

# Design of a Wireless Sensor Network Platform for Detecting Rare, Random, and Ephemeral Events

Prabal Dutta<sup>†‡</sup>, Mike Grimmer<sup>\*</sup>, Anish Arora<sup>†</sup>, Steven Bibyk<sup>†</sup>, and David Culler<sup>‡</sup>

<sup>†</sup>The Ohio State University  
Columbus, Ohio 43210  
{dutta.4, arora.9, bibyk.1}@osu.edu

<sup>\*</sup>Crossbow Technology  
San Jose, California 95134  
mgrimmer@xbow.com

<sup>‡</sup>University of California, Berkeley  
Berkeley, California 94720  
{prabal, culler}@cs.berkeley.edu

## Abstract—

We present the design of the eXtreme Scale Mote, a new sensor network platform for reliably detecting and classifying, and quickly reporting, rare, random, and ephemeral events in a large-scale, long-lived, and retaskable manner. This new mote was designed for the ExScal project which seeks to demonstrate a 10,000 node network capable of discriminating civilians, soldiers and vehicles, spread out over a 10km<sup>2</sup> area, with node lifetimes approaching 1,000 hours of continuous operation on two AA alkaline batteries. This application posed unique functional, usability, scalability, and robustness requirements which could not be met with existing hardware, and therefore motivated the design of a new platform. The detection and classification requirements are met using infrared, magnetic, and acoustic sensors. The infrared and acoustic sensors are designed for low-power continuous operation and include asynchronous processor wakeup circuitry. The usability and scalability requirements are met by minimizing the frequency and cost of human-in-the-loop operations during node deployment, activation, and verification through improvements in the user interface, packaging, and configurability of the platform. Recoverable retasking is addressed by using a grenade timer that periodically forces a system reset. The key contributions of this work are a specific design point and general design methods for building sensor network platforms to detect exceptional events.

## I. INTRODUCTION

Wireless sensor networks hold great promise as an enabling technology for a variety of applications, including data collection and event detection. While wireless multi-hop data collection applications have achieved operational lifetimes on the order of a year or more, we are unaware of lifetimes exceeding a few days or weeks for wireless multi-hop detection and classification of civilians, soldiers, and vehicles – our motivating application – using similar amounts of energy. A principal reason for this short lifetime is event detection requires sensors to be vigilant most of the time, whereas data collection generally allows sensors to be turned off. On the other hand, data collection requires frequent messaging to report measurements but event detection requires reporting only when an event occurs.

Considerable efforts have been invested toward lowering the power consumption of the networking and communications layers. These improvements benefit all applications but they benefit data collection applications more due to their greater levels of messaging. In contrast, far fewer ideas for architecting power-efficient event detection sensing and signal processing systems have been proposed. Our work was partially influenced by [1], [3], [4], which collectively identify tradeoffs in detection and communications, ideas for network management, importance of signal processing hierarchy, scalable network architectures, importance of aggressive power management at all levels, and differences between military surveillance and environmental monitoring which are representative of event detection and data collection, respectively.

We believe passive vigilance, which refers to a vague awareness of and limited attention to ambient conditions similar to what people

experience when sleeping, holds the key to extending the lifetime of event detection applications. An abrupt change in the ambient conditions due to a television turning off, a door opening, or another person moving nearby is often enough to wake up a person who might then exhibit more active awareness or even actions like answering the telephone, perhaps aided by a shot of adrenaline. During periods of passive vigilance, little energy is expended to processing ambient information and whatever processing that does occur happens in a crude manner. The key observation is that passive vigilance triggers active vigilance, which consumes more energy but is also capable of better discriminating real events from noise. This pattern, sometimes called the energy-quality tradeoff, is applicable to event detection using sensor networks: crude, low-energy detections in hardware triggering accurate, high-energy detection and classification algorithms, and communications. We emphasize our interest in events which occur rarely, happen at random times, and have short durations.

In addition to the challenging requirements of exceptional event detection, a key requirement of our application is to demonstrate a 10,000 node sensor network, spread out over a 10km<sup>2</sup> area, with node lifetimes approaching 1,000 hours of continuous operation on two AA alkaline batteries. The batteries fix the energy budget at 6Whr and average power consumption to 6mW. Unifying these diverse requirements, the central question underlying our work is:

*How does one engineer a wireless sensor network platform to reliably detect and classify, and quickly report, rare, random, and ephemeral events in a large-scale, long-lived, and retaskable manner?*

This question focused our research on reducing human-in-the loop operations, exploring low-power passive vigilance modes, investigating the conditions under which recoverable multi-hop wireless reprogramming is possible, and enhancing usability – an unusual combination considering the functional requirements of detection and classification. Our approach is independent of any particular intrusion detection application and may be applicable to an entire class of exceptional event detection problems including earthquake monitoring, traffic surveillance, and shooter localization. The remainder of this paper addresses in turn the various aspects of our central question.

