

Identifying and facilitating social interaction with a wearable wireless sensor network

Joseph A. Paradiso · Jonathan Gips ·
Mathew Laibowitz · Sajid Sadi · David Merrill ·
Ryan Aylward · Pattie Maes · Alex Pentland

Received: 6 October 2008 / Accepted: 9 March 2009
© Springer-Verlag London Limited 2009

Abstract We have designed a highly versatile badge system to facilitate a variety of interaction at large professional or social events and serve as a platform for conducting research into human dynamics. The badges are equipped with a large LED display, wireless infrared and radio frequency networking, and a host of sensors to collect data that we have used to develop features and algorithms aimed at classifying and predicting individual and group behavior. This paper overviews our badge system, describes the interactions and capabilities that it enabled for the wearers, and presents data collected over several large deployments. This data is analyzed to track and socially

classify the attendees, predict their interest in other people and demonstration installations, profile the restlessness of a crowd in an auditorium, and otherwise track the evolution and dynamics of the events at which the badges were run.

Keywords Electronic badge · Wearable sensing · Wearable computing · Social dynamics

1 Introduction

Electronic badges have a rich history that dates back circa 15 years to the dawn of ubiquitous computing. These first electronic nametags, pioneered by Olivetti Research as the “active badge” [1], were very simple platforms that periodically transmit 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. Other approaches used the badge as a dynamic display and as a facilitator for person–person interaction at large events. This direction was taken by two mid-90s 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.

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. A subsequent badge, designed primarily in collaboration between Georgia Tech and Charmed

J. A. Paradiso (✉) · M. Laibowitz · R. Aylward
MIT Media Lab, Responsive Environments Group,
Cambridge, MA 02139, USA
e-mail: joep@media.mit.edu

M. Laibowitz
e-mail: mat@media.mit.edu

R. Aylward
e-mail: aylward@media.mit.edu

J. Gips · A. Pentland
MIT Media Lab, Human Dynamics Group,
Cambridge, MA 02139, USA
e-mail: jgips@media.mit.edu

A. Pentland
e-mail: sandy@media.mit.edu

S. Sadi · D. Merrill · P. Maes
MIT Media Lab, Ambient Intelligence Group,
Cambridge, MA 02139, USA
e-mail: sajid@media.mit.edu

D. Merrill
e-mail: dmerrill@media.mit.edu

P. Maes
e-mail: pattie@media.mit.edu

Technologies [5] was a product aimed at facilitating conference events. In addition to similarly displaying affinity by flashing red/green according to preloaded profiles, the CharmBadge subsequently gave attendees a log of the other badged people and IR-tagged exhibits they appeared to be most interested in, determined primarily by integrated IR encounter time. A more recent product, the nTAG [6], was also designed to facilitate business meetings and conferences. It features a 128×64 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 <20 feet from them. The IntelliBadge [7] is also intended for conferences, but as it is only essentially 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. A current Swiss product called “SpotMe” [8] has many features for facilitating group interaction (e.g., locating people via RSSI zoning, messaging, etc.), but is designed as a handheld PDA rather than a wearable badge. These products target applications such as tracking people through a convention hall, detecting what booths they visited or were most interested in, and (especially for the nTAG), exchanging virtual business cards and encouraging inter-attendee interaction.

Badges and other wearable or mobile platforms are increasingly used to infer and facilitate social interaction using measurements based solely on proximity. Some of these systems (e.g., [9]) are minimal IR transceivers that harken back to Olivetti’s Active Badge, while others are simple RF beacons running on mobile devices [8, 10] or Bluetooth-enabled cell phones [11]. A recent badge called the “Life Thermoscope” from Hitachi Research was inspired by our work—it contains a suite of sensors, such as an accelerometer and a thermometer, and logged data from these badges is analyzed to infer relevant patterns in the user’s activity [12].

While badge platforms have evolved considerably since their inception at the beginning of the 90s, their recent manifestations have been aimed more at niche commercial markets rather than research. The system described in this article has been designed to support both directions. We have developed a very versatile platform, termed the “UbER-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, explore, 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/actuator network.

2 The UbER-badge system

A final-generation production badge is shown in Fig. 1. An early prototype system was introduced in [13] and details on the current badge’s hardware and software can be found in [14]. A complete UbER-badge (including frosted plastic faceplate) measures 11×12 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 h of continuous use. In quantities of 300, the cost of an assembled badge runs roughly US \$85, not including a case or front panel.

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 (its hard to read the LCD panels used with existing badges at any significant distance or outdoors). Accordingly, the badge was equipped with a 5×9 LED matrix driven by a dedicated controller capable of independently



Fig. 1 A UbER-Badge worn around a users’ neck with animated LED graphics

specifying the intensity of each LED. In addition, 4 brightness-controllable blue LEDs below the matrix provide another channel of 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.

Each badge can provide tactile feedback (typically felt on the neck through the badge's lanyard) via a pager-style vibrating motor. 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. Each badge also includes a 2-axis, $\pm 2\text{G}$ accelerometer to sense user motion and an ambient light sensor. Each badge is currently equipped with 2 MB of flash memory to enable the continuous recording of sensor data across the duration of a day-long event.

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 short lanyard also limited the amount of free swinging of the badge, improving the veracity of the accelerometer data.

