The Human Experience

To address Weiser's human-centered vision of ubiquitous computing, the authors focus on physical interaction, general application features, and theories of design and evaluation for this new mode of human–computer interaction.

Gregory D. Abowd and Elizabeth D. Mynatt
Georgia Institute of Technology

Tom Rodden
University of Nottingham

Weiser originated the term *ubiquitous computing*, creating a vision of people and environments augmented with computational resources that provide information and services when and where desired. Although this vision has excited many technologists, we must realize that the main motivation behind Weiser's vision is the impact ubicomp could have on the human experience: "Machines that fit the human environment instead of forcing humans to enter theirs will make using a computer as refreshing as a walk in the woods."[1] For the past decade, researchers have worked toward the implicit goal of assisting everyday life and not overwhelming it. Although the names applied to their research efforts vary (pervasive, wearable, augmented, invisible, disappearing, calm, and so forth), almost all share the goal that Weiser so eloquently characterized:

> Inspired by the social scientists, philosophers, and anthropologists at Parc, we have been trying to take a radical look at what computing and networking ought to be like. We believe that people live through their practices and tacit knowledge so that the most powerful things are those that are effectively invisible in use. This is a challenge that affects all of computer science. Our preliminary approach: Activate the world. Provide hundreds of wireless computing devices per person per office, of all scales (from 1-inch displays to wall-sized). This has required new work in operating systems, user interfaces, networks, wireless, displays, and many other areas. We call our work 'ubiquitous computing.' This is different from PDAs, dynabooks, or information at your fingertips. It is invisible, everywhere computing that does not live on a personal device of any sort, but is in the woodwork everywhere.[2]

To realize Weiser's vision, we must address several clear goals. First, the everyday practices of people must be understood and supported. Second, the world must be augmented through the provisioning of heterogeneous devices offering different forms of interactive experience. Finally, the networked devices must be orchestrated to provide for a holistic user experience. Here, we overview how these goals have affected research in three areas: the definition of the appropriate physical interaction experience, the discovery of general application features, and the evolution of theories for designing and evaluating the human experience in ubicomp.

Defining the appropriate physical interaction experience

Ubiquitous computing inspires application development that is "off the desktop." In addition to suggesting a freedom from well-defined interaction locales (such as the desktop), this vision assumes that physical interaction between humans and computation will be less like the current keyboard, mouse, and display paradigm and more like the way humans interact with the physical world. We speak, gesture, and write to communicate with other humans and alter physical artifacts. The drive for the correct ubicomp experience has resulted in a vari-

1536-1268/02/$17.00 © 2002 IEEE

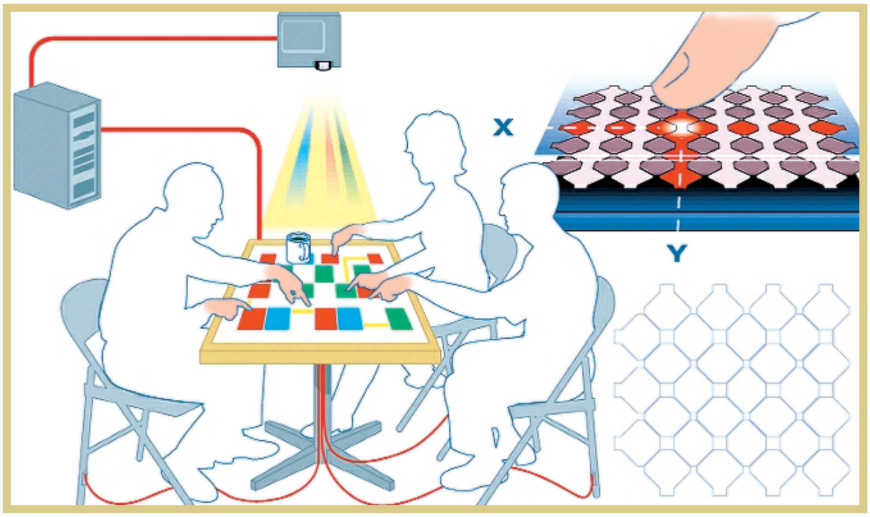

Figure 1. The DiamondTouch input technology from Mitsubishi Electric Research Lab uses capacitive coupling through humans to provide a large-scale input surface for multiple simultaneous users (see www.merl.com/projects/DiamondTouch for more details). Figure courtesy of MERL.

ety of important changes to the input, output, and interactions that define the human experience with computing.

We have traditionally treated input as explicit communication. However, the advance of sensing and recognition technologies has challenged us to provide more humanlike communications capabilities and effectively incorporate implicit actions into the subset of meaningful system input. Similarly, the communication from the environment to the user—the output—has become highly distributed and available in many form factors and modalities. The challenge is to coordinate across many output locations and modalities without overwhelming our limited attention spans. Finally, the relationship between input and output is important in ubicomp, because technology's invisible nature can be seen as a smooth integration between the physical and virtual worlds.

Toward implicit input

Input has moved beyond the explicit nature of textual input (from keyboards) and selection (from pointing devices) to a greater variety of data types. This has resulted in not only a greater variety of input technologies but also a shift from explicit means of human input to more implicit forms of input. In other words, our natural interactions with the physical environment provide sufficient input to a variety of attendant services, without any further user intervention. For example, walking into a space is enough to announce your presence and identity in that location.

