Pers Ubiquit Comput
DOI 10.1007/s00779-006-0117-0

ORIGINAL ARTICLE

Supporting the unremarkable:experiences with the obje Display Mirror

**Mark W. Newman · Nicolas Ducheneaut ·
W. Keith Edwards · Jana Z. Sedivy · Trevor F. Smith**

Received: 16 January 2006 / Accepted: 25 May 2006

Abstract Many believe that ubiquitous computing will succeed when it has faded into the background of everyday life and work—that is, when it has become mundane. This paper examines the potential for technology to enhance users' experience of their environments through the improvement of the unremarkable activities that comprise everyday experience. Based on a year-long longitudinal study, we describe how we designed, deployed, and evaluated technology to support and enhance a common but unremarkable practice: the act of connecting a portable computer to a shared display (e.g., VGA projector). We found that new capabilities of our technology introduced subtle but significant changes in the practices surrounding the sharing of information in meetings. However, we also met with substantial challenges in terms of deployment, adoption, and evaluation. We analyze and discuss these challenges in depth, in order to inform the design of future mundane, pervasive applications.

M. W. Newman (✉) · N. Ducheneaut ·
J. Z. Sedivy · T. F. Smith
Palo Alto Research Center, 3333 Coyote Hill Road,
Palo Alto, CA 94304, USA
e-mail: mnewman@parc.com

N. Ducheneaut
e-mail: nicolas@parc.com

J. Z. Sedivy
e-mail: janasedivy@yahoo.com

T. F. Smith
e-mail: trevor@trevor.smith.name

W. K. Edwards
GVU Center, Georgia Institute of Technology,
85 5th Street NW, Atlanta, GA 30332, USA
e-mail: keith@cc.gatech.edu

1 Introduction

In their paper "Unremarkable Computing," Tolmie et al. encouraged designers of ubiquitous computing technology to investigate the use of computing to support the mundane, everyday routines of life [29]. This agenda lines up with the original vision of ubiquitous and calm computing as expressed by Mark Weiser [31] and refined via ideas like "everyday" [22] and "ambient" [33] computing that have driven much research over the past several years.

Our work follows on a broader methodological and theoretical framework from the social sciences. Ethnomethodology [10] alerts us to the importance of studying the "detailed and observable practices that make up the incarnate production of ordinary social facts" [17]. These "ethnomethods" are, by definition, mundane and invisible (which is why they require such detailed observation to be uncovered or, alternatively, some form of "breaching experiment" to foreground them). Yet they are the foundations of the orderly functioning of society. Along the same lines, Suchman's influential book *Plans And Situated Actions* [28] emphasized the situated nature of human activity, and the importance of the context surrounding any instance of technology use. This context is, at first sight, also mundane and invisible. Thorough studies of, for instance, a person's work environment, are meant to reveal its effects.

In this paper we describe our attempts at studying, designing, and deploying ubiquitous computing technology specifically meant to support mundane practices. There are two sides to the mundane that are important for systems designers. On the one hand, achieving mundanity is a *goal* we have for much of the technology we design—we hope that our systems will eventually integrate so well into our users' lives that they will become ordinary and commonplace. At the same time, mundanity can be seen as a *topic for investigation* in the design of ubicomp technologies. In this latter case, we set out to understand, support, and improve activities that have *already become mundane*. In this paper, we hope to shed light on this latter, less studied aspect of the relationship between technology and the mundane.

To illustrate our focus on the application of ubicomp technologies to the mundane, we report on the design and deployment of a lightweight service to support and enhance the commonplace activity of connecting a portable computer to a shared display (e.g., VGA projector). Our experiences with this service provide insight into the strategies and methods that are required in order to tackle this peculiar design domain. However, our investigations into the unremarkable were not without difficulties. We encountered important challenges before, during, and after the deployment of our mundane technology. These challenges either differ entirely from those traditionally encountered during application design, or are amplified versions of more traditional problems. We discuss these difficulties, with the hope of informing and facilitating future work in this domain.

In the coming pages, we first describe our initial, 6-month study into device-oriented behaviors in a meeting room at our facility. Based on this study, we designed and deployed a service, the ''Obje Display Mirror,'' that was deployed in the same meeting room, and describe our observations regarding the adoption and use of the service after 6 months of (mostly) continuous availability. We then discuss what we learned from our experiences in the design, deployment, and evaluation of technology to support mundane activity. Finally, we provide our recommendations for other researchers interested in supporting unremarkable, everyday behavior.

2 Observing device-oriented practices in the workplace

The intent of our work is to advance the agenda of unremarkable computing by not only describing people's interactions with technology but by introducing technology that intervenes in those interactions, and by understanding the results of doing so. Thus we found it necessary to augment the ethnographic methods employed by Tolmie et al. with quantitative data collection via longitudinal sampling in an effort both provide concrete information about the nature and quantity of various interactions that were taking place, and to provide a before-and-after comparison of behavior among our study population.

2.1 Studying meeting rooms

We focused our observations on one meeting room in our facility that is used regularly by a variety of individuals from different parts of the organization including researchers, managers, and support staff. It is also commonly used to meet with external visitors. The room is equipped with a large oval conference table (14 seats), two whiteboards, a whiteboard capture camera, wireless and wired network connections, and a ceiling-mounted XGA projector.

Meeting rooms have been extensively studied in the CSCW literature [23, 30]. They have also been privileged experimental sites for ubicomp researchers [15, 27]. The reasons that CSCW and Ubicomp researchers have paid so much attention to meeting rooms is clear: they are sites that are rich in human, device, and computer interaction. Moreover, a less frequently acknowledged property of meeting rooms is that they are ''boring''—that is, they are the locus of many activities that are now considered mundane, such as presenting information to a group or collectively reviewing documents. Our interest in meeting rooms stems in part from this combination of technological richness and mundane activity.

Setting out, a particular goal in undertaking this project was to determine how the presence of physically situated, continuously available networked services such as screens, speakers, printers, audio and video capture devices, and so forth might influence the expectations and practices of users who were exposed to them over a long enough period to allow them to adopt them into everyday practice. Before designing any specific technology however, we had to understand current practices. We therefore designed a study of existing device-oriented behaviors in meeting rooms with three goals in mind: (1) to determine what types of activities were taking place during meetings that could be helped or improved by the addition of continuously available services; (2) to determine which devices and resources available in the meeting rooms would be most useful if made available as networked service; and (3) to learn what types of portable personal devices were being brought into the room that we could leverage as client devices as well as a source of additional services that could be combined with the ones embedded in the room.