The badge is equipped with an IR channel (875 nm modulated at 38 kHz) to support face-to-face and local communication. Badges can notice each other via the IR channel out to 3 m and across large angles (e.g., 60°). People facing one another in small groups can be reliably detected, and a group sitting around a table could be identified by transitively linking all badges that are noticed within a short interval. All badges broadcast a packet containing a unique ID code through their IR port to alert other nearby facing badges and squirts (see Fig. 2) of their presence. Although the average interval between IR pulses is 1 s, it varies by up to 25% from shot to shot to avoid persistent collisions. When a pair of badges (or a badge and a squirt) detects 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 s.

The badge also sports an RF section [15] to support higher bandwidth, non-line-of-sight communication across

larger distances. It is based around the Chipcon CC1010 radio chip, which contains a processor and RF transceiver set to run at 433 MHz with programmable transmission strength. The CC1010 software implements a peer-to-peer random access network using a carrier-sense method of media sharing and collision avoidance (CSMA). Using a simple wire monopole antenna, easily tucked behind the badge in its case, the badge's indoor RF range has been tested out to 100 m.

Infrared beacons called "squirts" were 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 2.5×5 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 m away) to inform 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 visible LEDs), an event that is logged in both the memory of the badge and squirt. As shown in Fig. 2, the squirts were generally fixed to a large placard that was posted near each demo in order to attract the attention of a nearby badge wearer in case he or she is interested enough in the project to register a bookmark.

Most data packets broadcast by the badge's RF system are not multihop routed across badges—instead, all badges directly radio a network of fixed base stations, each

Parasitic Mobility

Responsive Environments Group

To request more info on this demo:

Aim badge at the hot-spot
Do you see a green light?

Press either button on the badge.

Orange and Red Light?
The request is noted.

The UbER-Badge Demo Hot-Spot



Fig. 2 A demo placard with a "Squirt" (a compact IR tagging beacon) at lower right

plugged into the building's hardwired LAN. Six of these base stations were able to cover all three floors of our building and adequately cover other venues where we ran this system—participants were always within range of at least one base station.

The base stations enable PC-based kiosks and clients to query, command, synchronize and monitor the badges from anywhere on the Media Lab's network without the badges bearing the overhead of badge–badge routing. 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 base stations and sent to the kiosks and 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.

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 25 samples (4 Hz). Microphone signals are acquired 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 h of data recording time before the 2-MB flash memory on the badge fills up—certainly ample time to outlast a day-long event.

When not displaying messages or information, the badge display goes into a “pilot light” mode to indicate that the device is active, quiescently showing a dot that bounces around with the user's motion, driven by the accelerometer data.

3 Interactive applications

After testing and evaluating an earlier prototype badge design [13], 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 Lab consisted of two different environments, namely structured talks in an auditorium and an extended freeform “open house,” where participants could explore the Media Laboratory at will. The career fair was entirely an open house.

3.1 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 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. 3) or flashing a LED on the squirt. Analyzing the bookmark data also provides a very direct way to find groups that have common interests. Figure 4, for example, shows a scatter plot of the bookmarks that students gave to different vendors who were present at the high-school career fair—the distance between the plotted dots corresponds to the commonality of the bookmarks that


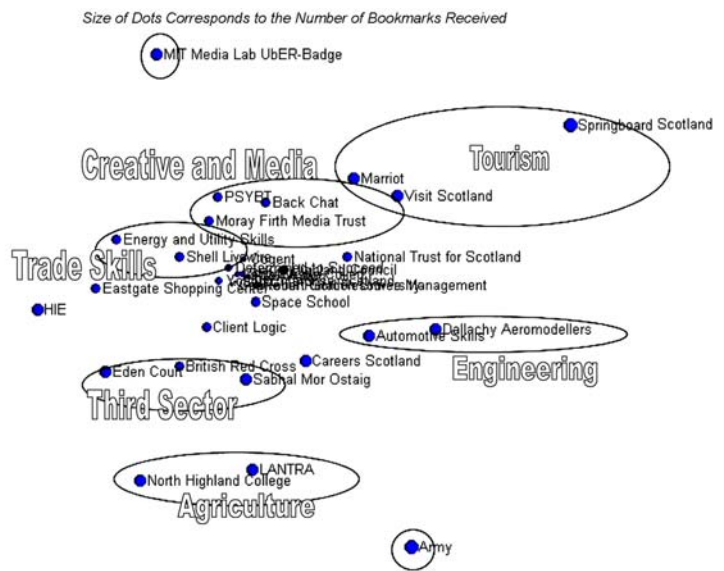


Fig. 3 The badge at *right* has bookmarked the badge at *left* by pushing a button—blue LEDs at the *bottom* of the badges signify an encounter in progress and the checkmark verifies the bookmark

Fig. 4 Distribution of vendor bookmarks at the high-school career fair, with their separation corresponding to the number of bookmarkers in common—vendors plotted closer together are bookmarked mainly by the same individuals, whereas widely separated vendors are bookmarked by primarily different groups



they received (vendors who tended to receive bookmarks from the same group of people are plotted closer together). Aside from a dense group of vendors near the center who received significantly fewer bookmarks, we tend to see vendors naturally group in topical clumps, as indicated in the hand annotation. Note that the Media Lab and Army received significant numbers of bookmarks, but were isolated islands opposite one another, indicating that they had few bookmarkers in common.

3.2 Displaying public messages

Computers running the badge management software (used by the meeting administrators) were able to command all badges to repeatedly scroll a cached 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. 5), etc.



