

Cordless phones operate with similar data rates at less than one tenth the power, but with a range reduced to 10–100 meters. On the order of 1 μJ/bit is common. The Bluetooth radio [22] is designed for short range, 1Mbps communication in a household or office environment. Transmit power is 1 mW, but the total radio power is still roughly 100 mW regardless of transmit power, due to radio circuit overhead. Regardless, the Bluetooth standard remains promising for civilian sensor networks with short-range communication costs (over tens of meters) near 100 nJ/bit in the 2.4 GHz band.

The fundamental ground-to-ground communication limit at 1kbps over 100m is 1μJ/bit. The path loss from (4) is 122 dB. For a noise factor and an SNR_{min}, both of 10dB (easily achievable with current amplifiers and signal processing), and an attenuation of 30 dB (corresponding to a $B = 1kbps$), (3) implies that the receive power must be no less than -124 dBm. Hence, transmission of 1mW meets the specification and represents 1μJ/bit. For the same bit rate and distance from the ground into the air, to an unmanned aerial vehicle (UAV) for example, the path loss exponent may drop as low as 2. Assuming a dipole antenna gain of 6dB on the UAV, the path loss is reduced to 66 dB. Further, assuming ideal values of unity for N_f and SNR_{min}, the receiver sensitivity drops to -144 dBm. Now $P_t = 100$pW, or 0.1 pJ/bit!

How far can this trend continue? As the current limit to low-power RF communication, the Cassini satellite communicates over a distance of 1.5×10^9 km at 115 kbps on an 8GHz carrier. With a receiver sensitivity of -130 dBm, the satellite can transmit at 20 W, resulting

in 10^{-21} J/bit. At 1kbps (requiring approximately 20 dB less sensitivity), this system seemingly violates the fundamental limits discussed above. However, the receiver is supercooled to reduce T and increase the sensitivity.

In summary, current RF ground communication over the 0–50m range requires approximately 100nJ/bit; over 1–10km, 50μJ/bit is required. In the physical limit of ground-to-air communication, 0.1pJ/bit is achievable at room temperature.

5.2 Optical Communication

Optical communication provides several advantages for communication from small devices with limited power. First of all, optical radiators, such as laser diodes, are small and have large antenna gain, as compared to the relatively large antennas needed for microwave communication. Secondly, media access control can be implemented using spatial division multiple access (SDMA), which just relies on the positional differences between different transmitters as seen with an optical imager. This is a simpler and lower energy technique than those used in RF communication, such as frequency, time, or code division multiple access (FDMA, TDMA, CDMA). The energy limits of optical communication are substantially lower than their RF counterparts as the energy emission can be directional and also as the path loss exponent is $n = 2$ when a line-of-sight path is available. With the same variables as before:

$$P_r = \frac{P_t G_{ant} A_{Rx}}{4\pi d^2}, G_{ant} \cong \frac{4\pi}{\theta^2}$$

where A_{Rx} is the area of the receiver in m^2 and θ is the divergence of the transmitter in radians. Table 4 summarizes predicted performance over a range of distances.

In summary, optical communication over the 0–50m range requires approximately 20pJ/bit; over 1–10 km, 10nJ/bit is required. However, line of sight and some method of orienting the optical beam are paramount.

6. Scenarios

In this section we consider a set of revealing applications for sensor networks. Based on the daily energy budget and the expenditures required to perform each individual task, an estimation of the overall capabilities of the network can be calculated. Each scenario involves a disparate set of constraints and performance criteria; we expound on applications for building environment monitoring, earthquake safety testing, and vehicle tracking. Each case is developed with the assumption that radio communication will be used. As explained earlier, this communication is more power-hungry than optical methods but is technologically closer to fruition in a deployable sensor network.

The following quantities are used from the previous sections:

1. Energy: lithium battery—2 J/mm^3; solar power outdoors—0.3 mW/mm^2; solar power indoors—0.3 μW/mm^2
2. Computation: 1 pJ/instruction
3. Sampling (10-bit): temperature, humidity, light—4 nJ/sample; magnetic, acceleration—2 μJ/sample; limit—375 pJ/sample
4. Communication (short range): RF—100 nJ/bit; limit—0.1 pJ/bit

Each scenario will first be developed using current achievable energy figures, then refined to explore the limiting behaviour.

