

A Sensor Network for Social Dynamics

Mathew Laibowitz
Responsive Environments Group
MIT Media Lab
Cambridge, MA 02139
mat@media.mit.edu

Jonathan Gips
Human Dynamics Group
MIT Media Lab
Cambridge, MA 02139
jgips@media.mit.edu

Ryan Aylward
Responsive Environments Group
MIT Media Lab
Cambridge, MA 02139
aylward@media.mit.edu

Alex (Sandy) Pentland
Human Dynamics Group
MIT Media Lab
Cambridge, MA 02139
sandy@media.mit.edu

Joseph A. Paradiso
Responsive Environments Group
MIT Media Lab
Cambridge, MA 02139
joep@media.mit.edu

ABSTRACT

This paper describes the design and architecture of the UbER-Badge, a wireless sensor node and wearable display designed to facilitate group interaction in large meetings and acquire a wide range of data for analyzing social dynamics. The platform design and its application suite are described. Data is presented that shows the social patterns developing across large events and indicates that certain individual characteristics (interest, affiliation) can be determined from the sensor data from deployments of this system with groups of over 100 people.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]

General Terms

Algorithms, Management, Design, Human Factors, Theory.

Keywords

Active badge, wearable sensors, behavior measurement.

1. INTRODUCTION AND PRIOR WORK

Electronic badges make ideal wireless transducer platforms for interactively facilitating social events and monitoring group behavior. As standard "nametag" badges tend to be pinned on the chest or hung about one's neck, they are appropriately placed for displaying simple messages to nearby people. Tending to face along the direction where the wearer's attention is focused, they are also well suited to broadcasting a line-of-sight ID code that enables proximate devices (including other badges) to be aware of individual presence. Accelerometers on a badge can detect characteristics of the user's motion, and an appropriately directed microphone can pick up the user's voice.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IPSN '06, April 19–21, 2006, Nashville, Tennessee, USA.

Copyright 2006 ACM 1-59593-334-4/06/0004...\$5.00.

Electronic badges have established a recent history in both research and products. The first electronic badges, pioneered over a decade ago by Olivetti Research [1], were very simple platforms that periodically transmitted a modulated InfraRed (IR) identification (ID) code to the vicinity, enabling people to be located by an infrastructure of embedded networked IR readers as they moved about a facility. Another approach is to use the badge as a dynamic display and as a facilitator for person-person interaction at large events. This is the direction taken by two mid-90's research projects at the MIT Media Lab - the "Thinking Tag" [2] (an electronic "icebreaker" that flashed red/green LED's according to agreement of proximate wearers on a series of provocative questions) and the "Meme Tag" [3], which featured a large LCD display that enabled users to selectively exchange brief catch phrases (or "memes") that were tracked as they propagated through large groups.



Figure 1. The final production version of the Uber-Badge – Front View.

Badge platforms have subsequently moved into the commercial world, with systems like the Matchstick and the Japanese Lovegetty [4] – similar to the Thinking Tags, these were designed as matchmakers for nightclub environments. In contrast, the

nTAG [5], designed to facilitate business meetings and conferences, features a 128x64 pixel, back-lit LCD display, a trio of navigation buttons, and both IR and quasipassive radio frequency (RF) backscatter communication – the IR is for line-of-sight communication with other badges and fixed beacons, while the backscatter system allows the badge to upload data to microwave beacons when it is less than 20 feet from them. The IntelliBadge [6] is also intended for conferences, but as it is only essentially only a hybrid inductive/RF ID tag, it is unable to display information or store state – the responsibility of tracking the tags is moved off the badge and onto the networked infrastructure of fixed readers. These products target applications such as tracking people through a convention hall, detecting what booths they visited or were most interested in, and (in the case of the nTAG), exchanging virtual business cards and encouraging inter-attende interaction.

Badges and other wearable or mobile platforms are being increasingly used to infer and facilitate social interaction using measurements based solely on proximity. Some of these systems (e.g., [7]) are minimal IR transceivers that harken back to Olivetti’s Active Badge, while others are simple RF beacons running on mobile devices [8] or Bluetooth-enabled cell phones [9].

We have developed a very versatile platform, termed the UbER (Ubiquitous Experimental Research) Badge, geared toward enabling a variety of user interactions at large events while collecting a rich suite of multimodal sensor data that can be used to analyze and respond to the structure and evolution of ongoing social dynamics. The UbER-Badge encompasses an extreme mix of capabilities not available in prior badge platforms, such as surveyed above. Featuring both line-of-sight IR and omnidirectional RF communication links, the badges form a large wearable sensor network.

