

Battery-Free Wireless Identification and Sensing

Wisps, a class of battery-free wireless sensors, scavenge power from ambient radio-frequency-identification readers to communicate sensed data over room-sized areas. A simple usage model enables a variety of distributed sensing applications.

Collections of tiny, inexpensive wireless sensor nodes capable of continuous, detailed, and unobtrusive measurement have attracted much attention in the past few years.¹ Prototypes exist for applications such as early detection of factory equipment failure, optimization of building energy use, habitat monitoring, microclimate monitoring, and monitoring structural integrity against earthquakes. Unfortunately, the very properties that make sensor nodes attractive for these applications—low cost, small size, wireless functioning, and timely, long-term radio communication of data—make powering them a challenge. Particularly challenging is the production of an energy source

that's tiny and has a long mean time to replacement yet can supply enough power for wireless communication of sensed data. Recent work has focused on batteries and ambient power scavenging to provide power for sensor networks (see the "Power Supply Options" sidebar for a more in-depth discussion of these power sources).

The Wireless Identification and Sensing Project (WISP) explores a third approach, based on passive radio-frequency-identification technology.² In traditional passive RFID systems, ambient high-power readers interrogate battery-free

devices, called *tags*, that modulate the interrogating signal to communicate a unique identifier to the reader. WISP aims to augment RFID tags with sensors so that tags can also send sensed data to the readers. We call these augmented tags *wisps*. Basing wisps on RFID has some immediate advantages. RFID tags communicate to ambient readers over distances of up to eight meters. So, a single US\$1,500 four-antenna reader currently available off the shelf should cover four rooms in a home. The tags can be read at nominal rates of up to 2,000 per second. They have the form factor of mailing labels, cost roughly \$.50 each, and can communicate even when obscured by many materials (conductive and water-rich materials are notable exceptions). Perhaps most important, business and manufacturing trends support the development and adoption of smaller, cheaper, and longer-range tags. Solutions compatible with RFID standards might therefore find quicker acceptance and see faster improvement than other solutions. The "Other Approaches to Reader-Powered Sensors" sidebar discusses other RFID-based approaches.

Many challenges remain, however. Many applications proposed for conventional sensor networks require instrumentation of large, often outdoor areas—clearly beyond the range of RFID-style readers. Some applications are also liberal in the density and size of tags needed: in many cases, even \$1 to \$10 solutions might not be too expensive, and relatively large devices might be acceptable. Furthermore, long-range

Matthai Philipose and Joshua R. Smith

Intel Research Seattle

Bing Jiang, Alexander Mamishev, and Sumit Roy

University of Washington

Kishore Sundara-Rajan

Intel Research Seattle and University of Washington

Power Supply Options

Recent distributed sensing work has explored two main power supply options: batteries¹⁻³ and ambient power scavenging.^{4,5} Battery technology is mature, extensively commercialized, and completely self-contained. However, given current energy density and shelf-life trends,⁵ even for relatively large batteries and conservative communication schedules, the mean time to replacement is only a year or two. The problem is aggravated significantly for batteries with more inconspicuous form factors. For deployments with hundreds of sensors, this means that a battery will need a replacement every few days, an unsustainable rate for many applications. Finally, systems based on battery-powered communication still cost more than US\$5 a unit, with no clear commercial or technical route to a solution that costs less than a dollar.

Ambient-power scavenging, which harvests energy (especially solar energy) from the surroundings for wireless sensing, can in principle supply power indefinitely.^{4,5} Implementations have provided low-duty-cycle communication at very small size under ambient conditions that are reasonable for many applications. However, challenges remain: reliance on ambient power constrains both where you can place the sensors and when you can use them. Furthermore, given these devices' use of power-harvesting components (such as

solar cells) and relatively sophisticated radios, it remains to be seen whether these devices can be manufactured at their target cost of tens of cents each. Almost as important, the commercial case for high-volume manufacture of these devices is still being made.