Fig. 5 Public messages scrolling across badges

3.3 Displaying personal messages

Badge kiosk PCs distributed around the building running GuideStar software (described in the following section) could be used to send a message to one particular badge via the RF basestation network. 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.

3.4 Finding people

Badge-wearers can be physically located via two techniques. One involves simply querying one of the GuideStar kiosks 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, a 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.

3.5 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 five groups that were defined by commonality of behavior (Fig. 6)—as anticipated, members of the same company often end up in the same group. 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 (Fig. 7). This was something of a digital





Fig. 6 A tree clustering badge wearers at a consortium meeting into five primary affinity groups based upon their encounter and bookmarking activity during the open house

Fig. 7 Screenshots of animated affinity icons



“T-shirt”—nearby participants would note that their icon was similar or different, often instigating conversation about their experiences during the day.

3.6 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—inhibiting direct transmissions for frequent pushes and sending only a summary push count after the pushing stops prevents the CSMA protocol from jamming when requesting votes from large audiences. The voting function also has a “public” mode, where the degree of agreement (dictated by repeated badge pushes) is indicated by the number of LED’s illuminated on the badge.

3.7 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 many very short research summary talks (4–8 min 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, could not see the badges, all were visible to the speaker). The time-keeping 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.

3.8 Badge network control and monitoring

An application was written to manage the above applications and interface with the badge network. The standard control window for this application is presented in Fig. 10. Badge commands generated by this application were sent via the radio of a slave badge connected to the host PC’s serial port, and could be also echoed by the other badge base stations distributed across the building’s computer network.

4 GuideStar interactive kiosks

The GuideStar system complements the intrinsic capabilities of the badge. While the badge is quite capable of displaying short messages and handling simple operations like bookmarking, it lacks the user interface (UI) affordances needed to manage more complex interactions. With its full-size computer display and UI, GuideStar also allows us to present a greater amount of on-demand, conveniently

Fig. 8 Timekeeping cues flashed by the badges—progressive warnings (*top*) and scrolled text when overrunning (*bottom*)




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

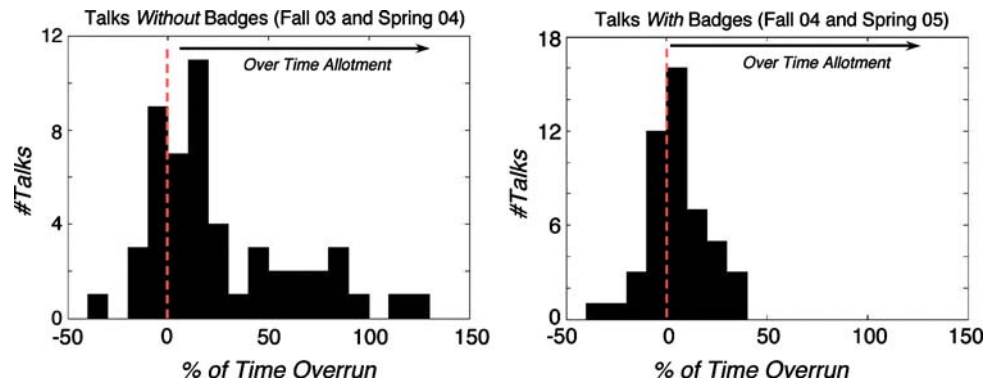
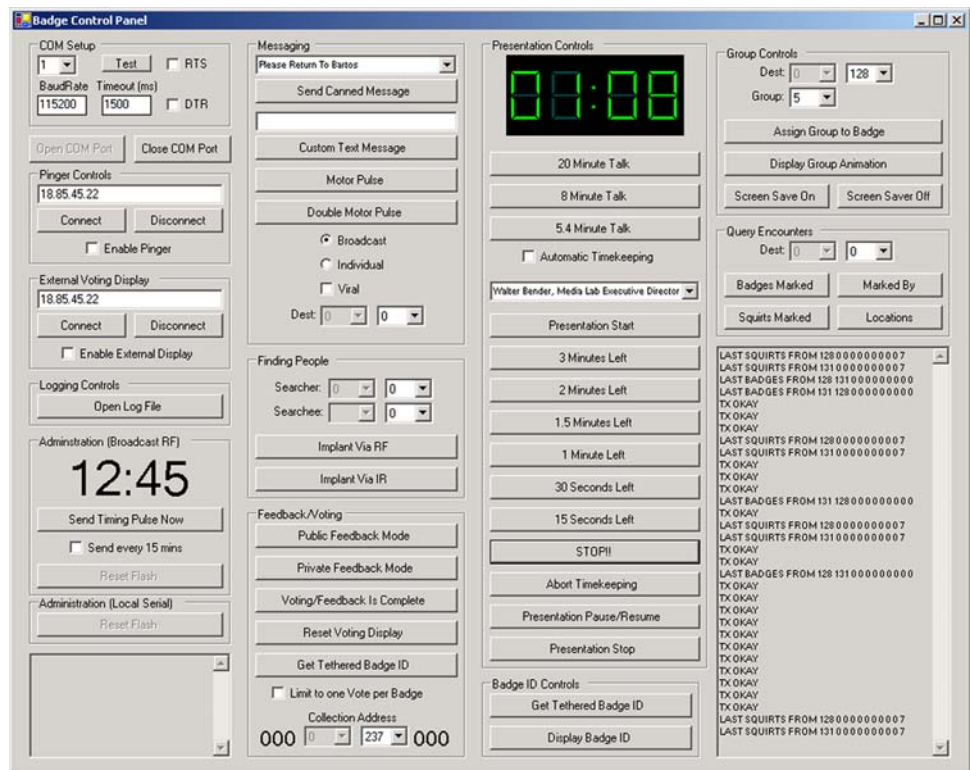


