# Challenges on Low-Power Platform Design for Real-World Wireless Sensing Applications

Pai H. Chou
Center for Embedded Computer Systems
University of California, Irvine, CA 92697-2625, USA

***Abstract*— Real-world wireless sensing applications pose a number of great challenges on low-power hardware/software platform designs, including a wide range of size, cost, power consumption, connectivity, performance, and flexibility requirements. Based on a classification of sensing functions, detection methods, timeliness of data, and characteristics of power supply, the platform may need to incorporate different features in order to operate in a low-power, energy-efficient manner. The design issues are highlighted in the context of a number of sensing systems ranging from high-performance, high-precision data acquisition wireless sensor node for civil engineering and an ultra-compact wireless sensor node for infant monitoring to a laser-based breast cancer detector.**

## I. Introduction

Real-world wireless sensing applications have received much attention recently due to their potential scale of deployment and impact on many interdisciplinary applications. A common challenge in many wireless sensor platforms is power. Because these wireless sensors are often deployed in deeply embedded or remote sites without access to a steady power supply, they must either operate entirely on batteries or harvest energy from the environment.

A sensing system contains much more than digital circuits. It also contains a variety of sensing devices, wireless interfaces, and even actuators and storage. In many cases, these peripheral devices dominate the system power consumption. The power problem is exacerbated by the increasing demand for even smaller, more compact form factors. Even though many electronic components have shrunk significantly, some components, particularly the battery and the antenna, are often the largest component in the system and also the most difficult to miniaturize. The battery, as the primary power supply, limits not only the energy capacity but also the instantaneous power level. The antenna size limits the gain on the RF power, which in turn affects the bit-error rate and communication range.

The challenges on low-power wireless sensor platforms cannot be addressed without first understanding the requirements of real-world applications. Sensing can be classified into

- event detection vs. data acquisition vs. data aggregation
- passive vs. active sensing
- data logging vs. real-time monitoring

Many techniques have been developed for low duty-cycle, passive sensing without hard real-time constraints. Duty cycling is the primary way to create power management opportunities, because the designer can then turn off the wireless module, the processor, or sensor devices temporarily to save power. However, many real-world applications demand high-rate data acquisition and real-time monitoring, possibly the most challenging combination to which many existing techniques do not apply. As a result, low power and energy efficiency must come from changes in the underlying architecture, rather than mainly from the policy level.

This paper first starts with case studies of real-world sensing applications. To address these challenges, we divide the problems into power consumption and power supply. We further divide power consumption into control, sensing, communication, and actuation. On the power supply side, we discuss techniques that enhance the efficiency of different types of power supplies, including batteries, solar panels, and energy storage. The platform as a whole will operate even more efficiently if the system designer can strike a balance between the consumption and supply.

## II. Case Studies

### A. Structural Health Monitoring

One popular application of wireless sensing is structural health monitoring in civil and mechanical systems. The purpose is to monitor and assess the structural integrity of buildings, airplanes, and many other structural bodies. Sensors have been installed for measuring vibration, strain, displacement, and other signals. Active sensing can be performed by measuring the bridge's response to artificial excitation, although most perform passive sensing by observing the bridge's natural movement or displacement without additional stimuli. Here, frequency-domain analysis requires around 200–1000 samples per second of tri-axial acceleration [1], while time-domain analysis requires an order of magnitude higher. High duty cycling means modern batteries cannot be expected to last for a long time especially if it must transmit data at a sustainable rate in real time. In case the sensor has no access to an AC power source, energy harvesting from solar, wind, and other ambient sources must be implemented.

### B. Wearable sensors

Sensor networks can be deployed not only in building structures but also on groups of humans, for whom the form factor is a primary concern. One application is to monitor the spontaneous motion of pre-term infants in terms of tri-axial acceleration of their limbs [2]. The purpose is to assess their growth in bones and muscle strength in response to assisted exercises. As these infants in incubators are already heavily wired for monitoring vital signs such as heart beat and temperature, wireless sensing is highly desirable if not essential. Another application is in interactive dance [3]. Here, the sensor nodes deployed throughout a dancer's body transmit not only acceleration and orientation data but also heartbeat and the angle of joint bending in real time to a host computer. It then synthesizes stage effects such as music, lighting, animation, sound, and even robotic motions according to a script.

