By Joshua R. Smith, Kenneth P. Fishkin, Bing Jiang, Alexander Mamishev, Matthai Philipose, Adam D. Rea, Sumit Roy, and Kishore Sundara-Rajan

The iBracelet and the Wireless Identification and Sensing Platform promise the ability to infer human activity directly from sensor readings.

RFID-BASED TECHNIQUES FOR HUMAN-ACTIVITY DETECTION

Many ubiquitous computing scenarios require an intelligent environment to infer what a person is doing or attempting to do. Historically, human-activity tracking techniques have focused on direct observation of people and their behavior—with cameras, worn accelerometers, or contact switches. A recent promising avenue [4, 7] is to supplement direct observation with an indirect approach, inferring people's actions from their effect on the environment, especially on the objects with which they interact.

Researchers have applied three main techniques to human-activity detection: computer vision, active sensor beacons [4], and passive RFID. Vision involves well-known robustness and scalability challenges. Active sensor beacons provide accurate object identification but require batteries, making them impractical for long-term dense deployment. RFID tags have the same object-identification accuracy as active beacons, with the advantage of being battery-free; however, unlike sensor beacons, they are unable to detect motion.

Figure 1. Medical iGlove (left) and iBracelet (right).

Despite this daunting RFID limitation for tracking human activity, we've been pursuing two very different approaches, both based on RFID. The iBracelet is a wrist-worn short-range RFID reader that detects object use via hand proximity. The Wireless Identification and Sensing Platform (WISP) is a family of long-range RFID tags augmented with sensors that detect object motion; they eliminate the need to wear something by moving from short-range tags matched with wearable readers to long-range motion-sensitive tags read by fixed infrastructure.

While both approaches modify and extend conventional RFID, neither requires batteries in the objects being tracked. WISPs deliver the motion-detection capabilities of active sensor beacons in the same battery-free form factor as RFID tags using line-powered readers. The iBracelet system uses just one battery to power its wrist-worn reader and yields information about who is using particular objects not directly available through the WISP approach. The wearable reader also gives the subjects being monitored more control over the system, since they can more conveniently disable it—by taking it off—than they have with the fixed-reader infrastructure.

iGlove. We created the iGlove in 2003 as part of our first effort to track object use with RFID. While the early prototype was too crude for true long-term deployment, it was usable and durable enough that we were able to enlist 14 volunteers to wear it while conducting a range of daily household tasks, averaging approximately 45 minutes per user. By tracking the objects they grasped, we were generally able to figure out which activities they performed and when [7].

For our next test of the concept, we wanted a set of users to wear it while performing a true workplace task. Collaborating with the University of Washington Medical School, we enlisted seven faculty volunteers to use their simulation lab to perform test procedures while wearing iGloves (see Figure 1, left). All reported the form factor to be acceptable, and we

WISPs deliver the motion-detection capabilities of active sensor beacons in the same battery-free form factor as RFID tags, using line-powered readers.

were able to collect a good set of sensor data [2]. However, the effort the doctors had to exert was greater than we anticipated, highlighting the need for a more robust sensing platform.

iBracelet. Our iGlove experience validates the concept of inferring actions from RFID generally and wearable readers specifically. However, any glove, no matter how well-designed, is acceptable to only a limited user population. Accordingly, we next focused on bracelets, which are aesthetically and ergonomically much preferred to gloves. We needed to extend the read range from the 1–2 cm of a palm antenna to the 10–15 cm of a bracelet while staying lightweight, low-power, and within regulatory limits on exposure to electromagnetic fields. We have been able to achieve these goals by using a tuned circular loop antenna worn coaxially around the arm. When worn by one of the authors (Matthai Philipose) in his home it was able to duplicate the functionality of the iGlove.

The initial functioning prototype is evolving into something truly wearable by a larger population. Figure 1, right, shows the iBracelet's current form factor. For the first prototype, read range was a challenge, as the antenna was typically further from the tags than for the iGlove. However, working with our colleague Stephen Hughes of Intel Dublin, we overcame these range limitations by improving the antenna and other passive components.

The iBracelet retains the desirable features of the iGlove while moving to a less-obtrusive, more broadly acceptable form factor. However, the wearable-reader approach still involves open questions about basic feasibility. It is not possible to know in advance what combination of size, aesthetics, and battery life will satisfy a picky consumer or if such a satisfactory combination even exists.

