Some Sensor Network Elements for Ubiquitous Computing

Waylon Brunette[1], Jonathan Lester[2]

[1]Dep't of Computer Science & Engineering
[2]Dep't of Electrical Engineering
University of Washington
Box 352350, Seattle, WA 98195 USA
{wrb, gaetano}@cs.washington.edu
jlester@u.washington.edu

Adam Rea[3], Gaetano Borriello[1,3]

[3]Intel Research Seattle
1100 NE 45th Street, Suite 600
Seattle, WA 98105 USA
{adam.d.rea, gaetano.borriello}@intel.com

Abstract— **Ubiquitous computing applications often use a user's context to automatically adjust their behavior to the situation. We have developed three types of wireless sensor nodes that can be worn, carried, or embedded in the environment that can provide interesting contextual information: a rich multi-sensor node to infer human activity that can be worn by a person or be part of a sensing environment, a portable wireless node for reading and writing short-range radio-frequency identification (RFID) tags, and a small wrist-watch-sized wireless display device that can serve as an easily accessible and 'glanceable' user interface. We report on our experiences in building these platforms and using them in some initial applications. We conclude with suggestions for future work in creating platforms that help blend sensor networks into ubiquitous computing.**

Keywords- Multi-Modal Sensors; Personal Area Networks; Ubiquitous Computing; passive RFID; Activity Inference; Wristwatch Display

I. Introduction

Monitoring and data collection have been the primary applications of sensor networks for domains as varied as military and agriculture [1, 2]. The research work has focused on the ad hoc networking aspects, namely, routing and power management. Ubiquitous computing makes extensive use of sensors to provide important context information so that applications can adjust their behavior based on what the user is likely to need next. These proactive applications are designed to interrupt the user as little as possible while assisting them throughout their day [3].

We view sensor networks as sets of nodes that are either on a person and or in the environment. As people move about, their personal sensor network can leverage the networks they encounter by exchanging information with them. New types of sensor nodes are needed that are appropriate for ubiquitous computing. These include ways of viewing information provided by the local environment's sensors, ways of providing data to local sensors, and using personal sensors in conjunction with the fixed sensors. In addition, we are investigating how to integrate support for the more dynamic interactions of ubiquitous computing into the sensor network domain and how this will impact power-management schemes, data-transfer, and routing protocols.

In this paper, we report on three platforms we have developed as well as some of the initial applications we have built using them. Our devices are built on a foundation of UCB and Intel sensor motes (mica2dots [4] and iMotes [5, 6]) so that we have the added flexibility of leveraging the sensing and communication infrastructure being developed by the sensor network community.

The three devices are a multi-modal sensor board (MSB), a portable short-range HF RFID reader/writer (iReader), and a sensor node in a wrist-watch form factor with an LCD display (DisplayMote). The MSB is a research tool that provides a super-set of sensors for human activity recognition which can be used in conjunction with sensors in the environment. We use it to experimentally determine the most appropriate sub-set of sensors for a given application. The iReader is a personal actuation mechanism that uses RFID tags embedded in objects to collect data and issue commands based on the tag that is read. The DisplayMote provides a general-purpose, lightweight user interface that can be adapted across a wide range of applications. We find the wrist-watch form-factor to be the most useful as it is easily worn and is easy to just glance at for information about the available sensor-based services in the user's proximity.

In our work, we are particularly interested in the intersection of sensor networks and ubiquitous computing. Although personal area networks can be built using the same radios and routing protocols, the more interesting space is how carried sensors interact with fixed sensors. Connections between devices will be short-lived with higher contention as people move through varied spaces. This will require new protocols for discovery and routing, as the computations and correlations necessary in the network will be determined by the mix of devices in the collection at any given time. These protocols should leverage powered sensors that may be part of the environment and minimize the power consumption of the mobile or battery-powered nodes. Moreover, the data from carried and local sensors can be correlated to reach interesting conclusions about where people are and what they are doing. For example, audio sensors can help assess spatial relationships between people. Barometric pressure sensors can assist a device in determining the floor of a building.

The remainder of this paper describes our three platforms in turn and provides an application for each.