2. UbER-BADGE HARDWARE

A production badge is shown in Fig. 1, and its block diagram is given in Fig. 2. Figure 3 shows a rear view of the main circuit card. A complete UbER-badge (including frosted plastic faceplate) measures 11 x 12 cm, and weighs about 170 g with all four AAA batteries installed. At an average current of about 100 mA, badges last for roughly 15 hours of continuous use. In quantities of 300, the cost of an assembled badge runs roughly US \$85, not including a case or front panel.

To maximize longevity, this device needs to consume low power, and since this platform is being made to support many different applications, it must be extremely agile. The Texas Instruments MSP430F149 fits these requirements, hence was used as the central microcontroller. In addition to the program memory within the MSP, the badge can access up to 256 MB of data flash for storing audio or user data (2 MB are currently installed). The entire flash memory can be offloaded within a few seconds via a fast USB adaptor.

The badge’s display was designed to show simple iconographic animations and bright scrolling text that can be easily read in any lighting condition across several meters (it’s hard to read the LCD panels used with existing badges at any significant distance or outdoors). Accordingly, the badge was equipped with a 5x9 LED matrix driven by a dedicated controller capable of independently

specifying the intensity of each LED. In addition, 4 brightness-controllable blue LEDs below the matrix are managed directly by the MSP430’s PWM to provide another visual output. To support applications that require the display of larger messages or more data, the badges are equipped with circuitry and connections to drive a large backlit alphanumeric LCD display that can be mounted beside or atop the LED matrix. Another connector is provided to support a narrow LCD that can be mounted on the top of the badge, allowing the viewer to see personal messages without rotating the badge to inspect the front. Although test badges were assembled with these additional displays, they were omitted from production units because of added cost, weight, and size.

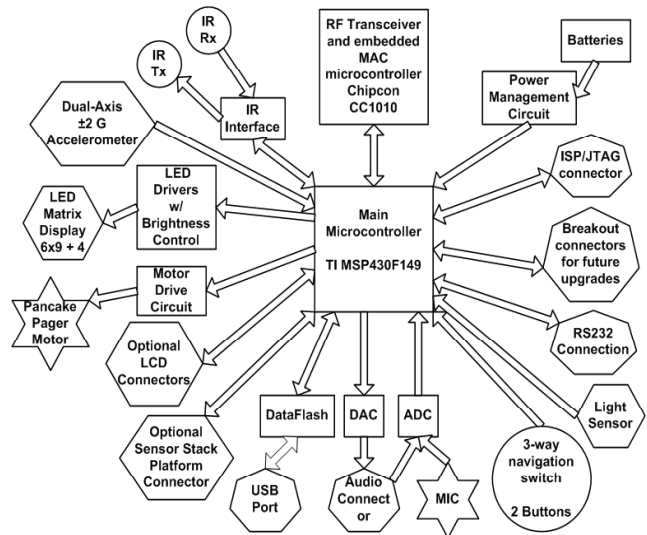


Figure 2. UbER-Badge Block Diagram

Each badge can provide tactile feedback (typically felt on the neck through the badge’s lanyard) via a pager-style vibrating motor, with force controlled by a PWM channel on the MSP430. A side mounted switch (providing up, down, and push-to-select) and a pair of easily accessible buttons on the lower front of the badge (see Fig. 1) are used for user input. An onboard microphone is connected to a 12-bit audio input, and a 12-bit monaural audio output is available at a headphone jack. This capability can be used to record, detect, process, generate, or transmit audio events. The headphone output, for example, can be used to provide contextual audio prompts that relate to the badge’s location in social or geographic space. Each badge also includes a 2-axis, ±2G accelerometer to sense user motion and an ambient light sensor. To support applications requiring additional sensing, compact cards from our modular Stack Sensor architecture [10] can be plugged into the badge at a set of dedicated headers; these cards currently include circuits that incorporate multiple modes of tactile sensing, 6-axis inertial measurement, digital imaging, and sonar proximity.

The badges were hung on conventional polyester lanyards with hooks allowing for easy put on and take off. The tops of the badges were typically suspended 13 cm below the user’s chin. This placed the badge close enough to the user’s face such that their own voice generally dominated the microphone signals without the badge’s proximity become annoying or cumbersome.