WISP. A WISP consists of passive RFID tags augmented beyond the basic identification capability of ordinary RFID to support sensing and general-purpose computation. Like ordinary passive RFID tags, WISPs lack an on-board power source; instead, they harvest their power from RFID readers. All the WISPs we have created are backward-compatible with existing RFID standards and hardware. They can be powered and read using unmodified, commercially available RFID readers employing standard protocols. We hope that the pre-existing ecosystem associated with the Electronic Product Code standard we are extending will advance the user-engagement process more quickly than might otherwise be possible. Most previous work in battery-free sensing employed a chipless approach [3, 5] that allowed only a small number of distinguishable sensors not compatible with the RFID infrastructure.The α-WISP [6] uses anti-parallel mercury switches as both one-bit accelerometers and as modulating elements to selectively enable or disable a first or second RFID tag. Whereas ordinary RFID tags consist of a single chip mounted on a single antenna, the α-WISP consists of an antenna with two chips and two mercury switches. The switches are configured so that when the object is level (in its rest configuration, if it has one) the first ID chip is connected to the antenna; therefore the first ID is detected by the reader. When the object is tilted, the first chip is disabled and the second enabled, so the second ID is detected by the reader.

Detecting either of a WISP's two IDs indicates that a physical object is present—the same information that would be provided by an ordinary RFID tag. Both ID values showing up within a second or so indicates that the object, in addition to being present in the reader's field of view, is also moving.

We use the term "ID modulation" to describe the process of encoding information in a choice of RFID-tag ID values. While ID modulation may appear to support only single-bit communication, by controlling the pattern of ID changes over time—through more complex hardware than the α-WISP—continuous sensor bitstreams can be embedded at a constrained rate in an ordinary RFID channel. ID modulation allows WISPs to communicate limited amounts of application data—typically sensor values—through standards-compliant RFID communication channels.

One disadvantage of the α-WISP is that it's a single-axis device lacking sensitivity in two of the three spatial directions. While this limited sensitivity is acceptable for some applications, for activity detection it could lead to undetected object use. In [8], we described a three-axis accelerometer WISP with one bit per axis of dynamic range. The system used ID modulation to reliably encode all three bits in the stream of RFID reads. Unfortunately, reliability came at a high price: The first version of our protocol took approximately 10 seconds to communicate the three bits and the necessary synchronization header information. While the three-axis sensor would seem to be an improvement for activity detection, the update rate is too slow to support the application.

ID modulation allows us to create RFID sensor systems that work today, despite the fact that RFID sensing standards are not likely to be finalized for at least several years. Even when RFID sensing standards are in place, the ID modulation approach might still be used, since it allows system integrators to build

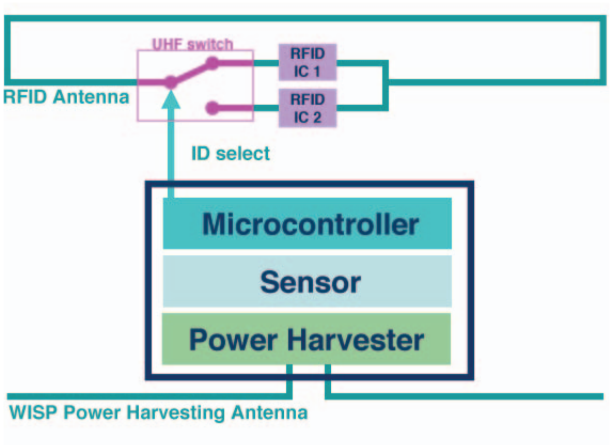

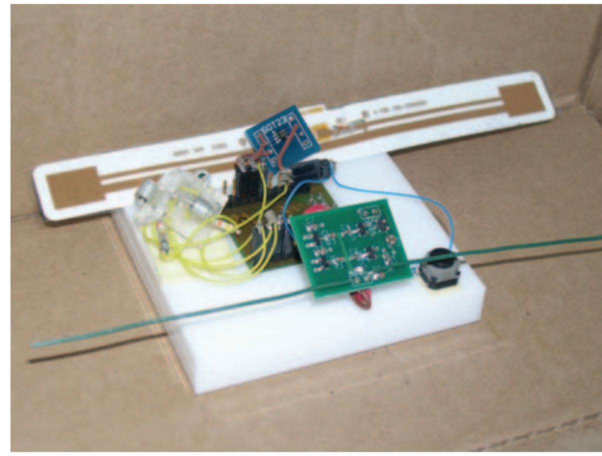

Figure 2. Block diagram of the WISP (left) and photograph of a functioning implementation (right). The WISP implementation uses separate antennae for power and for communication. The power-harvesting antenna is at the bottom of the diagram; in the photo, it is the pair of straight green wires emerging from the power-harvesting circuit board toward the bottom of the image. The microcontroller and sensor are above the harvester in the diagram. In the photo, the microcontroller is mounted on the reverse side of the circuit board in the center of the plastic mounting block. The assembly to the left of the microcontroller is the three-axis mercury tilt/acceleration sensor. At the top are the RFID chips and antenna used in the ID modulation process. The UHF RFID-select switch is mounted on the small circuit board in front of the RFID antenna.

solutions using the best (or least-expensive) available RFID tags and readers.