II. Multi-Modal Sensor Board

We built the multi-modal sensor board (MSB) to gather data simultaneously from a large set of sensors to help us better understand the usefulness of different sensor modalities and to determine which modalities will be most important in inferring human activities. The MSB is designed to attach to Intel's iMote, a Bluetooth/32-bit ARM7-

0-7803-9202-7/05/$20.00 (C) 2005 IEEE

based sensor, but can also communicate with handheld and desktop computers via either a USB or compact flash bridge. The sensor board contains the seven sensors listed in Table 1 and is shown in Figure 1 (along with an iMote and battery board).

The sensors on the MSB were selected for their general usefulness (as evidenced by related work in activity inference [7, 8, 9]), small footprint, low power consumption, and availability of digital interfaces (except for the microphone and phototransistor). Sensors such as these are already being incorporated into cell phones (like the Samsung SCH-S310 and Pantech & Curitel PH-S6500) and wrist-watches are likely to follow soon.

Each sensor sub-system is equipped with a low resistance analog switch between the sensor's power rail and the system power rail. This allows us to selectively shutdown a sensor completely; regardless of the disparate sleep modes the sensors themselves may support (the analog switches have a typical drain of 1nA, 1µA max). With all the sensors running continuously the sensor board consumes approximately 13mA (11.9µA in deep sleep) allowing us to run continuously for more than 12 hours on a 200mAh Li-Polymer battery. We designed a battery board (shown in Figure 1) with charging and protection circuitry that allows the system to run from USB rechargeable Li-Polymer batteries.

The MSB is built on a 6-layer PCB that includes an Atmel ATMega 128L microprocessor running at 7.3728MHz. The on-board microprocessor controls communication with each sensor, performs the sampling (either polled or streaming), and sends the data over a hi-speed UART. Just as importantly, the on-board microprocessor allows the sensor board to perform its sensing tasks independently of the platform to which it is attached. Allowing the MSB to attach to the iMote along with other devices, which can then be used as communications gateways as well as a platforms for in-network computation (currently we've attached the MSB to an iPAQ, a laptop, and wirelessly to a Bluetooth cell phone).

We designed the MSB with a focus on human activity inference. Some of the primitive activities we would like to recognize are walking, running, bicycling, going up/down stairs, going up/down in an elevator, driving a car, etc. In addition to recognizing these activities we would like to determine which sensors and features from those sensors are most useful for each activity [10]. Typically, machine learning algorithms don't use raw sensor data directly; instead they rely on higher-level features extracted from this data to reduce the complexity of the analysis.

The knowledge of which sensors and features are important will allow us to intelligently design future devices with the most appropriate sensors and computational power on-board. In addition, we are experimenting with different locations in the network to perform feature extraction and correlations, while taking into account the cost of communicating data between the network elements. A node in the infrastructure that is more computationally capable may be more appropriate for performing complex operations; however, it may be more appropriate to perform feature extraction on the sensor node itself to reduce the amount of communication bandwidth necessary. In addition feature extraction on the node itself may have positive privacy implications. For example, it is likely inappropriate for privacy and bandwidth considerations to transmit complete audio recordings, while features for detecting conversations, ambient sound intensity, frequency distribution of the sound, etc. would likely be more appropriate and acceptable (although not necessarily completely benign).

We are using the MSB to understand the relationships between the sensors, and how important these sensors are to one another. Sensors that may be important for recognizing one activity may not necessarily have to be together on the same device, person, or environment. This allows us to place sensors on different devices as well as taking

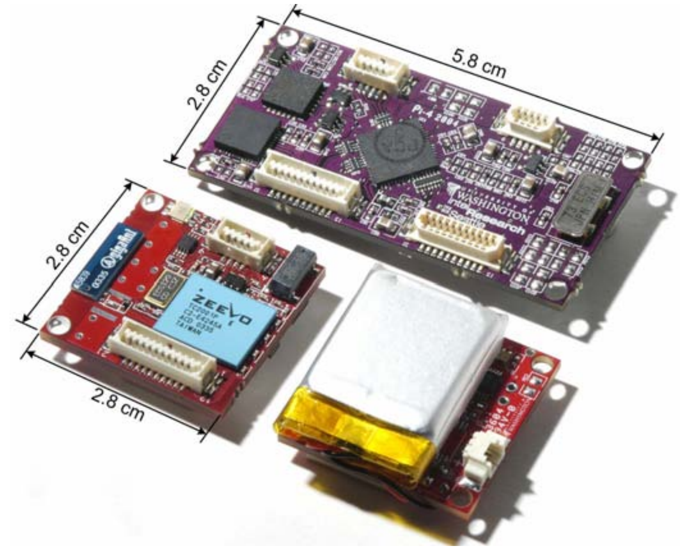

Figure1. Sensor board (top) along with an iMote (lower left) and a Li-Poly battery board (lower right).

advantage of a changing mix of sensors. For example, checking that a light sensor in a room saw the lights go on at the same time as a sensor on the user thereby providing evidence of the room the user is in. Moreover, sensors may not have to be on at the same time; an event detected by one sensor can activate another to obtain more specific data. For example, checking accelerometer data when there is a barometric pressure change, to determine if the change is due to an elevator, moving up/down, a stairway, or from weather fluctuations.