REFERENCES

1. M. Feldmeier and J. Paradiso, "Giveaway Wireless Sensors for Large-Group Interaction," *Extended Abstracts 2004 Conf. Human Factors in Computing Systems (CHI 2004)*, ACM Press, 2004, pp. 1291-1292.
2. J. Hill et al., "System Architecture Directions for Networked Sensors," *Proc. 9th Int'l Conf. Architectural Support of Programming Languages and Operating Systems (ASPLOS IX)*, ACM Press, 2000, pp. 93-104.
3. I. Munguia-Tapia et al., "MITes: Wireless Portable Sensors for Studying Behavior," *UbiComp 2004 Adjunct Proc.*, <http://ubicomp.org/ubicomp2004/adjunct/demos/tapia.pdf>.
4. J. Paradiso and M. Feldmeier, "A Compact, Wireless, Self-Powered Push-button Controller," *Proc. 3rd Int'l Conf. Ubiquitous Computing (UbiComp 2001)*, Springer-Verlag, 2001, pp. 299-304.
5. S. Roundy et al., *Energy Scavenging for Wireless Sensor Networks*, Kluwer Academic Publishers, 2003.

Other Approaches to Reader-Powered Sensors

Reversibility, ID space, and range distinguish approaches for acquiring and transmitting sensed data using scavenged signals from ambient readers.

Reversibility refers to whether the tag can be reused after sensing the phenomenon of interest. In some cases, because the phenomenon (commonly, exceeding a threshold temperature for perishables) irreversibly changes the tag structure, the tag lasts for only one measurement.

ID space refers to the number of distinct identifiers available to tags. To provide globally unique identification of the sensors and their associated objects, it's crucial to have at least 32 ID bits. The MIT Media Lab developed a series of devices that use the chipless RFID approach, in which the individual tags have distinct resonances that nearby readers can discern. Variations in magnetic-coupling strength can indicate tag movement.¹ More generally, the use of smart materials can change the tag's resonance frequency on the basis of ambient conditions, letting tags measure temperature, force, or humidity. Unfortunately, the number of distinguishable resonances available to these chipless approaches isn't sufficient to enable globally unique identification.²

Finally, to the best of our knowledge, all existing sensors that are

based on RFID (radio frequency identification) use near-field coupling and have ranges of a meter or less. The project most like WISP (Wireless Identification and Sensing Project—see the main article) is the University of Auburn's RFID sensor project, which uses (custom) passive RFID sensors coupled with biosensors to detect food spoilage. The project's most recent publication sketches a large (saucer-sized) inductively coupled short-range (10cm) battery-powered tag, with a stated goal of removing these constraints.³

REFERENCES

1. J. Paradiso, K. Hsiao, and A. Benbasat, "Tangible Music Interfaces Using Passive Magnetic Tags," *Proc. ACM Conf. Human Factors in Computing Systems: Special Workshop on New Interfaces for Musical Expression (CHI 2001)*, ACM Press, 2001.
2. R. Fletcher, *Low-Cost Electromagnetic Tagging: Design and Implementation*, PhD dissertation, MIT, 2001.
3. S. Nambi et al., "Radio Frequency Identification Sensors," *Proc. 7th World Multiconf. Systemics, Cybernetics, & Informatics (SCI 2003)*, Int'l Inst. of Informatics and Systemics, 2003, pp. 386-390.

RFID tags are susceptible to collision (limiting the density of readable tags), harvest very small amounts of power from the reader's signal (making it unclear whether we can power circuitry to acquire and transmit extra information from sensors), and are quite sensitive to environmental effects such as occlusion and presence of metal. Finally, it's unclear how to convert an RFID tag/reader system into a wisp transmitting sensed data in a manner

Our usage model for wisps enables a variety of applications related to detecting day-to-day human activities.

that preserves the RFID tags' form factor and is completely compatible with RFID protocols.

Nevertheless, we argue that distributed sensing based on wisps is useful and feasible. Our usage model for wisps enables a variety of applications related to detecting day-to-day human activities. Our working prototype wisp, a 1-bit accelerometer we call the α -wisp, is fully compatible with existing RFID protocols and in principle enables these applications. Measurements show that the α -wisp's performance is, for the most part, in the range required by the applications.

Usage model

We envision a simple usage model for wisps. Long-range RFID reader antennas are distributed around the area to be observed so that readers can provide blanket coverage of the area. These readers are typically connected directly to a power outlet in the wall, but they can be mobile. Next, we attach wisps to points of interest—for example, objects, people, and structures that either are in the space or will enter it. Finally, we set the readers to scan the space continuously. The wisps respond to the reader with their identifier

and sensed value. The ID lets the sensed value be associated with a unique sensor. The application can then act on the sensor snapshot of the space as appropriate.