## II. SCALABILITY

Immense scale is a distinguishing characteristic of the sensor networks vision. The DARPA Networked Embedded Systems Technology program, for example, envisions networks consisting of 100 to 100,000 simple computing nodes. However, networks approaching only the low end of this scale have been demonstrated to date [5], [6]. Our own experiences indicate that at large scales, logistical issues dominate deployment and configuration time, and that reducing the impact of these factors is critical to realizing the full scope of the sensor networks vision. Tennenhouse has suggested that “human

in the loop computing has its limits” and that “shrinking time constants and sheer numbers demand research into proactive modes of operation in which humans are *above* the loop”[7]. As a consequence, a key hardware design goal was to ensure that *human-in-the-loop* operations during deployment-time could be accomplished quickly and efficiently. More generally, scale considerations pervaded every aspect of the XSM design and scalability’s early presentation is a testament to the importance of considering it early, and often, in the design lifecycle.

The XSM, shown in Figure 1, provides three hardware features to support efficient deployment and configuration. First, the nodes are *highly-integrated* which helps minimize the amount of time required to place and orient them (and lowers their cost). A directional logo, which can optionally be colored, is embossed on the top of each node and makes orienting a node during deployment easy. Second, only *one-touch* is required to activate a node and initiate the process of discovering and joining the network. Third, after a node has been activated, it takes only *one-listen* of the buzzer to verify that it is operational. *One-glance* would have been nice but the LEDs are not exposed, but even if they had been, LEDs are not easily visible in bright sunlight. We found in bright sunlight that it was difficult to distinguish an LED that was turned on from one that was turned off, particularly for the small, low-power green, blue, and yellow LEDs we tested. Blinking red LEDs appeared most easily distinguishable but not without pausing for a second or two to see several on-off cycles. An unplanned benefit of using a telescoping antenna is that it can help speed up deployment by acting as a long handle which allows a node to be deployed on the ground without crouching and occasionally without even breaking a walking stride.



Fig. 1. The eXtreme Scale Mote. The XSM circuit board has a 3”x3” footprint and the enclosure has dimensions of 3.5”x3.5”x2.5”. The *one-touch* input (on/off switch) and *one-listen* output (buzzer) are mounted on the base next to the batteries. The XSM *platform* integrates an Atmel ATmega128L microcontroller, a Chipcon CC1000 radio operating at 433MHz, a 4Mbit serial flash memory, infrared, magnetic, acoustic, photo and temperature sensors, weatherproof packaging, a bootloader, and a “Golden Image” program that incorporates system management, network programming, and low-power listening.

Configuration is the other *human-in-the-loop* operation that can

be very time consuming. Our earlier experiences with manually configuring parameters on much smaller scale networks – on the order of 100 nodes – convinced us that any sort of individualized attention would be untenable with 10,000 nodes. On the other hand, the sheer size of the sensor field guarantees a heterogeneous environment. To support a zero-configuration *drop-and-go* approach, we provided digital potentiometers to adjust analog circuit parameters including amplifier gains and offsets, filter cutoffs, and detection thresholds under program control. Even though much of the work in realizing *drop-and-go* rests in software, appropriate hardware feedback and control is required.

### III. PASSIVE VIGILANCE

The functional problem the XSM addresses is hardware support for the detection and classification of *rare*, *random*, and *ephemeral* events in a *long-lived* manner. Rare implies a disproportionate amount of time is spent monitoring noise, which is where our efforts to reduce power should be targeted. Random implies continuous sensing is needed since it is difficult to predict event arrival times. Ephemeral implies either continuous sensing or frequent sampling for detection, and still faster sampling for noise reduction, is needed. Quantitatively, we are interested in detecting and classifying events that are separated in time on the order of hours to days (nominally, 10 events/day), follow Poisson distributions, and last between 1 and 10 seconds. The events of interest are parameter changes in ambient signals with spectra ranging from 1Hz to 5kHz. The lifetime requirement is 1,000 hours of continuous operation from two AA alkaline batteries which can deliver 6Whr or an average of 6mW. Qualitatively, these “events” correspond to the passage of civilians, soldiers, and vehicles.

#### A. Sensors for Detection and Classification

Pottie and Kaiser’s work on Wireless Integrated Network Sensors [3] strongly influenced our choice of a *dense network* of *multi-modal* sensors. Their work argued that, “due to the decay of signals with distance, shorter-range phenomenon (such as magnetics) can be used” and that “at short range...the environment is essentially homogeneous within the detection range” and finally “since higher SNR is obtained at short range, and we can use a variety of sensing modes that may be unavailable at distance, we are better able to choose a small feature set that distinguishes targets.”