Computer interfaces that support more natural human forms of communication (such as handwriting, speech, and gestures) are beginning to supplement or replace elements of the graphical user interface interaction paradigm. Long-standing research communities in computer vision and multimodal recognition technologies (mainly handwriting and speech) are driving the emerging area of perceptual interfaces. Pen-based interaction, unsuccessfully rushed to the market in the early 1990s, is also experiencing a resurgence. Large-scale touch-interactive surfaces, using technologies such as capacitive coupling, have made it possible to create multiperson interactive surfaces on tables and walls (see Figure 1). Recognition of freehand writing is improving, but more significantly, mass adoption has followed the introduction of less sophisticated and more robust recognition technologies, such as Grafiti (Parc's handwriting recognition system). We have even seen compelling examples of voice and pen input effectively used in applications without requiring any recognition.[3]

These recognition technologies exemplify how computers can interpret meaning from sensed signals of human activity. There are many other ways to infer information about people and environments by sensing a variety of other physical signals ("Connecting the Physical World with Pervasive Networks" in this issue directly addresses the advances in sensing of the physical world). However, sensing and interpreting human activity provides a more implicit notion of input to an interactive system. For example, many researchers have investigated[4,5] how we can incorporate simple sensors such as radio frequency identification, accelerometers, tilt sensors, capacitive coupling, and IR range finders into artifacts to increase the language a user can provide as input to control that artifact (see Figure 2).

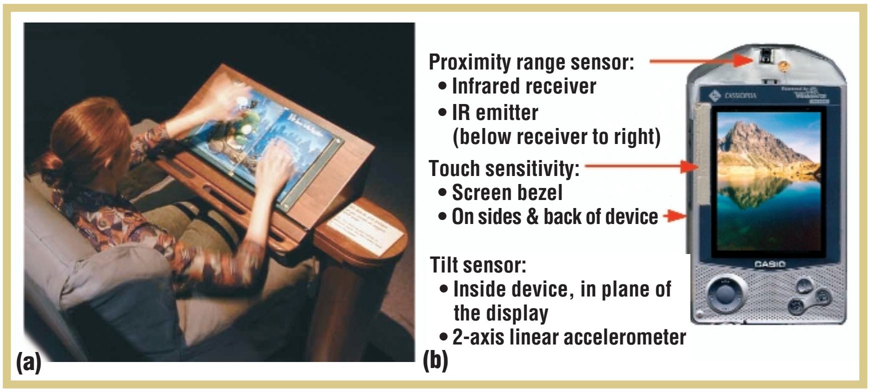

Figure 2. Two examples of simple sensing embedded into devices: (a) The Listen Reader from Xerox Parc uses electric field sensors located in the book binding to sense the proximity of the reader's hands and to control audio parameters.[4] RFID (radio frequency identification) tags embedded in each page allow fast, robust page identification. Picture courtesy of Xerox PARC. (b) An experimental PDA platform used at Microsoft Research to investigate how a variety of simple sensors can improve the interaction between a user and various handheld applications.[5] Figure courtesy of Ken Hinckley.

Figure 3. Form factors of ubiquitous computing. (a) The Stanford Interactive Mural—an example of a large-scale interactive display surface created by tiling multiple lower-resolution projectors. Figure courtesy of Françoise Guimbetrière. (b) The Water Lamp from Hiroshi Ishii's Tangible Media Group at the MIT Media Lab—an example of an ambient display. Light shines upward through a pan of water, which is actuated by digitally controlled solenoids that can tap the water and cause ripples. External information can drive the tapping of the solenoids. Figure courtesy of Hiroshi Ishii.

Invisibility of computing, from the human perspective, can start when we can determine an individual's identity, location, effect, or activity through his or her mere presence and natural interactions in an environment. The union of explicit and implicit input defines the context of interaction between the human and the environment.

Toward multiscale and distributed output

Integrating ubicomp capabilities into everyday life also requires novel output technologies and techniques. To start, the design of targeted information appliances, such as PDAs and future home technologies, requires addressing the technology's form, including its aesthetic appeal. Output is no longer exclusively in the form of self-contained desktop or laptop visual displays that demand our attention. A variety of scales of visual displays are being distributed throughout our environments. For example, many of us carry cell phones with small screens, and large electronic whiteboards—similar to the Liveboard that Weiser's group invented at Parc—are common in offices and classrooms today. More importantly, we are seeing multiple modalities of information sources that lie more at the periphery of our senses and provide qualitative, ambient forms of communication.

Weiser described the form factor of ubicomp technology in three scales —the inch, foot, and yard. The middle (foot) scale is similar to the standard laptop and desktop displays. We each have at least one of these devices, and we largely use them in stationary settings. A new generation of tablet-like portable pen-based computers—devices that rival the experimental Pad prototypes developed at Xerox Parc—is expected to hit the market this year. Pagers, cellular phones, and PDAs that incorporate handheld displays with relatively low resolution currently represent the small end of the scale (inch). We carry around an increasing number of these display devices at all times. High-resolution wall-sized displays now represent the large end of the scale (yard). Such displays are created by effectively stitching together multiple low-resolution projected displays, such as the Stanford Interactive Mural[6] (see Figure 3a) or the Princeton Display Wall.[7]

As these displays continue to proliferate in number and variety, two important trends have emerged. First, as Jun Rekimoto's "pick and drop" demonstration[8] initially motivated and the Stanford Interactive Room[9] further explored, we want to easily move information between separate displays and coordinate the interactions between multiple displays. Second, we want displays that are less demanding of our attention. We'll achieve Weiserian invisibility by designing output that provides for peripheral awareness of information out of the foreground of our conscious attention.

Researchers have explored the trend toward peripheral output for a particular class of displays—*ambient displays*. Such displays require minimal attention and cognitive effort and are thus more easily integrated into a persistent physical space. The artist Natalie Jeremijenko at Xerox Parc invented one of the first ambient displays, the Dangling String.[10] Using analog sensing of network traffic from the cabling in the ceiling, a motor drove the spin of a long string—and the more traffic, the faster the rotation. During high traffic periods, the whir of the string was faintly audible as well.

