Proactive Displays: Enhancing Awareness and Interactions in a Conference Context

Joseph F. McCarthy, David H. Nguyen, Al Mamunur Rashid, Suzanne Soroczak

Intel Research Seattle
1100 NE 45th Street, Suite 600
Seattle, WA 98105 USA
jmccarth@acm.org, dnguyen@erstwhile.org,
arashid@cs.umn.edu, suzka@u.washington.edu
http://www.proactivedisplays.org/

Abstract. We have designed a suite of *proactive display* applications that detect people in their vicinity and display content that is intended to promote greater awareness and interactions among those people. A conference provides an ideal context in which to explore the use of proactive displays, as attendees come together for the purpose of *mutual revelation*, eager both to learn more about others and what others are doing and to tell others about themselves and what they are doing. In this paper, we describe the design goals, technologies and algorithms used in a suite of proactive display applications that we deployed at a recent conference to aid and abet this desire for mutual revelation in the context of a paper presentation session, a demonstration and poster session, and an informal break area at a conference. We present results from our evaluation of the impact these applications had on attendees' experience of the conference and report on some lessons we learned from the deployment.

1 Introduction

Large electronic displays are becoming ubiquitous, as the technology advances and the costs decrease, showing up in an increasing variety of physical contexts, such as airports and train stations, billboards along the roads, retail stores and, with the growing popularity of high definition television (HDTV), people's homes. At the same time, sensing technologies are proliferating, from sophisticated multi-purpose sensors [Kahn, *et al.*, 1999, Gellersen, *et al.*, 2002] to rather simple radio frequency identification (RFID) tags and associated readers. We are exploring how these two trends may converge to create opportunities for *proactive displays* that can sense their context – nearby objects, people and/or activities – and respond with appropriate content.

Any proactive display application must address a number of research questions:

- What *contexts* are most amenable to the successful deployment of a proactive display?
- What kinds of information or *content* are best suited to the context(s) in which the displays are situated?
- What mechanisms are most appropriate for allowing users to *control*, influence or interact with the content in the context(s) of use?

Conferences provide an ideal setting in which to explore these questions. The conference context is often semi-public: many conferences are open to all, but the constraints of theme, time, place and cost typically result in a self-selected group of attendees with common interests who come together with the shared goal of *mutual revelation*: seeking to learn more about others and their work, as well as being open to opportunities to tell others about themselves and their own work. Furthermore, many conferences offer a variety of sub-contexts for different kinds interactions, including formal presentations such as keynotes, papers and panels; informal presentations such as may occur around demonstrations or posters; and the casual exchanges that typically take place during breaks and receptions.

Another advantage to designing for the conference context is the availability of content: conference attendees routinely reveal information about themselves – such as their names and the institutions with which they are affiliated – through conference registration forms filled out before or during the conference and badges they wear at the conference. Conference programs and proceedings offer another source of information about some of the attendees at the conference.

There are a range of possibilities for enabling people to control or influence the content shown on displays in various contexts within a conference. At one extreme, the displays could be interactive, allowing people to directly manipulate the information shown on the display, e.g., by interacting directly with the display or explicitly submitting content through some other device. At another extreme, the displays could be completely unresponsive, pushing content without regard to whoever happens to be nearby. A middle ground is to design applications that operate in a *proactive* mode, allowing the presence of those nearby to influence the content without requiring them to directly interact with or otherwise manipulate what is shown on the display.

In the design of the proactive displays applications for the conference context, our goals were to:

- Promote greater *awareness* and *interaction* opportunities within the community represented by attendees (aiding and abetting their desire for mutual revelation)
- Fit within the *common practices* at conferences as closely as possible (following the precept of *calm technology* [Weiser & Brown, 1997])
- Respect the *privacy* – notably, the varying desires for publicly revealing content – of the attendees

In pursuit of these goals, we designed, implemented and deployed a suite of three proactive display applications to sense and respond to the 500 attendees of a recent international conference. To fit within common practices, we used technology that would be minimally intrusive: radio frequency identification (RFID) tags that could be inserted unobtrusively into conference badge sleeves, detected by readers connected to large displays situated on the periphery in different conference subcontexts. In order to respect privacy, we required that attendees explicitly create profiles containing content that they wished to reveal on displays at the conference (such as name, affiliation and interests), and associate the profile with an RFID tag. The three applications responded to participants in different ways in different contexts to promote awareness and interactions at the conference:

- *AutoSpeakerID*: displayed the name, affiliation and photograph of someone asking a question at the microphone stand during the question and answer period following a paper or panel presentation, visually augmenting the common oral practice of introducing oneself before asking a question.
- *Ticket2Talk*: displayed images representing interests explicitly specified in the profile of someone in the coffee break area, providing a "ticket to talk" [Sacks, 1992] with that person about that interest while they are nearby.
- *Neighborhood Window*: displayed the name and picture of each participant gathered in a small lounge area, along with words and phrases found on those people's homepages (with links between people and their words); words and phrases shared by the group were highlighted, providing a sense of their shared interests.

The remainder of this paper will provide details about the common hardware and software infrastructure, more detailed descriptions of the proactive display applications, results from a survey we conducted of conference attendees on their experience with the applications, a summary of related work, and some of the ways in which our work might be extended in interesting ways.