We have collected 35 hours of annotated data using the MSB; recording various human activities performed by two volunteers over the course of several weeks. Figure 2 contains a sample trace of this data, from a person walking around the Allen Center at the University of Washington. The trace contains 5 minutes of recorded data with some basic activity labels shown at the bottom of the graph. One sensor that stands out from this trace is the barometric pressure sensor, which is capable of fine enough measurement to show differences due to the floor that a person is on. By looking at the transitions we can see the relative difference in pressure between single floors, with multiple floor changes very distinct (due to faster movement using an elevator).

From the raw sensor data we currently compute ≈650 features (both time domain and frequency domain) and use a modified version of AdaBoost, proposed by Viola and Jones [11], to automatically select the best features for recognizing a given activity. This feature selection method reduces our feature set by 60%; saving us a great deal of computation, by allowing us to only calculate 60% (250 of the 650) features. Our current classification system uses this subset of features as inputs to simple discriminative decision stump classifiers. The outputs of which are used to drive a Hidden Markov Model (HMM). The decision stump classifiers operate in an instantaneous fashion providing a classification for a small data window (currently ¼ of a second); while the HMM uses a sequence of these windows (currently 15 seconds) to provide temporal smoothness to our

TABLE I. MSB SENSORS AND ACHIEVABLE DATA RATES

Manufacturer	Part No.	Description	Sampling Rate
Panasonic	WM-61A	Electric Microphone	~ 16000 Hz
Osram	SFH-3410	Visible Light Phototransistor	~ 550 Hz
STMicro	LIS3L02DS	3-Axis Digital Accelerometer	~ 550 Hz
Honeywell	HMC6352	2-Axis Digital Compass	30 Hz
Intersema	MS5534AP	Digital Barometer / Temperature	15 Hz
TAOS	TSL2550	Digital Ambient (IR and Visible+IR) Light	5 Hz
Sensirion	SHT15	Digtial Humidity / Temperature	2 Hz