The Dangling String shares many features with subsequent efforts in ambient displays. A data source drives the abstract representation such that the user's peripheral perception can monitor the output. The data source is generally information of medium to low priority, but it is beneficial for the user to be aware of it, perhaps for some opportunistic action. Because these displays are meant to be persistently available in the environment, they are often designed to be aesthetically appealing and novel (see, for example, the Water Lamp in Figure 3b). Other examples of ambient displays include ambientROOM,[11] which projects information about colleagues as pinpoints of light on the wall; Audio Aura,[12] which encodes the arrival of incoming email as auditory cues in a mobile device; and Kandinsky,[13] which assembles images triggered from keywords in information bulletins into an aesthetically pleasing and intriguing collage.

Although using the visual channel dominates our experience with computing out-

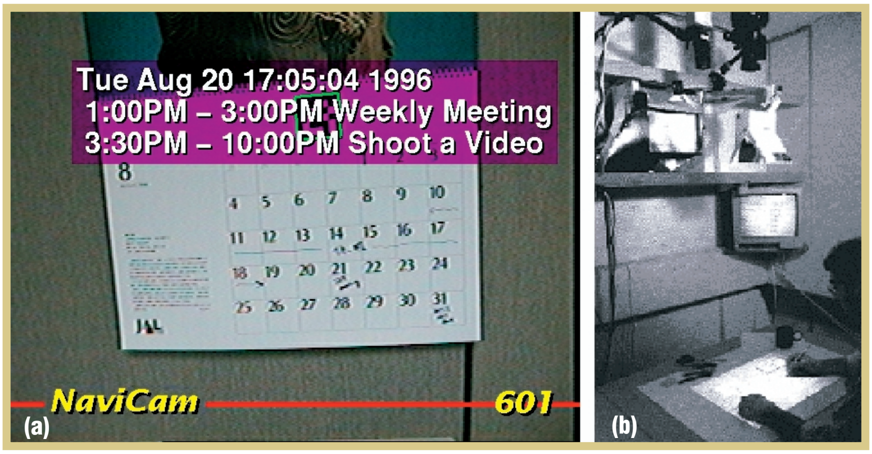

Figure 4. Augmented reality: (a) The NaviCam system, which recognizes 2D glyphs on objects and then superimposes additional information over that object.[16] Figure courtesy of Sony Computer Science Laboratories. (b) DigitalDesk prototype integrates physical and virtual desktop environments with the aid of projection and vision technology.[17]

put, examples such as Audio Aura demonstrate how other modes of output, such as actuation of small devices, can effectively communicate ambient information. With the introduction of simple programming tools for dealing with motors and other actuators, such as Phidgets,[14] mechanical actuation to drive distributed output devices will increase.

Seamless integration of physical and virtual worlds

An important feature of ubicomp technology is that it attempts to merge computational artifacts smoothly with the world of physical artifacts. Plenty of examples demonstrate how we can overlay electronic information on the real world, thus producing an augmented reality.[15] An example of such augmented reality is NaviCam,[16] a portable camera and TV that recognizes 2D glyphs placed on objects and can then superimpose relevant information over the object for display on a TV screen (see Figure 4a). This form of augmented reality only affects the output. When input and output are intermixed, as with the DigitalDesk[17] (see Figure 4b), we begin to approach the seamless integration of the physical and virtual worlds. Researchers have suggested techniques for using objects in the physical world to manipulate electronic artifacts, creating so-called *graspable*[18] or *tangible*[19] user interfaces. Sensors attached to devices provide ways to physically manipulate those devices to be interpreted appropriately by the applications they run (see Figure 2).[5,20]

Application themes

> Applications are of course the whole point of ubiquitous computing. —*Mark Weiser*[21]

Many applications-focused researchers in human–computer interaction (HCI) seek the Holy Grail of ubicomp, the killer application that will cause significant investment in the infrastructure and then enable a wide variety of ubicomp applications to flourish. It could be argued that person-to-person communication is such a killer app for ubicomp, because it has caused a large investment in environmental and personal infrastructure, moving us closer to a completely connected existence. Regardless of whether personal communication is the killer app, the vision of ubicomp from the human perspective is much more holistic. It is not the value of any single service that will make computing a disappearing technology. Rather, it is the combination of a large range of services, all of which are available when and as needed, and all of which work as desired without extraordinary human intervention. A major challenge for applications research is discovering an evolutionary path toward this idyllic interactive experience.

The brief history of ubicomp demonstrates three emergent features that appear across many applications. First, we must be able to use implicitly sensed context from the physical and electronic environment to determine a given service's correct behavior. Context-aware computing demonstrates promise for making our interactions with services more seamless and less distracting from our everyday activities. Applications can work well when properly informed about the context of their use. Second, we must provision automated services to easily capture and store memories of live experiences and serve them up for later use. Finally, we need continuously available services. As we move toward the infusion of ubicomp into our everyday lives, the services provided will need to become constantly available partners with the human users, always interrupted and easily resumed.

Context-aware computing

Two compelling early ubicomp demonstrations were the Olivetti Research Lab's Active Badge[22] and the Xerox ParcTab,[23] both location-aware appliances. These devices leverage a simple piece of context—user location—and provide valuable services (automatic call forwarding for a phone system and automatically updated maps of user locations in an office). These simple location-aware appliances are perhaps the first demonstration of linking implicit human activity with computational services that serve to augment general human activity.