2 Proactive Displays: Common Infrastructure & Applications

All three proactive display applications share a common infrastructure of digital profiles, physical tags, and kiosks for establishing associations between profiles and tags. In order to participate, conference attendees had to complete the following steps:

1. *Create*: establish a web-based digital profile with their name, affiliation and other information.
2. *Activate*: associate that profile with an RFID tag.
3. *Experience*: wear the RFID tag (in a plastic conference badge sleeve) when near proactive displays.

We designed this process so as to respect the privacy of attendees, requiring conference attendees who wanted to participate to *explicitly* opt-in at each step; while we tried to make each step as easy as possible, we didn't want to make any assumptions about any attendee's willingness to participate (hence a conflict between our second and third design goals). We also made it easy to opt-out: participants could stop participating at any time simply by removing their RFID tag.

More details on the profiles, RFID tags and their associated readers and antennas, and the applications we developed to allow creation and modification of profiles and association of profiles with tags, are provided below.

2.1 Profile Creation & Maintenance

The Proactive Displays web site provided the capability for creating or editing a profile before or during the conference. Each profile contains the following information:

- *E-mail address*: the participant's email address
- *Name*: the participant's name
- *Affiliation*: the organization with which the participant is affiliated
- *Photo*: the URL of a digital photograph of the participant
- *Ticket2Talk image*: the URL of an image of something that the participant would be happy to talk about at the conference
- *Ticket2Talk caption*: a short text description of the Ticket2Talk image
- *Homepage*: the URL of the participant's home page
- *Interests*: a brief list of phrases representing the participant's interests

In order to simplify the creation of a profile for some participants, we allowed people who had existing profiles on the conference web site's Community Directory to copy over information from those profiles, once they supplied the email address and password associated with a given Community Directory profile. The only required fields were e-mail address, name and affiliation: the former is required so we can send the participant a password with which to access their proactive display profile information later; the latter two simply represent the information people commonly display on their conference badges.

After a participant modifies the Homepage field, the URL is recursively crawled in the local domain only. The words and phrases specified in the Interests field (if any) are appended to the plain text that is extracted from the culled HTML documents. A probabilistic part-of-speech tagger [Mason, 2000] is used to identify noun phrases with the following regular expression pattern:

```
NP = [adjective]* [noun]+ [preposition NP]*
```

In keeping with our goal of preserving privacy, the resulting list of words and phrases is displayed to the participant, who may delete any (and all) of them; the participant must press an "Accept" button on the confirmation screen to add these words and phrases to his or her profile.

Profiles could be modified by a participant at any time by visiting the web site, entering the email address and password, and editing any field in the web-based profile form. Creation and modification of profiles was facilitated at the conference by having wireless Internet connectivity (802.11b, or WiFi) throughout most of the conference. However, many people created or modified profiles at one of our kiosk stations, which is also where they could "activate" the profiles by associating them with RFID tags (more on this in Section 2.3).

2.2 RFID Tags, Readers and Antennas

We investigated a variety of technologies that can be used to augment computer displays to enable them to sense and respond to people nearby. There is a long history of research into technologies for identifying and tracking people [Want, *et al*., 1992; Bahl, *et al*., 2000], and there are a growing number of companies that offer products and services using such technologies (e.g., Elpas, Savi, and Alien Technologies).

Since one of our goals was to fit within common practices, we wanted to use technology that was inexpensive and unobtrusive. Infrared, active RFID and WiFi solutions all involve tags or other devices that are expensive enough to require that they be turned in by participants after the conference, possibly requiring us to institute rigorous policies to reclaim the devices. Passive RFID tags are sufficiently inexpensive that they are, in effect, disposable. Furthermore, RFID tags are small enough – and in particular, thin enough – that they fit unobtrusively into conference badge sleeves that are worn as part of the common practice at many conferences. Other tags and devices, which require batteries, are heavier and bulkier, and thus somewhat more obtrusive.

We selected Alien Technology's 915 MHz passive RFID tag reader system, which has a read range of up to 5 meters. A reader and a pair of antennas cost approximately US$3000, and the feather-weight, paper-thin C-shaped tag costs US$0.50 and fits easily into a conference badge sleeve. Figure 1a shows an RFID tag reader and antenna; Figure 1b shows an RFID tag and a conference badge.

Fig. 1. (a) RFID tag reader (right) and antenna (left); (b) RFID tag and conference badge.

2.3 Kiosks

Two applications ran at the two kiosk stations we provided at the conference, one for profile creation and modification (as described in Section 2.1), the other to enable tag activation. In order to activate a tag, a participant had to access his or her profile, place a new RFID tag on the RFID antenna at the kiosk, and accept the terms of the privacy policy shown on the screen. The identification number for the tag on the antenna is thereby associated with that participant's profile in the database. The participant can then insert the tag into his or her badge sleeve so that RFID readers associated with the proactive display applications can sense the participant when he or she is nearby.

3 Proactive Display Applications for the Conference

We designed, implemented and deployed three applications at the conference: *AutoSpeakerID*, which displays the picture, name and affiliation of a person asking a question at the microphone during a question & answer period following a paper or panel presentation; *Ticket2Talk*, which displays explicitly specified content (a "ticket to talk" [Sacks, 1992]) for any single person as he or she approaches a proactive display in the coffee break area; and *Neighborhood Window*, which displays a visualization of implicit or "discovered" content (from explicitly provided homepage information) for a group of people who are in the neighborhood of a proactive display in a lounge area at the conference. These applications are described in more detail in the sections below.