The badge is equipped both with an IR channel to support face-to-face and local communication, and a RF channel to support higher bandwidth communication across larger distances. The IR system consists of a composite IR LED lamp with a 17-degree spread, and an IR receiver with integrated demodulator, photodiode, photodiode amp, and a Silicon Labs C8051F301 processor, which acts as a dedicated IR communication controller to buffer incoming and outgoing IR messages. A slightly quicker version of the Sony-IR protocol is used on the badges, with the IR modulated at 40kHz. In addition to controlling the IR communication, the F301 also manages a RS232 port on the badge, allowing connection to a PC. The badge's IR communication is sensitive out to 3 meters.




Figure 3. Rear of UbER-Badge, showing RF Card (top)

The badge's RF section [11] is based around the Chipcon CC1010, which contains a processor and RF transceiver with programmable transmission strength and set to run at 433 MHz. Similarly to the IR system, by using the Chipcon's integrated microcontroller, all of the communication processing and protocol is offloaded from the main application processor. The software that runs in the CC1010 implements a peer-to-peer random access network using a carrier-sense method of media sharing and collision avoidance. Using a simple wire monopole antenna, easily tucked behind the badge in its case, the RF range has been tested out to 100 meters.




Figure 4. The "Squirt" - a compact IR tagging beacon

Figure 4 shows an IR beacon called a "squirt" that was used to tag fixed locations, typically research demos running during the open house portion of the meetings at which the badges were used. The squirts are 2.5 x 5 cm in area and run off a pair of AA batteries for a week. They broadcast a byte of ID at over 1 Hz to nearby badges (up to roughly 2-6 meters away), thereby informing them of the squirt's proximity. As the location of all squirts is known, they serve to roughly localize the badges. Badge wearers can also "bookmark" the demo associated with a squirt by pushing a button on their badge when it is in IR range of the squirt (indicated by the squirt's LEDs), an event that is logged in both the memory of the badge and squirt.

The hardware of the badge is designed to shield an application designer from the low-level details of the communication protocols and at the same time allows the custom application to run directly on the main processor without any layers of abstraction. This is achieved by providing the badges pre-loaded with the communication code running on the two communication co-processors (CC1010's integrated processor and the C8051F301) and a compiled library with interrupt driven functions for the main application running on the MSP430 to use the communication channels. This pre-loaded code was written in C using the Keil C51 toolchain. A custom badge application is written in C with precompiled libraries containing functions for all the features of the hardware and communication channels using Rowley CrossStudio for MSP430.

3. IMPLEMENTATION

Most data packets broadcast by the badge's RF system aren't multihop routed across badges – instead, all badge radios directly talk to a network of fixed base stations, each made by connecting a Lantronix Xport Ethernet link plugged into the building's LAN to the serial port of a standard badge. Six of these base stations were able to cover the entire Media Lab and adequately cover other venues where we ran this system – participants were always within range of at least one base station.

The main role of the base stations was to extend the range of the badge's RF communication (without the badges bearing the overhead of multihop routing) and enable PC-based kiosks and clients to query and command the badges from anywhere on the Media Lab's network. Packets sent from the badges are received by a nearby base station, routed through the LAN, and, depending on the nature of the packet, rebroadcast by the other base stations. Commands from PC's running the badge control software are sent over the LAN and relayed by all base stations. This infrastructure was also used to broadcast period timestamp messages every 15 minutes, which were used to align the data between the badges for the post-analyses described in Section V (real-time badge operations are asynchronous). The base stations also continuously collect data from all badges. Each badge sends a data payload every minute containing the ID's of other badges and squirts that they encountered via their IR channel since the last payload was sent. This packet is received by the beacons and sent to the kiosks and any other real-time data processing servers. Badges can be coarsely localized by keeping track of the base stations that they see – most applications, however, use the most recent IR encounters with fixed squirts for this.

All badges broadcast a packet containing a unique ID code through their IR port to alert other nearby facing badges and

squirts of their presence. Although the average interval between IR pulses is 1 second, it varies by up to 25% from shot to shot to avoid persistent collisions. Badges can detect each other at up to 3 meters and across large angles (e.g., 60°). When a pair of badges (or a badge and a squirt) detect each other via IR, an “encounter” is defined between them. The length of this encounter is monitored – the encounter is declared over when the other badge or squirt is not detected for at least 30 seconds.




Figure 5. Accelerometer signals (top) and audio amplitude (bottom) for each badge plotted across an entire day-long consortium event.