6.1 Building Environment Monitoring

We are interested in monitoring environmental conditions at many locations within an office building. Each worker in the building is assumed to have individual preferences for temperature, humidity, and light level. Actuators controlling these levels over entire rooms and/or single workspaces are available. The role of the sensor network in this context is to measure the conditions for each worker and to communicate this information to a centralized computer for actuation of heating units, humidifiers, fans, and lights to best accommodate the current denizens of the workspace. Some actuators would serve the entire room; others serve only a single workspace. In each room, a typical number of sensor nodes would be 10, all communicating their data to a radio receiver within the confines of the same room. In an office building, perhaps 100 such rooms are controlled.

Human beings can discern temperature levels to a Celsius degree. In an operating environment that certainly spans fewer than one hundred such degrees, a 10-bit temperature measurement system proves sufficient to measure to within 0.1 degrees. Similarly, 10-bit light and humidity measurements provide more than adequate resolution to perform as required. The parameter in this system is the tightness of control; is it necessary (or even preferable) to transmit all information every second, or does it suffice to make a modicum of transmissions from each sensor per minute? The sampling period maps nearly linearly to the battery life of the device.

This scenario exemplifies perhaps the simplest class of systems: there is no need for network discovery, as all sensor nodes are within one radio hop of their target, sampling and communication is carried out periodically at a low rate, and replacement of sensors/batteries can be done manually upon failure. Each room can be controlled approximately independently of others, though some building-level optimization schemes may be employed.

For sampling of the three sensors every dT seconds, we expend $3 \cdot 4$ nJ for each set of samples and $3 \cdot 10 \cdot 100nJ = 3\mu J$ for each 30-bit radio transmission. Computational energy is negligible; we need only alternate between reading the sensors and queuing this information to the radio. With an estimated 10 sensors in the room and a conservative radio channel bandwidth of 10kbps, each sensor can be time-provisioned 1 kbps, in great excess of the 30 bits of data to transmit in a $dT = 1$ second case. Daily time synchronization within a room may be necessary to reset drifting clocks, but this would be energetically negligible as well. Assuming the network is fully functional all hours of the day, the daily energy requirement per node is $1/\text{dT} \cdot 86400\text{s}(12nJ + 3\mu J) = 86400(3\mu J) = 300\text{mJ}$. With

Table 4
Energy Required for Optical Communication

	Transmitter			Receiver			Total
	P_{total}	P_t	Bit Rate	A_{Rx}	P_r	P_{total}	E/bit
(a) 5m	100μW	10μW	5Mbps	$0.1mm^2$	10nW	50μW	20pJ
(b) 5km	50mW	5mW	5Mbps	$1cm^2$	10nW	50μW/pixel	10nJ
(c) 500km	50mW	5mW	2Mbps	$1m^2$	10nW	50μW	25nJ

Note. All cases assume a MEMS beam-steering laser system [23] as a transmitter.

In (a) the optical beam is received by a Smart Dust photodetector.

In (b), the receiver is an array of 1000 CMOS imaging pixels [24].

In (c), the receiver is a satellite.

The optical power is generated with 10% efficiency.

The receiver power is based on measured data.

battery power, each node would have a lifetime of about 1 week/mm^3. With a 900-mm^2-scale solar cell array, this system could survive indefinitely on power scavenged from indoor lighting.

With the same parameters in the physical limit, we require $30 \cdot 0.1$ pJ= 3pJ for communication of a 30-bit message and about 1 nJ to sample the three sensors. This results in a daily energy requirement of 100 μJ and nodes that could survive on 0.3 mm^2 solar cells.

6.2 Earthquake Scenario

In earthquake monitoring, we are interested in measuring the frequency of oscillation at several locations throughout a structure. Changes in resonant frequency in building members are possible indications of stress changes and impending failure following an earthquake. A similar paradigm to the environment monitoring is required here; all sensors are hand placed and in known positions, and are required to behave in an interactive and deterministic manner. We envision a situation where the sensor network collectively measures the building properties on a daily schedule and automatically assesses damage following a measured earthquake of a significant magnitude.