A number of practical considerations for sensor selection was presented in [5]. However, unlike our case, power consumption was not a factor in the other study, leading us to somewhat different conclusions on which sensors to select. Sensor selection for the XSM was principally influenced by three factors. First, a sensor ensemble that allowed discrimination between our target classes was essential and made multi-modal sensing desirable. Second, minimizing the power consumption, maximizing the class-conditional sensing range, and minimizing the power density (power consumed per unit area covered, measured in  $W/m^2$ ) allowed us to maximize node lifetime and minimize node density. Third, minimizing the probability of false alarm ( $P_{FA}$ ) was essential for scalability while maintaining a high probability of detection ( $P_D$ ) was essential for accuracy.

We focused on the “natural” aspects of a target class phenomenology rather than “accidental” ones and chose a multi-modal sensor suite to detect the corresponding signals. For example, humans (both civilians and soldiers) naturally emit thermal energy, but not necessarily acoustic energy. In addition, soldiers carry weapons made from ferrous materials. Vehicles contain substantial ferrous content and emit thermal and acoustic energy. Our motivation for multi-modal sensors followed from Pottie and Kaiser’s argument [3] that with only

one mode, it would be necessary to go deep into that mode's feature set, getting lower marginal returns for each feature.

After factoring in power, range, power density, and receiver operating characteristic (ROC) curves, we narrowed the sensing suite to acoustic, magnetic, and passive infrared. Table I shows the measured startup time and power consumption of the various XSM subsystems. The processor and radio, which are the same as the Mica2 mote due to reasons of risk, are included because they play an important role in passive vigilance, and the buzzer is included for completeness. Table II shows the typical sensor detection range, power density, and expected event duration as a function of sensor type and target class. The power density is  $power/(\pi \times range^2)$ . The expected event duration,  $\mathcal{E}[T_{event}]$ , is the expected covered path length, computed by averaging over all paths through a circle of radius  $range$ , divided by the maximum speed of the target class:  $\mathcal{E}[T_{event}] = \pi \times range / (2 \times v_{max})$ .

TABLE I

THE STARTUP LATENCY FROM THE OFF STATE TO THE LABEL STATE, POWER CONSUMPTION AT 3VDC, AND PERCENTAGE OF THE 6mW AVERAGE POWER BUDGET THAT WOULD BE CONSUMED, AS A FUNCTION OF THE GIVEN SUBSYSTEM AND STATE. LPL IS AN ACRONYM FOR LOW-POWER LISTEN.

Subsystem	State	Startup Time	Power	% of Budget
Acoustic	off	-n/a-	3 $\mu$ W	0.05 %
Acoustic	on	< 1 ms	1.73 mW	28.8 %
Magnetic	off	-n/a-	3 $\mu$ W	0.05 %
Magnetic	on	41 ms	19.4 mW	323 %
Infrared	off	-n/a-	3 $\mu$ W	0.05 %
Infrared	on	> 1000 ms	0.88 mW	14.7 %
Processor	sleep	-n/a-	30 $\mu$ W	0.5 %
Processor	active	0.2 ms	24 mW	400 %
Radio	off	-n/a-	3 $\mu$ W	0.05 %
Radio	receive	2.5 ms	24 mW	400 %
Radio	transmit	2.5 ms	48 mW	800 %
Radio	LPL	2.5 ms	411 $\mu$ W	6.84 %
Buzzer	off	-n/a-	3 $\mu$ W	0.05 %
Buzzer	on	-n/a-	45 mW	750 %

TABLE II

TYPICAL SENSOR DETECTION RANGE, POWER DENSITY, AND *expected* EVENT DURATION AS A FUNCTION OF SENSOR TYPE AND TARGET CLASS. MAXIMUM CIVILIAN, SOLDIER, AND VEHICLE SPEEDS OF 5, 10, AND 50 KMPH, RESPECTIVELY, ARE ASSUMED.

	Civilian	Soldier	Vehicle
<b>Infrared</b>			
Range	10 m – 15 m	10 m – 15 m	15 m – 25 m
Power Density	2.8 $\mu$ W/m <sup>2</sup>	2.8 $\mu$ W/m <sup>2</sup>	1.23 $\mu$ W/m <sup>2</sup>
Event Duration	11.3 s	5.7 s	1.7 s
<b>Magnetic</b>			
Range	-n/a-	1 m – 2 m	4 m – 8 m
Power Density	-n/a-	6,208 $\mu$ W/m <sup>2</sup>	388 $\mu$ W/m <sup>2</sup>
Event Duration	-n/a-	0.6 s	0.5 s
<b>Acoustic</b>			
Range	-n/a-	-n/a-	> 20 m
Power Density	-n/a-	-n/a-	< 1.38 $\mu$ W/m <sup>2</sup>
Event Duration	-n/a-	-n/a-	> 2.3 s

Based on Table II, the classification predicates are immediately obvious. By fusing simultaneous detections from these sensors, we can easily discriminate target classes:

$$civilian = infrared \wedge \neg magnetic \wedge \neg acoustic \quad (1)$$

$$soldier = infrared \wedge magnetic \wedge \neg acoustic \quad (2)$$

$$vehicle = infrared \wedge magnetic \wedge acoustic \quad (3)$$

where *infrared*, *magnetic*, and *acoustic*, refer to detections from the corresponding sensors. The details of computing detection thresholds and decision surfaces are not discussed in this paper but typical techniques are described in [8].