The badge continuously samples and logs signals from the accelerometer and microphone. Accelerometer readings are taken at 100 Hz, and an average absolute sample for each of the two dimensions (ACC_x, ACC_y) is computed and recorded every twenty-five samples (4 Hz). Microphone readings are taken at a rate of 8 KHz and averaged every eight samples, yielding a down-sampled rate of 1 KHz. These averaged readings are used to create two different parameters with different characteristics. The first is the average amplitude (AUD_{AMP}), and we calculate it by accumulating the absolute value of the averaged readings and dividing the sum by the frame size. The second measurement is the average difference between the 1 KHz averaged readings (AUD_{DIF}), yielding a high-pass response. Similar to the average amplitude, we accumulate the differences between successive averaged readings and divide by the frame size. The frame size

for our implementation was 256 samples producing a final audio feature-sampling rate of 3.91 Hz. The aforementioned sampling rates produce an upper bound of about 13.5 hours of data recording time before the 2 MB flash memory on the badge fills up – certainly ample time to outlast a day-long event.

Figure 5 shows sensor data from all badges accumulated throughout the entire day of the Media Lab event. The banded structure follows the timing of the event – buffet breakfast, first talk session, coffee break, second talk session, lunch, third talk session, then open house. The plots show that the environment during talks is quieter for most badges, and involve less motion dynamics, since people are seated and listening, as opposed to moving around and talking to one another. When not displaying messages or information, the badge display went into a “pilot light” mode, quiescently showing a dot that bounced around with the user’s motion, driven by the accelerometer data.

4. APPLICATIONS

After testing and evaluating an earlier prototype badge design [12], we fabricated 200 of the finalized UBER-Badges described here, together with a similar number of squirts, and deployed them at a couple of large research consortium meetings hosted at the Media Laboratory (in October 2004 and May 2005) and, in collaboration with one of our industrial partners, at a career fair for high school students in Scotland during September 2005. These events all involved on the order of 100 simultaneously badged individuals and 100 distributed, squirt-tagged demos. A variety of applications, as outlined below, ran on the badge system in order to facilitate many types of interactions between attendees. The events at the Media Lab consisted of two different environments, namely structured talks in an auditorium and an extended freeform “open house,” where participants could explore the Laboratory at will. The career fair was entirely an open house.




Figure 6. The badge at right has bookmarked the badge at left - blue LEDs signify an encounter in progress and the checkmark verifies the bookmark. Note that this early version of the badge lacked the easier-to-use front-mounted buttons and only used the side-mounted switch.

Bookmarking Demos and Exchanging Virtual Business Cards: Attendees were instructed to “bookmark” when they encountered either another badge wearer with whom they wished to exchange contact info or found an interesting Squirt-tagged demonstration that they desired to remember or investigate further. After the conference, bookmarks were downloaded from the badges and reported to the corresponding users, facilitating further contact with colleagues and deeper exploration of projects in which they

expressed interest. The bookmarking process was made to be very simple and intuitive. When a badge comes into IR range of another badge, the blue lights on both badges cycle – similarly, when a badge is within range of a squirt, LED’s on the squirt will glow. From this state, pushing any button on the badge will bookmark the ID of the other device. This process is visually verified by animating a checkmark on the other badge (Fig. 6) or flashing a LED on the squirt (Fig. 4).

Displaying public messages: Computers running the badge management software (used by the meeting administrators) were able to command all badges to repeatedly scroll a canned or custom text message. This was used to get attendees back into the auditorium for the next round of talks, inform them that food was being served during the Open House, tell the high school students that their bus was here (Fig. 7), etc.



Figure 7. Public broadcast messages scrolling across a set of badges.

Displaying personal messages: Badge kiosk PCs distributed around the building could be used to send a message to one particular badge. When the text is received, the badge’s vibrator pulses repeatedly to inform the wearer that a message is queued. The messages must be retrieved on another badge by getting within IR range and pushing any of the badge buttons. We opted not to allow users to retrieve messages on their own badge, since looking down at your badge is somewhat awkward and approaching a colleague for revealing your message tended to foster sporadic social mixing. Our original implementation of this feature used multihop communication over the IR channel, where the message would first “infect” the kiosk user’s badge, then propagate to other badges as they were encountered (which in turn propagated the message) until a timeout elapsed, ideally after the message reached the recipient’s badge. Our final implementation used the base station network to flood the building with a simultaneous broadcast, reaching the recipient with one transmission.

Finding People: Badge-wearers can be physically located via two techniques. One involves simply querying one of the badge kiosk PC’s with the wearer’s name – a location based on the most recent squirts noticed by the quarry’s badge is plotted on a building map. In another more interactive approach, the badge kiosk implants the ID of the quarry’s badge into the seeker’s badge. When the seeker presses a button within IR range of another badge, that badge displays a pattern that illuminates a number of LED’s in inverse proportion to the time elapsed since it last encountered the

quarry. If the seeker holds the button down, the request is sent through the radio, causing all badges in the vicinity to appropriately display. By following the trail of “brighter” badges, the seeker is led toward his goal.