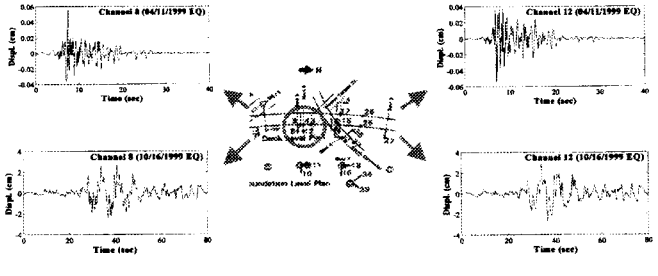

Figure 8. Response of Sylmar-I5/14 Interchange Bridge under temblor stress. The plots at the top are taken during a small temblor, and those at the bottom are taken during a large temblor. Plots on the right are for one span of the bridge; those on the left are for another. Responses show that the spans oscillate together during large deflections. Graphic courtesy of Steven Glaser.

Densely placed sensors are required to adequately measure structural response. If an insufficient number of nodes are placed on a structure, as shown in Fig. 8, the transition between decoupled movement and coupled movement of members as earthquake strength increases will not be seen. Sensor nodes in this scenario are equipped with three axes of accelerometers (to measure forces) and three axes of magnetometers (to measure displacement using the earth's field for reference), providing 6 degrees of measurement. Data from this system are sufficient to calculate the frequencies of interest, and will be transmitted via a radio channel either to another node or to a nearby receiver. All frequencies of interest lie below the hundreds of Hz. [25] provides a summary of structural damage survey methods.

A daily appraisal of the building requires each sensor to measure acceleration and displacement over a period of a few seconds. Again assuming a conservatively high 10 bits/degree of measurement, and a 100 Hz sampling rate, this requires 10 s $\cdot$ 100 Hz $\cdot$ (6 $\cdot 2\mu$J) = 10 mJ. There is a trade-off between on-board computation and communication at this point. The structural model at each sensor node is approximated by a 4-pole transfer function: either the node can transmit the entire sampled bitstream to a more powerful computer, or the entire system identification can be accomplished locally.

For local computation, the sampled 100 Hz signal is passed through an IIR filter before system identification is computed. In this paradigm, only a few bits identifying the system must be transmitted to the network. However, even if the entire stream of 10 bits $\cdot$ 6 sensors/sample $\cdot$ 1000 samples $\cdot$ 100 nJ = 6 mJ is transmitted, this is still less than the sampling energy. Either way, the energy required by this system is comparable to that of the office environment monitor even though sampling is done only once per day.

To achieve maximal lifetimes, the nodes will be run on a low duty cycle. The drawback of this approach is that nodes are likely dormant at the time of greatest interest, when an earthquake hits. There are two alternatives to solving this problem: self-awakening by the onboard accelerometer, or listening periodically for a beacon signal on the radio that could be provided by a single powerful transmitter for an entire municipality. A test is required every dT seconds, where dT is the maximum amount of delay we are willing to tolerate in the system. A self wake-up would require sampling from each accelerometer, an expenditure of 6μJ. A remote wake-up might require a 10 ms test reception on the radio, a cost of 1 mW $\cdot$ 10 ms = 10 μJ. Such a wake-up each second requires about 1 J per day, yielding 2 days/mm^3 of lithium battery.

In the physical limit, a single MSB sample of the accelerometer would suffice to detect an event, requiring only 37.5 pJ, or 3 μJ/day. The daily 1,000 measurements from each sensor requires 6,000 samples $\cdot$ (10 bits/sample $\cdot$ 0.1 pJ/bit + 375 pJ/sample) = 2 μJ. This is well within the scavengeable regime for a mm^2 sensor node.

6.3 Tracking Scenario

The third class of problems that we will consider has less definition than the previous two. The goal is to sense and track vehicles moving through a desert in a military setting. As a positive, the solar cells are operating in hours of full sunlight on most days, but the uncertainty of the network creates many complications. For example, the locations of the sensor nodes are unknown; in the extreme case, nodes may be deployed from an aircraft and scattered by the wind before settling on the ground. The nodes must establish an ad hoc network in highly variable transmission and reliability conditions.

Large vehicles can be detected with magnetometers, microphones, and accelerometers, the former with ranges of at least 10m and the latter with a longer range. Acoustic detection might be accomplished by recognizing the presence of certain characteristic frequency emissions from the vehicle requiring a sampling rate into the kHz range. Magnetic detection of large vehicles has been demonstrated at distances around 10m [26] from their distortion of the earth's field. With a simple 4-pole low-pass filter and threshold, vehicles can be detected at a 30 Hz sampling rate.

