

Editors: **Frédéric Thiesse** • frederic.thiesse@unisg.ch
Florian Michahelles • fmichahelles@ethz.ch

Building the Internet of Things Using RFID

The RFID Ecosystem Experience

At the University of Washington, the RFID Ecosystem creates a microcosm for the Internet of Things. The authors developed a suite of Web-based, user-level tools and applications designed to empower users by facilitating their understanding, management, and control of personal RFID data and privacy settings. They deployed these applications in the RFID Ecosystem and conducted a four-week user study to measure trends in adoption and utilization of the tools and applications as well as users' qualitative reactions.

**Evan Welbourne,
 Leilani Battle, Garret Cole,
 Kayla Gould, Kyle Rector,
 Samuel Raymer,
 Magdalena Balazinska,
 and Gaetano Borriello**
University of Washington

The rapid proliferation of passive RFID tags in the past decade has given rise to various concepts that integrate the physical world with the virtual one. One of the most popular is the Internet of Things (IoT), a vision in which the Internet extends into our everyday lives through a wireless network of uniquely identifiable objects. Given numerous predictions that we'll have hundreds of billions of RFID-tagged objects at approximately five cents per tag by 2015,¹ we're not only approaching such a world, we're on its doorstep.

In this type of RFID system, each physical object is accompanied by a rich, globally accessible virtual object that contains both current and historical information on that object's physi-

cal properties, origin, ownership, and sensory context (for example, the temperature at which a milk carton is being stored). When ubiquitous and available in real time, this information can dramatically streamline how we manufacture, distribute, manage, and recycle our goods. It can also transform the way we perform everyday activities by giving applications current and detailed knowledge about physical events. This "real-life" context can unlock the door to various business, environmental, personal, and social contexts hitherto inaccessible to Internet applications.

The incredible amount of information captured by a trillion RFID tags will have a tremendous impact on our lives. However, questions remain if we

are to use RFID in the IoT. How do we transform low-level RFID data into meaningful, high-level information? Can we design and build applications that are truly useful and not just novelties? If so, will their utility outweigh the potential loss of privacy, and how can we help users understand and control their privacy settings?

At the University of Washington, we're exploring these issues first-hand with a building-scale, community-oriented research infrastructure called the RFID Ecosystem (<http://rfid.cs.washington.edu>). This infrastructure creates a microcosm for the IoT in which we can investigate applications, systems, and social issues that are likely to emerge in a realistic, day-to-day setting. We've developed a suite of Web-based, user-level tools and applications for the IoT and deployed it in the RFID Ecosystem. We've also conducted a four-week user study to investigate patterns of adoption and utilization of our tools and applications as well as users' subjective reactions. We present the results of this study, focusing on tool and application usage.

The RFID Ecosystem

We built the RFID Ecosystem around an Electronic Product Code (EPC) Class-1 Generation-2 RFID deployment that spans all seven floors of our 8,000-square-meter computer science and engineering building (see Figure 1). The deployment includes 44 RFID readers (each equipped with up to four antennas for a total of 161) positioned at the building's entrances, on the stairwells, and throughout the corridors. Readers run embedded Linux and have wired or wireless Gigabit Ethernet over which they report their RFID data to a central server. Volunteers carry RFID tags as badges and attach tags to personal objects. Because most everyday objects aren't yet manufactured with tags embedded, and because manufacturer-assigned metadata might not be personally meaningful, we create the tag-object association manually. For this purpose, we created a special kiosk where users can select an RFID tag, physically attach it to an object, and create a corresponding association between the tag and that object.

All readers in our deployment run custom software that processes new RFID data before streaming it to the central server. This software continuously polls the reader hardware for newly detected RFID tags and generates



Figure 1. The RFID Ecosystem. RFID reader antennas are mounted on cable trays (upper left) and in custom-built wooden boxes (lower left). An RFID kiosk (upper right) lets users associate one of three types of tags (lower right) with a personal object.

one tag-read event (TRE) per tag per antenna per second, a tuple with the schema (`tag ID`, `antenna ID`, `time`). For example, if tag *A* is detected by reader antenna *X* at time stamp *t*, then the custom reader software will generate and send the following TRE to the server: (`tag A`, `antenna X`, `t`). Each reader also runs the network time protocol to synchronize its clock with the rest of the system.