Locating identifiable entities (usually people) is a common piece of context used in ubicomp application development. The most widespread applications have been GPS-based car navigation systems and handheld tour guide systems that vary the content displayed (video or audio) on a handheld unit, given the user's physical location in an exhibit area.[24,25] The Sentient Computing Project (see www.uk.research.att.com/spirit) demonstrates the most complex set of location-aware applications and provides the most complex indoor location-aware infrastructure and applications development (see Figure 5).

Of course, there is more to context than position (*where*) and identity (*who*). Although a complete definition of context remains an illusive research challenge, it is clear that in addition to who and where, context-awareness involves *when*, *what*, and *why*.

With the exception of using time as an index into a captured record or summarizing how long a person has been at a particular location, most context-driven applications are unaware of the passage of time. Of particular interest are the relative changes in time as an aid for interpreting human activity. For example, brief visits at an exhibit could indicate a general lack of interest. Additionally, when we can establish a baseline of behavior, action that violates a perceived pattern would be of particular interest. For example, a context-aware home might notice when an elderly person deviates from a typically active morning routine.

The interaction in current systems either assumes what the user is doing or leaves the question open. Perceiving and interpreting human activity is a difficult problem, but interaction with continuously worn, context-driven devices will likely need to incorporate interpretations of human activity to provide useful information. One strategy is to incorporate information about what a user is doing in the virtual realm. What application is he using? What information is he accessing? One example that has both positive and negative uses is cookies, which describe people's activity on the Web. Another way of interpreting the *what* of context is to view it as the focus of attention of one or more people during a live event. Knowing this focus of attention can inform a better capture of the event.

Even more challenging than perceiving what a person is doing is understanding why he is doing it. Sensing other forms of contextual information that could indicate a person's effective state, such as body temperature, heart rate, and galvanic skin response, might be a useful place to start.

Related to the definition of context is the question of how to represent context. This issue is important once we consider separating out an application's context-sensing portion from its context-aware behavior. Without good representations for context, applications developers are left to develop ad hoc and limited schemes for storing and manipulating this key information. The evolution of more sophisticated representations will enable a wider range of capabilities and a true separation of sensing context from the programmable reaction to that context.

An obvious challenge of context-aware computing is making it truly ubiquitous. Having certain context—in particular, positioning information—is useful. However, there are few truly ubiquitous, single-source context services. Positioning is a good example. GPS does not work indoors and is even suspect in some urban regions. Several indoor positioning schemes exist, with differing characteristics in terms of cost, range, granularity, and requirements for tagging, and no single solution is likely to ever meet all requirements.

The solution for obtaining ubiquitous context is to assemble contextual information from a combination of related context services. Such context fusion is similar in intent to the related and well-researched area of sensor fusion. Context fusion must seamlessly hand off sensing responsibility between boundaries of different context services. Negotiation and resolution strategies must integrate information from competing context services when more than one service concurrently provides the same piece of context. This fusion is also required because sensing technologies are not 100 percent reliable or deterministic. Combining measures from multiple sources could increase the confidence value for a particular interpretation. In short, context fusion assists in providing reliable ubiquitous context by combining services in parallel (to offset noise in the signal) and sequentially (to provide greater coverage). For example, we could combine information from speaker identification, face recognition, and gait recognition to improve identification of an individual in a home setting. Additionally, information from different sources might be available and more appropriate at different times.

Automated capture and access

Much of our life in business and academia is spent listening to and recording (more or less accurately) the events that surround us and then trying to remember the important information from those events. There is clear value, and potential danger, in using computational resources to augment the inefficiency of human record taking, especially when there are multiple streams of related information that are virtually impossible to manually capture as a whole. Tools that support automated capture of and access to live experiences can remove the burden of doing something at which humans struggle (such as recording), so they can focus attention on activities at which they succeed (indicating relationships, summarizing, and interpreting).

We define *capture and access* as the task of preserving a record of some live experience that is then reviewed at some point in the future. Vannevar Bush was perhaps the first to write about the benefits of a generalized capture and access system when he introduced the concept of the *memex*.[26] The memex was intended to store the artifacts that we come in contact with in our everyday lives and the associations that we create between them. Over the years, many researchers have worked toward this vision. As a result, many systems have been built to capture and access experiences in classrooms, meetings, and other live experiences.

Xerox Parc explored the earliest work on automated support for meeting capture. Over the past decade, numerous capture applications have been explored in a variety of environments (see the "Classroom 2000" sidebar for description of a classroom capture environment) in support of individuals or groups. A full review of automated capture appears elsewhere.[27]

Toward continuous interaction

Providing continuous interaction moves computing from a localized tool to a constant presence. A new thread of ubicomp research, *everyday computing*, promotes informal and unstructured activities typical of much of our everyday lives. Familiar examples are orchestrating daily routines,