We describe our study and the methods we used below. We then look at how the results of this study informed the design of one of our services: the Obje Display Mirror, which allows users to connect to shared meeting room displays (e.g., projectors and plasma screens) without using cables.

2.2 Observation methods

Our goals for this project made traditional ethnographic observations alone insufficient. Indeed, we needed to capture and analyze longitudinal data: the sporadic, interleaved and fine-grained nature of the activities we wanted to observe (e.g., connecting a laptop to a screen) made it impossible for us to dedicate the number of person-hours needed to have researchers present in meeting rooms continuously. These constraints led us to consider techniques that sample behavior or activity, rather than ones that demand constant attention. The classic sampling technique, the Experience Sampling Method (ESM) [8] has been used for a variety of purposes related to ubicomp [7, 13, 14]. This technique did not seem to be appropriate for us to use, because it depends on the subjects being able to articulate the relevant aspects of their experience at the time they are sampled. Since the activity we wished to study is often not the conscious aspects, this would not work.

Another related technique is Lag Sequential Analysis (LSA), which was applied to the design and evaluation of LabScape [6] in a way that closely resembled our approach. In LSA, a set of temporal interactions are captured and analyzed, either using video or in-person observations. The interactions are coded at the granularity of the *lag*, a short time duration (in the LabScape study, the lag was 1 min). A set of interesting event types are chosen, and for each lag the presence or absence of the event is noted. This provides a record of all significant events at the granularity of the lag. Most importantly, it provides a record of temporal relationship among events (e.g., precedence, co-occurrence, etc.) However, since we were seeking to capture and analyze a greater volume of data—spanning months rather than the hours or perhaps days that are conventionally covered by LSA—we had to find a way to meet our needs with a considerably smaller investment in terms of capture and coding effort.

Therefore, we elected instead to take snapshots of meeting rooms and their visitors once a minute from three angles. It turned out that this low sample rate of data (as compared to, say, streaming video), along with the assurance that we were capturing using low-resolution consumer webcams (640×480 resolution) and were not capturing any audio, was helpful in assuaging our users' concerns about loss of privacy. Data representing times when people were in the room was later coded using a tool that we built especially for this purpose (see Fig. 1). We supplemented these automated and coded observations with interviews and direct observations of meetings.

2.3 Study results and observations

The results of our study are divided into two phases, which are distinguished by the technology that was available in the meeting room at the time. The first phase (Phase 1a) represents the technology that was available before we began our study, as was described earlier. The second phase (Phase 1b) represents the period after a second display and dedicated public PC were added to the room. Two months into the study we added a 50'' plasma display and a public PC with wireless mouse and keyboard. The PC was connected to the internal network but left logged in using a local account so that anyone in the room could gain access to the Internet and those with access privileges on the corporate network could use the PC to log into their own account, access their home directories, check their email, etc. The PC, projector, plasma screen, and one additional VGA input cable were all connected together with a 2×2 VGA matrix switch that allowed either source (PC or VGA input) to be directed to either or both of the displays (plasma or projector). Our goal with the introduction of these additional devices was to add additional display options and client capabilities using the best off-the-shelf technology we could find.

We analyzed 47 meetings in detail: 27 in Phase 1a and 20 in Phase 1b. For each of these meetings, we tabulated the number of attendees, the number of personal computing devices (e.g., laptops and PDAs), the frequency of display and whiteboard use, and the use of non-electronic resources such as printouts and paper notebooks. We supplemented these detailed observations with five interviews with meeting room users, unstructured in-person observations of meetings, and numerous informal conversations with users.

The statistics we collected are presented in Tables 1 and 2. As stated earlier, we were especially interested in:

- What client devices would be available to room users that would allow them to configure and control other aspects of the environment?
- What existing room resources were the most commonly used and what were they used for?

Fig. 1 We developed a custom tool for reviewing and coding our observational data. Each *column* represents a time sample, and each *row* represents one of the cameras installed in the room

- What problems and frustrations did users encounter?
- What types of activities are carried out in meetings, and how are different technologies employed to support those activities?

The only client devices that were seen with any regularity were laptops. One PDA was seen, and it was used for less than 5 min by someone who was also using a laptop at the same time. If cellphones were present, they were hidden and did not emerge during the meetings. Our interviews and informal conversations with participants revealed that most cellphone users at our institution leave their phones in their offices. A minority (26–29%) of attendees were observed to use a personal device during meetings, though a majority of meetings (70–75%) had at least one attendee that was using a laptop. When we inquired about people's practices around carrying laptops, we found that the decision to carry or not to carry is based on a number of factors as well as personal preference. If someone was planning to present some information at a meeting in the form of a set of slides or a document, then they would certainly bring their laptop in order to do so. However, some individuals would bring a laptop if they thought there was a chance that they could use it to help the group discussion by retrieving and displaying relevant information (e.g.,

Table 1 Frequency of personal device usage in meetings during Phase 1 of the study

Phase	Meetings	Attendees	Personal devices	Meetings with at least one personal device
1a	27	157[a]	40 (26%)	19 (70%)
1b	20	127[a]	37 (29%)	15 (75%)

[a] Many of these were repeat users, though more than 40 unique participants were observed during the course of Phases 1a and 1b

Table 2 Frequency of room device usage in meetings during Phase 1 of the study

Phase	Meetings	Number of meetings where				
		Projector used	Printouts used	Whiteboard used	PC Used	Plasma used
1a	27	11 (41%)	11 (41%)	3 (11%)	NA	NA
1b	20	9 (45%)	11 (55%)	4 (20%)	1 (5%)	1 (5%)

web sites or documents from a shared group repository). One user, ''J,'' was well-known to play this role in a variety of groups, to the extent that other users said that they would leave their laptop behind if J was going to be there because they knew that he would be available to facilitate any information retrieval task. Still others would bring laptops if they thought that their full attention would not be required for the duration of the meeting, in order to get other things done such as responding to email or working on other projects.