We store all TRE data in a central SQL Server database (www.microsoft.com/SQL). This database also contains metadata about the deployment, including each antenna's latitude and longitude and a symbolic antenna name (for example, "front entrance," "4th floor stairwell," or "Room CSE 405"). We wrote software to transmit data between the readers and the database in Java using Apache's Multipurpose Infrastructure for Network Applications library for efficient, secure networking. This software implements various privacy policies^{2,3} and runs the Cascadia system⁴ to support application development and execution. Finally, our tools and applications are entirely Web-based, implemented with the Google Web Toolkit (<http://code.google.com/webtoolkit/>), and hosted with

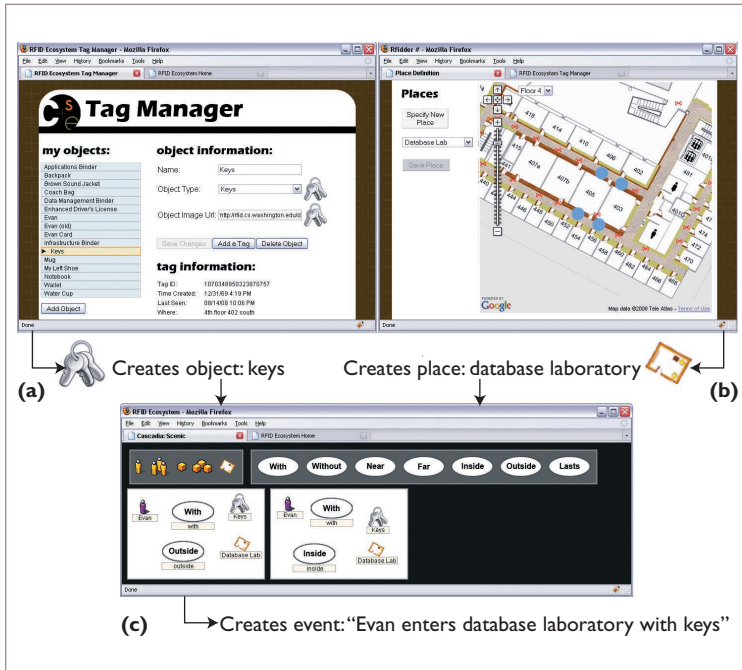


Figure 2. Metadata management tools. (a) The Tag Manager creates and manages virtual objects to which tags can be bound; (b) the Place Manager groups antennas into places; and (c) Scenic uses objects and places to specify how to transform a user's tag-read events (TREs) into higher-level events.

Apache and Tomcat (www.apache.org) on a separate server.

User-Level Tools

The RFID data our deployment supplies offer low-level location information in terms of tags and antennas. TREs – such as (tag A, antenna X, t) or (tag B, antenna Y, t + 1) – are helpful to IoT users only if a middleware can transform them into more meaningful, high-level information about events that applications and users can directly consume (for example, Ana is leaving the office with her purse). Furthermore, because such high-level events are personal and potentially sensitive, users must be able to precisely control all information disclosure to avoid privacy breaches. As such, we've developed several secure, Web-based tools that let users directly control how their RFID data is transformed and disclosed in the RFID Ecosystem.

Transforming Low-Level RFID Data

To support transforming TREs into higher-level events, we built tools that let users directly define metadata and associate it with tags and antennas. One such tool, the Tag Manager,

presents a highly interactive set of menus, tables, and Web forms for creating and managing metadata on a user's tags and personal objects. The Tag Manager interfaces with the RFID kiosk so that when users are at the kiosk, they can associate one or more physical tags with an object. For example, a new user can use the Tag Manager at the kiosk to register several personal tags. Later, the same user can access the Tag Manager from his laptop to delete objects and review or edit object metadata (such as name, type, image URL, or where the object's tags were last seen).

A second tool, the Place Manager, supports creating and editing high-level location information items, called *places*.⁵ A place in the RFID Ecosystem is a set of one or more RFID antennas with a label. For example, the two antennas in the corridor on either side of a user's office door might be grouped and labeled "my office." The Place Manager displays each RFID antenna's location as an icon in a Google Map mashup of the RFID Ecosystem deployment. Users can create or edit a place by clicking on antenna icons to select or deselect antennas and by entering the place label in a text box (see Figures 2a and 2b).