3.1 AutoSpeakerID

At the end of a paper presentation during a conference, people often approach a microphone stand to ask questions about the work described in the presentation. Most people in the audience know who the presenter is (via the presentation, the session chair's introduction and/or the conference schedule), but often don't know much about the person asking a question. A diligent session chair may remind the questioner to state his or her name and affiliation, but this is not always done, and even when encouraged to identify themselves, questioners' names or affiliations may not be heard clearly by others in the audience (especially if the questioner is hurrying to get to his or her question). This problem can be further exacerbated when the questioner's native language differs from that of a large segment of the audience.

Since conference attendees are normally prepared to state their name and affiliations, verbally, anytime they rise to ask a question during a paper (or panel) presentation, we designed AutoSpeakerID to visually augment this common practice by using a proactive display. We used an LCD projector and a 6' x 8' (183 cm x 244 cm) screen for AutoSpeakerID, in order to have a cost effective way to make the display large enough to be viewable by as many members of the audience as possible. Figure 2a shows the floorplan of the conference room, the position of the projectors and projection surfaces, and the position of the microphone stand, where the RFID antenna was mounted. The projectors and surfaces in the upper left and upper right corners were used to display the speakers' presentations. The projector and surface used for the AutoSpeakerID display, as well as the computer used to run the application, appear on the left further below the upper left corner projection surface.

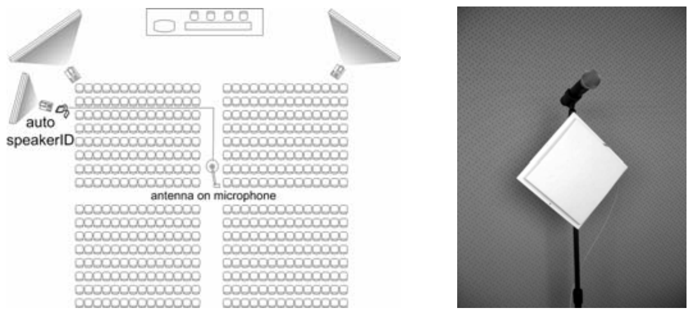

Fig. 2. (a) Floorplan of conference room; (b) closeup of antenna mounted on microphone stand

An RFID antenna mounted on the microphone stand (see Figure 2b for a closeup) was used to detect the RFID tag worn by the person approaching the microphone. AutoSpeakerID polls the RFID tag reader to which the antenna is connected every 250ms; upon sensing a tag, AutoSpeakerID displays the person's name and affiliation, along with a picture of the person (if provided), on the projection surface near the front of the room for as long as that person's tag is the sensed at the microphone stand. The projected images fade in or out gradually, over a period of two seconds, to reduce visual disruption. Figure 3a shows a screenshot from the application; Figure 3b shows AutoSpeakerID in the context of use.

Fig. 3. (a) AutoSpeakerID screenshot; (b) AutoSpeakerID in context of use

A corollary of our goal to fit within common practices was to minimize potential sources of disruption by the displays. We considered displaying more information about each speaker, e.g., their interests, involvement in the field or conference, email address or homepage. However, the more information we displayed, the greater the chance that people in the audience would focus on the displayed information rather than the question being asked. We decided that showing the name, affiliation and photo (if provided) achieved a reasonable balance between showing useful information without being too distracting.

The size and location of the screen – smaller than the main presentation screens, off to the side rather than in the front of the audience – were also designed to reduce disruptiveness. We did not want the AutoSpeakerID screen competing for attention with the screens used for the presentations. However, by placing the AutoSpeakerID screen on the front left side of the room, people in the front left quadrant of the room were forced to look left to see the name, affiliation and photo, and look right to see the person asking the question; people in the right rear side of the room could view both the screen and the questioner, but may have had difficulty reading the screen. There is always a delicate and dynamic balance between peripherality and in the design of any peripheral or ambient display application: if the display never attracts attention, it won't be useful; if it attracts too much attention, it won't be peripheral. Perhaps a second AutoSpeakerID screen on the right side of the room would have achieved the right balance in this case.

Our concern about disruptiveness also led us to look for ways to reduce spurious tag reads, e.g., when someone passes by the microphone stand (rather than approaching it to ask a question) or when someone asks a question and then steps to the side of the microphone stand (and thereby potentially establishing a clear line of sight to the tag worn by the next person in line). We created a signal-strength attenuator to reduce the power, and thereby the maximum range of the antenna, so that only RFID tags within at most 3 feet (1m) behind the microphone stand are detected.

In keeping with our goal to respect the privacy of attendees, those who did not wish to have their profile information displayed when they approached the microphone stand could opt out in a variety of ways: not creating a profile, not associating the profile with a tag, or not wearing the tag when approaching the microphone stand. Participants could control the amount of information that was displayed by limiting the information in their profiles (e.g., not including their photo). Furthermore, participants could control the information displayed in more subversive ways: wearing another person's tag or filling out profile information that was not entirely accurate.

The AutoSpeakerID display was viewable by the entire audience (a much larger group than was impacted by our other two applications), increasing our concern with potential negative effects beyond disruption, e.g., if highly offensive photographs were used. We thus implemented an on/off toggle control for the display, and had someone standing (or sitting) by at all times if in case we needed to censor content.

3.2 Ticket2Talk

A paper / panel presentation session is a rather formal context in which to deploy a proactive display. We also have proactive display applications that are designed for more informal contexts, such as a break area or a demo or poster session.