We had hoped that the introduction of the public PC and plasma screen would (a) allow some users to leave their laptops behind and/or (b) increase the incidence of serendipitous retrieval and display of information relevant to a discussion. It appears to have had neither effect for most users as it was only used once over the course of observations, and this instance was most likely due to the novelty of the installation. Other users expressed reservations or even distaste at the notion of using a public PC. One user even said that, for him, a computer was ''kind of like a toothbrush''—i.e., not something that you would share with other people. While this probably represents an extreme view, it was clear that among our target users the notion of a public client would not be likely to catch on in the near term. We concluded that users' laptops would be our best choice for any client software that would be used for configuring and controlling the room resources.

In addition to portable computing devices, a large portion of users used paper-based resources such as notebooks and printouts. A common pattern was that all meeting participants would show up with a printout of the same document, which had been emailed out to the group by one of the members a few minutes or hours beforehand. On some occasions, the document author would instead show up with a stack of copies of the document to hand out to other participants at the meeting. Printouts were used slightly more often than shared displays, and occasionally both modes of information sharing would be used in the same meeting. The prevalence of printouts can be attributed to a variety of factors, ranging from the simplicity and ubiquity of email as a document sharing medium to the fact that a personal printout is easy to peruse at one's own pace and annotate at will. We did not pursue the use of printed materials in meetings, though others have [16] and this continues to be an interesting area for future research.

The most commonly used room resource, by far, was the projector. Whiteboards were used on occasion and the whiteboard capture facility was used even less commonly. Clearly printers were used in advance of the meeting, as evidenced by the prevalence of printouts. On a couple of occasions a meeting attendee would leave for a few minutes and return with fresh printouts, indicating that a remote printer had probably been accessed by someone in the meeting. Other room resources such as a speakerphone and an overhead transparency projector were not used at all.

When asked to articulate problems and frustrations with this meeting room, a common complaint was that the meeting room equipment does not ''work right'' or is hard to get working. Everyone could tell a story of woe about getting their or a visitor's laptop to work with the projector because of some obscure incompatibility or setting on one or both of the devices. Another frustration was the difficulty of accessing information that one had not specifically prepared for sharing at the meeting. This frustration was most acute among attendees who had neglected to bring laptops, but could happen even when a laptop was available—for example, having a document on one's desktop computer and not being able to easily access it from the room. It should be noted that neither of these problems were seen as particularly severe, and in fact some users were hard pressed to think of any problems at all. By and large, people were reasonably satisfied with the room as it was currently configured but were open to trying new things if they didn't get in the way of what was already working.

2.4 Preparing to intervene

When we set out to design a set of services to deploy in and around our institution's meeting rooms, we took seriously the fact that the user population was largely satisfied with things as they were. Since our goal was not necessarily to make meetings more effective or efficient, but rather to explore how to design and deploy ubiquitous services that would be adopted and incorporated into the daily life of a work community, we were less concerned about the lack of a clear ''pain point'' and more concerned with ensuring that whatever technology we introduced formed a snug fit with existing practice. Therefore we made three decisions regarding the design of our service deployment:

1. Since laptops were fairly prevalent in meetings, and our users were reluctant to adopt shared computing devices, we would leverage users' existing laptops as client devices.
2. In order to facilitate adoption of our technology, we would begin with a small deployment of a single simple but high utility service. Since the most popular installed resource in the observed meeting

room was the projector, we determined that the first service to deploy would be a service that supported the use of shared displays.

3. The fact that similar meeting activities can be accomplished with varying technologies confirmed our intuition that it would be more fruitful to provide low level tools like display mirroring or document sharing that could be incorporated into a variety of activities rather than, say, a set of high-level integrated applications that would support specific activities such as brainstorming or group editing. Thus we reaffirmed our initial goal of deploying a set of loosely-coupled, flexible services that could be composed into a variety of applications by end-users.

3 The Display Mirror

The initial service offering that we created was called the Obje Display Mirror. This service allows any meeting participant with a networked laptop to mirror their laptop's screen to any public display that is running the service. In order to carry out the mirroring, the user visits a web site on our Intranet and clicks a link which downloads and runs a client application using Java Web Start. In most cases, this is a very simple operation and takes less than a minute the first time it runs and can take just a few seconds each successive time. The client application is shown in Fig. 2. Figure 2a shows the application as the user would first see it. After the user selects one of the screens and clicks "Connect," her laptop screen is mirrored to the selected screen and an additional control UI appears on her screen, as shown in Fig. 2b. When multiple users are connected, the control UI shows each connected user and allows the user to choose which users' screen to display on the projector. All connected users see the same control UI and are given the same ability to choose themselves or anyone else to take control of the shared display.

This service is not necessarily unique in terms of its core functionality. There are projectors [5] and network-to-VGA adapters [18] available for general sale that support some versions of direct screen mirroring. Well-known systems such as X11[25], VNC[24], and Windows Remote Desktop [21] allow some or all of one computer's screen to be mirrored to another computer's screen. The novel features of the Display Mirror are that (a) the conventional direction of screen mirroring is reversed—the user "pushes" their screen to a public, shared screen rather than "pulling" a remote computer's display to their local machine; (b) multiple users can mirror their computers to a single shared display simultaneously and easily control which one of them has the ability to control the shared screen; and (c) it is a stand-alone, platform independent application with a very narrow range of functionality and minimal user cost—a genre of technology we call *micro-applications* (this concept is discussed in more detail later in this paper). No screen mirroring capability was in widespread use among our target users at the time we were planning our deployment, and no such technology was deployed for public use in any meeting rooms.

3.1 Looking into the Display Mirror

Six months after the initial deployment, the service had achieved some success in terms of adoption. It had attracted a stable core of regular users and had gained a stable position in the toolbox of meeting room technologies that are employed by users. As we will report in this section, other successes were also observed: in particular, feedback we received indicated that the experience of using the service held unexpected benefits for both the direct users and the audience with which he or she was sharing data. Long-term usage data also revealed that a subtle but significant shift in meeting room data practices had taken place as a result of the Display Mirror deployment—namely that the instances and types of multiple user screen sharing increased.