To be wearable by a pre-term infant, such a wireless sensor node must occupy less than 1cm$^3$ and weigh under 3 grams. It is extremely challenging to make the wireless part consume power at a level friendly to batteries of this size. Most existing ultra-compact sensor nodes either fail to miniaturize their batteries and antennas and are thus not truly compact, or they operate at very power-inefficient points.

### C. Content Measurement by Active Sensing

Several classes of sensing applications involve the determination of certain content in gas, liquid, solids, and many objects. For humans, content measurement can determine the amount of fat and water in the issue, and this can aid the diagnosis of health problems and even detect certain types of cancer [4]. Content measurement often involves *active sensing*, i.e., emitting a signal in the form of light, electric current, electromagnetic wave, or ultrasound, and measuring the subject's response in terms of voltage drop, attenuation, phase shift, time delay of echo, or interference pattern.

Low power designs are challenging for active sensing systems because they must perform actuation in addition to sensing. Because of the physical limitation on these power levels, it is difficult to further reduce the instantaneous *power* levels of these systems much further, but the total *energy* consumption can be improved. If the actual detection time is reduced, then it means that the sensors, actuators, and amplifiers are turned on for a short duration and turned off as soon as possible. This will lead to much lower total energy consumption.

## III. Reducing Power Consumption

Power efficiency can be achieved on the consumption side by improving the sensing and actuation devices, system architecture, and RF communication. Higher level integration of these subsystems on a chip will also improve the power efficiency, and currently a number of integrated MCU+RF (microcontroller unit plus radio frequency) components in CMOS are available. However, a main challenge is the integration of technologies at very different feature sizes, especially with MEMS sensing devices. The design of sensors and actuators is domain specific and is outside the scope of this paper. This section discusses system architectures and RF communication by drawing lessons from existing designs.

### A. System Architectures

Wireless sensor platforms can be roughly divided into sensor *nodes* and sensor *computers*. The former, such as Mica motes [5], typically uses an 8-bit or 16-bit MCU with limited memory and thin runtime layer, if any. The latter, such as the Stargate, typically uses a 32-bit MCU capable of running embedded operating systems such as Linux or another RTOS, and the architecture is modeled after general purpose computers. The difference in power consumption is at least two orders of magnitude higher, and this gap will continue to exist. This paper focuses on sensor nodes.

Currently, most sensor nodes are based on single-MCU architectures, as depicted in Fig. 1(a). These 8/16-bit MCUs contain a small RAM and one or more flash memory banks for bootstrapping ROM, for the firmware, and possibly for data. To be low power, they support a much smaller instruction set, and the narrower bit width means fewer switching activities per operation. In addition, due to high leakage, RAM is usually kept on-chip and small, from under 100 bytes to 4–8K bytes. These resource constraints make it extremely difficult to support abstractions that programmers come to expect and demand from operating systems, including concurrency, timing control, memory management, and program update.

*1) Runtime Support:* Currently, virtually all sensor nodes implement runtime abstractions in software, just in very limited forms. For instance, TinyOS is designed to squeeze out overhead through tight cohesion with the application (i.e., compiled together as a single executable) and minimizing task switching overhead by supporting cooperative, event-triggered messages and hardware interrupts. It is modularized so that memory management and in-field program update can be configured and added as needed. Although synthesis of (Java) virtual machines has been proposed to reduce memory usage, the execution time increases by over two orders of magnitude, and thus the total energy increases. As another alternative, host-assisted scripting engines have been proposed [6]. Unlike Java-style virtual machines, applications written in scripts specify mainly the control flow while spending most of the time in native code. They are much smaller than the corresponding compiled version and thus require less energy to store and transmit. Script interpretation overhead can be further reduced by offloading complex parts to the host computer, including memory management and certain scheduling and networking tasks. This has been shown to be both faster and lower power than compiled code in several cases.