Affinity Group Display: During the Media Lab events, the encounter and bookmarking data that were continually offloaded from the badges were used to build affinity models that evolved as the day unfolded. Badge wearers were dynamically clustered into one of 5 groups that were defined by commonality of behavior (see Section V). After this model became somewhat stable by the end of the day (by the time the evening reception began), an animated icon corresponding to the wearer’s affinity group was displayed on the badges whenever a new encounter was detected. This was something of a digital “T-shirt” – nearby participants would note that their icon was similar or different, often instigating conversation about their experiences during the day.

Voting: The buttons on the badge can be used to take a poll of the badge wearers during presentations. Conventionally, the right (red) button indicates a disagreement, and the left (green) button indicates agreement. When voting mode is enabled, the button push flags are sent via the RF port.




Figure 8. Timekeeping cues flashed by the badges - progressive warnings (top) and scrolled text when overrunning (bottom).




Figure 9. Histograms for # of talks vs. their % time overrun for sessions without (left) and with (right) timekeeping badges.

Timekeeping: One of the most effective badge applications was as a distributed display used for keeping time in presentations. The Media Lab consortium meetings generally consist of circa 30 very short research summary talks (4-8 minutes in length) juxtaposed tightly back-to-back. Even though a large clock was visible to the speaker and session MC’s tried to intervene as speakers ran late, prior meetings had considerable problems holding time, as many enthusiastic researchers drifted over their allotment. Seeing your entire audience flash warnings to you (Fig. 8) in a darkened auditorium, however, is an experience that’s very difficult to ignore (while the audience, facing forward, couldn’t see the

badges, all were visible to the speaker). The timekeeping displays were triggered either autonomously or manually via radio broadcasts from the event administrator’s PC located in the auditorium. For the two recent meetings where the badges were used to flash timekeeping cues at the speakers, the sessions ran much more punctually. As the histograms of normalized talk duration in Fig. 9 attest, the badges worked well in eliminating the long tail of extreme stragglers.

5. INFERRING INTEREST AND AFFILIATIONS

One of the main limitations of today’s interactive badge systems is that their notion of human interest is set either by answering a few questions before the interaction, or is simply hardwired into the system design. This limits the range and flexibility of these systems, making them feel more like party games than serious social networking tools. Accordingly, we are developing automatic interest detectors that remove the restrictions imposed by use of preset questions and the requirement that users explicitly ‘bookmark’ interesting people/events. Instead we aspire to measure interest directly from normal human behavior. We are also developing an affiliation classifier that aims to infer relationships between subjects without any such explicit labels. A person should be able to pick up a badge, wear it, and have the system learn the group of people with whom he associates. If we can achieve both of these goals, then we can begin to group people by the pattern of interests they display, and make introductions based on these patterns, without requiring users to answer preset questions or input new data during the networking event. By learning the affiliations between people, we gain a social network that can be used to further guide the introductions.

Nalani Ambady and Robert Rosenthal [13] have shown that observers can accurately classify human attitudes (such as interest) from non-verbal behavior using observations as short as six seconds. The accuracy of such “thin slice” classifications are typically around 70%, corresponding to a correlation between observer prediction and measured response of about $r = 0.40$. Our initial experiments using a range of motion and sound features indicate that it is possible for computers to duplicate this human perceptual ability [14, 15]. We therefore set out to measure human interest levels and affiliations using the sensors and computation capacity of our badge platforms.

We created the interest detector described in this section by using the bookmarks recorded by the UbER-badges as labels for the sensor data. Individual models were created for both badge-to-badge and badge-to-demonstration encounters. Our affiliation detector draws upon company names as ground truth for its learning. The classifier infers dyadic (e.g., user-user) affiliation based on observations of face-to-face encounter duration as well as correlations in accelerometer-derived badge motions over time. The interest classifiers can run in real-time on the badge microprocessor alone, allowing classification of user interest during the course of the event. The affiliation classifier runs in real-time mostly on the badge, but requires using the badges’ RF link to a PC server (or a peer-peer badge network) in order to compare results between badges.