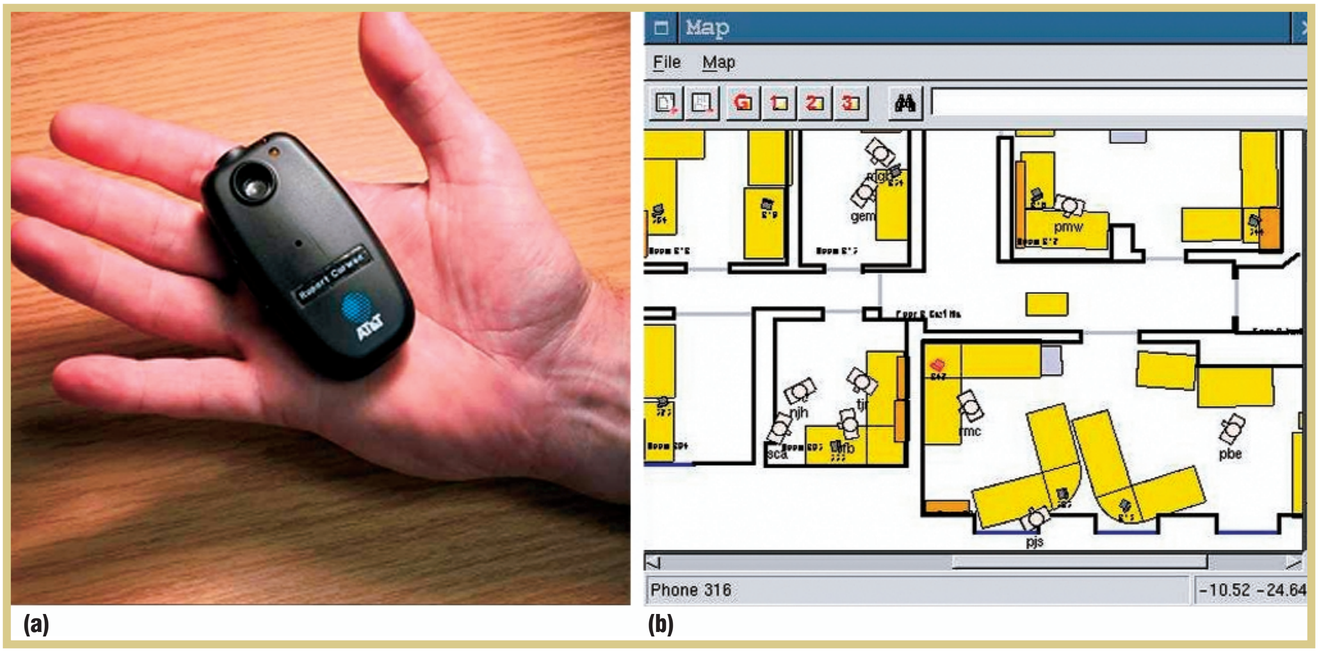

Figure 5. Indoor location systems: (a) AT&T Laboratories Bat device, which is part of a 3D ultrasonic indoor location system, and (b) A map of office worker locations that helps coworkers find each other and talk by phone.

communicating with family and friends, and managing information.

The focus on activities as opposed to tasks is a crucial departure from traditional HCI design. The majority of computer applications support well-defined tasks that have a marked beginning and end with multiple subtasks in between. Take word processing, for example. Word processing features are tuned for starting with a blank document (or template), entering text, formatting, printing, and saving. These applications are not well suited to the more general activity of writing, encompassing multiple versions of documents where text is reused and content evolves over time.

The emphasis on designing for continuously available interaction requires addressing these features of informal, daily activities:

- They rarely have a clear beginning or end, so the design cannot assume a common starting point or closure and thus requires greater flexibility and simplicity.
- Interruption is expected as users switch attention between competing concerns.
- Multiple activities operate concurrently and might need to be loosely coordinated.
- Time is an important discriminator in characterizing the ongoing relationship between people and computers.
- Associative models of information are needed, because information is reused from multiple perspectives.

Of course, activities and tasks are not unrelated to each other. Often an activity will comprise several tasks, but the activity itself is more than these component parts. For example, communication activities contain well-defined tasks such as reading a message or composing a reply. The interaction falters when the task refers to the larger activity: How does this new message relate to previous messages from this person? What other issues should be included in the reply? The challenge in designing for activities is encompassing these tasks in an environment that supports continuous interaction.

Theories of design and evaluation

Because the implementation of Weiser's ubicomp vision alters the relationship between humans and technology, we must revise our theories of HCI that inform design and evaluation. Traditional work in HCI has produced considerable human factors guidance for designing various kinds of computer interfaces (for example, graphical displays, direct manipulation interfaces, multimedia systems, and Web sites). Although widely used, most of these guidelines tend to focus on the needs and demands of designing desktop computer interfaces. During the past few years, ubicomp's emergence has led research communities to ask whether these current approaches are appropriate to the design of interfaces where interaction extends beyond the traditional monitor, keyboard, and mouse arranged on a desk.

A particular concern for ubicomp is developing support for designing and assessing systems that are appropriate when computing functionality becomes embedded in the surrounding environment, specific physical objects, or even objects that are carried. This movement away from the desktop with its well-understood and fixed arrangement of devices has been a catalyst for three broad research activities:

- The development of new models of interaction that incorporate the relationship of ubicomp with the physical world
- The emergence of methods that focus on gaining richer understandings of settings
- The development of approaches to assess ubicomp's utility

New models of interaction

The shift in focus from the desktop to the surrounding environment inherent in ubicomp mirrors previous work in HCI and computer-supported cooperative work. As the computer has increasingly spread throughout organizations, researchers have had to shift their focus from a single machine engaging with an individual to a broader set of organizational and social arrangements and the cooperative interaction inherent in these arrangements. This shift has seen the development of new models of interaction to support the design process in broader organizational settings. Many of these models apply to ubicomp with its emphasis on integrating numerous devices in one setting.

Traditionally, the Model Human Processor theory of human cognition and behavior has informed HCI research and evaluation efforts.[28] This model focuses on internal cognition driven by the cooperation of three independent units of sensory, cognitive, and motor activity, where each unit maintains its own working store of information. As the application of computers has broadened, designers have turned to models that consider the nature of the relationship between the internal cognitive processes and the outside world. Designing for a balance between "knowledge in the world" versus "knowledge in the head" is now a common maxim in the design community.[29,30] The ubicomp community is currently exploring three main models of cognition as guides for future design and evaluation.