3.1.1 Initial experiences with the Display Mirror

In the first few weeks of the deployment, we sat in on three meetings in which we asked one or more participants to use the Display Mirror to present any information they were planning to share. Afterwards, we asked each participant to fill out a short questionnaire and we conducted a focus group to collect initial reactions to using the technology. Since not all participants directly experienced using the Display Mirror client, we refer to those who did as the "primary" and those who did not as "auxiliary" users. All of the "primary" users felt that the process of connecting with the Display Mirror was as fast or faster than the conventional way of connecting. In addition, the experience of using the Display Mirror was felt to be more "natural" in certain ways. All said that they very much liked being able to connect wirelessly to the projector. First, the awkwardness of dealing with the physical cabling was eliminated (for instance, partici-

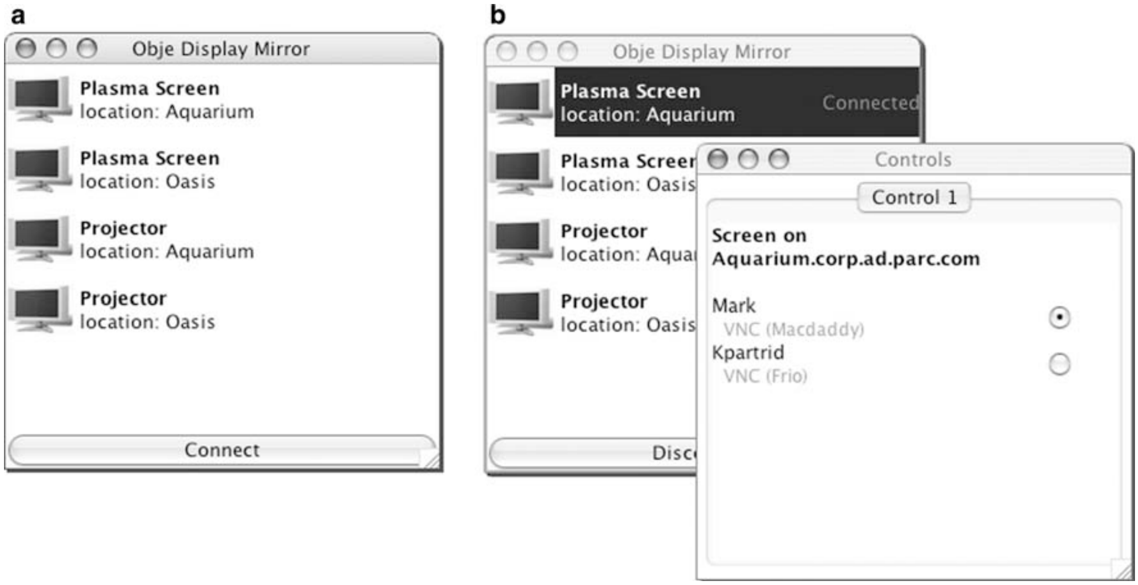

Fig. 2 The Display Mirror before and after a connection is made. First, **a** the user sees a list of available screens. After connection, a control UI **b** is presented *to all connected users* of the selected screen, allowing control of the display to be shared

pants did not have to crawl under the table to retrieve the VGA cable). Second, this newfound flexibility allowed them to sit in parts of the room that had previously been unusable for them because they were too far away from the cable.

The largest subset of ''auxiliary'' users was essentially unaware that a different technology was being used. Others remarked that the setup seemed faster and easier. One participant remarked upon the absence of ''physical thrashing about,'' which suggests, in this participant's words, that the new technology seemed ''calmer'' than the previous way of doing things.

However, we also became aware of the Display Mirror's downsides. The increased latency of the connection is an issue for certain tasks, such as group editing of a text document. The lag between when the primary users update the document and when it is reflected to the rest of the group is irritating to some users, though some find this to be unproblematic even though they are aware of the lag. For other tasks, such as showing a slide presentation and scribing notes the performance was felt to be adequate.

3.1.2 Sustained usage and experience

We analyzed device usage data from two time periods after the deployment of the Display Mirror. Phase 2a spans the 4 weeks immediately after initial deployment, and Phase 2b represents 6 weeks spanning the fifth and sixth months of the deployment.

We believe that Phase 2b represents a stable adoption state, in which novelty effects have worn off and users have had a chance to determine whether or not the service fits in with their work practices or not. To be sure, the fact that we are ourselves users of the meeting room (though our own usage data has been expunged from all statistics reported in this paper), and the fact that our colleagues are at least somewhat aware of our project's goals and that we are observing their behavior has some effect on their motivation to adopt or abandon the technology. However, we believe that the steady state adoption figures reflect that the service was somewhat successful at integrating into users' work practices. A significant portion of the rooms' users are people with whom none of our team members have any regular contact and in several cases whom none of us have ever met. Also, our assessment is that our population of users is, by and large, eager to try new technology (thus representing an ''early adopter'' mentality) but is also quite picky about what they actually adopt on a long-term basis (cf. the earlier quote about a computing environment being ''like a toothbrush''). Thus we feel confident that any users who were continuing to use the Display Mirror after 6 months had decided that it was a good fit for their practices, and in fact conversations with several of the core users confirmed this belief.

Tables 3 and 4 tell the story of the Display Mirror's effect on practice. In Phase 2a, almost no usage was recorded. This is interesting because it appears to demonstrate an anti-novelty effect—rather a reluc-

Table 3 Personal device and display usage statistics for Phase 2a and 2b

Phase	Meetings	Attendees	Personal devices	Meetings with at least one personal device	Number of meetings where			
					Projector used	Plasma used	Any display used	Display Mirror used
2a	14	82	28 (34%)	10 (71%)	6 (43%)	2 (14%)	8 (57%)	1 (7%)
2b	33	220	88 (40%)	27 (82%)	14 (42%)	4 (12%)	14 (42%)	8 (24%)

Phase 2a spanned the 4 weeks immediately after the deployment of the Display Mirror. Phase 2b represents 6 weeks of observations after the service had been available for 6 months

Table 4 Adoption figures for the Display Mirror (ODM)

Phase	Projector connections	Plasma connections	All display connections	Projector connections using ODM	Plasma connections using ODM	All display connections using ODM
2a	8	2	10	1 (13%)	0 (0%)	1 (10%)
2b	22	7	29	9 (41%)	7 (100%)	16 (55%)

In Phase 2b, the ODM was used for 55% of all display connections

tance to try an unproven technology. Therefore we used the new service on a regular basis in order to gradually introduce other meeting attendees to its capabilities and demonstrating its use. Eventually it appears to have caught on: by Phase 2b, the Display Mirror was used in over half of the meetings in which shared displays were used, and in nearly a quarter of the meetings overall. Notably, the number of meetings in which one or more shared displays were used did not change from Phases 1a and 1b, it was simply the case that the Display Mirror replaced the VGA cable as the means of connecting to shared displays in a portion of the meetings.