One such application is Ticket2Talk, which runs on a large plasma display in portrait mode orientation, and cycles through visual content explicitly contributed (via earlier profile creation or modification) by participants who are near the display, representing each person's "ticket to talk": a topic about which the person would be happy to talk with someone while at the conference. The ticket to talk image may represent a professional interest (e.g., a research project the participant is working on or the cover of a recently published book), or a more personal interest (e.g., a picture of a favorite pet, vacation spot or piece of art).

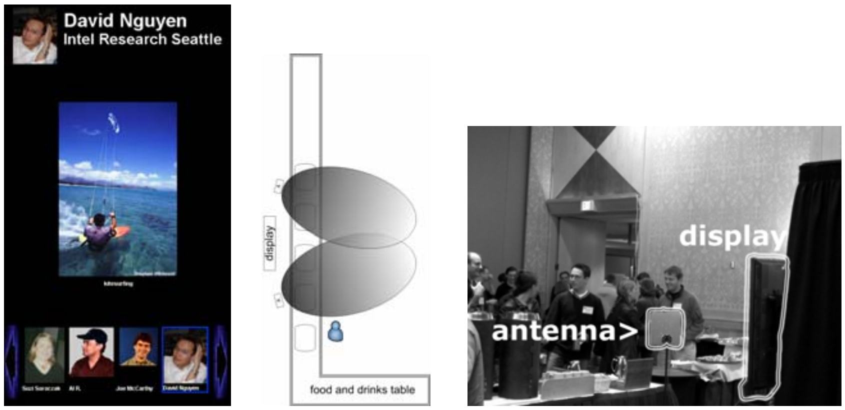

Fig. 4. Ticket2Talk (a) screenshot, (b) layout, and (c) context of use

Figure 4a shows a screenshot of Ticket2Talk. The ticket to talk image is displayed in the central region of the screen, a caption describing the image appears immediately below it (early testing revealed that it wasn't always obvious what the image was intended to represent), the picture, name and affiliation of the participant who posted the ticket to talk appearing at the top (similar to AutoSpeakerID), and a collection of thumbnail pictures and names of other people whose RFID tags have been detected near the display appearing in a row at the bottom. This queue of people is limited to four, as we considered this a reasonable compromise for providing some feedback to others in the area – if their thumbnail image is shown, their tag has been detected – and maintaining sufficient image quality of the thumbnail pictures. Arrows on either side of the queue appear whenever there are more people detected in the area than can be shown on the screen.

We designed a queue management algorithm to balance freshness with fairness. Each image is selected for display based on a priority determined by both the recency of the participant's tag being detected – higher priority for more recently sighted tags (freshness) – and the recency of the participant's ticket having been shown – lower priority for more recently displayed tickets (fairness). Images are displayed for a preset period of time; there is also a timeout period after which point a tag is considered no longer nearby, resulting in the removal of the associated profile from the queue. For the conference, we used a display period of 5 seconds and a timeout period of 60 seconds. In order to prevent disruptive "thrashing" – frequent changes to the contents being displayed – we recalculated the priority only after each person listed in the queue has had their ticket shown, e.g., if there were four people in the queue, we would recalculate priorities after 20 seconds.

At the conference, the Ticket2Talk proactive display was deployed behind a table used for a coffee urn and pastries during breaks. Figure 4b shows the layout of the display (middle of left side), antennas (either side of the display) and the approximate tag read zones (ovals in front of the antennas); Figure 4c shows Ticket2Talk in the context of use. We designed the application so that the sequencing of "tickets" shown on the display would correspond to the serial nature of the movement of people through the line, providing each person who comes through the line – and who has chosen to participate – an opportunity to both learn more about those nearby in the line and allow those same people to learn more about him or her.

The goal of this application (and Neighborhood Window) was to provide opportunities for conversation for attendees. However, as part of our goal to respect privacy and differing comfort levels, we also wanted to ensure *plausible ignorability*, i.e., no one should feel compelled to strike up conversation with a fellow attendee who happens to be nearby. By cycling through content, one can simply notice the stream of tickets, without acting on any particular one. Even when opportunities for direct conversation were not taken, the Ticket2Talk could still contribute to raising the level of awareness about other attendees' interests – helping people learn things about their colleagues that they may later choose to act on (e.g., at a demonstration or poster session, or the conference reception).

3.3 Neighborhood Window

A demonstration and poster session provides another context in which to explore the utility of proactive displays at a conference. Attendees often mill about such a session, forming ad-hoc groups as they cluster around a demonstration or poster of interest, or simply congregate off to the side to rest. The Neighborhood Window application displays a visualization of interests of the group of people in its vicinity, based on the collection of words and phrases found on their respective homepages, depicting both the unique and shared interests among that group.

While we could have run the Ticket2Talk application on a second display in the demonstration and poster session, we wanted to take advantage of this context to explore other dimensions of proactive display applications (and people's experience with them). Neighborhood Window utilizes *implicit* or latent profile information that can be attained through participants' explicit profiles, and generates visualizations of this content based on the *group* that is nearby. This contrasts with the Ticket2Talk application's utilization of explicit information about one person (at a time).

The Neighborhood Window uses a package of open-source, Java-based, graph visualization interfaces called TouchGraph, that provides support for spring-layout and focus+context techniques (http://sourceforge.net/projects/touchgraph). The nodes in the Neighborhood Window graph represent people (their pictures and names) as well as words and phrases; the links connect people to the words and phrases associated with them (via their homepages and/or list of interests specified in their profiles). Figure 5 shows a screenshot of Neighborhood Window.