### B. Achieving Long-life

The important thing to note in Table I is that only the acoustic and infrared subsystems do not exceed, across all operating states, the average system power budget of 6mW. Duty-cycling and hierarchy are two techniques which have been proposed to address this problem. However, there are some subtle constraints when applying these techniques to the detection of rare, random, and ephemeral events, and their quick reporting, as discussed below.

1) *Duty-cycling*: Duty-cycling lowers the *average* power consumed by a subsystem by cycling its power on and off. The duty-cycle period,  $T_{DC}$  is the sum of the on-time,  $T_{on}$ , and the off-time,  $T_{off}$ . The duty-cycle,  $DC$ , is the ratio of  $T_{on}$  and  $T_{DC}$ . The range of  $T_{on}$  and  $T_{off}$  values are constrained by several factors. The startup latency,  $T_{startup}$ , of a subsystem is the amount of time required for the subsystem to stabilize after power is applied. The sample duration,  $T_{sample}$ , is the number of samples,  $N$ , less one, required to achieve an acceptable signal-to-noise ratio, divided by the sampling frequency,  $f_s$ . Finally,  $\mathcal{E}[T_{event}]$  is the expected event duration, as described previously. The values of  $T_{on}$  and  $T_{off}$  are constrained by Equation 7 and Equation 8, respectively. The factor of two in Equation 8 captures the need to sample an *entire* event. The average power after duty-cycling is given by Equation 9.

$$T_{DC} = T_{on} + T_{off} \quad (4)$$

$$DC = \frac{T_{on}}{T_{on} + T_{off}} \quad (5)$$

$$T_{sample} = (N - 1) / f_s \quad (6)$$

$$T_{on} \geq T_{startup} + T_{sample} \quad (7)$$

$$T_{off} \leq \mathcal{E}[T_{event}] - 2 \times T_{on} \quad (8)$$

$$Power_{avg} = \frac{1}{T_{DC}} \int_{T_{DC}} Power(t) dt \quad (9)$$

There are two important cases in which duty-cycling does not solve the power budget problem. First, when  $T_{off}$  is negative, the constraints cannot be met because an event is *too* ephemeral, and the subsystem must be powered continuously. Second, when  $Power_{avg}$  is close to or in excess of the power budget, duty-cycling does not lower the average power to an acceptable level.

The infrared sensor, with parameters  $N = 1$ ,  $f_s = 32$  Hz,  $T_{startup} = 1,000$  ms, and  $\mathcal{E}[T_{event}] = 1,700$  ms, fails to meet the constraint of Equation 8. Fortunately, since the infrared subsystem only consumes 0.88 mW, it can be powered continuously. The magnetic sensor, with parameters  $N = 1$ ,  $f_s = 128$  Hz,  $T_{startup} = 41$  ms, and  $\mathcal{E}[T_{event}] = 500$  ms, results in an average power of 4.3 mW, which does not leave a sufficient margin of error or power for all other system processes. This *is* a problem and we will return to it in Section III-B.2. The acoustic sensor, with parameters  $N = 256$ ,  $f_s = 8,192$  Hz,  $T_{startup} = 1$  ms, and  $\mathcal{E}[T_{event}] = 2,300$  ms, meets the constraint of Equation 8 and results in an average power of 370  $\mu$ W.

Since the events of interest are rare and the radio consumes 400% of the power budget when listening and 800% of the power budget when transmitting, the radio is a candidate for duty-cycling.

The only complication is that reporting latency is an important application quality-of-service metric that needs to be balanced with system lifetime. We use an implementation of low-power listening [9] available in the TinyOS distribution. Low-power listening extends message preambles on transmit and duty-cycles the radio on receive, trading off energy and bandwidth against communications latency. Figure 2 shows the power profile of the receiver during the duty-cycle on-time. Integrating the power over this period gives an energy usage of  $441\mu\text{J}$ , and an average power consumption of  $411\mu\text{W}$ , given our acceptable duty-cycle period of 1.07s.

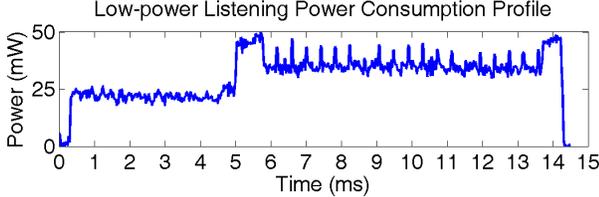


Fig. 2. Power consumption profile of low-power listening. The approximately 18 spikes between 6ms and 14ms correspond to the processor communicating with the radio over SPI bus to check for the presence of a message preamble.

Low-power listening shifts most of the power burden from the receiver to the transmitter, which is appropriate given how rare our events are and since the transmitter only pays a price when a message is (rarely) sent. The worst-case reporting latency is the product of the duty-cycle period and the network diameter or depth.