Activity theory is the oldest of the three, building on Lev Vygotsky's work.[31] The closest to traditional theories, activity theory recognizes concepts such as goals (objects), actions, and operations. However, both goals and actions are fluid, based on the world's changing physical state instead of more fixed, a priori plans. Additionally, although operations require little to no explicit attention, such as an expert driver driving home, the operation can shift to an action based on changing circumstances such as difficult traffic and weather conditions. Activity theory also emphasizes the transformational properties of artifacts that implicitly carry knowledge and traditions, such as musical instruments, cars, and other tools. The user's behavior is shaped by the capabilities implicit in the tool itself.[32] Ubicomp's efforts informed by activity theory, therefore, focus on the transformational properties of artifacts and the fluid execution of actions and operations.

Situated action emphasizes the improvisational aspects of human behavior and deemphasizes a priori plans that the person simply executes. In this model, knowledge in the world continually shapes the ongoing interpretation and execution of a task. For example, a downhill skier constantly adjusts her behavior given the changing physical terrain, the presence of other people, and the signals from her own body. Any external plan, such as "ski to the bottom of the hill without falling," is vague and driven by the context of the ski resort itself (based on the terrain, snow conditions, other skiers, and so forth).[33] Ubicomp's efforts informed by a situated action also emphasize improvisational behavior and would not require, nor anticipate, the user to follow a predefined script. The system would aim to add knowledge to the world that could effectively assist in shaping the user's action, hence an emphasis on continuously updated peripheral displays. Additionally, evaluating this system would require watching authentic human behavior and would discount post-task interviews in which people attempt to explain their actions as rationalizations of behavior that is not necessarily rationale.

Distributed cognition also de-emphasizes internal human cognition, but in this case, it turns to a systems perspective where humans are just part of a larger system. This theory focuses on the collaborative process, where multiple people use multiple objects to achieve a larger systems goal, such as naval crewmembers using numerous tools to navigate a ship into port.[29] Of all three theories, distributed cognition pays the greatest attention to knowledge in the world because much of the information needed to accomplish a system's goal is encoded in the individual objects. Cognition occurs as people translate this information to achieve one part of the larger task. Ubicomp efforts informed by distributed cognition focus on designing for a larger system goal in contrast to using an individual appliance. These efforts emphasize how information is encoded in objects and how different users translate or transcribe that information.

Gaining a richer understanding of settings

Considerable debate exists in the social sciences about the nature of cognition and the observable world of everyday practices. Lucy Suchman first highlighted the need to gain a rich understanding of the everyday world to inform IT development.[33] In contrast to developing abstract models, many researchers focus on gaining rich understandings of particular settings and conveying these understandings to the design process. Weiser emphasized the importance of understanding these everyday practices to inform ubicomp research: "We believe that *people live through their practices* and tacit knowledge so that the most powerful things are those that are *effectively invisible in use* [emphasis ours]."[2]

The challenge for ubicomp designers is to uncover the very practices through which people live and to make these invisible practices visible and available to the developers of ubicomp environments. Ethnography has emerged as a primary approach to address the need to gain rich understandings of a particular setting and the everyday practices that encompass these settings.

Ethnographic studies have their roots in anthropology and sociology and focus on uncovering everyday practices as they are understood within a particular community. Ethnography relies on an observer going into the field and learning the ropes through questioning, listening, watching, talking, and so forth, with practitioners. The fieldworker's task is to immerse himself into the setting and its activities with a view to describing these as the skillful and socially organized accomplishments of that setting's inhabitants.

In the context of ubicomp, the ethnographic investigation's goal is to provide these descriptions and analysis of everyday life to the IT designers and developers so that ubicomp environments seamlessly mesh with the everyday practices that encapsulate the goals, attitudes, social relationships, knowledge, and language of the

Classroom 2000

An influential case study in deploying and evaluating a ubicomp application is the Classroom 2000 system, developed at Georgia Institute of Technology (see Figure A).[1] The project began in July 1995 with the intent of producing a system that would capture as much of the classroom experience as possible to facilitate later review by both students and teachers. In many lectures, students have their heads down, furiously writing down what they hear and see for future reference. Although some of this writing activity is useful as a processing cue for the student, we wanted to offer students the opportunity to lift their heads occasionally and engage in the lecture experience. The capture system aimed to relieve some of the note-taking burden.

To quickly test this hypothesis's feasibility, we implemented an environment for capture in six months and used it to capture an entire course. We learned some valuable lessons during this first extended experience. The initial experiments included student note-taking devices that clearly distracted the students. We abandoned support for individual student note-taking, only to resume it two years later when the technology caught up.

To understand this capture system's impact on teaching and learning, more students and teachers had to use it in a wider variety of courses. This required a significant engineering effort to create a robust and reliable capture system that, by early 1997, could support multiple classes simultaneously. During a three-year experimental period ending in mid-2000, over 100 courses were supported for 30 different instructors. In what will hopefully serve as a model for longitudinal study of ubicomp systems, Jason Brotherton reports in his thesis on the extensive quantitative analysis that reveals how such an automated capture and access system affects the educational experience once it is incorporated into the everyday experience.[2] As a direct result of these deeper evaluations, we know that the system encourages 60 percent of its users to modify their in-class note-taking behavior. However, not all of this modified behavior is for the better. Taking no notes, for example, is not a good learning practice to reinforce. We need to facilitate more content-based retrieval and synchronized playback of the lecture experience. These insights have motivated further research efforts and established a long-term research project, eClass, that stands as a model for ubicomp research and automated capture and access.