*2) Multi-core Architecture:* Alternative software abstractions are unlikely to reduce power much further without changing the underlying system architecture. One candidate architecture for further lowering the power is one that integrates multiple MCU cores on the same chip. It has been shown that for structural health monitoring, distributing the tasks onto two or more MCUs can result in better performance, timing determinism, and low power at the same time [1]. For instance, one MCU can handle sensing and power management, and the other MCU can be dedicated to networking tasks without interfering with each other. When integrated on the same chip, additional cores can be dedicated to other tasks. Multiple cores can be voltage/frequency-scaled to run at a much more efficient point than a single core. They can also be more power manageable, as idle cores can be powered off.

A multi-core MCU architecture can use more hardware to provide various runtime services, including true concurrency, that single-core MCUs would have to implement as software abstractions. One or more cores can be dedicated to real-time scheduling, power management, and possibly firmware update for other cores. One master core controls the other slave cores to perform sensing and communication tasks. Depending on the duty cycling requirements, different cores may be dedicated to different sets of I/O devices running at different rates and voltages. Although inter-core communication can incur additional overhead, having multiple smaller RAM modules will allow power management at a finer-grain level than previous MCUs that use one larger RAM. This can also lead to lower power at the system level [1].

*3) Lookup-based computation:* Recent development in larger flash memory components is opening up new opportunities in power management. Traditionally, many lower-power MCUs implement only simple instructions, and operations such as multiplication, division, and other tasks must be emulated in software. With a large flash memory, it will be possible to make trade-offs between computation and storage. That is, results for many arbitrarily complex operations can actually be computed before deployment, so that at run time it can be converted into primarily a lookup with little or no computation. For instance, 8-bit multiplication tables can be stored in a 64K-entry table, or in a 16-bit address space. This approach also can free up precious RAM for both code and data.

### B. Radio Communication

Radios have been miniaturized and integrated onto several MCUs, including Nordic nRF24E1 [7], Chipcon (now TI) CC2430 [8], Microchip rfPIC12F675F [9], and others. However, higher-level integration does not help reduce power in the radio transceiver, which remains the single largest power consumer in many sensor nodes. In addition, the antenna gain can make a tremendous difference and can trade size for transceiver power. Additional power savings can be achieved though either better MAC protocols or hardware MAC. The challenge will be to select the right set of features that will enable improvements in other areas.

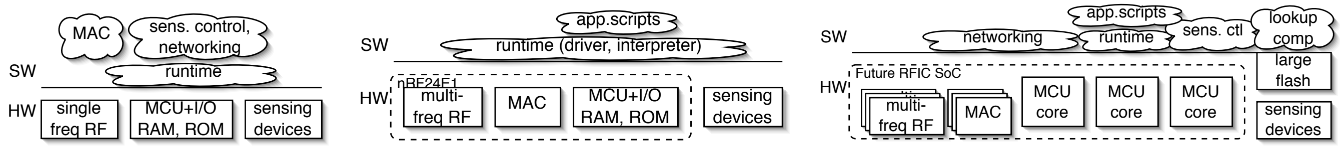


(a) (b) (c)

Fig. 1. System hardware/software architectures of sensor nodes: (a) Mica-2 class running TinyOS on an MCU, which controls a single-frequency RF with software MAC; (b) Eco node with nRF24E1 (integrated MCU, multi-channel RF, and simple hardware MAC), running a scripting engine; (c) proposed low-power architecture with multiple cores dedicated to networking, sensing control, and runtime; multiple RF, hardware coding support, and large flash.

*1) Antenna Efficiency:* The antenna can make a great difference in power efficiency. At a given level of transmission power, the RF power efficiency (and therefore bit-error rate) can vary by several orders of magnitude simply by changing the antenna type. For example, conventional quarter-wavelength antennas for the 2.4GHz band can have a gain of 6–8dBi, but such an antenna is significantly larger (10–100x the volume) than the entire sensor node itself. The antenna size also depends on the wavelength of the radio frequency. The Mica and $\mu$Part [10] sensors use lower ISM and UHF bands, respectively, and they both use wires several times the length of the sensor node itself as their antenna. However, for ultra-compact sensor nodes where such a protruding antenna is not an option, chip antennas may be the only option, though they have a gain ranging from 0–2dBi [11]. A main challenge therefore is to design new miniature antennas with significantly higher gain without increasing transmission power.