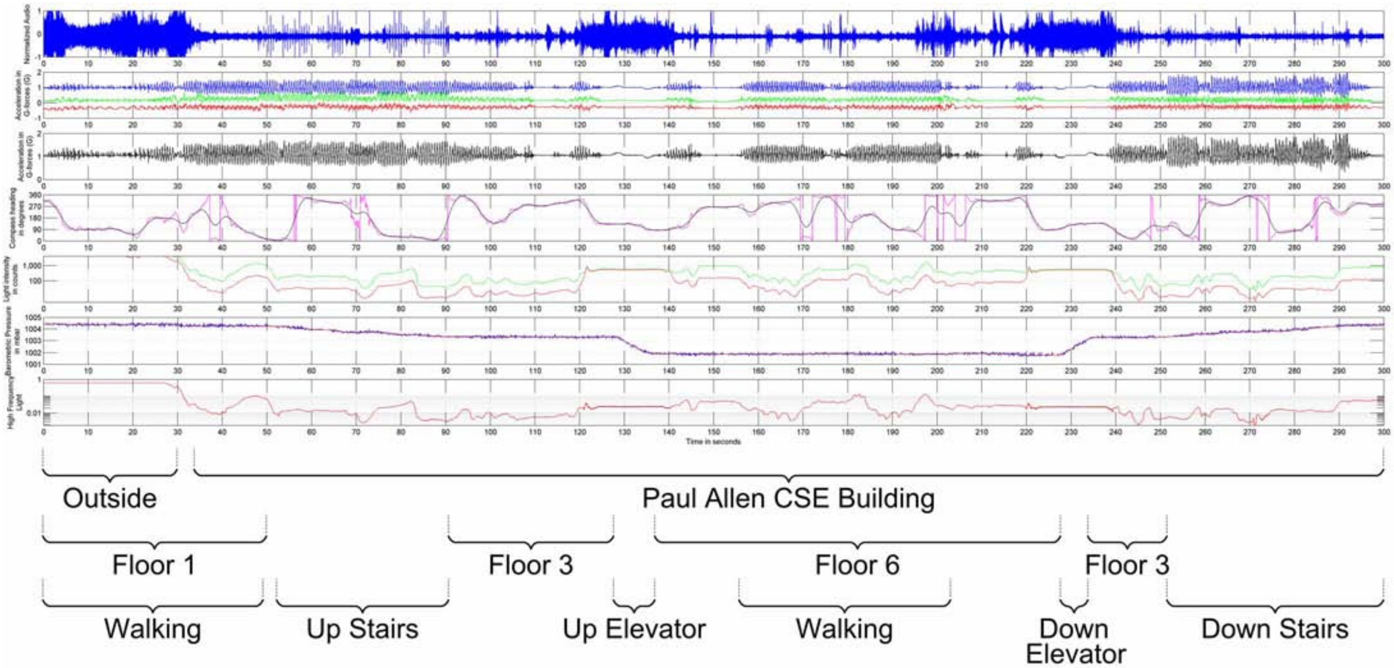

Figure 2. Example trace of sensor data collected from a person walking around the Allen Center at the University of Washington. A rough ground truth is shown below the graph describing the locations (building and floor) and the activities performed by the person.

classification results. The current system is able to achieve a 92% overall accuracy for 10 classes of activities including: walking, jogging, standing, sitting, riding an elevator, and climbing stairs. We are currently working towards implementing our system on a cell phone platform to perform real-time classification.

A. Activity and Place

One application of particular interest is the relationship between activities and place. The semantics of a place (Joe's office, the coffee shop at the corner, etc.) are more important for most users than their coordinates. If we know the activities that take place at a location we can develop a more meaningful understanding of that location. For example, if you have a location that you drive to, walk a little bit, and sit in a more or less fixed orientation for a duration between 10 and 60 minutes is likely to be a restaurant while a location that you walk to, walk around in and are engaged in high-impact activities is more likely to be a gym. This is particularly important when you have a location system that works indoors but has limited accuracy (e.g., the Wi-Fi based PlaceLab location system with its approximately 20 - 40m error bound [12]) and correlated sensor data can be used to distinguish between a gym and the café next door.

III. PORTABLE WIRELESS RFID READER/WRITER

To allow users to interact with the growing number of passive RFID tags in the environment, we created a small handheld RFID reader that can be used as a personal actuation device. We worked with SkyeTek Engineering, manufacturers of multi-protocol RFID readers, to develop a unit compatible with the UCB mica2dot mote that can both read and write 13.56 MHz RFID tags [13]. The low-power iReader is shown in Figure 3 and has a read range of only a few inches. Its short range allows users to directly interact with specific objects rather than all the tags around them. With a click of one button the user can read a tag; the other button allows the user to write a pre-loaded buffer onto the tag.

We designed the iReader as a small, mobile RFID reader with wireless communication and control capabilities. The iReader can read and write information into an assortment of passive RFID tags that have a globally unique ID and solid-state storage space for writing additional information (usually on the order of 2 – 4K bits). It can also use its wireless capability (via a mica2dot) to communicate with other devices in a user's personal area network (PAN). The iReader uses a rechargeable Li-Polymer battery with an accompanying USB charger and can run for more than a month when used an average of 50 times per day.

An envisioned usage model for RFID tags is in smart spaces and location sensing. Tagged objects can contain many forms of information such as part history, schematics, or even pointers to product manuals. In addition, tags can be used for physical access control or can even contain code that is executed upon a tag read, allowing the iReader to be used as an actuator. Tags can also be associated with auxiliary data contained in databases in the infrastructure making for a wide range of possible applications. A simple example is a virtual light switch – a user reads the tag on the switch using their key-fob iReader, the tag's ID as well as the iReader's ID (the owner's identity) is communicated to the infrastructure which can check for access privileges and turn the light on, if appropriate. This permits cheap enhancements to the environment (a tag on a wall sticker) while enhancing functionality (controlled access).