REFERENCES

1. G.D. Abowd, "Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment," *IBM Systems J.*, vol. 38, no. 4, Oct. 1999, pp. 508–530.

2. J.A. Brotherton, *eClass: Building, Observing, and Understanding the Impact of Capture and Access in an Educational Domain*, doctoral dissertation, College of Computing and GVU Center, Georgia Inst. of Technology, 2001; www.jasonbrotherton.com/brothert/thesis/Thesis.pdf.

Figure A. Classroom 2000 (1) in operation, with an extended electronic whiteboard; (2) the automatically generated lecture notes, which include slides presented with teacher annotations and Web pages visited during the lecture, all organized in a timeline presentation allowing random access to streaming audio or video recordings of the lecture.

intended setting. These techniques have been applied to inform design of social communications devices for the home[34] and to enhance the social connection between senior citizens and their extended families.[35]

Perhaps a more intriguing method for conveying the nature of settings to developers of future technologies has emerged from an art and design tradition. Bill Gaver and colleagues have explored the use of cultural probes to collect information from settings to inspire the development of new digital devices.[36] As part of a broader research project, the design group at the Royal College of Art in London is undertaking a cultural probe study of domestic environments in which volunteers are given packages of evocative materials (such as cameras, telephone pads, visitor books, listening glasses, and dream recorders) designed to elicit inspirational data. Unlike ethnography, which focuses on the everyday and routine nature of the setting, cultural probes seek to uncover the emotional, unusual, and even spiritual to inspire designers.

Assessment of use

We have considered models and theories to understand users and some of the methods researchers have used to uncover user needs and desires. However, we must also assess the utility of ubicomp solutions. Researchers have only recently begun to address the development of assessment and evaluation techniques that meet ubicomp's demands. One reason for this relatively slow development is the gradual evolution of ubiquitous technology and applications. To understand ubicomp's impact on everyday life, we navigate a delicate balance between predicting how novel technologies will serve a real human need and observing authentic use and subsequent coevolution of human activities and novel technologies.

Formative and summative evaluation of ubicomp systems is difficult and represents a real challenge for the ubicomp community. With the notable exception of the work at Xerox Parc on the use of the Tivoli capture system and at Georgia Tech with the Classroom 2000/eClass system (see the related sidebar), there has been surprisingly little research published from an evaluation or end-user perspective in the ubicomp community. We must address significant challenges to develop appropriate assessment methods and techniques.

The need for new measures

The shift away from the desktop inherent in the ubicomp vision also represents a shift away from the office and the managed structuring of work inherent within these environments. Much of our understanding of work has developed from Fordist and Taylorist principles on the structuring of activities and tasks (human behavior can be decomposed into structured tasks). Evaluation in HCI reflects these roots and is often predicated on notions of task and the measurement of performance and efficiency in meeting these goals and tasks.

However, it is not clear that these measures can apply universally across activities when we move away from structured and paid work to other activities. For example, it is unclear how we might assess the domestic devices suggested by the Royal College of Art[37] or the broad range of devices to emerge from Philips' vision of the Future.[38] This shift away from the world of work means that there is still the question of how to apply qualitative or quantitative evaluation methods. Answering this question requires researchers to consider new representations of human activity and to consider how to undertake assessment that broadens from existing task-oriented approaches. Although many researchers have investigated the use of observational and semistructured interviews, the lack of deployment of ubiquitous environments has hampered many of these activities.

The technology used to create ubicomp systems is often on the cutting edge; creating reliable and robust systems that support some activity on a continuous basis is difficult. Consequently, a good portion of reported ubicomp applications work remains at the level of demonstrational prototypes that are not designed to be robust. Deeper empirical evaluation results cannot be obtained through controlled studies in traditional, contained usability laboratory. Rather, the requirement is for real use of a system, deployed in an authentic setting.

Many researchers are seeking to roll out ubiquitous devices into a range of settings, such as museums, outdoor city centers, and the home. These researchers are creating "living laboratories" for ubicomp research by creating testbeds that support advanced research and development as well as use by a targeted user community. By pushing on the deployment of more living laboratories for ubicomp research, the science and practice of HCI evaluation will mature. It is our hope that *IEEE Pervasive Computing* will become a forum for the dissemination of information about other living laboratories in the coming years.

As Weiser said,

> The most profound technologies are those that disappear.[1]

And the most prophetic visions are those that are succinct.

REFERENCES

1. M. Weiser, " The Computer of the 21st Century," *Scientific American*, vol. 265, no. 3, Sept. 1991, pp. 66–75 (reprinted in this issue, see pp. 19–25).
2. M. Weiser, "The World Is Not a Desktop," *ACM Interactions*, vol. 1, no. 1, Jan. 1994, pp. 7–8.
3. D. Hindus and C. Schmandt, "Ubiquitous Audio: Capturing Spontaneous Collaboration," *Proc. ACM Conf. Computer Supported Cooperative Work* (CSCW 92), ACM Press, New York, 1992, pp. 210–217.
4. M. Back et al., "Listen Reader: An Electronically Augmented Paper-Based Book," *Proc. 2001 ACM Conf. Human Factors in Computing Systems* (CHI 2001), ACM Press, New York, 2001, pp. 23–29.
5. K. Hinckley et al., "Sensing Techniques for Mobile Interaction," *Proc. 13th Ann. ACM Symp. User Interface Software and Technology* (UIST 2000), ACM Press, New York, 2000, pp. 91–100.
6. G. Humphreys and P. Hanrahan, "A Distributed Graphics System for Large Tiled Displays," *Proc. IEEE Visualization*, IEEE Press, Piscataway, N.J.,1999.
7. Y.D. Chen et al., "Early Experiences and Challenges in Building and Using a Scalable Display Wall System," *IEEE Computer Graphics & Applications*, vol. 20, no. 4, July/Aug. 2000, pp. 671–680.
8. J. Rekimoto, "Pick-and-Drop: A Direct Manipulation Technique for Multiple Computer Environments," *Proc. ACM Symp. User Interface Software and Technology* (UIST 97), ACM Press, New York, 1997, pp. 31–39.
9. A. Fox et al., "Integrating Information Appliances into an Interactive Workspace," *IEEE Computer Graphics & Applications*, vol. 20, no. 3, May/June 2000, pp. 54–65.
10. M. Weiser and J.S. Brown, "Designing Calm Technology," *PowerGrid J.*, vol. 1.01, July 1996; www.ubiq.com/hypertext/weiser/calmtech/calmtech.htm.
11. H. Ishii et al., "ambientROOM: Integrating Ambient Media with Architectural Space," *Proc. 1998 ACM Conf. Human Factors in Computing Systems* (CHI 98), *Companion Proc.*, ACM Press, New York, 1998, pp. 173–174.