2) *Energy-Quality Hierarchy*: When duty-cycling fails to meet the system power budget, arranging sensors in a hierarchy with one sensor triggering another allows power consumption to be further lowered. We advocate an approach that uses a trigger network composed of continuously-operating low-power sensors (infrared) with asynchronous wakeup (*wake*), multi-modal (infrared, acoustic, and magnetic) sensor signal processing (*pir*, *mic*, and *mag*, respectively), data fusion algorithm (*fuse*), and radio transmission (*tx*). One possible trigger network is shown in Figure 3. Let  $X \rightarrow Y$  denote  $X$  triggers  $Y$  and let  $p(H_1^Y | H_1^X)$  be the conditional probability of  $Y$  given  $X$  where  $H_1^X$  is true for both detections and false alarms under  $X$ .

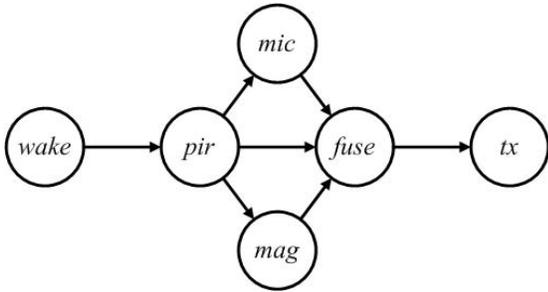


Fig. 3. Detection trigger network.  $X \rightarrow Y$  denotes  $X$  triggers  $Y$ .

Given our trigger network and classification predicates, the upper bound on our transmission probability is given by

$$p(H_1^{tx} | H_1^{fuse})p(H_1^{fuse} | H_1^{pir})p(H_1^{pir} | H_1^{wake}) \quad (10)$$

and is critically dependent on  $p(H_1^{pir} | H_1^{wake})$ . Assume that each time *pir* triggers, *mic*, *mag*, and *fuse* are all turned on for a duration of  $\max\{\mathcal{E}[T_{event}]\}$ , and then a message is transmitted. Then, for each *pir* trigger (a generalization is beyond the scope of this paper)

$$E = \max\{\mathcal{E}[T_{event}]\} \times \sum_{s \in S} Power_s + T_{DC}^{tx} \times Power_{tx} \quad (11)$$

units of energy is consumed where  $S = \{pir, mic, mag, fuse\}$ . We chose  $\max\{\mathcal{E}[T_{event}]\}$ , or the maximum expected event duration across *all* sensors, target classes, and trajectories, because an event could last up to this length of time. Since we do not know in advance whether the target is a civilian, soldier, or vehicle, and we do not know the prior probabilities of these classes, we cannot use the  $\mathcal{E}[T_{event}]$ . The maximum event duration provides an upper bound for how long a sensor has to remain in active vigilance.

Assuming  $\max\{\mathcal{E}[T_{event}]\} = 60$  s and  $T_{DC}^{tx} = 1$  s,  $E = 2,760$  mJ +  $48$  mJ =  $2,808$  mJ per *pir* event. Assuming that the wakeup-enabled infrared sensor ( $0.88$  mW  $\times$  1,000 hrs =  $3,168$  J) and low-power listen ( $411$   $\mu$ W  $\times$  1,000 hrs =  $1,479$  J) are always on, our energy budget of 6 Whr =  $21,600$  J is reduced to  $16,953$  J, which allows for more than 6,000 detections ( $16,953$  J/ $2,808$  J/detection), or more than 6 detections per hour. This exceeds by 14 times the expected event rate of 10 events per day, assuming *pir* can suppress false alarms from *wake*, which we believe it can. The remaining energy can be allocated to other middleware functions. We added wakeup circuitry to the infrared and acoustic sensors, and decreased the startup time of the magnetic sensor, to support a triggering hierarchy.

#### IV. RECOVERABLE RETASKING

A goal of the ExScal project is to create a large-scale testbed that can be used by researchers to develop and test middleware algorithms at heretofore unprecedented scale. Since manually reprogramming a large network repeatedly is impossible, a multi-hop wireless reprogramming capability was required. However, since many of the algorithms to be tested at scale are experimental, it is possible that latent bugs may exist, the algorithms may fail at scale, or that emergent pathological (Byzantine) behavior may emerge and lock up the network. Since the XSM uses the Atmel ATmega128L which, like most inexpensive and low-power microcontrollers, does not provide privileged instructions or memory protection, it is possible for an application to take control of the hardware and, in the worst case, render it remotely inaccessible.

The specter of manually reprogramming 10,000 dead nodes convinced us that a mechanism was needed to ensure trusted code eventually regained control of the hardware. To address this need, we implemented a hardware grenade timer [10] which periodically resets the system, whether it is healthy or not, and returns control to the trusted computing base (TCB). The TCB includes a bootloader that invokes a ‘‘Golden Image’’ which includes the Sensor Network Management Service (SNMS) [11] and the Deluge wireless multi-hop network reprogramming service [12]. The TCB is also responsible for configuring and arming the grenade timer. The details of our implementation of the grenade timer are presented in Table III and Figure 4.