The α -wisps are intended to measure the acceleration of the objects to which they are affixed. When interrogated, α -wisps transmit a 1 along with an ID if the object is moving or otherwise out of its rest configuration. If the object is in the rest configuration, the α -wisp either transmits a 0 and its ID or doesn't respond at

all, depending on the application.

Applications

The high-density limited-range wisp usage model is particularly well-suited to at least one rich, important family of applications—that is, indoor human activity tracking. The WISP project is part of the System for Human Activity Recognition and Prediction. SHARP models human activities probabilistically in terms of the sequence and duration of the objects' use during those activities. Given a trace of objects used, an inference engine tracks the likelihood of various activities in progress. Early results have shown that given these observations, the inference engine can tell with high certainty and in considerable detail what physical activity a person is performing.

SHARP relies on getting traces of objects actively being used in the activities it's tracking. Currently, the project relies on hand-worn RFID readers to report touch (and therefore use) of RFID-tagged objects. Wisps promise to remove the need for the wearable device, which is unsuitable for many applications. Motion (in particular, a change in orienta-

tion) is typically a good indicator that an object is in use. Tagging objects of interest with the α -wisp will potentially let us detect object usage unobtrusively and in great detail.

Caregiving for elders

Caring for elders is an emerging challenge in many societies.³ A central challenge in elder care is facilitating their independent living while ensuring that they're going about their daily lives with sufficient competence. To monitor an elder's activities, caregivers must spend substantial time near the individual. Unfortunately, given current social and economic trends, neither family nor professional caregivers can meet the required time commitment, motivating the need for a semiautomated solution.

We're investigating two applications in this space. The electronic Activities of Daily Living⁴ (eADL) form attempts to automatically fill in a state-mandated form requiring professional caregivers to rate their elderly clients' ability to perform various activities, such as housework, making snacks, and grooming. Professional caregivers would use this form as a starting point when visiting elders, rather than spending time attempting to recreate the information. A second application, the CareNet Display, provides family members with up-to-date information on whether elder relatives completed key activities.⁵ Family members can coordinate elder care using information from the display.

To get the trace in either application, we tag a variety of household objects with α -wisps and place RFID antennas in room corners. Our experience over the past year has indicated that sensing the vast majority of ADLs will probably require tagging about a thousand objects over the entire house (objects can be as small as a box of tea or a spatula). At less than \$0.50 a tag, a thousand tags per house seems acceptable for multiple-year

elder care support. Our experiments show that even partial deployments are useful. Typically, fewer than 10 objects are in active use at any time.

Another key aspect of this kind of home-based deployment is the health risk of having an always-on RFID reader in the space. Current standards require that users be more than 0.5 meter from the reader, implying that readers will have to be securely mounted high on walls or otherwise sufficiently removed from users.

Best-known method management

Most organizations have recognized experts to perform key tasks. The techniques the experts use are termed *best-known methods*. Managing BKMs, which includes capturing them, training workers to use them, and enabling their use on the work floor, is a strategic advantage for organizations. SHARP is particularly interested in BKMs related to manual activities such as operating machinery, cleaning, caring for patients, and assembling parts.

We're investigating two applications in this space. In one application, we monitor novice anesthesiologists to evaluate how well they perform relative to experts.⁶ Currently, onsite instructors perform these evaluations in an ad hoc manner. The instructors are receptive to augmenting the evaluation process with detailed sensors both because the US government is pushing for more quantitative grading criteria and because they feel that detailed information could improve training.

The second application monitors workers in a semiconductor fabrication plant to reduce system downtime due to human error. In these plants, the machines are so highly optimized for performance that human error is increasingly the dominant cause of downtime.

These applications involve tagging various tools, containers, and machines with wisp to track the progress of the anes-

thesiologists and plant workers, and positioning readers over designated work areas. Our preliminary studies suggest that the presence of large metal structures in constricted areas and requirements for electromagnetic shielding complicate wisp deployments in these milieus. On the other hand, because the activities are much more structured and the performers more cooperative than in the home-sensing case, we can carefully place readers and tags. We would likely tag tens to

The most common autonomous action studied so far is prompting users. Our particular interest is in activity-based reminders.

hundreds of potentially small objects (such as screwdrivers or syringes) in each work area.

Personal activity-based prompts

Proactivity is a stated goal of pervasive computing. In the proactive usage model, the computer need not wait for explicit interactive commands from the user. Rather, it acts autonomously on the basis of implicit triggers associated with user activities. The most common autonomous action studied so far is prompting users. Our particular interest is in activity-based reminders (for example, your phone might remind you to take your pills when you sit down to a meal). We envision that simply tagging the set of objects relevant to an activity should enable reminders for that activity.