The most interesting result of our deployment, however, is depicted in Table 5. We observed an increase in the number and type of multiple user display events in Phase 2b when compared to all previous phases. We define a ‘‘multiple user display event’’ as occurring whenever more than one individual connects his or her personal device to a shared display during the course of a single meeting. Examples of such events include serial display events (Alice shows her slides on the projector, then Bob takes over and does the same), overlapping events (while Alice is showing her slides on the projector, Bob shows a web page on the plasma screen), and interleaved events (Alice shows something on a display, then Bob takes over the same display, then Alice resumes control). Before the introduction of the Display Mirror, the only type of events that were observed were serial display events. Overlapping events require at least two screens, so they could not have been observed in Phase 1a, but could have been observed after that. Interleaved events could have been observed at any time, but the awkwardness of unplugging and replugging VGA cables between laptops could account for the fact that such events were not in fact observed.

While we did not seek to explicitly characterize the situations in which the new display capabilities were used, it became clear through informal observations and conversations with users that people were appreciative of the new practices that had been made possible. For example, one pattern that was observed after the Display Mirror deployment was the quick display by a meeting participant of a web page to the rest of the participants that was relevant to the current discussion. This type of thing was possible before the Display Mirror but was almost never observed until after it was available. Display Mirror users who experienced this new capability reported that it was a useful aid to meeting discussions, and that the low overhead involved in displaying information to other meeting participants was the key to enabling the new practice. Another new practice that emerged after the Display Mirror deployment was for meeting participants to display different types of information on the two displays, for example displaying an agenda on the plasma display while displaying a series of slide presentations on the projector. This practice was also possible before the Display Mirror was deployed (though only after the second display was available) but was not observed until after the Display Mirror deployment.

We view the emergence of these new use patterns as the most promising outcome of our deployment thus far, as we believe it underscores the potential of small, subtle changes in mundane aspects of interactions with technology to effect subtle but significant impacts in the execution of higher-level human tasks and activities.

Table 5 Types of multiple display events with and without the Display Mirror

Phase	Total Meetings	Meetings w/ display use	Meetings w/ multiple events	Serial display events	Interleaved display events	Overlap display events
1a	27	11 (41%)	3 (11%)	3 (11%)	0 (0%)	n/a
1b	20	10 (50%)	2 (10%)	2 (10%)	0 (0%)	0 (0%)
2a	14	8 (57%)	1 (13%)	2 (14%)	0 (0%)	0 (0%)
2b	33	14 (42%)	9 (57%)	2 (6%)	2 (6%)	4 (12%)

Two new types of multiple user display events were observed. Also multiple user display increased overall

3.1.3 Challenges faced

By and large, we were able to achieve our goals with the Display Mirror. We introduced a technological intervention in the form of a lightweight, continuously available networked service, and this service was adopted by our target users, incorporated by them into everyday practice, and used to improve and extend existing practices in new ways. However, these gains were achieved at some cost.

The effort required to collect, manage, code, and analyze a year's worth of data, even given our efforts to streamline the amount of data that was collected and the analysis process, was greater than expected. Much of this effort was expended in automating the data collection, migration, and backups, as well as in building and improving the analysis tool shown in Fig. 1. The coding itself was not overwhelming—a 1-hour meeting would take approximately 30 min to code. We coded in pairs so that we could more easily notice and discuss patterns and interpretations of the data. Since a typical week would contain 5–10 meetings, the coding amounted to 6–12 person-hours per week of data (3–6 pair-hours, which includes time for loading data into the tool and searching for meeting start and end times). This is considerably less than reported by Consolvo et al. in their use of LSA [5], though our data requirements differ in that we were fairly familiar with the activities being studied and we were able to articulate in advance what types of events we were looking for.

The effort required to produce and maintain a deployment was also quite substantial. Even a fairly modest deployment involving 10–15 client devices and four display mirror services in two meeting rooms required approximately 50% of one engineer's time and 10–25% of another's over the course of the 6 months of deployment reported in this paper. Part of this was due to the fact that even a small deployment encounters many of the problems that a larger one would encounter: client platform compatibility problems, problems with the corporate network configuration, and issues with the stability of the service nodes (which were running Windows XP) to name a few.

In essence, the high degree of effort to create and maintain the data collection infrastructure as well as the deployment itself, represent poorly amortized costs. We anticipate that these costs would not be incurred to such a large degree for the deployment of the 2nd, 3rd and Nth service. However, for a single service with such limited utility, we regard the costs as having been unreasonably high. The observation that longitudinal studies and robust, sustained, networked deployments are expensive should not be news to anybody. What is worth noting, however, is that such techniques appear to be essential for the design and evaluation of ubicomp technologies. This indicates that the search for improved data collection and analysis techniques and better deployment platforms continues to be critical to the success of ubicomp research.

Another challenge that we faced appears to be more fundamental to the goal of our research. It was somewhat difficult to achieve and sustain adoption of the Display Mirror over a long period of time. We do not believe that this was as much due to limitations of the technology as it was due to its very invisibility. Connecting one's laptop to a shared display, while a common activity in aggregate, is but a sporadic, occasional activity for any one individual. It was very rare for anyone in our study to connect to the projector more than once per week. Thus it is easy to forget that a service such as the Display Mirror even exists between opportunities to use it. It is more likely that a user will default to the ingrained habit of reaching for the VGA cable than that she will consider the various options for connecting and select the one that provides the greatest utility. The very nature of mundane tasks is such that they are not foregrounded in the user's conscious mind, thus making it difficult to replace the old habit with a new one. We found it necessary to temporarily increase the visibility of the display connection activity through advertisements in the meeting room, and public and personal reminders in order to gain any traction whatsoever with our intended users.

4 Discussion

4.1 The art of the boring

In her article on the ethnography of infrastructure, Star called on researchers to ''study boring things'' [26]. She proposed that computers are frequently less of an ''information highway'' and closer to ''symbolic sewers.'' A such, she argues that we need to pay more attention to ''the plugs, settings, sizes, and other profoundly mundane aspects of cyberspace.''