*2) Media Access Controller:* MAC protocols for sensor networks have been some hybrid between CSMA and TDMA, under the assumption of one frequency channel. CSMA has the advantage of low overhead when the utilization is low (i.e., the probability of collision is low), and it is the basis for protocols like ZigBee. However, during high utilization, TDMA is more power efficient, because it reduces or eliminates collisions by limiting communication to take place during assigned time slots only. In a hybrid scheme, TDMA may define the time slots in which CSMA is used to arbitrate the multiple accesses. To handle both low-duty and high-duty communication, Z-MAC enables slot stealing adaptively so that the protocol behaves like TDMA during high contention, and CSMA during low contention [12]. Today's single-frequency transceivers have low hardware complexity, and their MAC protocol is usually simple enough that it can be implemented in software, but they are limited to modem-like speeds and is practical mainly for low duty-cycle communications.

Further power reduction can be achieved by increasing the data rate of the radio and reducing collision. The former will require hardware MAC that turns on the radio for a short duration, and it may need to be used in conjunction with low-power listening [13]. It may also perform error-correctable coding to be more tolerant to bit errors. Making use of additional frequency channels is another way to reduce collision, and this must be supported by both the MAC and the transceiver. One key difference between TDMA and FDMA is that TDMA nodes can snoop outside their time slots, but FDMA nodes cannot receive outside their selected frequency. To address this problem, a second receiver can be added for this purpose. Currently one such radio is the nRF2401, also used in the Eco node. It supports 125 frequencies channels in 1MHz increments in the 2.4GHz band. It is also integrated as part of the nRF24E1 MCU on the same chip, making it suitable for these advanced compact sensor node designs. This makes available much more bandwidth and can be effective in reducing collision even in high bandwidth communications. At the same time, it gives users the option to access the radio either through

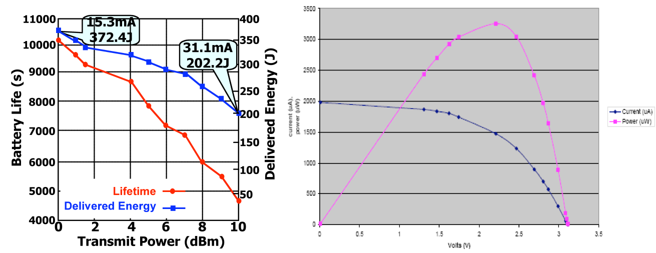


(a) (b)

Fig. 2. (a) Tx power vs. battery life, (b) max. power point of a solar panel.

its ShockBurst MAC or in direct mode without MAC. Although slightly higher in hardware and software complexity, the efficient use of bandwidth and time will be able to achieve significantly higher power efficiency at the system level.

For future transceivers in sensing systems, the challenge will be to find a balance between MAC complexity and power. One candidate architecture that combines these features is shown in Fig. 1(c), which integrates multiple cores that can be dedicated to different tasks, multiple multi-frequency transceivers with significantly broader bandwidth, and hardware MAC for accelerated data rate and error-correctable coding.

## IV. Considerations for Power Supply

Balancing the power consumption and supply will be critical to designing a power-efficient sensor platform. As most sensor nodes are powered by batteries with limited energy capacities, they will need to be replaced or recharged either manually or by energy harvesting sources. The key issues here are battery awareness and impedance matching between supply and load. The improvements made on the supply side will benefit both passive and active sensors.

### A. Batteries

Miniaturization poses a particularly challenging problem to the power supply, because it is difficult to shrink batteries further while still providing sufficient power. For instance, in the Eco node, which fits in a 1cm$^3$ form factor, the 40mAH lithium-polymer battery takes up more than half of the volume. Even though for infant monitoring applications the energy capacity is sufficient with duty cycling, the greater problem is the power.