Fig. 10 Control panel window from a PC-based application that interfaces with the badge network and manages all badge system functions



located information customized to individual users, thus providing benefit to the user beyond the convenience of real-time bookmarking and voting. The system was implemented as a series of kiosks (networked PC's connected to a badge IR transceiver mounted next to the monitor) strategically placed at highly trafficked "cross-road" locations created by the structure of the venue. Each is capable of automatically identifying and greeting the user by name. This simple and widely used metaphor for ownership allows the user to immediately recognize their interaction with the system. The GuideStar kiosks provide two explicit functions to the user, as well as an implicit ability to display the distribution of users within the venue (Fig. 11). Firstly, they allow the user to locate another user graphically, simply by typing in the first few letters of their name. While the information received from the badges

about encountered squirts is generally too sparse for accurate tracking, it does provide the location where the quarry was most recently noticed (Fig. 12), which turned out to be useful information for individuals wanting to find others. Another function of the GuideStar system is to allow a user to enter text messages that can be sent to a particular badge, as mentioned in the previous section. Lastly, the system can provide users personalized recommendations for particular demos to visit during the open house, according to their evolving interests.

In order to produce the personalized recommendations, we use a combination of implicit and explicit user input. When a user first requests a recommendation, he or she is presented with a short game to explicitly select a few keywords or phrases to represent their interests. In the game, the user sees several keywords fly outward from the



Fig. 11 GuideStar personalized startup screen, showing the most recent visitor distribution at demos across the building

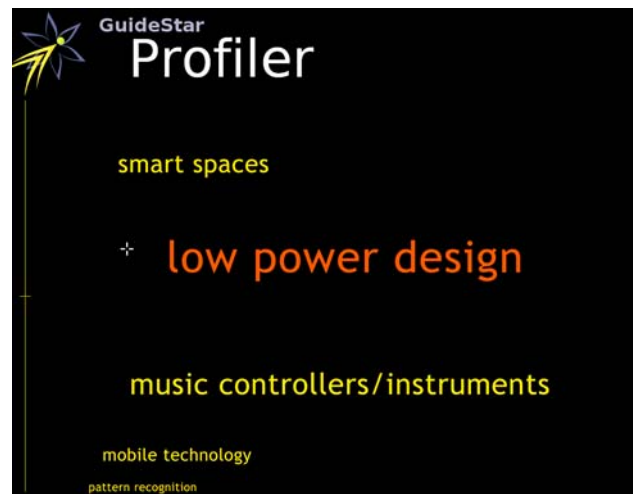


Fig. 13 The GuideStar Profiler game screen. Words fly from *right to left*, and text closest to the mouse position is made larger. Here, the user has selected “low power design”




Fig. 12 The GuideStar People-Finder screen, showing the user’s location (*left*) and the most recently noted location of the selected quarry (*right*)

center of the GuideStar display (Fig. 13). The goal of the user is to hit the word that most attracts their attention by moving the mouse cursor toward it. Feedback is provided by dynamically zooming into the word that the user moves towards. Since the words fly off the screen within a few

seconds, there is implicit time pressure to select a word quickly, which is intended to bias users toward known words. The game is designed with the express goal of being immediately rewarding and enjoyable, while providing seed information for the recommendation system. Due to the design of the game, it can only provide levels of cross-validated information similar to a formal Likert-type questionnaire when played for a several minutes. Since the users only played for very brief intervals (typically 30 s), we chose to use the game in a coarser context, presenting only high-level words within the search tree that the user could select in exclusion to others, and not displaying lower-level, more niche-centric words. This design paralleled the demands of the events in which the badge system was used, where decisions had to be made about how to invest limited time towards a large number of options (e.g., which out of many demos or groups to visit during the 4 h of open house). Preliminary testing also suggests that the time for the user to react to a word may be an indicator for how strongly the user is attached to a word. However, the presence of multiple stimuli on one screen precludes the use of this metric without a baseline measure. Nonetheless, using these initial seed keywords as binary selections, we are able to bootstrap the recommendation system immediately, thus providing the user instant output that improves over time as the system receives implicit input from the user’s behavior while exploring the venue.

Implicit user behavior data is collected whenever the badge is used to bookmark a squirt associated with a demonstration. Each demo is represented by a weighted set of keywords. When the user bookmarks a demo, its weights are impressed upon the weight from the initial seeding generated by the game. While the initial seed keywords are given a higher weight in the beginning, they can be

overridden over time by demo bookmarks representing different keywords. Since the system is meant for use within a short period, old keywords are not aged. However, this would become necessary should the system be used over longer periods, in order to prevent stale selections from improperly biasing recommendation calculation.