Our experience in deploying and evaluating the Display Mirror resonates with many of Star's comments. A service that allows users to bypass the VGA cable when connecting to shared displays is probably not ubicomp's ''killer app''. The most widespread activities we observed in meeting rooms (e.g., connecting a laptop to a screen using a VGA cable) are indeed profoundly mundane; the patterns of interconnection between devices we coded could easily be categorized as boring. Yet it is exactly these mundane activities, prevalent yet ignored, that many ubiquitous computing systems could be best suited to support. This, in turn, requires alternative methodological and design approaches.

1. Longitudinal observations highlight background activities. Mundane practices pose several methodological difficulties. As they are diffuse and often pushed to the background, they require long and repeated observation in order to be uncovered. Traditional interview techniques and laboratory studies, for instance, cannot foreground these activities that, for the most part, are not attended to by their participants. Instead, longitudinal observations are required. But these cannot be entirely automated, as only careful qualitative analysis of the data will progressively reveal these widespread tasks unconsciously carried out by the users. While we have described an approach to mitigate the problem by blending automatic data collection, interviews, and qualitative data coding, it remains that longitudinal studies are inherently costly in terms of manpower. The observe/design/evaluate cycle that is foundational to HCI research might be significantly longer when designing for the support of mundane activities.

2. Alternative means of assessing ''improvement'' are needed. By definition, mundane practices have reached a point where they are simple or integrated enough that they disappear from consciousness. While they can sometimes come to the foreground in the case of novice users or outliers [19], the main user population simply forgets about their cost and implications. As such, it is extremely difficult to justify alternative ways of carrying out mundane tasks using traditional success metrics. The time required to complete a task, for instance, may very well not be significantly lower in ubicomp-supported scenarios such as ours than in the standard case, since the user cost for the latter is already extremely low. However other, less obvious, benefits may accrue to the users and other members of the users' environment. For instance, *the beneficiaries of these ubicomp systems might not be their direct users*. In the case of the Display Mirror, displaying information on a projector is not only useful to the laptop's user: the other meeting participants are also affected. As we learned, the audience was able to articulate benefits of the Display Mirror that we had not anticipated: now that public displays' users were unwired, meeting participants described a less ''chaotic'', more ''fluid'' meeting experience [32]. Furthermore, *benefits to the direct users may go beyond human-computer interaction*. Users of the Display Mirror identified increased mobility within the room as a benefit. By being untethered and therefore able to use the space in the meeting rooms more flexibly (e.g., not to be forced to seat in the ''presenter's spot''), their experience of the meeting was improved. Space is a very important social resource and symbol, used to signify status and roles [12]. ''Untethered'' computing gives control of the space back the users. While this has nothing to do with human-computer interaction per se, it is certainly a benefit. Ubicomp researchers need to consider the global benefits of their system, beyond the confines of a single user interacting with a machine, or even a whole network of machines.

3. Sustained adoption of mundane technology is especially challenging. It has often been proposed that a mark of success for ubicomp systems is when they ''blend in'' their environment. However, when activities are already ''blended in,'' it creates a significant challenge with regards to driving adoption of a new technology.

While we initially tried as best we could not to be disruptive when deploying our infrastructure, it quickly became clear that this was not likely to bear fruit in a reasonable time frame. Mundane activities are deeply entrenched. If new technology to support them remains invisible, it simply won't be adopted. Therefore, to drive the adoption of our system, ''infrastructural inversion'' [4] was necessary: we had to temporarily foreground the backstage elements of our users' work practices, for instance by attending meeting ourselves and repeatedly pointing to our use of the new technology. Without such insistence and somewhat ''heavy handed'' behavior, nobody would have known that a new technology was in use—after all, the end result was no different from earlier meetings (some information appeared on a public display). While this shows that the Display Mirror was truly transparent, it obviously did not favor its ultimate adoption.

Therefore, unlike naturalistic ethnographies of systems' deployments, ubicomp researchers may need to be forceful and directly intervene in order to show people the new possibilities when dealing with mundane activities. Systems such as the Display Mirror deal with the ''taken for granted'' part of computing systems, and human practices in this area are extremely inertial. This is unlike entirely new devices and/or applications that support new (foreground) tasks, whose novelty or strangeness makes them inherently visible. A paradox of mundane ubicomp systems design and evaluation, in our experience, is that one needs to be simultaneously forceful and gentle, that is, to highlight new ways (intervention) of interacting that will end up being as mundane as what they are meant to replace.

An application or system designed to support mundane tasks may face additional challenges when compared with systems that are designed to be used in a focused way to accomplish a set of conscious tasks. This is especially true when the application is a replacement for an existing system that is still available for use. In the case of the Obje Display Mirror, the demands for ease-of-use and robustness were especially high, since the users could easily abandon our system in favor of the previous, familiar, and still available VGA cable. If any difficulties were encountered during use, users would simply abandon the Display Mirror and revert to earlier practices. Once this reversion had taken place, they were reluctant to re-try the Display Mirror without another explicit intervention. In systems such as the Display Mirror, the tasks being supported are, by their nature, in the background and not the focus of conscious attention, and as a result users are extremely intolerant of any glitch that forces them to pay attention to an otherwise unconscious activity.

4.2 Ubiquitous computing in the real world

Some of the lessons we learned in this project would apply equally well to other types of information technology, but some appear to be if not unique to, at least more pronounced in, ubiquitous computing environments.

Longitudinal studies of users' practices before and after the adoption of a new technology are almost always considered desirable, but generally viewed to be expensive and difficult. A common practice in user-centered design is to conduct small, iterative evaluations with progressively refined prototypes. This allows designers to identify critical shortcomings in functionality that will prevent users from accomplishing tasks that the system is being designed to support. However, given the background nature of the tasks we sought to support in the Obje Display Mirror, which we believe are of a similar type to tasks supported by other ubicomp systems, such design and evaluation methods are not likely to yield useful results. Much more so than in conventional desktop or client-server applications, ubicomp applications are explicitly designed to be absorbed into everyday practice, and so the design and evaluation methods need to be more closely interwoven into the background of the practices and environments being supported. In other words, longitudinal methods that are desirable in desktop systems become essential in ubicomp.