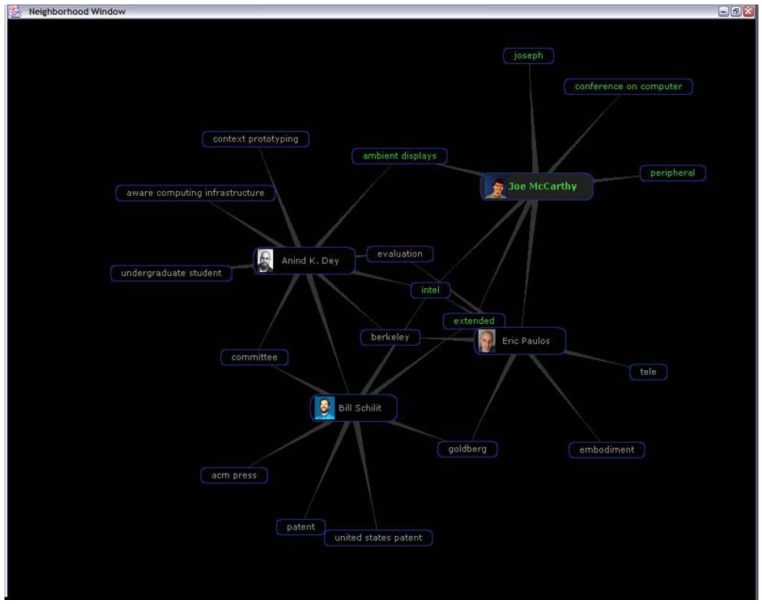

Fig. 5. Neighborhood Window screenshot

Neighborhood Window maintains a vector representation of the words and phrases contained in all the profiles (after passing through the Porter stemmer [Porter, 1980] as well as a stop-words pruner). Each participant's profile is represented by an instance of the vector where each element's value indicates the number of occurrences of a word or phrase within that participant's homepage.

These vectors are used to select the words to be associated with each nearby participant. In order to reduce screen clutter, we limit the number of participants who are included in the TouchGraph visualization (*N*) to four. We further limited the number of words associated with each participant in two different categories:

- *unique words* that are relatively unique to that participant's profile (based on standard term frequency and inverse document frequency metrics [Salton, 1988]) , which we limit to *N*, and
- *shared words* that were also found in the profiles of other nearby participants, which we limit to 2 * (*N* – 1).

As with Ticket2Talk, there are times when more than N participants are detected near the display. We use the same queue management algorithm that we used for Ticket2Talk to determine the priority with which a participant is highlighted. Highlighting increases the brightness of all the nodes and links associated with a participant for a time period (again, we used 5 seconds); to emphasize the participant's possible shared interests with other people nearby, we also highlight the links between the words shared by the highlighted person and the other people with whom those words are shared, seeking to provide prospective topics of conversation based on these shared interests.

Figure 6a shows the layout of Neighborhood Window. The display is at the top, with antennas on either side (with ovals depicting the approximate tag read zones). A loveseat is on the left hand side. Figure 6b shows Neighborhood Window in its context of use.

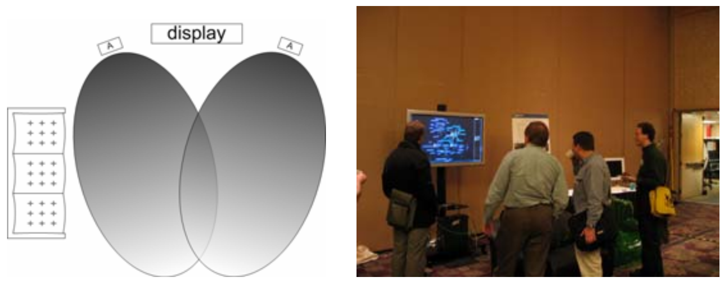

Fig. 6. Neighborhood Window (a) layout and (b) context of use

4 Evaluation & Lessons Learned

Approximately 500 people attended the conference. 229 (46%) of the attendees created profiles and 201 (40%) activated those profiles (associated RFID tags with their profiles). By design, everyone had to provide a name and email address for their profiles, but the other fields were optional: 226 (99% of people with profiles) provided their affiliations, 173 (76%) provided a list of their interests, 169 (74%) provided a homepage URL, 145 (63%) provided a photo to represent themselves, and 107 (47%) provided a Ticket2Talk image.

We conducted a web-based survey during the four weeks following the conference, a link to which was emailed to many of the attendees of the conference, to help us understand how the proactive display applications affected the experience of the conference by attendees. More specifically, we asked people to tell us about any experiences with the proactive displays and whether they thought the impact of each of the displays was positive, negative or had no impact at all.

Of the 500 conference attendees, 94 responded to the survey (a 19% response rate). A majority of the respondents (68%) reported active participation in the proactive displays deployment by creating a profile and wearing an RFID tag during the conference. Thus, survey respondents were biased toward active participation relative to the level of participation by the overall conference population.

For each proactive display application, the survey specifically asked if the respondent felt the application had a positive or negative impact on the conference. Table 1 shows the summary of responses.

	Positive	*Negative*	*Mixed / No Impact*	*No Response*
AutoSpeakerID	70	6	13	6
Ticket2Talk	40	3	33	19
Neighborhood Window	21	2	39	33

Table 1. Responses on the kind of impact each of the three proactive display applications had on the experience of attendees at the conference.