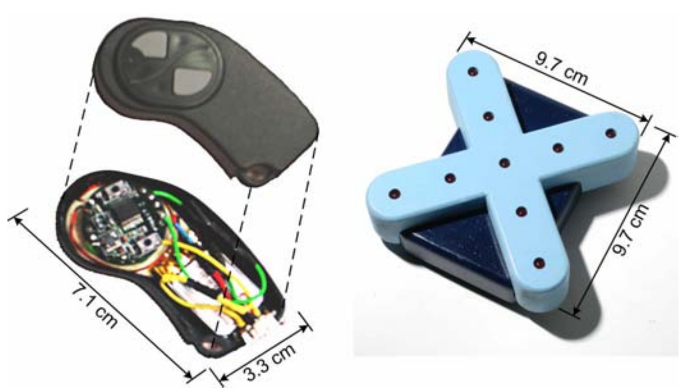

Figure 3. The iReader (left) and eXspot (right).

The iReader can also be used as an out-of-band connection mechanism. It can read a tag placed in a conference room that contains data about an office building's wireless network (SSID and WEP key). The read data is sent to a laptop, PDA, or any other device so that it can bootstrap itself into the wireless network. A Bluetooth capable device can be augmented with an RFID tag containing its MAC address allowing the discovery process to take less time and giving the user direct control of which devices to communicate with rather than having to deal with the multitude of devices that may be in the vicinity. Conferences rooms' RFID tags can contain all the data necessary to configure a laptop for that location. In addition connectivity information such as SSID and WEP key, it can contain the names of local printers and/or projectors available in the room. Not only does it allow for a convenient method of configuration, it can also be used as a security mechanism, since only people within the short read range of the tag can possibly collect the information. These actuation events don't have to be limited to computer interactions. As already mentioned, RFID tags can be placed anywhere and be used as widgets to trigger events such as virtual switches to turn lights on and off. This allows for extremely dynamic environments where widgets can be reprogrammed and reconfigured.

The iReader was also designed for applications that augment objects and spaces with information. These applications allow users to access and control data that is associated with a particular item. A prime example of this individual control of data is for associating repair histories with a specific device. For example, past repairs and scheduled maintenance for an elevator can be annotated at the elevator itself by writing into the elevator's RFID tag as well as in the centralized database. This means that the information needed on the worksite is always there and repair people do not have to rely on a network connection to get the most crucial information. Another strong advantage to having inexpensive, lightweight RFID reader/writers is the ability for people to create personalized content. Business cards are imprinted with a variety of static information (e.g. name, title, e-mail address, etc.). With an embedded writable RFID tag, business cards can contain active content that allows them to become a malleable document full of additional information that can be varied depending on who is the intended recipient. For example, it would be appropriate to embed the URL for a work homepage within the card to give to a work colleague but also to be able point a friend to a site of pictures from last week's golf outing using the same business card. With the reprogrammable memory available with RFID, business cards can now contain sounds, product descriptions, or any other data or pointer that can fit in the memory available on the tag.

A. The eXspot

A version of the iReader is undergoing trial deployments at the Exploratorium, an interactive science museum in San Francisco. The Exploratorium is creating a system to link visitors' expressed interests while in the museum to their online educational content. The Exploratorium wants to create a post-visit learning experience by using data about the visitor's interests to present relevant information (e.g., activities and online exhibits) to the visitor at a later time through the web.

The iReader was an ideal platform for the Exploratorium; because they needed a sensor that was wireless, portable, and could act as an actuation device to capture users' interests. Users get a card containing an RFID tag when they enter the museum. They can swipe the card by readers at various exhibits to indicate their interest in that particular exhibit. The packaging, battery, and antenna of the iReader needed to be adapted to accommodate the high duty-cycle of reads and the desire to create a highly visible artifact for visitors. The modified iReader is called an 'eXspot' and is shown in Figure 3 [14].

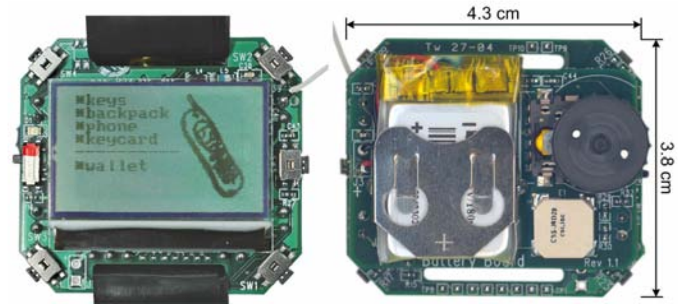

Figure 4. Display mote top and bottom.