*1) Rate Capacity and Recovery Effects:* As a non-ideal power source, each battery has a nominal *rated current* based on which the specified battery capacity is computed. When the current draw is higher than this rated current, the battery efficiency decreases, and this is called the *rate capacity* effect. Given sufficient time, the battery can recover part of its lost efficiency, and this is called *rate recovery* effect [14].

*2) Battery Effects vs. RF Power:* For miniature ($<$1cm$^3$) sensor nodes that perform real-time monitoring at non-trivial rates, the power consumed by the RF has possibly the greatest effect on battery efficiency. Lithium coin batteries of this size are rated for $\sim$0.2mA, while the CC1000 used in the Mica motes draws 26.7mA (3V) at 10dBm. That is over two orders of magnitude above the rated current. As a result, the battery operates at only 10–25% the specified capacity, depending on the duty cycle. If the RF transmission power is set to $-$20dBm, then the transceiver draws 5.3mA, and the battery gains much more efficiency, but at the same time the transmission range decreases and the bit-error rate increases. The need to re-transmit or relay packets may lead to even higher energy that may more than offset any recovered efficiency. Based on measurement, the battery efficiency of the 560mAh CR2354 lithium battery at 10dBm is 45% lower than at 0dBm; however, the bit error rate at 0dBm varies between 26–300% higher than 10dBm, depending on the distance and packet size. When the battery effect is combined, then the number of packets that can be delivered successfully can vary by 400% [15].

The problem is further exacerbated by constraints on the antenna size, because smaller antennas have much lower gain, as discussed in §III-B.1. One has the option of using a higher gain antenna to achieve the same effect as increasing RF amplifier power. However, the antenna size usually increases exponentially, and it may be more effective to use the volume on a larger battery instead of an antenna, as a larger battery will have a larger rated current, which then enables higher transmission power with much less rate capacity effect.

A power-efficient platform for wireless sensing should optimize for the trade-offs among the battery capacity, the bit error rate, the antenna size, the communication frequency, and the packet size.

### B. Ambient Power Sources

An increasing number of sensor nodes starts harvesting energy from the ambient sources, including the sun, wind, water flow, heat differential, and vibration. Such a system is usually built with an energy storage such as a rechargeable battery [16]. However, batteries have a limited number of recharge cycles. More recently, supercapacitors are starting to be considered in these systems to minimize discharging from the battery [17]. However, sensor systems that leverage ambient power sources fail to perform *maximum power point tracking* (MPPT) of these generators. In photovoltaic cells and wind mills, the impedance of these sources is a function of the ambient power level and the current being drawn. Harvested power is maximized if the supply and consumer sides are impedance matched. Although MPPT circuits and capacitor charging circuits exist separately, combining them directly will not work. This is because a supercapacitor can appear as a short circuit to a solar panel, causing the charger to limit the current at the other very inefficient point. To address this problem, two designs have been proposed. One is a digitally controlled feed-forward PFM converter that performs MPPT based on open-circuit voltage [18]. The other is based on a PWM switching regulator controlled by entirely analog MPPT circuitry, so that the MPPT charging can continue autonomously [19].

A challenge here is the integration of not only power consumers but also all the circuitry needed on the power supply side, including supercapacitor-like structures and MPPT circuitry on the same chip or in the same package.

## V. Conclusions

Real-world wireless sensing applications are posing many exciting new challenges on low-power platform designs. Higher level integration is only part of the answer; understanding all the issues involved and optimizing the design together as an entire system will be key to success for the next generation of miniature sensor platforms. We call for the tighter integration of multiple MCU cores for finer-grain power management and better timing determinism; significantly larger flash memory for lookup-based computation; multi-channel RF transceivers for better bandwidth utilization; hardware MAC for faster (lower total energy) communication; hardware for error-correctable coding to reduce the number of re-transmissions; digital or analog circuitry for MPPT; and possibly integration of supercapacitor-based energy storage device on the same chip or same package. In addition, improvements in miniature antenna designs will enable the same RF maintain the same performance level but at a much lower power level. This in turn is beneficial to the battery efficiency. Only then will the sensor platform design take the next quantum leap.