We are encouraged by the responses, which indicate a positive impact by the applications. 70 of the 94 of survey respondents (74%) reported that AutoSpeakerID had a positive impact on their experience at the conference. The smaller proportion of respondents who reported positive or negative experiences with Ticket2Talk or Neighborhood Window likely reflects the settings of each application: while AutoSpeakerID was in the main presentation room, Ticket2Talk was next to the coffee break area near the main foyer, and Neighborhood Window was set back far into the ballroom used for demonstrations. For all three applications, we are happy to note that the ratio of those who assessed their impact as positive to those who rated their impact as negative impact is 10:1.

One of the free response survey questions was “If you omitted any information from your profile, please tell us why.” The most oft-cited reasons were inaccessibility of content (we required URLs for photos & tickets), the hassles of waiting in line for the kiosks (we had two kiosks and registration took anywhere from 2-10 minutes each) and confusion about how the information would be used. Based on these responses, we believe that with more advance notice to describe the applications and encourage pre-registration, support for uploading images on-site, and more kiosks, we would have had a much higher participation rate among the attendees.

Conference attendees came to the microphone stand to ask a question 70 times during the paper and panel sessions during the second and third days of the conference.[1] During 26 of those questions (37%), AutoSpeakerID displayed information associated with the tag worn by the questioner; this level of response by the application corresponds to the 40% of attendees with activated profiles. Several survey respondents reported on the utility of having names and affiliations appear on the AutoSpeakerID display, variously commenting on how it aided their notetaking, real-time searches and followup contacts. Among those who assessed the impact of AutoSpeakerID as negative, the most oft-cited reason was the distraction it created, especially when the system didn’t function correctly. We believe this was, at times, the result of our use of a signal strength attenuator to reduce false positive tag reads; the side effect of reducing false positives was reducing true positives (people’s tags weren’t always detected immediately when they approached the microphone stand).

Three of the questioners had information in their profiles that might be characterized as “inaccurate” – including one person with a mythical affiliation, a person who listed their name as “I’m the real <name>” and a person whose profile contained the

[1] AutoSpeakerID was not running on the first day, when a single keynote and single paper session were held.

name, photo and affiliation of the well-known CEO of a large software company (who was not at the conference). Interestingly, of the 18 respondents who specifically commented on this "gaming" of the system, 16 (89%) described this type of use of the system in positive terms (one person reported enjoying "the little performance pieces"). Fortunately, despite some playful gaming, nothing offensive ever appeared on the AutoSpeakerID display, and thus the "kill switch" we implemented to immediately blank the screen in the event of any offensive content being displayed was never used.

Due to the respective settings of the displays, we believe fewer people encountered Ticket2Talk and Neighborhood Window than AutoSpeakerID, and thus we have less data on specific experiences with those applications. However, several respondents reported having conversations with new acquaintances based on content shown on one or the other display. Other respondents reported learning new things about old acquaintances. There were reports of people putting names to faces or otherwise better recognizing who they were actually standing next to (or even talking with) due to something appearing on the screen relating to one or both participants. Slightly more than half of the "tickets" appear to have been related to non-professional interests (travel, sports, hobbies, family, pets) with the remainder reflecting professional interests (photos, posters or logos of projects or research groups). Respondents reported specific instances of both types of content helping to facilitate conversations.

Although we didn't explicitly ask people to compare Ticket2Talk and Neighborhood Window in the survey, several respondents noted the similarity in their perceived goals and effects. Some people preferred the relative simplicity of Ticket2Talk, while others preferred the more complex (though sometimes noisy) words and links highlighted by Neighborhood Window. We had a few reports of conversations that were initiated around Neighborhood Window based on words that were inexplicably chosen to represent people's unique or shared interests.

5 Related Work

The three proactive display applications deployed at the conference combined sensing technologies with large displays to proactively reveal information about people in the vicinity of the displays, with the primary goal of increasing interactions and awareness among people in the shared context of the conference, while fitting into common practices and respecting people's privacy. This work is related to research others have done to explore the use technologies for use by people attending various types of events, as well as research into the use of large displays in group settings.

Other researchers are exploring the use of technology to enhance interactions among people who are collocated during certain events. Several of these have used handheld computers or personal digital assistants (PDAs) to allow people to represent their interests and experiences. Woodruff, *et al.*, [2001] have explored the use of PDAs to encourage conversations among small groups during museum visits; Sumi & Mase [2002] have used PDAs to support both face-to-face interactions and recommendations of things to see at a conference or laboratory open house; the SpotMe

Conference Navigator (http://www.spotme.ch), runs on a PDA that people can use to detect the proximity of other devices used by conference attendees with similar interest profiles. We believe that each of these systems can help enhance interactions among the people who use them, and they offer the advantage of having a [small] display anywhere a user – and his or her device – is located. However the requirement that participants carry special devices which, while increasingly common in museums and some laboratories, are not in widespread use at most conferences, may entail significant cost for the equipment and its management. Furthermore, the use of a handheld device may, at times, distract or even detract from more natural face-to-face interactions in a conference context: people may pay more attention to their handheld computers than to the people around them.