Currently, the project has moved from a purely sensing system to including an actuation mechanism as well. This has enriched the possibilities for personalizing the post-visit web site. For example, a user can swipe their RFID card at the infrared camera exhibit and trigger two cameras – one thermal imaging camera and one for visible light – to take a picture of their group. The images are then associated with their RFID tag and linked to their personalized web page. When the user logs into the website the two photos are overlaid on top of each other so that the user can use a slider bar to adjust their relative transparency. This allows users to investigate which objects were giving off heat in their picture and explore the relationships between infrared energy and the physical objects in the picture. The Exploratorium is finding that the eXspot is a versatile device that will lead to new web interactions and learning environments where information gathered during the visit is used for learning after the visit and prolong the exposure to scientific concepts.

IV. DISPLAY MOTE

The driving motivation for the DisplayMote (shown in Figure 4) is to make a wrist-watch-sized platform with lightweight I/O capabilities to enable short messages from other devices to be displayed to the user, and for the user to give feedback to these devices. Messages to the DisplayMote are alarms, reminders, or lists of available services in the environment around the user. Messages from the DisplayMote are confirmations, selections, and actuations. Our goal is to create a fully programmable device that allows a user to provide input to and receive output from their own personal area network or a sensor network in their environment.

The DisplayMote is based on a UC Berkeley mote with the addition of a graphical 64x128 pixel LCD, a 2-axis accelerometer, a jog dial, a buzzer and five buttons. The mote design was extended while maintaining as much hardware and TinyOS compatibility as possible. This lowers the barriers to entry for others looking to use the DisplayMote to create a specialty I/O device. In a typical usage model the DisplayMote consumes approximately 14.5mA and is capable of running for more than 10 hours on a single 200mAh Li-Polymer battery. Each component on the DisplayMote is equipped with an FET allowing us to turn off the power to the various sub-systems and extend battery lifetime.

Minimizing the size of the DisplayMote was a key aspect to the design; the final form factor is a single PCB design, which is roughly the size of a wristwatch. The integrated 64x128 graphical LCD enables the use of text and graphical icons to communicate information to the user, thereby minimizing the need of more cumbersome devices such as PDAs or laptops. The buttons and accelerometer can be used for many different applications, including an accelerometer-based text entry method for very small devices derived from TiltType [15]. The device uses the push buttons and accelerometer to create a mouse, keyboard, and menuing system based on a tilt-and-click input modality. The DisplayMote is also designed for use as a remote control, a reminding device, and as an input device for kiosks, and digital public displays. It can also be use as a remote terminal to connect to specialized devices such as the Intel Personal Server [16] that provide large amounts of personal storage and computational power but have no integrated display of their own.

The DisplayMote enables a user to act as a "human sensor" and become a member of the sensor network. By showing messages and giving the user a means of input, a human can interact in real time with sensors using this lightweight user interface. This capability is useful in many ways. First, it can be used to check the deployment of sensor networks to ensure that each node is properly configured and working. Second, a user can collect data from nearby sensors to "see" what services are available nearby – for example, that there is a data projection service or wireless network available in the room.

A. Reminding Application

An interesting class of applications is one that uses the DisplayMote to alert a user when something in their environment requires their attention. One application we have developed uses the DisplayMote to issue reminders to a user when they leave behind important objects [17]. We use passive RFID tags attached to objects and long-range RFID readers installed in the environment that broadcast all the tags they read to their immediate vicinity over the mote radio. As the user goes through his or her day, the tag reads collected from the broadcasts by the readers in the environment are used to determine which items are currently on the user and which items are missing. The system then generates alarms on the DisplayMote to inform the user that he may have forgotten something.

The DisplayMote has worked out well as a generic platform for this application because the form factor is familiar to something that people already are used to and normally carry. The output capabilities of the screen, buzzer, and LED have worked well in signaling the user and communicating exactly what is missing. The DisplayMote buttons have enabled the user to give feedback to the system, such as to ignore a reminder. We plan on using this capability to make our system learn a user's habits in terms of which objects they normally have in their possession. Feedback from the buttons will provide an easy way to enable supervised learning for which reminders are appropriate.