As mentioned earlier, each demo known to the GuideStar system is represented by a set of weighted keywords, with each weight in $(0, 1)$. Each demo can therefore be represented as an n -component feature vector, where n is the number of keywords, and the set of all demos can be represented as an $m \times n$ feature matrix, where m is the number of squirt-tagged demos. The user preferences are likewise represented as an n -component personalization vector, indicating the weight given to each keyword. The seed keywords chosen via the game are given a weight of 5—determined according to expected number of demo bookmarks between interactions with GuideStar—and each demo bookmark causes the feature vector for that demo to be added to the personalization vector. The product of the personalization vector with the feature matrix produces an m -component array of relevance ranking. Any already-seen entries (either previously suggested or bookmarked) are initially excluded. An adaptive threshold based on the largest cross-user relevance rating seen so far is applied to produce suggestions. If no unseen entries exceed the threshold, the search is expanded by increasing the relevance of other entries in the same cluster as the seen entries as a function of the relevance of the already-seen entry. This is a self-cumulative effect, which over time expands the search to the entire cluster. Since data from the badges are received over their RF link in batches, recommendations for the user are computed asynchronously ahead of time in order to reduce delay for individual interactions.

5 Identifying social dynamics and behavior

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 or during the event (e.g., the GuideStar game introduced above), or is simply hardwired into the system design. This limits the range and flexibility of these systems, sometimes 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 or she 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 suggestions 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 identify a social network that can further guide the recommendations.

Nalani Ambady and Robert Rosenthal [16] have shown that observers can accurately classify human attitudes (such as interest) from non-verbal behavior, using observations as short as 6 s. The accuracy of such “thin slice” classifications are typically around 70%, corresponding to a correlation between observer prediction and measured response of about $r = 0.40$. Initial experiments using a range of motion and sound features indicate that it is possible for computers to duplicate this human perceptual ability [17, 18]. 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. 3,703 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 flash memory 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 mean (μAUD_{AMP} , μAUD_{DIF} , μAUD_{SUB}) and standard deviations (σAUD_{AMP} , σAUD_{DIF} , σAUD_{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 (μACC_x , μACC_y , μACC_{SUB}) and standard deviations (σACC_x , σACC_y , σACC_{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.

More detail on the algorithms and procedures developed for deriving socially relevant variables from the UBER-Badge data is given in [19].

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 15 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.

Using the six highest ranked badge encounter features (σACC_y , σACC_x , μAUD_{AMP} , σAUD_{AMP} , μAUD_{DIF} , σAUD_{SUB}), 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 (σACC_{SUB} , σAUD_{AMP} , σAUD_{DIF} , μAUD_{SUB} , σAUD_{SUB}) 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 14 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 presented in [14], 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 ($r = 0.4681$, $p \ll 0.001$). 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.

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 [20]. 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

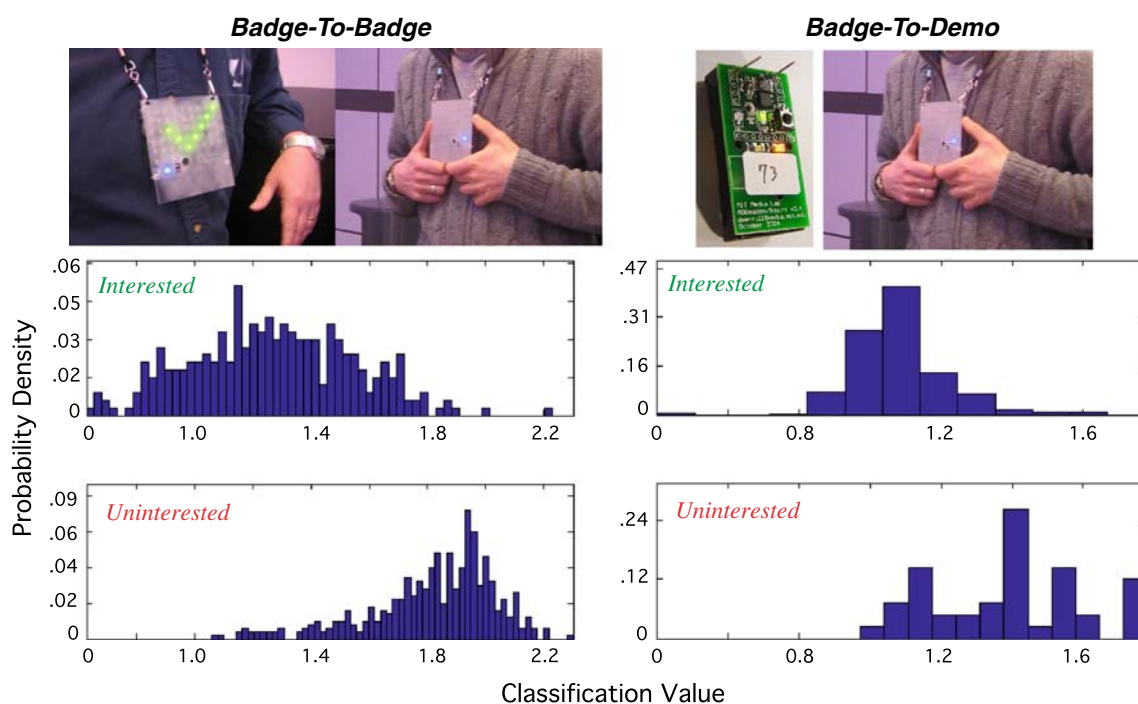


Fig. 14 Performance of interest detectors for badge-to-badge (left) and badge-to-demo (right) encounters

($r = 0.3981$, $p \ll 0.001$), producing 69.28% prediction accuracy [14].