Design requirements

Our early experiences with the applications suggest some constraints on α -wisp design:

- To minimize the density of RFID readers (which serve as wisp readers), wisps should have a read-range comparable to that of the underlying RFID system.

- Unobtrusively tagging mundane objects requires very small tags. We should be able to tag a spatula with little change in its usability, for example.
- Because we might tag hundreds of objects, tags should be very cheap. A dollar is (barely) acceptable for health-care applications; \$10 isn't.
- Tens to hundreds of tagged objects can be close together in front of an antenna, although objects in use tend to be a further from other objects. The

reader should be able to spot the few wisps moving at any given time regardless of the number of tags in the space.

- Most objects tend to move for a few (up to 10) seconds when in use, with substantial change in orientation during the process. Readers should be able to detect orientation changes robustly within this window.

α -Wisp design

The α -wisp uses ID modulation to communicate its sensor data. Rather than the obvious approach of encoding sensor data in special bits reserved in the protocol, we encode the sensor information in correlations in the pattern of ID reads. More specifically, we associate each object with two IDs instead of just one. We mount two off-the-shelf RFID ICs on each object. The sensing subsystem controls which of the two IDs is returned by short-circuiting one or the other of the ID chips. If the first ID is detected, it indicates that the tagged object is present and its sensor is in state one; if the second ID is detected, we know that the object is present and has sensor state two. This is equivalent to allocating one bit of ID data to sensing because it causes the ID space's




Figure 1. a-Wisp serial-switch configuration: (a) schematic, (b) photograph, (c) deployment on a coffee cup.

size to shrink by one bit. Future wisps will use more sophisticated form of ID modulation in which multiple bits of sensor data are encoded in patterns of ID change over time.

ID modulation lets us add new capabilities, such as sensing, while maintaining compatibility with existing protocols, standards, and reader hardware. Building the new capabilities as an additional application layer on top of the unmodified preexisting lower layers lets application-layer capabilities grow and develop independently of the lower layers. The investment in these lower layers can be preserved, while still allowing innovation at the application layer. By contrast, introducing a new protocol with bit positions reserved for sensor data would require writing off previous investments in reader hardware as well as new investments in replacement reader hardware and standards development.

One-bit accelerometer

We've created an experimental one-bit accelerometer that can be read by ID modulation. It requires no additional power supply and thus is a battery-free,

wirelessly readable, uniquely identified one-bit inertial sensor. It's compatible with the increasingly important 915-MHz Electronic Product Code (EPC) standard for RFID. Because the sensor information is represented at the application layer, the system is backward compatible with the existing protocol.

We implemented the experimental sensor using mercury switches. A single mercury switch can serve as a one-bit, one-axis accelerometer. Suppose the switch is initially far from any gravitational field. The switch has a particular orientation, an axis along which the mercury is free to move. When the switch housing accelerates in one direction, the mercury appears (in the switch housing frame of reference) to move in the opposite direction because of its inertia. When the mercury is at one end of its container, the switch is closed; when the mercury is at the other end, the switch is open. So, the switch will be in one state (closed, say) when the acceleration vector's projection onto the switch orientation vector is positive. The switch will be in the opposite state (open, in this example) when the sign of the acceleration in the


switch-orientation direction is negative. Because of gravity, accelerometers are also good tilt sensors: when the object's orientation changes, the three components of the gravitational force vector change in the object's frame of reference. Thus, even a one-bit accelerometer functions effectively as a tilt sensor.

We can use the switch's state to modulate the information returned by the tag. If the switch is connected in parallel with the RFID IC, when the switch is open, the tag will return its ID normally. When the switch is closed, the tag will be shorted out and won't return an ID. If the switch is connected in series between the RFID chip and the antenna, the ID reads normally when the switch is closed. Both of these configurations correspond to on-off keying (OOK) modulation. (Here, "modulation" refers not to the low-level physical-layer communication protocol but to modulation of the entire tag ID presence.)

For many applications, OOK modulation is a disadvantage, because distinguishing "object present, acceleration state 0" from "object absent" isn't possible. But using two antiparallel switches and two RFID ICs, we have implemented a *binary-code-shift keying* scheme that doesn't have this problem. Under positive acceleration, the first ID is returned; under negative acceleration, the second ID is returned. If neither ID is returned, the object is absent.