Following on the above point, the inadequacy of standard usability metrics such as task performance time and error rate for characterizing the acceptability of a software system has been acknowledged in both the HCI community (e.g., [3]) and ubicomp community (e.g., [2, 6]). However, this point is worth making again, as in cases where the tasks being supported are not the explicit focus of attention, usability metrics and methods are even less adequate as they provide little information about how the system being designed will affect the larger context in which it will be used. Another way to say this is that in ubicomp, the larger context of use may be even more important than in other domains, and so the design and evaluation methods used must be selected or modified to take larger context into work.

The difficulty of getting users to adopt a new technology, especially when a viable and established alternative exists, is well known in multiple domains (e.g., [11]). The issues in ubicomp are probably not much different. One slight but important difference, however, comes back to the issue that the tasks being supported may not be conscious or explicit. In this case, the necessity to foreground the task in order to force the user to become aware of a new technology, and to choose between the existing and new methods of accomplishing the task, may cause an unwelcome interruption in the flow of some other activity. This implies that the methods used to foster adoption need to be as unobtrusive as possible and need to fit in well with existing practice. In our project, we attempted to do this by attaching advertisements to the devices whose functionality we were attempting to replace or enhance. Even this proved to be too subtle, however, and we had to more forceful methods—namely social pressure. On the other hand, the fact that our system was designed for public use was helpful from the perspective of disseminating information about its availability, usefulness, and usage instructions. Another approach that has been taken to increasing adoption of ubicomp technologies is reported in [20], in which the authors introduced a number of different applications aimed at increasing the adoption of an enabling technology—in this case, wireless location tracking badges. However, we did not wish to introduce a number of different applications that used the Display Mirror, because this likely would have caused Display Mirror usage to become more of an explicit goal rather than a means to an existing work-oriented goal.

Creating and maintaining robust, available, easy-to-use services is always challenging, but certain factors in ubicomp may make this even more challenging. For one thing, the expectation of reliability may be higher, since in at least some cases the practices being replaced are ones being carried out with more reliable tools (e.g., analog or physical). For another thing, ubicomp systems often stress standard assumptions about device, system, and application configurations. They may, for example, rely on the coordination of multiple distributed processes, and may depend on assumptions about network topologies and security policies that are sufficiently different from existing application models (e.g., client-server or peer-to-peer) that standard configurations will not support them. In our case, we had two such considerations that interfered with the functioning of our application—one was the fact that we relied on multicast (mDNS) for the Display Mirror clients to discover available display services and the other was that we depended on being able to open sockets on otherwise unused ports in order to stream the video data from the laptop to the display service. In the case of multicast, we discovered that our application employed a pattern that had not been encountered by our Networking Support department—namely the ability to pass multicast traffic back and forth between our wired and wireless networks—and we had to raise the issue through several levels of management in order to get the wireless bridges configured properly to support our application. Regarding the need to open unused ports, it is our company's policy to install and configure a personal firewall on all laptops. Some of these firewalls silently block traffic on all ports unless the user explicitly enables it using an advanced control panel. Other firewalls would allow traffic on these ports but only after the user had agreed to a series of rather cryptic dialog boxes. In the common case where the firewall was configured incorrectly to allow our application to run, the application would fail silently and the user would assume that our service was down.

There are certainly other classes of applications where mysterious and seemingly unrelated system settings can impede proper functioning of the application, but it seems that for the time being ubicomp applications may run up against these problems more frequently than most—at least until the coordination patterns and configurations required to support them become commonly understood.

4.3 Challenge: micro-applications to support micro-tasks

We believe that our experiences can inform a range of existing and forthcoming ubicomp applications that focus on support for infrequent, mundane tasks, especially those that form constituent parts of a variety of foreground tasks. In our observations, for example, connecting a laptop to a projector was a subtask of a variety of larger activities, such as "giving a presentation," "scribing notes," and "sharing information" (e.g., by showing a web page to a group).

It has been previously noted that task-oriented application design and evaluation may not be appropriate for many ubicomp scenarios, because ubiquitous computing is at its best when it fades into the routines of life—that is, when it supports ongoing activities through continuously available, yet sporadically accessed services [2]. Many of the mundane activities we observed do not fit well into the standard HCI notion of a "task." Still, they are activities that users perform, and perform quite often, and which—despite the fact that they are not foreground activities—still to some degree determine the experience that users have in a space. We call these activities *micro-tasks*, because they are rarely explicitly attended to by those who perform them, are generally short-lived, and are typically merely a step in the process of accomplishing some larger tasks (such as giving a practice talk for colleagues). With the Display Mirror, we have taken a look at the micro-task of appropriating a public display in order to share information with others in a meeting room, and we have designed a single, small, limited-functionality service to improve the experience of carrying out that micro-task. Correspondingly, we propose that a service like the Display Mirror should be thought of as a *micro-application*.

While the term "application" can take on many meanings at many granularities, it has frequently been used by technologists and researchers to mean a substantial, user-visible collection of functionality, along with some user interface, designed to allow the user to accomplish some set of related tasks during a focused interaction. Such a meaning does not seem to be applicable to a tool designed to support a single, simple function that is not a part of a consciously attended action on the part of the user. Shifting from applications to micro-applications, particularly when we imagine environments in which myriad micro-applications co-exist, converge, and even compete for the users' attention, causes us to rethink much of what we know about application design, deployment, and evaluation.

While the "micro-application" terminology is our own, we believe that similar conceptualizations exist throughout the research literature. Abowd and Mynatt [1], for instance, describe the "informal and unstructured activities typical of much of our everyday lives," and distinguish these from the dialog styles that are normally a focus of HCI research. Others [9] have argued that small, lightweight applications are an appropriate unit of analysis in ubicomp settings. From purely an evaluative standpoint, separating functionality into such micro-apps may make it easier to tease apart the effects of the application and the role of the infrastructure, by limiting confounding factors.

Our experiences with the Obje Display Mirror provide some insight into the methods and approaches that will be effective for the design and evaluation of micro-applications in ubiciomp environments, but we believe that there is a good deal of further investigation that can be done in this area.

5 Conclusion

In this paper we have presented an approach for investigating and intervening technologically into mundane practices. Methodologically, we have argued that designing for mundane practices requires longitudinal observations, prior, during, and after deployment. During deployment ubicomp designers should expect to intervene gently in the space they are trying to support in order to temporarily foreground activities that began, and will hopefully end, as background activities. The need for intervention stems from both the background nature of the task being supported and the inertia of practices surrounding mundane activities.