In full sunlight we expect 30% $\cdot$ 1mW/mm^2 for 8–16 hours per day. In this domain, it is instructive to think in terms of power instead of energy: we are afforded approximately 0.15 mW on average over the day-night cycle. This permits several tasks to be accomplished in a second at each node. In one second, a node can take 80 samples, transmit 150 bits, or make 150 million computations.

A dense network will profit from introducing heterogeneous modes of operation. Nodes will "specialize" as either sensor nodes or communication relays back to the central observer. Distributed algorithms exist for assigning relay nodes that ensure each sensor node is one hop away from the link of relays ([27], for example). Instead of expending energy to poll sensors, these nodes will use their radio receivers. Sensor nodes rarely, if ever, are required to receive data. With a 1 mW radio, the relay nodes can listen at a 15% duty cycle or can transmit up to 150 bits to the downstream relay node. As shown in Fig. 9, in the event of a failure, network rediscovery costs each node about C (the connectivity of the node) receive events and 1 transmission, a negligible fraction of the daily costs providing that failures are not excessively frequent.

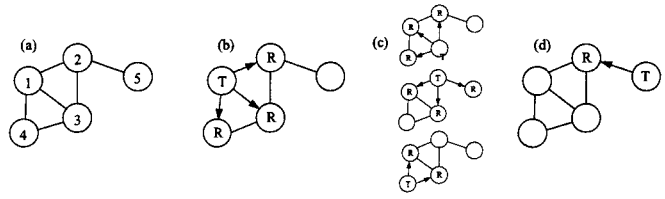

Figure 9. Illustration of a simple flooding network generator: (a) depicts node names and bidirectional broadcast connectivity by edges. In (b), node 1 initiates the network building with a single broadcast transmission; nodes 2, 3, and 4 must all receive the broadcast. In the second round (c), all nodes newly added to the network each transmit a message. In the final round (d), only node 5 is newly added and transmits a message. Each node transmits once and receives a message from every neighbour.

This scenario was recently implemented on a sensor network consisting of eight Rene nodes (as depicted in Fig. 3) to detect military vehicles passing along a desert roadway. Sensor nodes were dropped from a UAV, then configured themselves into an ad hoc network. Each node sampled its two-axis magnetometer at 30 Hz and communicated a 30-byte message to the network when it detected a vehicle. Once five observations were recorded, the network computed a least-squares solution of the vehicle's velocity. The estimated velocity is logged by the nodes and transmitted up to the UAV on a return trip. As implemented, the nodes drew approximately 10 mA from a 3V lithium coin cell. With an energy content of 540 mA $\cdot$ hr, the nodes could monitor the roadway for over two days. An optimized solution, assuming one hundred vehicles per day, would require the following energy values:

- 30 Hz $\cdot$ 86000 s/day $\cdot$ 375 pJ = 1 mJ for sampling
- 30 bytes/message $\cdot$ 0.8 pJ/byte $\cdot$ 100 messages/day = 2 nJ for communication to UAV
- 30 bytes/message $\cdot$ 800 pJ/byte $\cdot$ 100 messages/day = 3 μJ for communication among nodes
- 1000 instructions/sample $\cdot$ 1 pJ/instruction = 1 nJ for computation

The sampling energy dwarfs the other requirements. Solar cells on mm^2 surfaces are sufficient to achieve these design goals.

7. Conclusion

With current technology, sensor networks capable of monitoring environments are possible. Node control is based

primarily on energy requirements for sensor sampling, computation, and communication. In a typical implementation, a single sample requires 1 nJ – 1 μJ, a computation requires roughly 1 pJ, and each bit communicated by RF requires 100 nJ – 50 μJ. Cubic millimeter scale nodes scavenging photovoltaic energy will be budgeted about 0.3 J/day outdoors or 0.3 mJ/day indoors, sufficient for useful network operation. The most sensitive parameter is the distance between nodes, as for ground communication, the communication energy consumption rises with the fourth power of distance. It is hence advantageous to reduce internode distances as much as possible by increasing the density of the network.

Three sensor network scenarios were developed. In the environmental monitoring scenario, energy expenditure (and hence useful lifetime) was limited by communication costs. In earthquake damage monitoring, communication and sensing costs were of the same magnitude. In tracking applications, the sensor sampling overwhelms communication energy. To reduce communication requirements, distributed compression methods may be used without information transfer among nodes as proposed in [28].