Devices for promoting interactions can be worn rather than carried. nTAGs (http://www.ntag.com, see also Borovoy, *et al.*, [1998]) include infrared and radio frequency communication capabilities, as well as a small display and buttons for interaction. These devices have also been deployed at conferences, with the goal of creating conversation opportunities and raising mutual awareness among the people attending the conference). We believe that the use of large, situated displays that react to RFID tags embedded in ordinary conference badges worn by attendees fits more closely into existing practices at conferences. Also, showing content that may spark conversations on a peripheral display leaves more room for plausible ignorability – it is easier to glance at (and ignore) a display on the periphery than to ignore content shown on a display worn on the person in front of you – and thus engenders different types of social expectations, interactions, and reactions, among the conference attendees.

There is a growing body of research into the use of shared displays situated in various contexts to support different types of tasks (cf O'Hara, *et al.* [2003], for a comprehensive collection of such applications). Most closely related to our work has been research into the use of large, interactive displays to enhance the awareness and interactions among people gathered together for an event. Opinionizer [Brignull & Rogers, 2003] shows the opinions, and responses to those opinions, typed in at a keyboard near the display by people attending a party. PlasmaPlace [Churchill, *et al.*, 2004] shows content relating to a conference and its attendees, allowing users to navigate an online community directory for the conference using a trackball or touchscreen. An interesting variation of a shared display is AgentSalon [Sumi & Mase, 2002], where the display shows interactions among animated characters representing nearby users (who signal their presence via their handheld PalmGuide devices).

. All of these applications help to initiate conversations among those near the respective displays. However, the requirement that people directly and visibly interact with the displays may limit the range of people willing to step forward to do so. Our applications operate in *proactive* mode, so that the content is responsive to those nearby without requiring their direct input (while they are there). We believe that this proactive model may encourage broader participation in certain contexts; for example, having a picture appear on a large display simply because one happens to be near the display may be more socially acceptable than having to explicitly post or retrieve a picture on that display (in real-time) if the people near the display don't already know each other fairly well.

The Intellibadge system [Cox, *et al.*, 2003] allowed for people to search for conference-related information, with the additional capability of showing visualizations of aggregate information collected through active radio frequency (RF) tags worn by approximately 20% of the attendees of the SC 2002 conference. As an example, one application showed the distribution of interests among the people attending each parallel session (e.g., the number of compiler people vs. middleware people, etc.). However, these applications did not actually respond directly to people wearing the tags near the large displays. Our work explores applications that directly react to the small number of people in the vicinity of the displays, rather than showing more general, aggregate data regarding the overall conference population.

GroupCast [McCarthy, *et al.*, 2001] is an application that runs on a large display that does respond to the people nearby based on the detection of infrared badges worn by those people, and profiles previously created by them. However, GroupCast was deployed in a corporate environment where all the passersby were members of the same company (indeed, most were members of the same research group within the organization), and had profiles for only 20 people. Villar, *et al*., [2003] have also created a system that enables displays to respond to people nearby based on wearable *pendles*, small wireless devices that can store information and detect gestures. As with other systems based on handheld computers or wearable displays, this system affords users more control over the information displayed, but it has the added advantage of being potentially less distracting in face-to-face interactions (once users master the gestures). However, like GroupCast, it has thus far been limited to a lab environment, and evaluation of its impact has been largely anecdotal. In our conference deployment of proactive displays, we targeted a less restricted context, with a much larger number of people, from multiple organizations, disciplines and geographical areas; we have also collected and reported data from our evaluation of this large-scale deployment.

6 Future Work

We are encouraged by the success of our deployment of proactive displays at the conference, but recognize that we are only scratching the surface of what is a very rich area for exploration. Our future explorations in this area can be categorized into three primary areas: new processes we would want to incorporate into any subsequent deployment, new features we may incorporate into the existing applications, and new deployment opportunities.

The participation at the recent conference deployment was higher than we had anticipated. We thought that only about 10-20% of the attendees would be willing to participate; 50% participated, and we believe the vast majority of attendees would have participated if we had been better prepared. For any future deployment, we would want to have more kiosks available, advertise the web site URL more effectively (so more people could create / modify profiles away from the kiosks) and make sure all attendees better understand the applications and how they can participate.

We also have ideas about modifications and new features we might want to add to the applications. One improvement would be to incorporate better feedback mechanisms to let people know when their tags have been detected. Other examples of potential changes include better tuning of the parameters (read ranges, timer values, queue management weights), allowing for multiple tickets in Ticket2Talk, modifications of the algorithm to select words for Neighborhood Window (possibly using Decision Theory to select words that have more information value and relevance to a participant), and the accommodation of other sensing technologies (e.g., Bluetooth, perhaps via Personal Servers [Want, *et al*., 2002]).

Despite these shortcomings, several people have suggested – at the conference, in the surveys and in email – that these applications be deployed at future conferences. We would like to be able to support this, ideally through some kind of customizable toolkit that would let others manage subsequent deployments. We would also like to investigate other kinds of venues in that may also be well-suited to applications that aid and abet mutual revelation, e.g., cafés, bars, or other "third places" (cf. Oldenburg [1999]). We fully expect that new design challenges would have to be addressed as we move to other types of venues, e.g., perhaps more anonymization would be required to promote more sharing of information in some of these other venues, and correspondingly greater care would be required to balance privacy and accountability.

7 Conclusion

We have designed, implemented and deployed a suite of proactive display applications intended to enhance the conference experience for attendees by providing conversation opportunities and fostering greater awareness among the community. This paper describes the three applications – AutoSpeakerID, Ticket2Talk and Neighborhood Window – along with our design goals, the common infrastructure we built to support these applications, and the results and lessons we learned from our deployment. We look forward to sharing these applications, and more of the results of our experiences with them, with the rest of the research community. We welcome new opportunities to work together on – and reap the benefits from – ongoing investigations into the use of technology to create, maintain or enhance interactions and relationships with each other in the physical world.