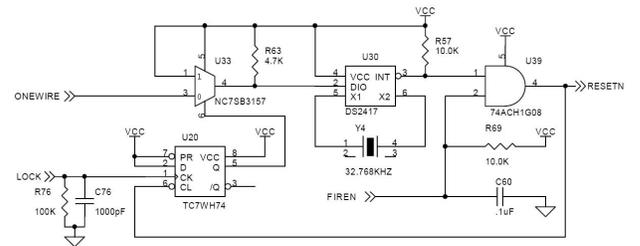


Fig. 4. The XSM grenade timer circuit.

Trusted code can set the timeout interval and start the timer by sending commands over the ONEWIRE bus. Further access to the

TABLE III  
GRENADE TIMER FEATURES.

Feature	Description
<b>Adjustable Timeout Interval</b>	The amount of time that the grenade timer “fizzes” before forcibly resetting the node (1, 4, 32, 64, 2048, 4096, 65536, or 131072 seconds).
<b>One-Shot Lock-out</b>	Once the grenade timer is started, it cannot be stopped and the “fizz” time cannot be changed.
<b>Asynchronous Trigger</b>	The grenade timer may be triggered by either trusted or untrusted software at any time.
<b>Alternate Uses</b>	As long as the grenade timer has <i>not</i> been started, the real-time clock used to implement it remains accessible to application code.

grenade timer can be disabled by asserting LOCK high. The timer may be asynchronously triggered at any time by asserting FIREN low. As long as LOCK has not been asserted, the DS2417 real-time clock (RTC) at the heart of the circuit remains accessible over the ONEWIRE bus. A quirk of the RTC causes occasional spurious interrupts, which reset the XSM, during certain register writes. Since these writes only occur immediately after a reset, stuttering resets are not unusual.

## V. PACKAGING

Packaging primarily protects the electronics from the environment but also constrains, and is constrained by, circuit board footprint, antenna placement, sensor positioning, and logistical issues. The packaging is shown in Figure 1.

The enclosure is molded from a white plastic material that reflects sunlight, shades the electronics, and avoids the overheating problem observed in earlier designs [5], [13]. Twenty-four acoustic apertures expose the microphone sensor to the environment. These apertures are covered with water-resistant windscreens which reduce wind noise. Four windows expose the passive infrared sensors to the environment. These windows are covered with an infrared-transparent film which reduces false alarms due to air thermals. An easily accessible on/off switch and a piezo buzzer support the *one-touch* and *one-listen* goals, respectively. A reset switch is accessible through a pinhole opening on one side of the enclosure. The antenna protrudes through the top of the package and is located in the center of the circuit board to provide a more isotropic ground plane than the Mica family of motes [14] and reduce the “degree-of-irregularity” in the radiation pattern [15]. The circuit board is mounted to the top cover using brass screws. By doing so, the height of the antenna is raised by about one inch, which nominally reduces ground absorption of radio transmissions and increases communications range. Since the sensitivity of the magnetic sensor decreases when placed near ferromagnetic materials, the choice of non-ferrous brass screws to attach the circuit board to the enclosure is important, as is maximizing the distance between the sensor and the batteries which contain ferrous materials.

## VI. DISCUSSION

In this section, we highlight some of the key lessons that we learned and tradeoffs we made during the design of the XSM platform. Perhaps the most important lesson is that scale impacts every aspect of the platform lifecycle including concept, design, verification, manufacturing, shipment, testing, deployment, configuration, operation, retasking, and reclamation. There is no substitute for having a broad audience critique the design and even that may be

insufficient to find all the problems before testing at moderate scale. We suggest short and frequent hardware iterations. Unfortunately, we discovered, a broad audience can also increase risk aversion and confine innovation. In particular, we observed an unwillingness to support innovation unless it was critical to functionality or the shadow cost was unacceptably high. The passive infrared circuit is one example of essential functionality while the grenade timer is an example of an unacceptably high shadow cost. In contrast, a proposal to lower the power consumption of the magnetic sensing circuit by nearly a factor of 20 was dismissed due to risk. While such behavior might be expected of the mainstream market, its presence in a research endeavor dampened the potential platform innovation.

Three generations of XSM circuit boards were evaluated during this work. The first generation investigated some esoteric ideas like digitally-clocked switched-capacitor lowpass filters and analog wakeup circuits. The second generation was a heavily-modified design that added many new features like digitally-controlled analog filtering and grenade timer circuits, and addressed the problems of the first design. The third generation corrected an interference problem caused by the digital control of analog filtering circuits. Fulford-Jones et al. reported similar experiences with radio transmissions causing interference in analog circuitry [16]. The interference was unexpected in both cases, we suspect, due to poor test coverage. In our case, the design worked correctly when only one sensor channel was used since calibration and sampling operations were mutually exclusive. When multiple channels were tested simultaneously, the interference and crosstalk became noticeable, and the need for additional shielding and decoupling became obvious. Other researchers using the platform still report crosstalk between some subsystems. The key lesson is that concurrency, in all its forms, must be tested to ensure adequate coverage.