π-WISP for Activity Detection

The three-axis sensor WISP requires a much more complex and general-purpose hardware platform than the α-WISP. We call this more general-purpose platform the π-WISP; its sensing and modulating functions, which are performed by the mercury switches in the α-WISP, are cleanly split. The initial π-WISP consists of a power-harvesting circuit, an ultra-low-power microcontroller, a three-axis-by-one-bit mercury-switch-based accelerometer, and an electronically (rather than inertially) controlled 2:1 multiplexer that allows the microcontroller to connect a first or second RFID integrated circuit to the antenna. Figure 2 outlines the π-WISP's components (schematic on the left, photograph on the right). Our custom-designed power-harvesting unit (at the bottom of the photograph) consists of four cascaded voltage doublers based on zero-bias Schottky diodes, which convert the low-voltage RF signal from the reader to a larger DC voltage.

The brain of the π-WISP is a Texas Instruments MSP430F1121 ultra-low-power microcontroller, which in its low-power operating mode nominally draws as little as 160μA at 2.2V with a 1MHz processor clock speed, in standby mode only 0.7μA, and in off mode (with RAM retention) only 0.1μA.

The π-WISP's sensor consists of three orthogonally mounted mercury switches that function as a three-axis accelerometer with one bit of dynamic range per axis.

The modulation, or multiplexing, is performed by a NEC UPG152TA SPDT GaAs switch capable of switching high-frequency signals (up to 2.5GHz) and offering insertion loss that is low for a semiconductor device. However, the loss is not as low as the all-metal mercury switch used in the α-WISP. While the α-WISP can be read at 10 feet or more, the π-WISP's range is limited by losses in the switch to approximately four feet. The microcontroller itself continues to run off harvested power at five to six feet. In the terminology of RFID, the π-WISP is "reverse link limited," or limited by the ability of the reader to detect the tag's backscattered signal, whereas optimized RFID systems are typically "forward link limited," or limited by the supply of power from the reader to the tag. Future versions of the WISP will not rely on the GaAs switch and thus will have improved read range.

By selecting the first or second tag IDs—ID1 or ID2—the microcontroller can transmit a single bit of application information, with a sacrifice of a bit of ID space. By generating the appropriate switching pattern over time, multiple bits can be sent from WISP to RFID reader using this apparatus. To hardware and middleware unaware of the additional structure we have now introduced, the output of the RFID reader looks like an ordinary time series of RFID read events. But a decoder that is aware of the additional structure can extract the sensor data by examining the entire time series—not just single read events—and then decoding the embedded sensor data.

This bitstream communication method relies on the WISP microcontroller to add the necessary redundancy and synchronization header. We are now exploring another application of the WISP's computational power, allowing greater sensing capabilities than the α-WISP, without the update rate sacrifices that bitstream ID modulation would require. The new aggregation and communication scheme achieves a sensor frame update rate of 6Hz vs. 0.1Hz for our prior bitstream-coding scheme. But note that the two schemes are not directly comparable on the basis of rate alone; rather, they may be viewed as two distinct points on the rate-distortion trade-off curve. The earlier scheme sacrificed communication rate for reliability or distortion; the current scheme has a higher communication rate, along with a higher cost in terms of distortion. For activity detection, the higher distortion is acceptable, and the higher rate is desirable.