*Acknowledgments:* This work is sponsored by NSF CCR-0205712, CMS-0509018, and NSF CAREER CNS-0448668.

## References

[1] H. Chung, C. Park, Q. Xie, P. H. Chou, and M. Shinozuka, "DuraNode: wireless network sensor for structural safety monitoring," in *SPIE Nondestructive Detection and Measurement for Homeland Security III Conference*, Mar. 2005.

[2] C. Park, J. Liu, and P. H. Chou, "Eco: an ultra-compact low-power wireless sensor node for real-time motion monitoring," in *Proc. IPSN*, May 2005.

[3] C. Park, P. H. Chou, and Y. Sun, "A wearable wireless sensor platform for interactive art performance," in *Proc. PerCom*, March 13-17 2006.

[4] K.-S. No and P. H. Chou, "Mini-FDPM and heterodyne Mini-FDPM: Handheld non-invasive breast cancer detectors based on frequency domain photon migration," *IEEE Trans. Circuits and Systems*, no. 11, Nov. 2005.

[5] "MPR400/410/420 MICA2 Mote," http://www.xbow.com/Products/productsdetails.aspx?sid=72.

[6] J. Hahn, Q. Xie, and P. H. Chou, "Rappit: A framework for the synthesis of host-assisted light-weight scripting engines for adaptive embedded systems," in *Proc. (CODES+ISSS)*, Sep. 2005, pp. 315–320.

[7] "nRF24E1 transceiver / MCU / ADC," http://www.nvlsi.com/index.cfm?obj=product&act=display&pro=79.

[8] Chipcon, "CC2430 Product Information," http://www.chipcon.com/index.cfm?kat_id=2&subkat_id=12&dok_id=240, 2005.

[9] Microchip, "rfPIC12F675F," http://ww1.microchip.com/downloads/en/devicedoc/70091a.pdf, 2003.

[10] M. Beigl, C. Decker, A. Krohn, T. Riedel, and T. Zimmer, "μParts: Low cost sensor networks at scale," in *Ubicomp 2005*, 2005. [Online]. Available: http://particle.teco.edu/upart/whatareuparts/krohn-muparts-demo.pdf

[11] "AN3216 multilayer chip antenna for 2.4GHz wireless communication," http://www.rainsun.com/Products/LTCC/AN3216/AN3216.pdf, 2004.

[12] I. Rhee, A. Warrier, M. Aia, and J. Min, "Z-MAC: a hybrid MAC for wireless sensor networks," in *SenSys*, 2005, pp. 90–101.

[13] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access," in *SenSys*, Nov. 2004.

[14] K. Lahiri, A. Raghunathan, S. Dey, and D. Panigrahi, "Battery-driven system design: A new frontier in low power design," in *Proc. Intl. Conf. on VLSI Design/ASP-DAC*, January 2002, pp. 261–267. [Online]. Available: http://esdat.ucsd.edu/~klahiri/papers/vlsi02.pdf

[15] C. Park, K. Lahiri, and A. Raghunathan, "Battery discharge charateristics of wireless sensor nodes: An experimental analysis," in *IEEE Conf. on Sensor and Ad-hoc Communications and Networks (SECON)*, Sep. 2005.

[16] A. Kansal, D. Potter, and M. B. Srivastava, "Performance aware tasking for environmentally powered sensor networks," in *SIGMETRICS/Performance'04*, 2004, pp. 223–234.

[17] X. Jiang, J. Polastre, and D. Culler, "Perpetual environmentally powered sensor networks," in *Proc. IPSN*, May 2005.

[18] F. Simjee and P. H. Chou, "Everlast: Long-life, supercapacitor-operated wireless sensor node," University of California, Irvine, Tech. Rep., 2005.

[19] C. Park and P. H. Chou, "AmbiMax: Efficient, autonomous energy harvesting platform for multiple-supply wireless sensor nodes," University of California, Irvine, Tech. Rep., 2005.