The fall consortium meeting resulted in a data set that included 113 badges and 76 Squirts. Unfortunately, due to a combination

of hardware and software problems, a sizeable (but random) part of the full sensor data was lost. We corrected these problems for the spring meeting and successfully collected data from 84 badges and 73 Squirts that were deployed. After validating the data, we isolated sections of the sensor data that pertained to the badge-to-badge encounters (‘badge encounters’) and the badge-to-demo encounters (‘Squirt encounters’). Within each of these categories, we further divided segments into two groups: 1) those that received bookmarks and 2) those that did not. Our data sample included 311 bookmarked badge encounters and 320 bookmarked squirt encounters vs. 3703 non-bookmarked badge encounters and 400 non-bookmarked squirt encounters.

Two types of preprocessing were performed on the measurements that are used in the feature vectors. First, the sensor data recorded to dataflash was normalized on a per badge basis. This allowed variation in badge hardware to be controlled. Second, the encounter data from the IR was propagated between all badges, to minimize the possibility of an incorrectly labeled encounter in the training dataset. We also verified that the act of making a bookmark was not skewing the accelerometer features by testing our model on the badge-to-badge encounters that received bookmarks. These badges did not need to be handled in order to receive a bookmark and showed a similar classification distribution to the bookmarked encounters.

Using this sensor and interaction data, we created a 15 dimensional feature vector for every encounter. The average amplitude (AUD_{AMP}) and average difference (AUD_{DIF}) samples were subtracted to create a third audio measurement (AUD_{SUB}). For each encounter, the means ($\mu_{AUD_{AMP}}, \mu_{AUD_{DIF}}, \mu_{AUD_{SUB}}$) and standard deviations ($\sigma_{AUD_{AMP}}, \sigma_{AUD_{DIF}}, \sigma_{AUD_{SUB}}$) of these measurements were used as audio features. In a similar manner to the audio measurements, the accelerometer measurements (ACC_x, ACC_y) were subtracted to create a third accelerometer measurement (ACC_{SUB}). For each encounter, the means ($\mu_{ACC_x}, \mu_{ACC_y}, \mu_{ACC_{SUB}}$) and standard deviations ($\sigma_{ACC_x}, \sigma_{ACC_y}, \sigma_{ACC_{SUB}}$) of these measurements were used as audio features.

The remaining three features were derived from the IR data and represented the number of other encounters that occurred during the primary encounter (IR_{COUNT}), the sum of the lengths of all the encounters that occurred during the encounter (IR_{SUM}), and the length of the specific encounter being considered (IR_{LEN}).

In addition to the per-encounter features, we created a symmetric adjacency matrix that contains the sums of the durations that each dyad of badges spends within IR range of each other. These sums were accumulated for the course of the entire spring event.

5.1 Interest Detection

We analyzed the encounter data set with the goal of creating two classifiers: one that would predict bookmarking of badge-to-badge (badge) encounters and another that would predict bookmarking of badge-to-Squirt (Squirt) encounters. We found strong correlations between the features and an encounter being bookmarked for both the badge and Squirt encounters. Badge encounters showed a significant correlation between accelerometer features and bookmarks, primarily in the standard deviation features. Squirt encounters showed a very different set

of correlations. Audio features exhibited a negative correlation with receiving a bookmark but accelerometers showed no significant correlation at all. This may indicate that, for demo (squirt) bookmarks, the interested badge wearer is quietly reading or observing the demo before taking a bookmark.

From the original set of fifteen encounter features, we picked the most-correlated features, and constructed a predictor function using simple linear regression. Cross-validation was performed using a “leave-twenty-percent-out” method, and decision boundaries were selected such that the difference between classification accuracy for the bookmarked and non-bookmarked encounters was minimized.




Figure 10. Performance of interest detectors are shown for badge-to-badge (top) and badge-to-demo (bottom) encounters.

Using the six highest ranked badge encounter features (σ_{ACC} , σ_{ACC} , μ_{AUDAMP} , σ_{AUDAMP} , $\mu_{AUDDIFF}$, σ_{AUDSUB}), our linear combination model classifies 86.2% of badge-to-badge encounters correctly with a cross-validation accuracy of 85.5%. Accuracy was very similar at both the spring and fall conferences. The performance of the top five Squirt encounter features (σ_{ACC} , σ_{AUDAMP} , $\sigma_{AUDDIFF}$, μ_{AUDSUB} , σ_{AUDSUB}) was almost as good with a classification accuracy of 78.4% and cross-validation accuracy

of 78.3%. Accuracy was very similar at both the spring and fall Media Lab meetings. Figure 10 shows the classification distributions for both classifiers combining both datasets.

5.2 Affiliation Detection

We analyzed the encounter data set with the goal of determining what behaviors were useful predictors of affiliation. We found two factors (which we term cumulative time and influence), which can be used independently or in combination.