It was determined that computation is significantly less expensive than either communication or sensing. Whenever possible, the brunt of the network's work should be done locally to minimize the communication costs of sending unnecessary information. Memory and hardware limitations may impose constraints on the types of computation that are possible on individual sensor nodes. It is essential that nodes be powered down whenever possible to conserve energy. More robust and efficient wake-up algorithms would facilitate the progression to the theoretical energy consumption levels.

Acknowledgements

This work was supported by the DARPA/MTO MEMS program and a Howard Hughes Doctoral fellowship. The authors thank Michael Scott and Carl Chang for valuable discussions.

References

[1] Taiwan Semiconductor Manufacturing Company, http://www.tsmc.com.

[2] J.M. Kahn, R.H. Katz, & K.S.J. Pister, Emerging challenges: Mobile networking for smart dust, *J. Commun. and Networks,2*(3), 2000, 188–196.

[3] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, & K.S.J. Pister, System architecture directions for networked sensors, *Operating Systems Review, 34*, 2000, 93–104.

[4] G. Asada, M. Dong, T.S. Lin, F. Newberg, G. Pottie, W.J. Kaiser, & H.O. Marcy, Wireless integrated network sensors: Low power systems on a chip, ESSCIRC '98, *Proc. 24th European Solid-State Circuits Conf.*, The Hague, Netherlands, September 1998, 9–16.

[5] J.H. Harreld, W. Dong, & B. Dunn, Ambient pressure synthesis of aerogel-like vanadium oxide and molybdenum oxide, *Materials Research Bulletin, 33*(4), 1998, 561–7.

[6] http://www.powercache.com.

[7] S. Meninger, T.O. Mur-Miranda, R. Amirtharajah, A. Chandrakasan, & J. Lang, Vibration-to-electric energy conversion, *Proc. 1999 Int. Symp. on Low Power Electronics and Design*, San Diego, CA, USA, 1999, 48–53.

[8] D. Teasdale, W. Lindsay, V. Milanovic, K.S.J. Pister, & C. Fernandez-Pello, Thrust and electrical power from solid propellant microrockets, *Proc. 14th Annual Int. Conf. on Microelectromechanical Systems (MEMS 2001)*, Interlaken, Switzerland, 2001, 606–610.

[9] K.W. Mai, T. Mori, B.S. Amrutur, R. Ho, B. Wilburn, M.A. Horowitz, I. Fukushi, T. Izawa, & S. Mitarai, Low-power SRAM design using half-swing pulse-mode techniques, *IEEE J. Solid-State Circuits, 33*(11), 1998, 1659–1671.

[10] M. Margala, Low-power SRAM circuit design, *Records of the 1999 IEEE Int. Workshop on Memory Technology, Design and Testing*, San Jose, CA, 1999, 115–122.

[11] P. Laramie, *Instruction level power analysis and low power design methodology of a core processor*, Master of Science Research Project, University of California at Berkeley, 1998.

[12] C. Piguet, J.-M. Masgonty, C. Arm, S. Durand, T. Schneider, F. Rampogna, C. Scarnera, C. Iseli, J.-P. Bardyn, R. Pache, & E. Dijkstra, Low-power design of 8-b embedded CoolRisc microcontroller cores, *IEEE J. Solid-State Circuits*, July 1997, 1067–1078.

[13] R. Amirtharaja, S. Meninger, J.O. Mur-Miranda, A. Chandrakasan, & J. Lang, A micropower programmable DSP powered using a MEMS-based vibration-to-electric energy converter, *Digest of the 2000 IEEE Int. Solid-State Cicuits Conf.*, San Francisco, CA, Feb. 7–9, 2000, 362–363.

[14] B.M. Baas, A low-power, high-performance, 1024-point FFT processor, *IEEE J. Solid-State Circuits, 34*(3), 1999, 380–387.

[15] P.R. Gray & R.G. Meyer, *Analysis and design of integrated circuits*, 3rd ed. (New York: Wiley, 1993), Chapter 11.

[16] LM20, 2.4V, 10mA Temperature Sensor, *National Semiconductors Datasheet*, http://www.national.com.

[17] ADXL202 Dual Axis Accelerometer, *Analog Devices Datasheet*, http://www.analog.com.

[18] R.H. Walden, Analog-to-digital converter survey analysis, *IEEE J. Selected Areas in Comm., 17*(4), 1999, 539–550.