Once the Tag and Place Managers define metadata that binds tags to objects and antennas to places, respectively, applications and other system components can use that data to generate higher-level information that's personalized and more directly meaningful to users. An additional third tool, Scenic, lets users specify what higher-level events they would like to have extracted from their TREs (see Figure 2c). Scenic uses an iconic visual language and a storyboard metaphor to describe how people and objects enact an event through a sequence of movements between places. Specifically, the Scenic interface lets users drag and drop icons representing people, objects, places, and basic relationships (such as inside, outside, near, or far) onto one or more panels, or "scenes," in a storyboard to specify a movement sequence corresponding to an event. Thus, to specify an event, users simply "tell the event's story," scene by scene. For instance, a user could drag icons representing himself, his keys, and a lab into one scene with icons representing the "with" and "outside" relationships to indicate that he's outside the database lab with his keys. Then, he creates a second scene with the same icons, re-

placing “outside” with “inside” to indicate that he’s entered the lab with his keys. Further details on Scenic’s interface and implementation are available elsewhere.⁴

Controlling Privacy

RFID security and privacy present many challenges, and potential solutions, from hardware and wireless protocol security to the management, regulation, and sharing of collected RFID data.^{6,7} Our privacy work in the RFID Ecosystem has focused on controlling access to collected RFID data.^{2,3} As such, we accomplish privacy control in the ecosystem chiefly through personal data auditing and by enforcing novel access-control policies. Two tools let users directly interact with their personal RFID data and with the access-control framework that governs data disclosure. Both can operate in conjunction with our *physical access control* (PAC) policy. We discuss PAC in other work,² but the main idea is that users’ access is constrained to events that occurred only when and where they were physically present. In this way, their personal RFID data can serve as a detailed log of events they might have observed in person throughout the day. Whether a system employs PAC or not, users can review or delete their data with the Data Browser, or extend the data they choose to share with the Access Control Interface (see Figure 3).

The Data Browser lets users review all TREs collected on their tags in an interactive, table-based interface. The table displays only TREs that occur in a user-specified time window (the default window is the current day so far). Each table row shows human-readable information about an individual TRE, including the tagged object’s name, the reading antenna’s symbolic location, and the time stamp. The Data Browser sorts rows in chronological order, but users can also sort based on object or location by clicking that field in the header row. Users can delete their data by selecting and deleting specific rows. To make large deletions easier, the Data Browser has a deletion menu with which users can quickly delete all data collected over the past 30 minutes, an entire day, or some other user-specified time range.

In addition to directly managing their personal RFID data, users can use the Access Control Interface to control what data the RFID Ecosystem automatically stores and discloses about their tags. This interface features a set

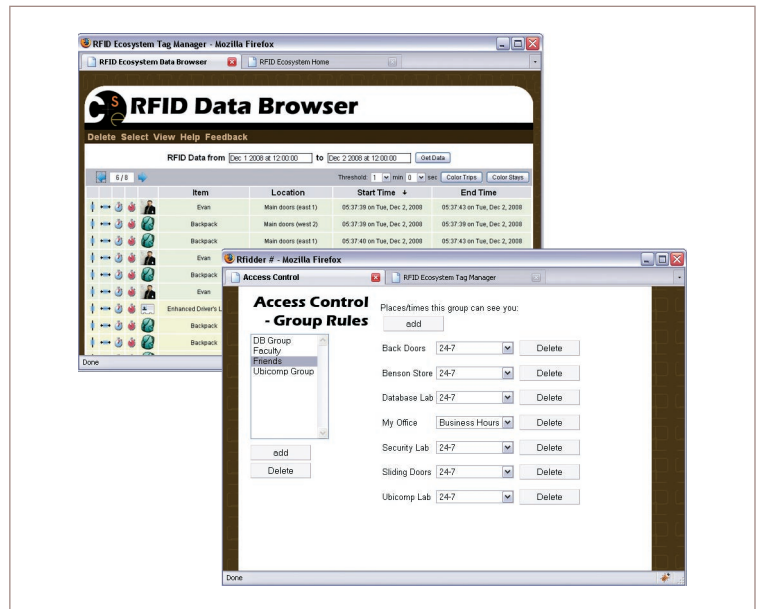


Figure 3. The Data Browser and Access Control Interface. With these tools, users can directly interact with their personal RFID data and with the access-control framework that governs data disclosure.

of Web forms that let users extend the data set they share with others by defining additional circumstances in which the system can disclose their data. For example, a user could define a rule that says “Professors can see when I’m in my office during business hours” or “My friends can see my location at any time.” Specifically, the interface lets users define friend groups to which they can apply various access-control rules concerning places and time ranges. Thus, all people in a particular group have access to the data that PAC (if it’s in use) or a user’s access-control rules allow. The key advantage to this type of access control is that it increases the amount of data available to applications in a way that’s tied to physical events, a method users can more easily understand.