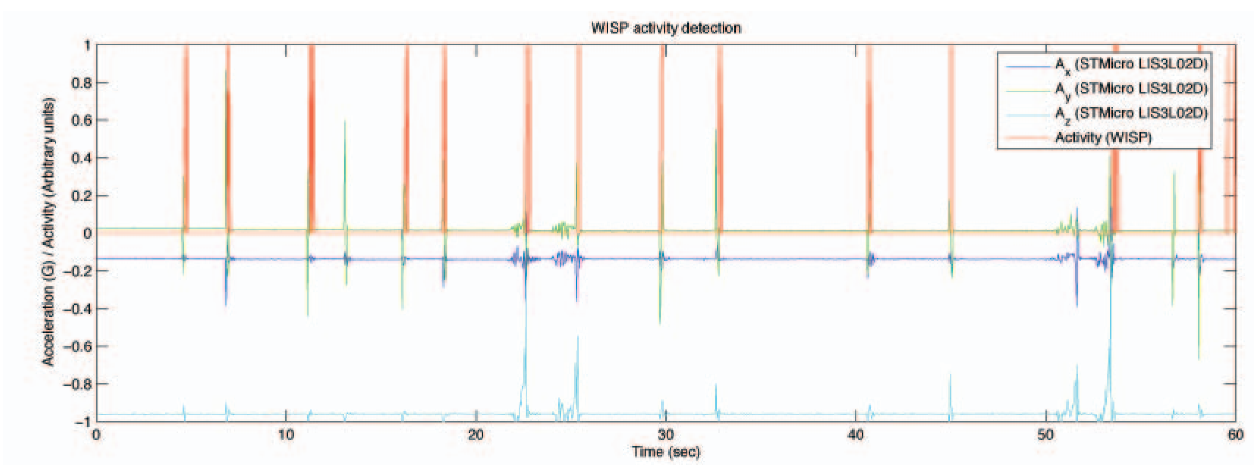

Figure 3. WISP activity detection. We subjected a plain cardboard box with a WISP mounted inside and a commercial three-axis accelerometer (an STMicroelectronics LISL02D) on the outside to shocks from various directions. The commercial accelerometer provides the "ground truth" against which we evaluated the WISP. The traces, labeled A_x, A_y, and A_z, show the shocks, as measured by the LISL02D. The red trace, labeled "Activity (WISP)," represents activity detected by the WISP. Of the 16 substantial shocks visible in the continuous accelerometer measurements, the WISP successfully detected and communicated 14 of them to the reader.

Data Aggregation

To operate at the higher communication rate, the WISP's microcontroller polls the three sensors, stores their values in RAM, and polls the sensors again. If any of the three values has changed since the previous reading, the microcontroller causes the ID returned by the WISP to toggle to its other state. The host analyzing the data from the WISP determines that the object has moved or been subjected to shock when it observes that the ID has toggled from one value to the other.

This aggregation method is just one of many compression schemes that might be considered for sensing applications; others include averaging, integration, differentiation, and thresholding. More general approaches (such as vector quantization and model-based coding schemes) are also worth considering. Just as lossy compression/coding schemes for images, audio, and video represent a rich area of research, lossy compression schemes tuned to other forms of sensor data should be as well. Since many of the constraints on WISPs differ from previous systems (such as sensor networks), it is reasonable to expect that the most appropriate sensor data coding schemes may be different for WISPs than for sensor networks.

In future work, we will examine in detail the rate-distortion trade-off curve for WISP channels, hopefully enabling us to make principled comparisons among sensor-data coding schemes.

Response of Sensors

Figure 3 shows a WISP performing object-activity detection. We mounted a WISP on a test object (a cardboard box) and attached a commercially available, wired three-axis accelerometer to the box. The figure shows the response of the sensors when we subjected the box to various small external impulses. We generated the x and y impulses by manually tapping the box on one side at a time. The z-axis impulses were generated by picking up and dropping the box from a height of about 2 cm. We subjected the box to 16 external shock-type impulses from all sides over a period of 60 seconds. These impulses appear in the figure on either the A_x, A_y, A_z trace or on a combination, depending on the direction of the shock.

The microcontroller toggles the ID it returns whenever the new sensor state disagrees with the previous state. Thus, it is collecting three bits of sensor data, performing a particular aggregation computation, and communicating back a one-bit result. The advantage of communicating back a one-bit result is that it requires no header overhead. The absence of header data allows for a much faster update rate; for example, using an Alien Technologies reader in Verify mode, we were able to poll the sensor and ID at up to 24Hz. The maximum frequency at which the WISP changes its ID is 6Hz, so this is the factor that limits the update rate.

At the host PC, the WISP's output signal is converted by software to an activity signal. Whenever the currently received ID changes, the software sets the activity level high. Simple failures to read an ID are common in the RFID channel, but they have little effect on the ability to detect activity, as long as the RFID read rate is faster than the activity timescale.

Of the 16 significant impulses visible in the figure, the WISP failed to detect two. Although the three-axis-by-one-bit accelerometer is more nearly isotropic than the single-axis version of the sensor, it is still far from uniform. The missed impulses may have been in a direction of low sensitivity. It is also possible that the WISP really did detect the impulses and changed its ID, but the reader may not have successfully read these ID values. The system's rate and distortion appear to be within the bounds required to support activity detection.