As seen in Fig. 11, cumulative time spent face-to-face with someone as measured by IR encounters has a medium correlation with whether two people are affiliated or not ($r = 0.4681$, $p < 0.001$). Using this feature alone, a simple threshold model will achieve 88.7% accuracy in determining whether two badge-wearing attendees at our consortium meeting are from the same corporation or not.




Figure 11. Histograms of the total amount of time that affiliated and unaffiliated dyads of people spend face-to-face with each other.

We could also determine affiliations from correlations in wearer activity. To accomplish this we employed the influence model, a partially coupled Hidden Markov Model that can be used to learn “influence values” across multiple chains [16]. We modeled each badge as a Markov chain with two hidden states (moving, not moving) whose observations were accelerometer motion features. Using expectation maximization, we learned the parameters of this model, including the influence values. We found the influence values across two badges correlate with their wearers being from the same corporation ($r = 0.3981$, $p < 0.001$), producing 69.28% prediction accuracy (Fig. 12).

Combining the cumulative time and influence predictors using a simple polynomial regression model produces a predictor with 93.0% accuracy and cross-validation accuracy of 92.7%.

Thinking further along these lines, a feature that has potential bearing on affiliation detection is correlated motion, as inferred from similarities in accelerometer data across a pair of badges. This feature is attractive, as it doesn’t require an IR line-of-sight, and could perhaps be measured by accelerometers embedded in

canonical platforms such as mobile phones kept in the user’s pocket rather than an IR transceiver that needs to be mounted on a visible platform like a badge. The most basic dyadic motion feature that we have been using is the correlation across an “energy feature” calculated independently for two individuals’ accelerometers. This energy feature is the standard deviation of the magnitude of the 2D accelerometer over a two second period.




Figure 12. Histograms of log influence values for affiliated and unaffiliated dyads of people.

We have recently observed this relationship in the Scottish high school data set (dominated by groups of young people walking through the career fair), where the natural logarithm of time spent face-to-face (determined by the IR system) between two people had a medium correlation ($r = 0.55$, $p < 0.001$) with the correlation in their dyadic energy feature (Fig. 13).




Figure 13. Dyadic motion feature (from correlated accelerometer activity) vs. log of time spent face-face.

6. CONCLUSIONS

The UBER-Badges proved to be robust, and held up well to the rigors of long meetings. The vast majority of participants kept

them on throughout the events, regardless of their 170 g weight, and follow-up questioning of the meeting participants indicated that most were happy to wear them just for the added benefit of convenient bookmarking of demos and colleagues. Although badges are generally designed to be primarily seen by people other than the wearer, some individuals expressed discomfort at having messages for other people come across their badge. The distributed display provided by the badges worn in the auditorium was very effective at keeping speakers to their allotted times despite very jammed meeting agendas. Although populating the badges with full-color RGB LED’s was prohibitively expensive at the time of this project, using 2-color RG LED’s may have been feasible, and (even if all LED’s switched color together) this could have provided a useful parameter space in which to explore expressing immediate agreement/disagreement, yes/no, like/dislike etc.

The asynchronous CSMA RF protocol used by the badges performed well in all instances except for voting, where people who hit buttons repeatedly could jam the airspace. This could be accounted for by throttling back the transmissions – e.g., transmitting vote updates only every second or two, and sending a “button push” count as opposed to repeatedly sending single-push messages.

The IR localization update rate was somewhat slow because of the narrow field of view on the squirt IR transceivers. Combining IR and RF localization schemes could provide better performance – e.g., when badges see several RF base stations, a RSSI fingerprinting or interpolation technique could serve to coarsely locate them, with refinement provided by acquired squirts.

Our sensor analysis shows that we can automatically generate bookmarks that approximate the decisions made by UBER-Badge wearers with 80% accuracy, without taking into account personal characteristics, history, or other prior knowledge. Similarly, we can infer affiliations of the wearers with greater than 90% accuracy, again without prior knowledge. A next step is to work towards having the badges begin proactively suggesting things of interest to the wearer, as the model starts to correlate their behavior patterns and associates them with other individuals and groups.

Ultimately, such a badge platform could be a wireless peripheral to a mobile phone – the badge would provide a wearable display, which may become a fashionable piece of apparel outside of industrial conventions, while the phone provides computational power and networks to the local infrastructure. Although the badge is well-suited to line-of-sight sensors such as IR transceivers, we have indications that some social association can be extracted from correlation motion cues extracted from an accelerometer that could be perhaps carried in one’s pocket or purse.