RFID-Based Web Applications

To illustrate what’s possible in the IoT with RFID and our user-level tools, we developed and deployed several Web-based applications. These applications combine TREs with metadata on objects, places, and events in accordance with user-defined access-control rules to offer services to users. The applications we present rely chiefly on place-level location information, unlike other recent RFID applications that use more fine-grained locations or information about people’s and objects’ close-range interaction.^{8,9} Although several location-based Web

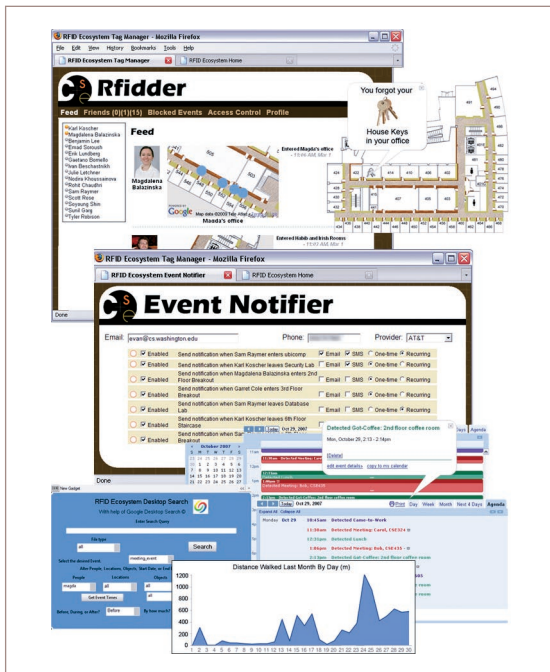


Figure 4. RFID Web applications. We developed the Rfidder, Event Notifier, Object Search, Digital Diary, and Event-Based Desktop Search applications.

infrastructures and applications are emerging (see <http://fireeagle.yahoo.net> or <http://plazes.com>), we feel that RFID-enabled information on indoor location, object location, and events will provide significant incremental functionality. Figure 4 shows our applications.

A Search Engine for Things

Perhaps the simplest Web-based RFID application is a search engine for things. We implemented a Web interface that lets users view the last recorded location for their tagged objects or search for a particular object's location. The Event Notifier, a more proactive extension of this application, leverages user-defined events to notify users when the last recorded object location matches some condition. For example, the application can send an SMS reminder to users when they leave the building without a particular item that they've designated as important (such as a laptop power cable).

Social Applications

Some of the most popular Web services offer information and updates on activities in our social networks; feeds from sites such as Twitter (<http://twitter.com>) and Plazes are some examples. To explore this space, we developed

an application called Rfidder that uses events about people and places to give users real-time updates in their social networks. Users' Rfidder interfaces display a feed of events that their friends have defined ("Fiona has entered her office," or "Raj is taking a coffee break"). Users can control their friend lists as well as what events are disclosed to which friends by defining access-control rules with the Access Control Interface. Rfidder also integrates with Twitter to provide greater utility and social networking capability.

Personal Trends

Historical queries about object and event data let users study trends in their activities over time. This can be extremely useful for applications that support long-term activities such as business projects and collaborations. We built a Digital Diary application that records and displays events in a Google Calendar (<http://code.google.com/apis/calendar>) for later perusal. This way, users can look back over their diaries to see how and with whom they've spent their time. We've also added support for plotting historical trends using the Google Charts API. Users' charts can succinctly display where, how, and with whom or what they've spent their time over some arbitrary period. For example, Fiona can have her diary record how often she enters the building with her bicycle helmet and later use that data to visualize how often she biked to work during the past month or year.

Event-Based Desktop Search

The log of events that applications such as the Digital Diary collect can also enable search-based applications that leverage a user's memory of events in the physical world. To demonstrate this, we implemented an event-based search plug-in for Google Desktop Search (<http://desktop.google.com>). Our plug-in retrieves digital documents created or modified around the time of some physical event that the user remembers. If Ana remembers that she visited a Web site during a meeting with her advisor a few weeks ago but doesn't remember exactly what that site was, she can search by the event "advisor meeting" to select all Web sites visited during meetings with her advisor.