Conclusion

RFID-based sensing of object use provides computers a view of human activity that is unprecedented in detail and breadth. This family of techniques can help solve the longstanding problem of how to infer human activity. While the iBracelet and the WISP both use RFID for human-activity inference, each represents a quite different solution, appropriate in different circumstances. The WISP approach appears promising for a variety of sensing problems, in addition to human-activity inferencing, in which battery-free operation or compatibility with RFID infrastructure is key. The iBracelet appears most promising for industrial or enterprise-context human-activity inference applications in which a wearable reader is not a burdensome requirement. The WISP approach appears to be suited to more casual or consumer scenarios in which a wearable device is not acceptable. As researchers gain more experience in human-activity inference [1], it will be interesting to discover which combination of approaches turns out to be most suitable for real-world applications and environments. **C**

References

1. Dishman, E. Inventing wellness systems for aging in place. *Computer 37,* 5 (May 2004), 34–41.
2. Fishkin, K., Consolvo, S., Rode, J., Ross, B., Smith, I., and Souter, K. Ubiquitous computing support for skills assessment in medical school. Presented at the Second International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (Nottingham, England, Sept. 6–7, 2004); www.pervasivehealthcare.org/ubicomp2004/papers.
3. Fletcher, R. *Low-Cost Electromagnetic Tagging: Design and Implementation.* Ph.D. dissertation, MIT, Cambridge, MA, 2002; www.media.mit.edu/physics/publications/theses.
4. Munguia Tapia, E., Marmasse, N., Intille, S., and Larson, K. MITes: Wireless portable sensors for studying behavior. In *Extended Abstracts of the Sixth International Conference on Ubiquitous Computing* (Nottingham, England, Sept. 7–10, 2004); ubicomp.org/ubicomp2004/adjunct/demos.
5. Paradiso, J., Pardue, L., Hsiao, K.-Y., and Benbasat, A. Electromagnetic tagging for electronic music interfaces. *Journal of New Music Research 32,* 4 (Dec. 2003), 395–409.
6. Philipose, M., Smith, J., Jiang, B., Mamishev, A., Roy, S., and Sundara-Rajan, K. Battery-free wireless identification and sensing. *IEEE Pervasive Computing Magazine 4,* 1 (Jan.–Mar. 2005), 10–18.
7. Philipose, M., Fishkin, K., Patterson, D., Perkowitz, M., Hahnel, D., Fox, D., and Kautz, H. Inferring activities from interactions with objects. *IEEE Pervasive Computing Magazine 3,* 4 (Oct.–Dec. 2004), 50–57.
8. Smith, J., Jiang, B., Roy, S., Philipose, M., Sundara-Rajan, K., and Mamishev, A. ID modulation: Embedding sensor data in an RFID timeseries. In *Proceedings of The Seventh Information Hiding Workshop* (Barcelona, Spain, June 6–8, 2005). Springer-Verlag, Berlin, 2005.

Joshua R. Smith (joshua.r.smith@intel.com) is a member of the research staff at Intel Research Seattle, Seattle, WA, and an affiliate professor of electrical engineering at the University of Washington, Seattle.

Kenneth P. Fishkin (kenneth.p.fishkin@intel.com) is a researcher at Intel Research Seattle, Seattle, WA, and an affiliate professor of computer science at the University of Washington, Seattle, WA.

Bing Jiang (bjiang@ee.washington.edu) is a Ph.D. candidate in the Department of Electrical Engineering at the University of Washington, Seattle, WA.

Alexander Mamishev (mamishev@ee.washington.edu) is an assistant professor of electrical engineering and director of the Sensors, Energy, and Automation Laboratory at the University of Washington, Seattle, WA.

Matthai Philipose (matthai.philipose@intel.com) is a member of the research staff at Intel Research Seattle, Seattle, WA.

Adam Rea (Adam.d.rea@Intel.com) is a hardware engineer at Intel Research Seattle, Seattle, WA.

Sumit Roy (roy@ee.washington.edu) is a professor of electrical engineering in the Department of Electrical Engineering at the University of Washington, Seattle, WA.

Kishore Sundara-Rajan (kishore@ee.washington.edu) is a Ph.D. candidate in the Department of Electrical Engineering at the University of Washington, Seattle, WA, and an intern at Intel Research Seattle, Seattle, WA.