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 does not 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 2-s period.

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 [14].

Figure 15 shows a 2D distribution of badge-wearers at the consortium meeting open house separated by a metric

derived from their dyadic energy feature—points that are close together in this diagram have high correlation in their motion. It can be seen that when just using this energy feature alone, badge wearers from the same companies tend to cluster together, indicating that they tend to move in synchrony and attesting to the effectiveness of motion features in determining affiliation.

5.3 Group restlessness during presentations

It is widely assumed that seated audience members squirm and fidget increasingly as a lengthy presentation progresses, and perhaps their degree of restlessness can change whether they are interested vs. bored. As our badges feature onboard accelerometers, and our participants endured several hours of rapid-fire lectures in our auditorium, we have examined the data from our spring meeting to look for correlations of this sort.

Figure 16 shows accelerometer data from all badges accumulated throughout the entire day of a 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 tends to involve less motion dynamics, since people are seated and listening, as opposed to moving around and talking to one another (the collective audio data show a similar, although somewhat less distinct, segmentation, as the sound environment is different during each phase of the

Fig. 15 Distribution of badge wearers for the consortium meetings across an affiliation metric derived solely from their dyadic energy feature

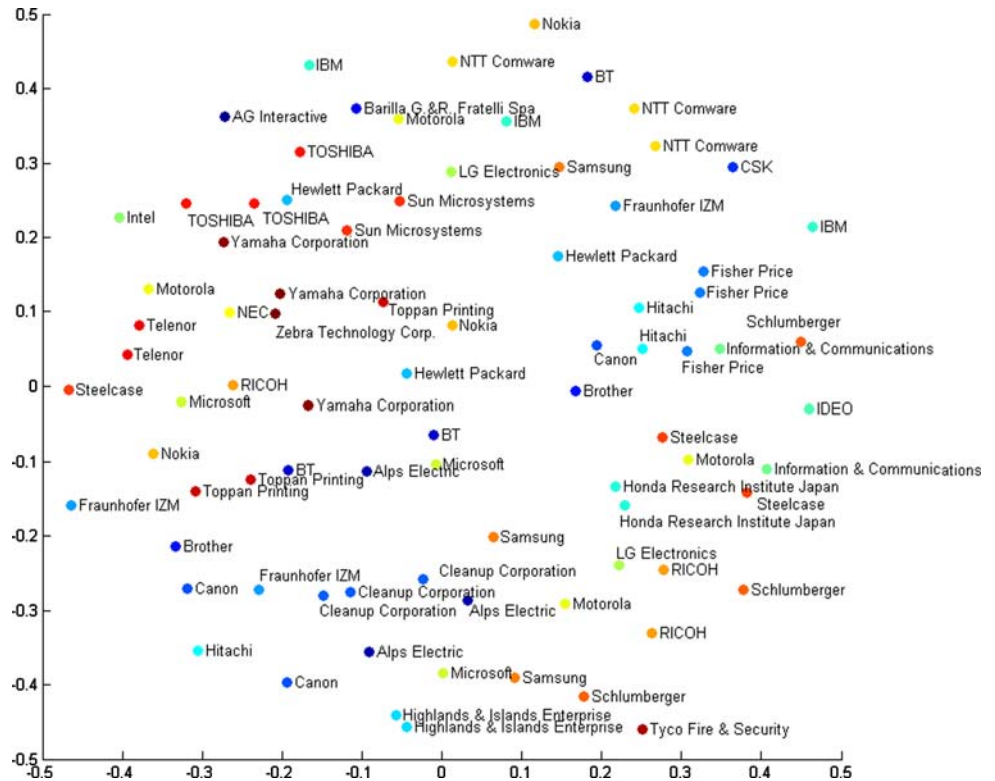



Fig. 16 Accelerometer signals for each badge plotted across an entire day-long consortium event



meeting). We have examined the accelerometer data during the three talk sessions to look for meaningful trends. As participants during this meeting were asked to vote (push a button on their badge) if they wanted more detail sent to them about the presentation that they were hearing, we had a set of labels that indicated interest in particular talks. Although more analysis could be performed, we saw no significant correlation between our fidget feature (the variance in accelerometer data integrated across each talk and

normalized by the talk duration) with votes of interest. At this stage in our analysis, we also do not see any significant indication of an increase in fidgeting with the net amount of time people were continuously sitting in the auditorium (sessions would last between 1 and 2 h). As plotted in Fig. 17, we did, however, see a medium-low, but still significant correlation ($r = 0.42, p < 0.001$) between the duration of each individual presentation and the fidget feature, suggesting that people tended to become more

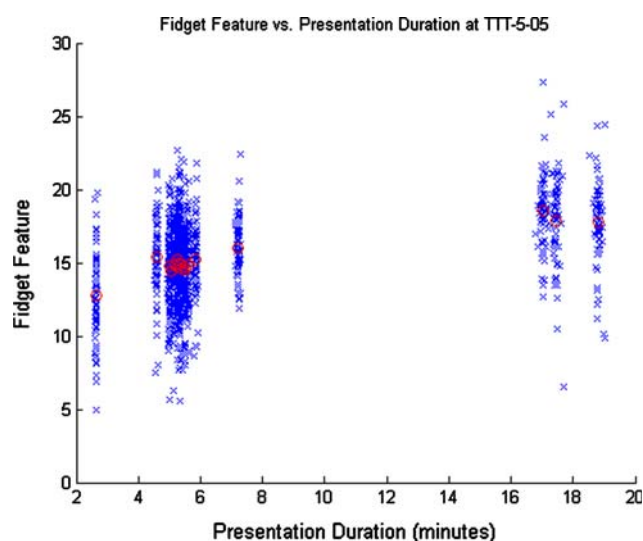


Fig. 17 Fidgeting distributions (average is denoted by red circles) vs. the length of individual presentations for all sessions in the auditorium

restless as individual talks dragged on (the two morning sessions consisted of circa fifteen 5-min talks each, and the afternoon session consisted of three 20-min talks).