Usage Patterns

To better understand how people might use our




Figure 5. Results of the four-week user study. We measured (a) a cumulative distribution function (CDF) of operations per participant, (b) tool and application usage over time, (c) the number of operations performed through each tool and application, and (d) how many times participants performed each type of operation.

tools and applications in the IoT, we conducted a four-week user study with the RFID Ecosystem. We recruited 67 participants, including 33 undergraduates, 30 graduates, two faculty members, and two staff who occupy the computer science building on a daily basis. We offered participants US\$50 as an additional incentive to use our applications. Each participant wore a tag as a personal badge and attached tags to personal objects using the Tag Manager with the RFID kiosk. Participants tagged a total of 324 objects in 19 different categories (such as wallets, bags, jackets, mugs, books, power cables, and laptops). Further details on the study, including measurements on data rates and volumes, tag-read rates, and system performance, are available elsewhere.¹⁰

Throughout the four weeks, participants carried their tags every day and had access to the Tag Manager, Place Manager, Data Browser, and Access Control Interface as well as the Object Search and Rfidder applications. Note that we

integrated the Object Search application with the Tag Manager tool in this study to provide a centralized point of access for object-related information. In addition to RFID data, regular survey feedback, and a detailed exit survey, we collected detailed logs of each participant's interaction with tools and applications. Figure 5 provides a high-level summary of how participants interacted with their RFID data. In particular, the plots show operation counts, where an operation represents one database query (selection, insert, update, deletion, and so on) a participant conducted on his or her data. Participants perform operations implicitly when interacting with any tool or application.

The first plot shows the cumulative distribution function (CDF) of the total operation count over all 67 participants. The CDF indicates that roughly 40 percent of participants performed fewer than 100 operations, which reflects minimal system use. In the exit survey, these users explained that the system

wasn't detecting their tags well enough for the applications to be useful. The other 60 percent of users didn't have such severe problems. Indeed, the average read rate for a personal badge was 0.4, and we found a correlation with coefficient .37 between a participant's number of operations and the number of times his or her personal badge was detected.¹⁰ Figure 5b shows how much participants used each tool and application each day relative to that tool or application's total usage throughout the study. We immediately see spikes in roughly the middle of each week (the busiest days) as well as a gradual decline in usage over the entire study. This decline might be in part due to winter exams in the last two weeks or to the US Thanksgiving holiday on the last three days; this might also be the effect of fading

applications an average of 3.2 (standard deviation 0.7) on a 5-point scale (where 5 is most useful and 1 isn't useful) – most of them (37 of 40) felt that the applications were at least fun and novel. Those giving a ranking less than 3 often cited a lack of tagged objects or participating friends as a reason for low utility. Most participants also expressed a desire to be pushed the location information, rather than have to pull it from a Web page. The data also shows that graduate students defined 20 of the 26 places defined in the study; this is likely because graduate students occupy a larger region of the building (five additional floors) than undergraduates.

The second histogram also reveals some interesting results about privacy control. Perhaps surprisingly, a participant deleted collected RFID data in only one instance. Instead, users focused on defining access-control rules to manage personal privacy. The user who deleted the data explained in the exit survey that he did so simply as an experiment to verify that the tool worked. Furthermore, most users (50 of 67) said that they had very few privacy concerns involving data collected in our controlled experiment. However, 51 users said they would be more concerned (4 on a 5-point scale, where 5 is extremely concerned and 1 is unconcerned) if their employer had this data, and 56 said they would be extremely concerned if their government had it. Other deletion operations occurred during metadata management and included minor operations such as deleting an object that the system couldn't effectively track (for example, a metallic laptop or mobile phone).

Metadata associated with tags, antennas, and events must be personalized and carefully controlled to create a safe, meaningful user experience.

novelty. The plot also shows that participants used the Data Browser, Tag Manager, and Rfidder more constantly throughout the study, and the Place Manager and Access Control Interface more often in the beginning of the study, when they were likely to have started defining policies through which friends could view their activities in Rfidder.

Figure 5c displays the total number of operations performed through each tool or application for each participant category. The columns for Access Control Interface and Place Manager count only metadata updates, insertions, and deletions, omitting the initial selection of existing metadata that occurs each time a user loads these tools. Similarly, Figure 5d shows the operation counts for each operation type. The histograms illustrate that participants interacted most frequently with their data by issuing object and person location queries via the Tag Manager and Rfidder. Here, the exit survey revealed that – although the 60 percent of participants without severe tag-detection problems only rated these appli-

Building applications with RFID data in the IoT is challenging, not just because TREs provide only low-level information but also because the metadata associated with tags, antennas, and events must be personalized and carefully controlled to create a safe, meaningful user experience. Our Web-based tools aim to empower users by letting them manage metadata and control privacy.