We had high hopes for the low-power wakeup circuits used with the infrared and acoustic sensors. Unfortunately, these circuits did not live up to our early expectations. With a sufficiently high detection threshold, the infrared wakeup circuit can achieve vanishingly-small false alarm rates, but detection range suffers. Greater range is possible by increasing the sensitivity, but the false alarm rate increases due to environment-specific air thermal movements. In the worst cases, false alarms can occur as frequently as every few seconds. Fortunately, it is possible to discriminate false alarms from true detections in software by analyzing signal parameters and the peak-to-peak interval. In contrast, problems with the acoustic wakeup makes this circuit practically useless. In the original XSM design, the threshold detection circuit included non-linear and storage elements which provided some hysteresis and noise immunity. We determined this was unnecessary for the infrared case, but mistakenly removed these circuit elements from both the infrared and acoustic wakeup circuits. Hence, the acoustic wakeup circuit has an unacceptably high false alarm rate.

Low-power wakeup sensors and radios represent promising areas of innovation since they extend the event-driven hierarchy into hardware. The state-of-the-art in low-power signal processing, event detection, and wakeup triggering are highly-integrated VLSI solutions. In [17], Dong et al., present a single-chip micropower VLSI spectrum analyzer implemented in  $0.8\mu$  HPCMOS using 45,000 transistors. The system operates with a  $1\mu$ A drain current at a 3V supply bias at a 200 samples/second processing rate. In [18], Goldberg et al., describe a low-power VLSI wakeup detector for use in an acoustic surveillance sensor network. In [19], Gu and Stankovic propose using ultra low-power wakeup radios as a power management technique for sensor networks.

Robust network reprogramming is a much deeper problem than simply recoverable-retasking. While our grenade timer implementation guarantees that a trusted bootloader regains control of the hardware after a forced periodic reset, that is only the first step. What happens after the reset is essential to recoverability, which we addressed using SNMS, Deluge, and the Golden Image, in the XSM design. However, our current approach does not address several issues that are important for secure and robust network reprogramming. For example, if during reprogramming, the general set of network protocols are upgraded, then a node may become isolated since after a forced reset, its neighbors might operate on a new and incompatible version of the network protocol. We believe considerable work remains in securing the bootloader and the programs which it calls, collectively the trusted computing base, from a variety of accidental and malicious threats [2]. This work is unrelated to hardware but is constrained by the available computational and storage resources.

We encountered several problems with the enclosure design. A pinhole used to access a user-switch had to be plugged since it allowed standing water to seep into the enclosure. The telescoping antenna used in the XSM is chrome-plated and has ferro-magnetic properties. The battery holder is attached to the enclosure base but the batteries themselves were not secured and many became dislodged during transit and damaged internal components. Battery replacement requires removing four screws and makes battery changes time-consuming. Even worse, manual reprogramming requires removing eight screws to access the 51-pin expansion connector. Fortunately, this operation is required only when reprogramming the bootloader – something that has been necessary only once. Selecting the winning enclosure design was difficult. We considered nearly a dozen different, sometimes conflicting, packaging requirements and five different designs before selecting the one presented here. To our surprise, no single design addressed all of our requirements and each design was better in some regards and worse in others. Our final choice was, by no means, an obvious winner. The key lesson is that package design is highly constraint-driven and packaging must be co-designed and co-evaluated with the electronics and logistics. Sufficient time needs to be allocated for these coupled activities.

## VII. CONCLUSIONS

This work presents a novel embedded systems platform and architecture, design methods for exceptional event detection in a power-conscious manner, and experiences in designing sensor nodes for large-scale operation. The XSM is the first fully-integrated mote-class wireless sensor node that provides hardware support for low-power passive vigilance, large-scale operation, and recoverable-retasking. Our passive vigilance architecture, used to detect and classify *rare*, *random*, and *ephemeral* events, is directly applicable to a wide variety of event detection applications. While considerable work remains in improving the detection rates and lowering the false alarm rates of low-power wakeup sensors, our work demonstrates the feasibility and promise of this architectural approach. Support for large-scale operation is achieved through a *highly-integrated* platform that allows *one-touch* activation and *one-listen* verification, all of which minimize *human-in-the-loop* operations. Recoverable-retasking is achieved through the use of a hardware grenade timer and is the first implementation of its kind of which we are aware. The XSM presents a novel and timely addition to the evolution of wireless sensor network platforms. Additional information about the eXtreme Scale Mote is available at: <http://www.cs.berkeley.edu/~prabal/projects/xsm>

## ACKNOWLEDGMENT

We would like to acknowledge the contributions of Sandip Bapat, Hui Cao, Benjamin Coifman, Robert Cuenin, Emre Ertin, William Ferreira, Lin Gu, Tien He, Wei Hong, Mike Horton, Jonathan Hui, Santosh Kumar, Kenneth Parker, Joseph Polastre, Vijay Raghavan, Rajiv Ramnath, Al Sciarretta, Cory Sharp, Robert Szweczyk, Gilman Tolle, and Norm Whitaker. We would also like to thank the anonymous reviewers and Frank Vernon, our shepherd. This work was supported by the Defense Advanced Research Projects Agency and a National Science Foundation Graduate Research Fellowship.