V. CONCLUSION

Our goal is to enable the blending of personal area and environment-based sensor networks and understand how this, more generally, changes the requirements of sensor network services and protocols. We've created a set of devices that utilize standard wireless communication platforms to connect sensors and users. We leverage the low-power radio and sensor network protocol work already in progress on the UC Berkeley mote platform and Intel's iMotes to create general-purpose I/O devices (the iReader and DisplayMote) and activity inference capabilities (MSB). This set of programmable sensor building blocks for ubiquitous computing share a common programming language and low power communication protocol, making it easier for application developers to integrate them into a wide variety of new applications.

VI. FUTURE WORK

We have two major directions for our future work. The first direction seeks to develop approaches to dynamically allocating computations in blended sensor networks for activity inference. Deciding how to best use available bandwidth and power capabilities is an important area of research that should have general application in extending the lifetime of sensor networks.

The second direction is to develop interesting applications that exploit the availability of environmental sensors while having user interfaces that stretch across what the user is carrying (wristwatch, cell phone display, etc.) and what is available in the environment (large displays, speakers, etc.).

ACKNOWLEDGMENTS

We would like to thank Roy Want's team, who took our design and produced a prototype of the DisplayMote. Additionally, we would like to thank Ken Smith of Intel Research Seattle for his help with packaging solutions and Intel's iMote team for their invaluable assistance. We could not have created the mica2dot compatible RFID reader so quickly without the folks at SkyeTek Engineering. We would like to thank Kurt Partridge and Saurav Chatterjee for advice in building the DisplayMote. And finally we would like to thank our two volunteers, Megan Karalus and Jodi Shi, for their work in collecting and annotating our activity data.

ADDITIONAL CONTENT

Additional multimedia content is available online at:
http://www.cs.washington.edu/ubicomp/ubidevices

REFERENCES

[1] Burrell, J., Brooke, T., and Beckwith, R., "Vineyard computing: sensor networks in agricultural production", IEEE Pervasive Computing, pp 38–45, vol. 3, no. 1, 2004.

[2] Simon, G., Maróti, M., and Lédeczi, A., "Sensor Network-Based Countersniper System", Conference On Embedded Networked Sensor Systems, pp 1–12, 2004.

[3] Tennenhouse, D, "Proactive computing", Communications of the ACM, pp 43–50, vol. 43, no. 5, 2000.

[4] Hill, J., et al., "System architecture directions for networked sensors." Proc. 9th Int'l Conf. on Architectural Support for Programming Languages and Operating Systems pp. 93-104, 2000.

[5] Hill, J., et al., "The platforms enabling wireless sensor networks", Communications of the ACM, pp 41–46 , vol 47, no. 6, 2004.

[6] Culler, D., and Mulder, H., "Smart Sensors to Network the World", Scientific American, pp 84–91, 2004.

[7] Lukowicz, P., et al., "WearNET: A Distributed Multi-sensor System for Context Aware Wearables", Proceedings of the 4th international conference on Ubiquitous Computing, pp 361–370, 2002.

[8] Mainwaring, A., et al., "Wireless Sensor Networks for Habitat Monitoring", Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications, pp 88–97, 2002.

[9] Laerhoven, K., and Cakmakci, O., "What Shall We Teach Our Pants?", Fourth International Symposium on Wearable Computers, 2000.

[10] Choudhury, T., and Pentland, A., "Sensing and Modeling Human Networks using the Sociometer" Seventh IEEE International Symposium on Wearable Computers 2003, pp 216-222

[11] P. Viola, and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple Features," Computer Vision and Pattern Recognition, 2001.

[12] LaMarca, A., et al., "PlaceLab: Device Positioning Using radio Beacons in the Wild", Intel Research Tech Report IRS-TR-04-016, 2004.

[13] SkyeTekM1-mini http://www.skyetek.com/readers_Mini.html

[14] Hsi, S., et al., "eXspot: A Wireless RFID Transceiver for Recording and Extending Museum Visits", Demonstration at The Sixth International Conference on Ubiquitous Computing, 2004.

[15] K. Partridge, et al., "TiltType: Accelerometer-Supported Text-Entry for Very Small Devices," Proc. International Conference of User Interface Software and Technology 2002, pp 201-204.

[16] R. Want, et al, "The Personal Server: Changing the Way We Think about Ubiquitous Computing," Proc. International Conference of Ubiquitous Computing 2002, pp 194-209.

[17] G. Borriello, et al., "Reminding About Tagged Objects Using Passive RFIDs," Proc. Ubicomp 2004, pp 36-53.