6 Conclusions and ongoing work

The UBeR-Badge proved itself to be a useful platform for acquiring data for analyzing social dynamics, while providing a set of convenient user features that made wearing the badge worthwhile for event participants. The badge design was 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, indicating that people will willingly take part in such experiments if they derive direct benefit from the process. Although the LED array limited displays to simple scrolling text and coarse animation, its high brightness proved to be very effective, as it was easy to notice and read the badges from large distances (e.g., across a room). A display made of bi-colored LEDs (rather than the fixed yellow or green LEDs used in the current badges) would have still been economically feasible and quite useful, e.g., for readily displaying agreement/disagreement as in the Thinking Tags [2] and other modes of user status.

Even though everyday nametag badges are generally designed to be primarily seen by people other than the wearer, some individuals expressed discomfort at having messages for other people display on 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 our analysis was preliminary, we see a significant correlation with the restlessness of a seated audience with the duration of the presentation that they are observing.

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 common RSSI fingerprinting or interpolation technique could serve to coarsely locate them [e.g., 21], with refinement provided by acquired IR squirts.

Although our GuideStar badge kiosk system worked well, its usage during the consortium meeting was limited, as there were not enough kiosks scattered around the building, making GuideStar encounters less frequent than anticipated.

Our sensor analysis shows that we can automatically generate bookmarks that approximate the decisions made by UBeR-Badge wearers with over 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 toward 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. Indeed, cell phone manufacturers have explored prototypes of modular phones, where a thin wireless display can detach, enabling it to be worn or carried away from the body of the phone [22]. 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 correlated motion cues derived from an accelerometer that could be perhaps carried in one's pocket or purse.

We have subsequently designed a smaller next-generation badge platform called the *Communicator Badge* [23] (Fig. 18) that sacrifices the display and operates quiescently at much lower power by using analog signal processing and passive wakeup techniques [24]. In addition to supporting our own low-overhead RF networking scheme, these badges also are Bluetooth enabled, enabling them to interact with common platforms such as mobile phones. These badges were designed to also support functions useful for the wearers (e.g., push-to-talk person-to-person direct communication, as they host an onboard speaker) to

Fig. 18 Recent badges from the MIT Media Lab to facilitate research involving social interaction. The *Communicator Badge* (left) designed by Olguin in collaboration with several of us [23] and the *uBadge* (right) designed by Laibowitz [28]



encourage them to be worn and kept on (a recent repackaging of this badge has eliminated this feature to make the package roughly 50% smaller). They have been used in recent social networking studies by the MIT Media Lab's Human Dynamics Group that explore interpersonal dynamics occurring at various stages of projects evolving at large companies [25]. To accomplish this, badges are deployed to all the employees within participating groups at the commercial organization for periods of 2–6 weeks. As in the work presented in this paper, the primary data recorded by these badges are conversational patterns, tone of voice, body motion, and proximity. To this face-to-face information, we can add email exchanges, and patterns of other digital communications. This produces a fairly complete profile of communications within the organization, but does not include any content, thus avoiding some of the more difficult privacy problems that can arise in such pervasive deployments. These communication patterns are then compared to productivity metrics conventionally used by these businesses, including measures of creative productivity. The results have been dramatic, showing that variations in patterns of communication can account for up to 40% of variation in productivity [26]—this work has won the 'Breakthrough Idea of 2009' award by Harvard Business Review [27]. As a consequence, several companies are now beginning to use these techniques to improve results in call centers and in programming teams.

We are now finalizing an even smaller badge with more sensing capability, an OLED display, and onboard Zigbee radio—this *μBadge* (Fig. 18) [28] will be used in a framework that automatically labels and assembles media clips according to queries based on social context inferred from the badge signals [29].

Acknowledgments We thank many of our colleagues at the MIT Media Lab who participated in the UbER-Badge project, namely Mark Feldmeier for badge mechanics, Matt Malinowski helping with badge construction, Lisa Lieberon for logistics, Missy Corley and Michael Broxton for fabrication help, Wen Dong for help with audio analysis, Walter Bender for analyzing and disseminating bookmark data to event attendees, and 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.