Based on our study results, we feel that RFID-based personal object and friend tracking are promising, basic services for the IoT that our tools can quickly enable. One key problem we must overcome is achieving a sufficient den-

sity of tags and users. Another problem is finding techniques that improve or compensate for low tag-read rates – we’re currently exploring using stricter tag-mounting strategies as well as probabilistic data management.^{4,10} We also conclude that although context-aware access control seems to be a useful, easily understood abstraction for managing location privacy, more evaluation is needed to determine whether it meets users’ needs when privacy concerns are magnified. Finally, because 41 out of 67 users expressed an interest in using personal trending applications such as the Digital Diary, we’re studying this application in another, longer user study. □

References

1. Complete RFID Analysis and Forecasts, 2008–2018, www.idtechex.com/research/reports/.
2. T. Kriplean et al., “Physical Access Control for Captured RFID Data,” *IEEE Pervasive Computing*, vol. 6, no. 4, 2007, pp. 48–55.
3. V. Rastogi et al., “Access Control over Uncertain Data,” *Proc. 34th Int’l Conf. Very Large Databases (VLDB 08)*, VLDB Endowment, 2008, pp. 821–832.
4. E. Welbourne et al., “Cascadia: A System for Specifying, Detecting, and Managing RFID Events,” *Proc. 6th Int’l Conf. Mobile Systems, Applications and Services (MobiSys 08)*, ACM Press, 2008, pp. 281–294.
5. J. Hightower et al., “Learning and Recognizing the Places We Go,” *Proc. 7th Int’l Conf. Ubiquitous Computing (UbiComp 05)*, vol. 3660, Springer, 2005, pp. 159–176, 2005.
6. A. Juels, “RFID Security and Privacy: A Research Survey,” *IEEE J. Selected Areas in Comm.*, vol. 24, Feb. 2006, pp. 381–394.
7. M. Langheinrich, “A Survey of RFID Privacy Approaches,” *Personal and Ubiquitous Computing*, Springer, Oct. 2008.
8. A. Nemmaluri et al., “Sherlock: Automatically Locating Objects for Humans,” *Proc. 6th Int’l Conf. Mobile Systems, Applications and Services (MobiSys 08)*, ACM Press, 2008, pp. 187–198.
9. M. Philipose et al., “Inferring Activities from Interactions with Objects,” *IEEE Pervasive Computing*, vol. 3, no. 4, 2004, pp. 10–17.
10. E. Welbourne et al., “Longitudinal Study of a Building-Wide RFID Ecosystem,” to appear in *Proc. 7th Int’l Conf. Mobile Systems, Applications, and Services (Mobicomp)*, June 2009.

Evan Welbourne is a PhD student in the University of Washington’s Department of Computer Science and En-

gineering. His interests are in sensor networks and data management for mobile and Web-based computing. He is the lead graduate student on the RFID Ecosystem project. Contact him at evan@cs.washington.edu.

Leilani Battle is pursuing a BS in computer science and engineering at the University of Washington. She worked on the RFID Ecosystem project during summer 2008 under the sponsorship of Intel’s Research Experience for Undergraduates program. Contact her at leibatt@cs.washington.edu.

Garret Cole is a software engineer at Microsoft Research. His primary interests include sensor networks and databases. Cole has a BS in computer science and engineering from the University of Washington. Contact him at riftdgarret@gmail.com.

Kayla Gould will graduate in June 2009 from Western Oregon University with a BS in computer science. She worked on the RFID Ecosystem project at the University of Washington during summer 2008 under the sponsorship of CRA-W’s Distributed Mentor Program. Contact her at kaylajgould@gmail.com.

Kyle Rector is a senior in electrical engineering and computer science at Oregon State University. She worked on the RFID Ecosystem project at the University of Washington during summer 2008 under the sponsorship of CRA-W’s Distributed Mentor Program. Contact her at rectorky@eecs.oregonstate.edu.

Samuel Raymer is pursuing a BS in electrical engineering at the University of Washington. He is an undergraduate researcher on the RFID Ecosystem project and a member of the IEEE Computer Society. Contact him at samwr@cs.washington.edu.

Magdalena Balazinska is an assistant professor in the University of Washington’s Department of Computer Science and Engineering. Her interests are broadly in databases and distributed systems with a focus on distributed stream processing, sensor and scientific data management, and cloud computing. Contact her at magda@cs.washington.edu.

Gaetano Borriello is the Noe Professor of computer science and engineering at the University of Washington. His interests are in the impact technology can have on solving important social and health problems. His recent work has been in sensing for healthcare and mobile systems for data collection in the developing world. Contact him at gaetano@cs.washington.edu.