[19] J. Rabaey, M.J. Ammer, J.L. da Silver Jr., D. Patel, & S. Roundy, PicoRadio supports ad hoc ultra-low power wireless networking, *IEEE Computer Magazine*, July 2000, 42–48.

[20] W.C.Y. Lee, *Mobile cellular telecommunications: Analog and digital systems* (New York: McGraw-Hill, 1995).

[21] European Telecommunications Standards Institute, *European digital cellular telecommunications system (phase 2): Radio transmission and reception (GSM 05.05)*, Technical report ETSI, December 1995, Sophia Antipolis, France.

[22] http://www.bluetooth.com Local area radio.

[23] M. Last & K.S.J. Pister, 2-DOF actuated micromirror designed for large DC deflection, *MOEMS '99*, Mainz, Germany, 1999.

[24] B. Leibowitz, B.E. Boser, & K.S.J. Pister, CMOS "smart pixel" for free-space optical communication, *Proc. SPIE - The International Society for Optical Engineering, 4306A (Electronic Imaging '01)*, San Jose, CA, 2001, 308–318.

[25] S.W. Doebling, C.R. Farrar, M.B. Prime, & D.W. Shevitz, Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: A literature review, *Los Alamos National Laboratory report LA-13070-MS*, Los Alamos, NM, April 1996.

[26] M.J. Caruso & L.S. Withanawasam, Vehicle detection and compass applications using AMR magnetic sensors, *Proc. 1999 Sensors Expo*, Baltimore, MD, 1999, 477–489.

[27] J.D. McLurkin, *Distributed algorithms for wireless sensor networks*, Master of Science Research Project, University of California at Berkeley, 1999.

[28] S.S. Pradhan & K. Ramchandran, Geometric proof of rate distortion function of Gaussian sources with side information at the decoder, *Proc. IEEE Int. Symp. on Information Theory (ISIT)*, 2000, 351.

Biographies

Lance Doherty received his B.Sc. and the Gold Medal in Engineering Physics from the University of Alberta, Canada, in 1998 and his M.Sc. in electrical engineering from the University of California, Berkeley in 2000. He is currently a Ph.D. candidate in the Berkeley Sensor and Actuator Center at UC Berkeley with research interests in wireless sensor networks, control systems, and microsystems.

Brett A. Warneke received his B.Sc. in electrical engineering from the California Institute of Technology in 1994 and his M.Sc. in electrical engineering from the University of California, Los Angeles in 1996. Since 1995 he has been a Howard Hughes doctoral fellow sponsored by HRL Laboratories in Malibu, CA. He is currently a Ph.D. candidate in electrical engineering and computer science at the Berkeley Sensor and Actuator Center, University of California, Berkeley. His research interests include ultra-low-power integrated circuits, CMOS micromachining, distributed sensor networks, and RF MEMS.

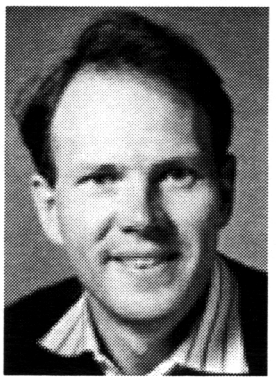

Bernhard E. Boser received the diploma in electrical engineering from the Swiss Federal Institute of Technology in 1984 and the M.Sc. and Ph.D. from Stanford University in 1985 and 1988. From 1988 he was a member of the technical staff in the Adaptive Systems Department at AT&T Bell Laboratories. In 1992 he joined the faculty of the Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley, where he also serves as a director of the Berkeley Sensor and Actuator Center. Dr. Boser's research is in the area of analog and mixed signal circuits, with special emphasis on micromechanical sensors and actuators.

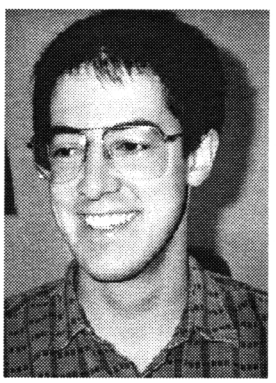

Kristofer S.J. Pister received his B.A. in applied physics from the University of California, San Diego in 1982, and his M.Sc. and Ph.D. in electrical engineering from University of California, Berkeley in 1989 and 1992. From 1992 to 1997 he was an Assistant Professor of Electrical Engineering at UCLA. In 1996 he joined the Department of Electrical Engineering and Computer Sciences at UC Berkeley as an Associate Professor.