Acknowledgements

The authors wish to gratefully acknowledge the contributions of a number of wonderful people to the planning, design, configuration, implementation, deployment and/or evaluation of the proactive display applications and associated hardware at the conference: Ken Anderson, Gaetano Borriello, Waylon Brunette, Sunny Consolvo, Anind Dey, James Gurganus, Michael Ham, Sabrina Hsueh, John LaMont, Sean Lanksbury, Jonathan Lester, David McDonald, Eric Paulos, Trevor Pering, Pauline Powledge, Adam Rea, Bill Schilit, and Ken Smith.

References

1. Bahl, P., and V. N. Padmanabhan, Radar: An In-Building RF-Based User Location and Tracking System. In *Proc. of IEEE Infocom 2000*, 775-784.
2. Borovoy, R., F. Martin, S. Vemuri, M. Resnick, B. Silverman and C. Hancock. 1998. Meme Tags and Community Mirrors: Moving from Conferences to Collaboration. In *Proc. of the ACM 1998 Conf. on Computer Supported Cooperative Work (CSCW '98)*, ACM Press, 159-168.
3. Brignull, H., and Y. Rogers. 2003. Enticing People to Interact with Large Public Displays in Public Spaces. In *Proc. of the IFIP Int'l. Conf. on Human-Computer Interaction (INTERACT 2003)*, 17-24.
4. Churchill, E., A. Girgensohn, L. Nelson and A. Lee. 2004. Blending Digital and Physical Spaces for Ubiquitous Community Participation. In *Comm. of the ACM*, Feb. 2004, 39-44.
5. Cox, D., V. Kindratenko and D. Pointer. 2003. IntelliBadge: Towards Providing Location-Aware Value-Added Services at Academic Conferences. In *Proc. of the Fifth Int'l. Conf. on Ubiquitous Computing (UbiComp 2003)*, LNCS 2864, Springer–Verlag, 264-280.
6. Gellersen, H-W., A. Schmidt and M. Beigl. 2002. Multi-Sensor Context-Awareness in Mobile Devices and Smart Artefacts. *Mobile Networks and Applications*, 7 (5), Oct. 2002, 341-351.
7. Kahn, J. M., R. H. Katz and K. S. J. Pister. 1999. Next Century Challenges: Mobile Networking for "Smart Dust". In *Proc. of the Fifth Annual ACM/IEEE Int'l. Conf. on Mobile Computing and Networking*, 271-278.
8. Mason, O.. 2000. *Programming for Corpus Linguistics: How to do Text Analysis in Java*. Edinburgh: Edinburgh University Press.
9. McCarthy, J. F., T. J. Costa and E. S. Liongosari. 2001. UniCast, OutCast & GroupCast: Three Steps toward Ubiquitous Peripheral Displays. In *Proc. of the Int'l. Conf. on Ubiquitous Computing (UbiComp 2001)*, LNCS 2201, Springer–Verlag, 332-345.
10. O'Hara, K., M. Perry, E. Churchill and D. Russell. 2003. *Public & Situated Displays: Social & Interactional Aspects of Shared Display Technologies.* Kluwer.
11. Oldenburg, R. 1999. *The Great Good Place: Cafes, Coffee Shops, Bookstores, Bars, Hair Salons, and Other Hangouts at the Heart of a Community.* Marlowe & Company.
12. Porter, M. F. 1980. An Algorithm for Suffix Stripping. *Program*, 14 (3), 130-137.
13. Sacks, H. 1992. *Lectures on Conversation*. Basil Blackwell, Oxford.
14. Salton, G., and C. Buckley. 1988. Term-Weighting Approaches in Automatic Retrieval. *Information Processing & Management*, 24 (5), 513-523.
15. Sumi, Y., and K. Mase. 2002. Supporting the Awareness of Shared Interests and Experiences in Communities. *Int. J. Human-Computer Studies*, 56 (1), 127-146.
16. Villar, N., A. Schmidt, G. Kortuem and H-W. Gellersen. 2003. Interacting with Proactive Community Displays. *Computers & Graphics Magazine*,. 27 (6), Dec. 2003, 849-857.
17. Want, R., A. Hopper, V. Falcao, and J. Gibbons. 1992. The Active Badge Location System,. *ACM Transactions on Information Systems*, 10 (1), Jan. 1992, 91-102.
18. Want, R., T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. 2002. The Personal Server: Changing the Way We Think About Ubiquitous Computing. In *Proc. of UbiComp 2002: Fourth Int'l. Conf. on Ubiquitous Computing*, LNCS 2498, Springer–Verlag, 194-209.
19. Weiser, M., and J. S. Brown. 1997. The Coming Age of Calm Technology. In Peter J. Denning & Robert M. Metcalfe (Eds), *Beyond Calculation: The Next Fifty Years of Computing*. Springer – Verlag, 75-85.
20. Woodruff, A., M. H. Szymanski, P. M. Aoki and A. Hurst. 2001. The Conversational Role of Electronic Guidebooks. In *Proc. of the Int'l. Conf. on Ubiquitous Computing (UbiComp 2001)*, LNCS 2201, Springer–Verlag, 332-345.