Binary-code-shift keyed ID modulation can be implemented in the serial-switch or the parallel-switch configuration. The serial-switch configuration (which we used for the α -wisp we describe in the "Performance results" section) consists of two RFID IC chips attached to one antenna with a mercury switch in series with each chip, as Figure 1a shows. The parallel-switch configuration uses two IDs and two antennas, with the switches

Figure 2. The dependence of read rate on read distance. Data is for an α -wisp serial-switch configuration. Read rate is the fraction of reads in which the reader detected a single stationary tag with no occlusion in typical indoor conditions. α -Wisps are detectable up to 3.8 meters from the reader. Read rates are lower than for unmodified RFID tags.



short-circuiting one of the IDs. By mounting the mercury switches in a geometrically antiparallel configuration, we ensure that when one switch is open, the other is closed. This way, only one RFID IC is enabled at any given time. Which ID is returned depends on the object’s acceleration: parallel to switch one corresponds to ID 1; parallel to switch two corresponds to ID 2.

Variants

We have built additional α -wisp variants. One lets us use just one switch but two tags and two chips. We apply an unmodified tag to the object as a presence detector, and use the second tag (with the switch) to encode the sensor data via OOK. Other extensions are possible. We can wire more than two IC-switch pairs in parallel to sense additional states or use more than one switch with each IC. Wiring the switches in parallel with each other (and as a group in series with the chip) causes the ID to return when any switch is triggered. Wiring the switches in series with each other causes the ID to return only when all the switches are triggered.

A fundamental trade-off of RFID-style communication is between the number of tags sharing the channel to the reader and the rate at which each tag can communicate with the reader. In general, the more tags in the reader’s view field, the lower the communication rate between each tag and the reader. When using RFID technology to implement dense sensor networks (as with wisps), you need to manage this trade-off in a way that meets the sensor deployment’s overall performance requirements. Existing

RFID technology provides built-in methods to mitigate this problem. For instance, clients of RFID readers can put particular tags to sleep so that they don’t use the channel for an interval of time. In a sensor network context, clients can let tags remain quiet unless the sensor has significant data to report. For example, we’ve implemented a variant of the α -wisp that shorts the RFID antenna when the object is in its rest position, so only objects in use—that is, those being moved—consume reader channel resources.

Performance results

We fabricated a serial switch α -wisp using Alien Technologies’ ALL-9250 EPC-compliant UHF RFID tags and Signal Systems International’s Series 5008pc mercury switches. We oriented the switches to be approximately antiparallel to one another. Figure 1b shows a similar α -wisp.

The RFID tags are 1 cm by 10 cm, and the switches are roughly 5 mm in length and diameter. As Figure 1c shows, an α -wisp fits comfortably on a coffee cup. Furthermore, the tags cost roughly \$.40 each, and the switches cost \$.50. A one-switch, one-tag α -wisp therefore costs

\$.90; double for the two-switch, two-tag version used in our tests. Although for per-item RFID tagging to be viable in supply chain applications, tag cost needs to drop much lower, a \$1 cost for tagging select items is reasonable for the healthcare and BKM applications we’re considering.

Figure 2 shows RFID read rates as a function of distance. The reader polled the tag under computer control at approximately 6 Hz (6 polls per second) for about 20 seconds. The read rate is the ratio of tag responses to read attempts. The top curve is the read rate for an unmodified RFID tag; the lower curves are rates for α -wisps in two different orientations. We generated the green curve (α -wisp, down) by placing the cup upside-down, which activates RFID IC 2—the extra IC we added to the antenna. The red curve represents responses from the IC that the manufacturer originally included with the antenna.

Despite the modifications, the tag continues to be readable (at a degraded rate) to three meters. At two meters, the modified tag is readable at nearly its full rate. We conducted this test under somewhat ideal circumstances (a metal-free




Figure 3. The effects of a worst-case dense tag deployment. The fraction of deployed tags that the reader can read is inversely proportional to the number of tags deployed. The tags are spaced uniformly in a 0.6-meter-wide paper envelope, read over 30-second windows, at a 0.9-meter range.

environment and stationary tag), so the particular range values should be viewed as upper bounds for what's achievable with this α -wisp implementation. However, we didn't optimize the implementation itself for range and made no special effort to minimize the impact of attaching the sensor to the tag. Furthermore, unmodified RFID tags have exhibited these ranges even in nonideal conditions (such as assembly lines and warehouses). We expect a more careful wisp implementation to have nearly the same communication range as the accompanying RFID tag. Wisps should therefore be able to leverage the trend in range improvement of commercial UHF RFID tags. Recently released tags, for instance, have a nominal range of 6 to 10 meters. At these ranges, we could cover four rooms (with one antenna each) using a single \$1,500 four-antenna reader.