References

1. Want R, Hopper A, Falcao V, Gibbons J (1992) The active badge location system. *ACM Trans Inf Syst* 10(1):91–102
2. Borovoy R, McDonald M, Martin F, Resnick M (1996) Things that blink: computationally augmented nametags. *IBM Syst J* 35(3–4):488–495
3. Borovoy R et al (1998) Meme tags and community mirrors: moving from conferences to collaboration. In: *Proceedings of the ACM 1998 conference on computer supported cooperative work*. ACM Press, New York, pp 159–168
4. McCrone J (2000) You buzzing at me? *New Scientist*, no. 2221, January 15, pp 20–23
5. <http://www.imtc.gatech.edu/projects/technology/CharmBadge.html>. Also see: <http://www.charmed.com>
6. Lewis DE (2004) Breaking the ice 2.0. *Boston Globe*, vol 266, no. 138, November 15, pp C1, C14
7. Cox D, Kindratenko V, Pointer D (2003) IntelliBadge: towards providing location-aware value-added services at academic conferences. In: Dey AK et al (eds) *UbiComp 2003*. Springer (LNCS), Berlin, pp 264–280
8. Fund J (2003) In the fray: people spotters—European Gizmo tells who's who. *Wall St J*, vol 241, no. 16, January 23, p D8
9. Bradford Paley W, Han JY (2004) Trace encounters: a social networks visualization at *Ars Electronica 2004*. In: Stocker G,

- Schopf C (eds) *Timeshift* (Proceedings of Ars Electronica 2004). Hatje Cantz Verlag, Ostfildern-Ruit, pp 252–255
10. Terry M, Mynatt ED, Ryall K, Leigh D (2002) Social net: using patterns of physical proximity over time to infer shared interests. In: *Extended abstracts of the ACM conference on human factors in computing systems (CHI 2002)*. Proceedings of CHI 2002, Minneapolis, pp 817–817
 11. Eagle N, Pentland A (2005) Social serendipity: mobilizing social software. *IEEE Pervasive Comput* 4(2):28–34
 12. Yano K, Sato N, Wakisaka Y, Tsuji S, Ohkubo N, Hayakawa M, Moriwaki N (2008) Life thermoscope: integrated microelectronics for visualizing hidden life rhythm. In: *Proceedings of the IEEE solid-state circuits conference, ISSCC 2008, San Francisco, 3–7 February*, pp 136–137
 13. Laibowitz M, Paradiso JA (2004) 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, Wien, pp 363–368
 14. Laibowitz M, Gips J, Aylward R, Pentland A, Paradiso JA (2006) A sensor network for social dynamics. *Proceedings of the fifth international conference on information processing in sensor networks (IPSN 06)*, Nashville, pp 483–491, April 19–21
 15. Laibowitz M, Paradiso J (2004) Wearable wireless transceivers. *Circuit Cellar* 163:28–39
 16. Ambady N, Rosenthal R (1992) Thin slices of expressive behavior as predictors of interpersonal consequences: a meta-analysis. *Psychol Bull* 111(2):256–274
 17. Madan A, Caneel R, Pentland A (2004) Towards socially intelligent wearable networks. In: *Proceedings of the student colloquium, the 8th IEEE international symposium on wearable computers (ISWC 2004)*. Institute for Computer Systems and Networks, UMIT, Austria; Electronics Laboratory, ETH Zurich, pp 30–32, October 31–November 3
 18. Pentland A (2004) Social dynamics: signals and behavior. In: *Triesch J, Jebara T (eds) Proceedings of the third international conference on developmental learning (ICDL'04)*. Salk Institute, San Diego. UCSD Institute for Neural Computation, pp 263–267
 19. Gips JP (2006) Social motion: mobile networking through sensing human behavior. M.S. thesis, MIT Media Laboratory
 20. Pentland A, Choudhury T, Eagle N, Singh P (2005) Human dynamics: computation for organizations. *Pattern Recognit Lett* 26:503–511
 21. Gwon Y, Jain R, Kawahara T (2004) Robust indoor location estimation of stationary and mobile users. In: *Proceedings of IEEE INFOCOM 2004, vol 2, Hong Kong*, pp 1032–1043
 22. Ferenczi PM (2006) Super prototypes—what's next for mobility—Motorola goes modular. *Laptop*, vol 24, no. 8, January, p 98
 23. Olguín Olguín D, Paradiso JA, Pentland A (2006) Wearable communicator badge: designing a new platform for revealing organizational dynamics. *Proceedings of student colloquium proposals of the 10th international IEEE symposium on wearable computing (ISWC)*, Montreux, pp 4–6, October 11–14
 24. Malinowski M, Moskwa M, Feldmeier M, Laibowitz M, Paradiso JA (2007) CargoNet: a low-cost micropower sensor node exploiting quasi-passive wakeup for adaptive asynchronous monitoring of exceptional events. In: *Proceedings of the fifth ACM conference on embedded networked sensor systems (SenSys'07)*, Sydney, pp 145–159, November 6–9
 25. Olguín Olguín D, Waber B, Kim T, Mohan A, Ara K, Pentland A (2009) Sensible organizations: technology and methodology for automatically measuring organizational behavior. *IEEE Trans Syst Man Cybern B Cybern* 39(1):43–55
 26. Pentland A (2008) *Honest signals: how they shape our world*. MIT Press, Cambridge
 27. Farrall D (2009) How social networks network best. *Harvard Business Review*, vol 87, issue 2, February, p 37
 28. Laibowitz M, Gong N-W, Paradiso JA (2009) Wearable sensing for dynamic management of dense ubiquitous media. In: *Proceedings of 6th international workshop on wearable and implantable body sensor networks (BSN 09)*. IEEE CS Press, pp 3–8
 29. Laibowitz M (2009) Distributed narrative extraction using wearable and imaging sensor networks. Ph.D. thesis, MIT Media Laboratory