## REFERENCES

- [1] Gregory J. Pottie, "Wireless sensor networks," in *IEEE Information Theory Workshop Proceedings*, June 1998.
- [2] William A. Arbaugh, David J. Farber and Jonathan M. Smith, "A Secure and Reliable Bootstrap Architecture," in *Proceedings 1997 IEEE Symposium on Security and Privacy*, pp. 65–71, May 1997.
- [3] G.J. Pottie and W.J. Kaiser, "Wireless integrated network sensors," *Communications of the ACM*, vol. 43, no. 5, pp. 51–58, May 2000.
- [4] B. West, P. Flikkema, T. Sisk, and G. Koch, "Wireless sensor networks for dense spatio-temporal monitoring of the environment: A case for integrated circuit, system, and network design," in *2001 IEEE CAS Workshop on Wireless Communications and Networking*, Notre Dame, Indiana, Aug. 2001.
- [5] A. Arora, P. Dutta, S. Bapat, V. Kulathumani, H. Zhang, V. Naik, V. Mittal, H. Cao, M. Gouda, Y. Choi, T. Herman, S. Kulkarni, U. Arumugam, M. Nesterenko, A. Vora, and M. Miyashita, "A line in the sand: A wireless sensor network for target detection, classification, and tracking," *Computer Networks Journal*, Oct. 2004.
- [6] T. He, S. Krishnamurthy, J. Stankovic, T. Abdelzaher, L. Luo, T. Yan, L. Gu, J. Hui, and B. Krogh, "Energy-efficient surveillance systems using wireless sensor networks," in *Mobisys 2004*, June 2004.
- [7] David Tennenhouse, "Proactive computing," *Communications of the ACM*, vol. 43, no. 5, pp. 43–50, May 2000.
- [8] Richard O. Duda, Peter E. Hart, and David G. Stork, *Pattern Classification*, John Wiley & Sons, Inc., second edition, 2001.
- [9] Joseph Polastre, Jason Hill, and David Culler, "Versatile low power media access for wireless sensor networks," in *SenSys'04*, 2004.
- [10] Frank Stajano and Ross Anderson, "The grenade timer: Fortifying the watchdog timer against malicious mobile code," in *Proceedings of 7th International Workshop on Mobile Multimedia Communications (MoMuC 2000)*, Waseda, Tokyo, Japan, Oct. 2000.
- [11] Gilman Tolle and David Culler, "Design of an application-cooperative management system for wireless sensor networks," in *2nd European Workshop on Wireless Sensor Networks (in submission)*, Istanbul, Turkey, Jan. 2005.
- [12] Jonathan W. Hui and David Culler, "The dynamic behavior of a data dissemination protocol for network programming at scale," in *The 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys'04)*, 2004.
- [13] Joseph Polastre, "Design and implementation of wireless sensor networks for habitat monitoring," M.S. thesis, U.C. Berkeley, 2003.
- [14] Jason Hill and David Culler, "Mica: A wireless platform for deeply embedded networks," *IEEE Micro*, vol. 22, no. 6, pp. 12–24, 2002.
- [15] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic, "Impact of radio irregularity on wireless sensor networks," in *Mobisys 2004*, June 2004.
- [16] Thaddeus R.F. Fulford-Jones, Gu-Yeon Wei, and Matt Welsh, "A portable, low-power, wireless two-lead ekg systems," in *IEEE Sensor and Ad Hoc Communications and Networks (SECON)*, Sept. 2004.
- [17] M.J. Dong, G. Yung, and W.J. Kaiser, "Low power signal processing architectures for network microsensors," in *Proceedings of 1997 International Symposium on Low Power Electronics and Design*, Monterey, CA, USA, Aug. 1997, pp. 173–177.
- [18] David H. Goldberg, Andreas G. Andreou, Pedro Julian, Philippe O. Pouliquen, Laurence Riddle, and Rich Rosasco, "A wake-up detector for an acoustic surveillance sensor network: algorithm and vlsi implementation," in *Third International Symposium on Information Processing in Sensor Networks (IPSN 2004)*, 2004.
- [19] Lin Gu and Jack Stankovic, "Radio triggered wake-up capability for sensor networks," in *Real-Time Applications Symposium*, May 2004.