Figure 3 shows the potential impact of a worst-case dense deployment on tag detection rate. The curve's x -axis is the number of unmodified RFID tags evenly spaced along a 0.6-meter-long paper envelope. The antenna reading the tags is 0.9 meters away from them. The y -axis is the fraction of placed tags that the reader detects in a 30-second window. In the

absence of special measures, the fraction of tags detected clearly falls off steeply with tag density. We emphasize that this figure illustrates a worst-case scenario that is nevertheless quite possible in a wisp deployment. We expect improvements both from using different reader configurations and from more common lower-density deployments. However, in applications in which very few sensors are simultaneously active, optimally disabling inactive sensors can be key in enabling a high-density deployment.

Figure 4 shows the α -wisp's operation. We mounted the wisp and a commercial accelerometer on a coffee cup, which we tilted periodically by motor at 0.18 Hz. The RFID reader provides wireless communication and power for the α -wisp; a wired USB connection provides communication and power for the STMicroelectronics LIS3L02D accelerometer. The blue trace shows the acceleration in the x -direction measured by the accelerometer incorporated into the Intel Research/University of Washington iMote sensor board. The superimposed red trace is the α -wisp's output. The red trace is clearly a one-bit quantization of the blue trace. The quantization threshold is nonzero because of the angle at which

we mounted the α -wisp on the coffee cup. The y -axis label for the red trace is arbitrary: the +0.8 G and -0.8 G levels shown actually correspond to ID 1 and ID 2.

The thickening of the red trace visible at some transitions is due to the sensor switch's *bouncing*—that is, as it transitions, the switch state can oscillate back and forth before settling to its final new value. The host application could “debounce” the α -wisp output using a simple software-filtering operation of the kind commonly used to debounce computer keyboard input.

Low-to-high transitions started at $0.14 \text{ G} \pm 0.08 \text{ G}$ (8.2 degrees \pm 4.4 degrees) and ended at $0.28 \text{ G} \pm 0.09 \text{ G}$ (16.5 degrees \pm 5.2 degrees). (In other words, we're reporting the mean and standard deviation of the distribution of acceleration values measured before and after a low-to-high transition.) High-to-low transitions started at $0.23 \text{ G} \pm 0.06 \text{ G}$ (13.0 degrees \pm 3.5 degrees) and ended at $0.10 \text{ G} \pm 0.07 \text{ G}$ (5.7 degrees \pm 4.3 degrees). We can't observe the actual point at which the one-bit sensor changes; we can only see the acceleration values immediately before and immediately after the transition. So, the true low-to-high threshold is probably around 12.3 degrees (averaging the values before and after the transition), and the true high-to-low threshold is probably around 9.3 degrees. The three-degree difference is hysteresis and is likely due to imperfect alignment of the (nominally) antiparallel mercury switches. The spread in transition values, approximately 0.1 G or 5 degrees, is the sensor's effective angular-resolution limit.

The measurement time resolution—that is, the rate at which the PC polled




Figure 4. An α -Wisp parasitically powered one-bit accelerometer in operation. Comparison of data returned by the α -wisp (the red trace) with continuous acceleration data provided by an STMicroelectronics LIS3L02D three-axis accelerometer (the blue trace).

the reader—is approximately 0.17 seconds. The RFID interrogation process appears to be the factor limiting the time resolution. Some of the apparent acceleration/angle noise is presumably due to the time quantization. Angular resolution could likely be improved somewhat by filtering out switch bounces, though at some cost in time resolution.

A natural question is whether we can generalize the working α -wisp sensor to enable more sophisticated sensing capabilities. The answer is that we can generalize it and are working to do so.

The mercury switches in the α -wisp support two logically distinct functions:

sensing (the one-bit accelerometer) and communication (ID modulation). In our next generation of wisps, these functions will be cleanly split. Power-harvesting circuitry will drive an ultra-low-power microcontroller with an analog-to-digital converter that we can connect to any compatible sensor, subject to power constraints. The microcontroller will perform ID modulation using an electronic switch, which, unlike the mercury switch, won't play a role in sensing. By actively modulating the pattern of IDs returned over time, the microcontroller will be able to send arbitrary multibit sensor data, at a limited rate, through a conventional, standards-compliant RFID reader channel.