We believe that our experiences with the Display Mirror point towards the possibility of experiencing the promise of ubiquitous computing on a small scale. By diffusing technology in the environment and supporting mundane activities, we have seen how practice can be modified in subtle but significant ways. We have also seen how "calm" and "natural" experiences can be delivered through the enhancement of mundane activities. Perhaps it is through the accumulation of small, incomplete benefits through the proliferation of micro-applications, rather than through the sudden, disruptive introduction of a single revolutionary technology, that the eventual promise of ubiquitous computing will eventually be fulfilled.

Acknowledgments This research was supported in part by the NIST Advanced Technology Program Award #3H3052. Kurt Partridge, Bo Begole, and Ame Elliott provided helpful comments on drafts of this paper. Beki Grinter provided helpful guidance during the formative stages of this research.

References

1. Abowd GD, Mynatt ED (2000) Charting past, present, and future research in ubiquitous computing. ACM Trans Comput Hum Interact 7(1):29–58
2. Abowd GD, Mynatt ED, Rodden T (2002) The human experience. IEEE Pervasive Comput Mobile Ubiquitous Syst 1:48–57
3. Beyer H, Holtzblatt K (1998) Contextual design: defining customer-centered systems. Morgan Kaufmann, San Francisco
4. Bowker G (1994) Information mythology and infrastructure. In: Bud-Frierman (ed) Information acumen: the understanding and use of knowledge in modern business. Routledge, London , pp 231–247
5. Colin D (2003) NEC's Wireless MT1065. http://www.projectorcentral.com/wireless_nec_mt1065.htm Accessed: 2005
6. Consolvo S, Arnstein L, Franza BR (2002) User study techniques in the design and evaluation of a ubicomp environment. In: Proceedinghs of the Ubicomp 2002. Goteborg, Sweden, pp 73–90
7. Consolvo S, Walker M (2003) Using the experience sampling method to evaluate ubicomp applications. IEEE Pervasive Comput Mobile Ubiquitous Syst 2:24–31
8. Csikszentmihalyi M, Larson R (1987) Validity and reliability of the experience-sampling method. J Nervous Mental Dis 175(9):526–536
9. Edwards WK, Bellotti V, Dey AK, Newman MW (2003) Stuck in the middle: the challenges of user-centered design and evaluation for infrastructure. In: Proceedings of the CHI 2003. Fort Lauderdale, FL USA, pp 297–304
10. Garfinkel H (1967) Studies in ethnomethodology. Prentice-Hall, Englewood Cliffs
11. Grudin J (1989) Why groupware applications fail: problems in design and evaluation. Office Technol People 4(3):245–264
12. Hall ET (1966) The hidden dimension. Anchor Books, New York
13. Hudson SE, Fogarty J, Atkeson CG, Avrahami D, Forlizzi J, Kiesler S, Lee JC, Yang J (2003) Predicting human interruptibility with sensors: a wizard of oz feasibility study. In: Proceedings of the CHI 2003. Fort Lauderdale, FL USA pp. 257–264
14. Intille SS, Rondoni J, Kukla C, Anacona I, Bao L (2003) A context-aware experience sampling tool. In: Proceedings of the CHI 2003: Extended Abstracts. ACM Press Fort Lauderdale, FL USA, pp 972–973

15. Johanson, B, Fox A, Winograd T (2002) The interactive workspaces project: experiences with ubiquitous computing rooms. IEEE Pervasive Comput Mobile Ubiquitous Syst 1
16. Landay J, Davis RC (1999) Making sharing pervasive: ubiquitous computing for shared note taking. IBM Syst J 38(4):531–550
17. Lynch M, Livingstone E, Garfield E (1983) Temporal order in laboratory work. In: Knoll-Centina, Mulkay (eds) Science observed. Sage, London
18. Mac News Network: WiJET.Video: 802.11g wireless display adapter. http://www.macnn.com/articles/04/01/08/wijet.video.802.11g/ Accessed: 2005 (2004)
19. Mainwaring S, Chang MF, Anderson K (2004) Infrastructures and their discontents: implications for ubicomp. In: Proceedings of the Ubicomp 2004. Springer, Nottingham, pp 418–432
20. Mansley K, Beresford A, Scott D (2004) The carrot approach: encouraging use of location systems. In: Proceedings of the Ubicomp 2004 Nottingham, UK, pp 366–383
21. Microsoft Corporation (2005) working remotely with Windows XP. http://www.microsoft.com/windowsxp/using/mobility/default.mspx Accessed: 2005
22. Mynatt ED everyday computing lab (2004) http://www.cc.gatech.edu/fce/ecl/
23. Nunamaker JF, Dennis AR, Valencich JS (1991) Electronic meeting systems to support group work. Commun ACM 34(7):40–61
24. RealVNC (2002) http://www.realvnc.com Accessed: 2004
25. Scheier RW, Gettys J (1986) The x window system. ACM Transactions on Graphics, 2
26. Star SL (1999) The ethnography of infrastructure. Am Behav Sci 43(3):377–391
27. Streitz NA, Geißler J, Holmer T, Konomi Si, Müller-Tomfelde C, Reischl W, Rexroth P, Seitz P, Steinmetz R (1999) i-LAND: an interactive landscape for creativity and innovation. In: Proceedings of the CHI '99. Pittsburgh, Pennsylvania, USA, pp 120–127
28. Suchman L (1987) Plans and situated actions. Cambridge University Press, Cambridge
29. Tolmie P, Pycock J, Diggins T, MacLean A, Karsteny A (2002) Unremarkable Computing. In: Proceedings of the CHI '02. ACM Press, Minneapolis, MN, USA, pp 399–406
30. Turoff M (1991) Computer-mediated communication requirements for group support. J Organ Comput 1(1):85–113
31. Weiser M (1994) The world is not a desktop. Interactions 1:7–8
32. Weiser M, Brown JS (1996) Designing calm technology. PowerGrid J 1.01(1)
33. Wisneski C, Ishii H, Dahley A, Gorbett M, Brave S, Ullmer B, Yarin P (1998) Ambient displays: turning architectural space into an interface between people and digital information. In: Proceedings of the first international workshop on cooperative buildings (CoBuild 98). Springer-Verlag, Darmstadt, Germany, pp 22–32