7. ACKNOWLEDGMENTS

We thank many people at the MIT Media Lab who participated in the UBER-Badge project, namely David Merrill for developing the Squirt IR beacons, Sajid Sadi for the badge kiosks, Mark Feldmeier for badge mechanics, Matt Malinowski for graphing talk data and helping with badge construction, Lisa Lieberman for logistics, Missy Corley and Michael Broxton for fabrication help, and Wen Dong, Ron Caneel, and Anmol Madan for their

comments. We also thank our industrial partners, namely Highlands and Islands for the collaboration in Scotland and British Telecom and Ricoh for their partial sponsorship of the badge fabrication. We thank MEMSIC for donating the badge accelerometers. This project was supported by the Things That Think Consortium and the other sponsors of the MIT Media Lab.

8. REFERENCES

- [1] Want, R., Hopper, A., Falcao, V., Gibbons, J., "The Active Badge Location System," *ACM Transactions on Information Systems*, Vol. 10, No. 1, Jan. 1992, pp. 91-102.
- [2] Borovoy, M. McDonald, F. Martin, M. Resnick, "Things That Blink: Computationally Augmented Nametags," *IBM Systems Journal*, Vol 35, Nos. 3 & 4, 1996.
- [3] Borovoy, R., et al, "Meme tags and community mirrors: Moving from conferences to collaboration," in *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work*, New York: ACM Press, p. 159, 1998.
- [4] McCrone, J., "You buzzing at me?," *New Scientist*, January 15, 2000, pp 20-23.
- [5] Lewis, D.E., "Breaking the Ice 2.0," *Boston Globe*, November 15, 2004, pp. C1,C14.
- [6] Cox, D., Kindratenko, V., and Pointer, D., "IntelliBadge: Towards Providing Location-Aware Value-Added Services at Academic Conferences," in A.K. Dey et al. (Eds.), *UbiComp 2003*, Springer Verlag LNCS, pp. 264-280, 2003.
- [7] Bradford Paley, W. and Han, J.Y., "Trace Encounters: A Social Networks Visualization at Ars Electronica 2004," in in Stocker, G., and Schopf, C. (eds), *Timeshift (Proc. of Ars Electronica 2004)*, Hatje Cantz Verlag, Ostfildern-Ruit, Germany, 2004, pp. 252-255.
- [8] Terry, M., Mynatt, E.D., Ryall, K., Leigh, D., "Social Net: Using Patterns of Physical Proximity Over Time to Infer Shared Interests," in *Proc. of CHI 2002, Extended Abstracts*, pp. 817-817.
- [9] Eagle, N. and Pentland, A., " Mobile Matchmaking: Proximity Sensing and Cueing , To Appear: *IEEE Pervasive Computing*, Special Issue on Smart Phones (April 2005).
- [10] Benbasat A.Y. and Paradiso, J.A., "A Compact Modular Wireless Sensor Platform," in *Proceedings of the 2005 Symposium on Information Processing in Sensor Networks (IPSN)*, Los Angeles, CA, April 25-27, 2005, pp. 410-415.
- [11] Laibowitz, M. and Paradiso, J., "Wearable Wireless Transceivers," *Circuit Cellar*, Vol. 163, Feb. 2004, pp. 28-39.
- [12] Laibowitz, M., and Paradiso, J.A., "The UbER-Badge, A Versatile Platform at the Juncture Between Wearable and Social Computing," in Fersha, A., Hortner, H., Kostis, G. (eds), *Advances in Pervasive Computing*, Oesterreichische Computer Gesellschaft, 2004, pp.363-368.
- [13] Ambady, N., & Rosenthal, R. (1992). "Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis". *Psychological Bulletin*, 111(2), 256-274.
- [14] Madan, A., Caneel, R., and Pentland, A. " Towards Socially Intelligent Wearable Networks". in the *Proceedings of the Student Colloquium, the 8'th IEEE International Symposium on Wearable Computers (ISWC 2004)*, Published by by the Institute for Computer Systems and Networks, UMIT, Austria and by the Electronics Laboratory, ETH Zurich, Switzerland, October 31 - November 3, 2004, pp. 30-32.
- [15] Pentland, A. (2004, October). "Social Dynamics: Signals and Behavior". *International Conference on Developmental Learning*, Salk Institute, San Diego, CA.
- [16] Pentland, A., Choudhury, T., Eagle, N., and Singh, P. (2005) *Human Dynamics: Computation for Organizations*, *Pattern Recognition Letters*, Vol. 26, pp. 503-511.