Just as in the single-bit sensor data

case, multibit data communicated by ID modulation requires only two IDs. The pattern, spread over time, of switching between these two ID values will represent the multibit sensor information. So, moving from single- to multibit sensor data won't require sacrificing additional bits of ID space. We will give up just one bit of ID space and gain the ability to send multiple bits of sensor data through the reader channel, at a constrained rate. The generalized platform we envisage could support multibit measurement of parameters such as temperature, light, strain, and acceleration. ■

ACKNOWLEDGMENTS

Thanks to Jonathan Lester for help with his iMote Sensor Board.



Matthai Philipose is a member of the research staff at Intel Research Seattle. His research interests include statistical reasoning and programming languages. He received his MS in computer science from the University of Washington, Seattle. Contact him at Intel Research Seattle, 1100 NE 65th St., 6th floor, Seattle, WA 98105; matthai.philipose@intel.com.



Joshua R. Smith is a staff research scientist at Intel Research Seattle. His research focuses on the detection of weak signals for sensing or communications purposes. He has also published in the field of digital watermarking and invented a technology for identifying physical objects using human-imperceptible random variations in surface characteristics. He has a PhD from the MIT Media Lab's Physics and Media group. He is a member of the IEEE and the AAAS. Contact him at Intel Research Seattle, 1100 NE 45th St., Seattle, WA 98105; joshua.r.smith@intel.com.



Bing Jiang is a PhD candidate at the University of Washington. His research interests include radio-frequency-identification technology, sensors, robotics, and power. He received his MS degree in electrical engineering from the University of Washington, Seattle. Contact him at Dept. of Electrical Engineering, Box 352500, Univ. of Washington, Seattle, WA 98195; bjiang@ee.washington.edu.



Alexander Mamishev is an assistant professor of electrical engineering and director of the Sensors, Energy, and Automation Laboratory (SEAL) at the University of Washington. His research interests include sensor design and integration, dielectrometry, electric power applications, bioengineering, and robotics. He received his PhD in electrical engineering from the Massachusetts Institute of Technology. He's a member of the IEEE. Contact him at Dept. of Electrical Engineering, Box 352500, Univ. of Washington, Seattle, WA 98195; mamishev@ee.washington.edu.



Sumit Roy is a professor of electrical engineering at the University of Washington. His research interests include analysis and design of communication systems and networks, with a topical emphasis on next-generation mobile and wireless networks. He received his PhD in electrical engineering from the University of California, Santa Barbara. He is a member of the IEEE. Contact him at Dept. of Electrical Engineering, Box 352500, Univ. of Washington, Seattle, WA 98195; roy@ee.washington.edu.



Kishore Sundara-Rajan is a PhD candidate at the University of Washington and an intern at Intel Research Seattle. His research interests include sensor design and integration, on-chip power scavenging, RFID, and dielectric spectroscopy. He received his MS in electrical engineering from the University of Washington, Seattle. He's a student member of the IEEE. Contact him at Dept. of Electrical Engineering, Box 352500, Univ. of Washington, Seattle, WA 98195; kishore@ee.washington.edu.

REFERENCES

1. D.E. Culler and H. Mulder, "Smart Sensors to Network the World," *Scientific American*, June 2004, pp. 85–91.
2. K. Finkenzerler, *RFID Handbook*, 2nd ed., John Wiley & Sons, 2003.
3. E. Dishman, "Inventing Wellness Systems for Aging in Place," *Computer*, vol. 37, no. 5, 2004, pp. 34–41.
4. M. Philipose et al., "Inferring ADLs from Interactions with Objects," *IEEE Pervasive Computing*, vol. 3, no. 4, 2004, pp. 50–57.
5. S. Consolvo et al., "Technology for Care Networks of Elders," *IEEE Pervasive Computing*, vol. 3, no. 2, 2004, pp. 22–29.
6. K. Fishkin et al., "Ubiquitous Computing Support for Skills Assessment in Medical School," *Proc. Ubichealth 2004: The 3rd Int'l Workshop Ubiquitous Computing for Pervasive Healthcare Applications*, www.pervasivehealthcare.com/ubicomp2004/papers/final_papers/fishkin.pdf.