12. E.D. Mynatt et al., “Designing Audio Aura,” *Proc. 1998 ACM Conf. Human Factors in Computing Systems* (CHI 98), ACM Press, New York, 1998, pp. 566–573.

13. J. Fogarty, J. Forlizzi, and S.E. Hudson, “Aesthetic Information Collages: Generating Decorative Displays that Contain Information,” *Proc. 14th Ann. ACM Symp. User Interface Software and Technology* (UIST 2001), ACM Press, New York, 2001, pp. 141–150.

14. S. Greenberg and C. Fitchett, “Phidgets: Easy Development of Physical Interfaces through Physical Widgets,” *Proc. 14th Ann. ACM Symp. User Interface Software and Technology* (UIST 2001), ACM Press, New York, 2001, pp. 209–218.

15. S. Feiner, B. MacIntyre, and D. Seligmann, “Knowledge-Based Augmented Reality,” *Comm. ACM,* vol. 36, no. 7, July 1993, pp. 53–62.

16. J. Rekimoto and N. Katashi, “The World through the Computer: Computer Augmented Interaction with Real World Environments,” *Proc. ACM Symp. User Interface Software and Technology* (UIST 95), ACM Press, New York, 1995, pp. 29–36.

17. P. Wellner, “Interacting with Paper on the Digital Desk,” *Comm. ACM* , vol. 36, no. 7, July 1993, pp. 87–96.

18. G.W. Fitzmaurice, H. Ishii, and W. Buxton, “Bricks: Laying the Foundations for Graspable User Interfaces,” *Proc. 1995 ACM Conf. Human Factors in Computing Systems* (CHI 95), ACM Press, New York, 1995, pp. 442–449.

19. H. Ishii and B. Ullmer, “Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms,” *Proc. ACM Conf. Human Factors in Computing Systems* (CHI 97), ACM Press, New York, 1997, pp. 234–241.

20. B.J. Harrison et al., “Squeeze Me, Hold Me, Tilt Me! An Exploration of Manipulative User Interfaces,” *Proc. 1998 ACM Conf. Human Factors in Computing Systems* (CHI 98), ACM Press, New York, 1998, pp. 17–24.

21. M. Weiser, “Some Computer Science Issues in Ubiquitous Computing,” *Comm. ACM*, vol. 36, no. 7, July 1993, pp. 75–84.

22. R. Want et al., “The Active Badge Location System,” *ACM Trans. Information Systems*, vol. 10, no. 1, 1992, pp. 91–102.

23. R. Want et al., *The PARCTab Ubiquitous Computing Experiment*,” tech. report CSL-95-1, Xerox Palo Alto Research Center, 1995.

24. G.D. Abowd et al., “Cyberguide: A Mobile Context-Aware Tour Guide,” *ACM Wireless Networks*, vol. 3, 1997, pp. 421–433.

25. K. Cheverst, K. Mitchell, and N. Davies, “Design of an Object Model for a Context Sensitive Tourist Guide,” *Proc. Interactive Applications of Mobile Computing* (IMC 98), Neuer Hochschulschriftverlag, 1998; www.egd.igd.fhg.de/~imc98/proceedings.html.

26. V. Bush, “As We May Think,” *Atlantic Monthly*, vol. 176, no. 1, July 1945, pp. 101–108.

27. K.N. Truong, G.D. Abowd, and J.A. Brotherton, “Who, What, When, Where, How: Design Issues of Capture and Access Applications,” *Ubicomp2001: Ubiquitous Computing*, Lecture Notes in Computer Science,, vol. 2201, Springer-Verlag, New York, 2001, pp. 209–224.

28. S. Card, T. Moran, and A. Newell, *The Psychology of Human-Computer Interaction*, Lawrence Erlbaum, Mahwah, N.J., 1983.

29. E. Hutchins, *Cognition in the Wild*, MIT Press, Cambridge, Mass., 1995.

30. D.E. Norman, *The Design of Everyday Things*, Doubleday, New York, 1990.

31. L. Vygotsky, “The Instrumental Method in Psychology,” *The Concept of Activity in Soviet Psychology*, J. Wertsch, ed., Sharpe, Armonk, New York, 1981.

32. B. Nardi, ed., *Context and Consciousness: Activity Theory and Human-Computer Interaction*, MIT Press, Cambridge, Mass., 1996.

33. L. Suchman, *Plans and Situated Actions. The Problem of Human-Machine Communication*, Cambridge Univ. Press, Cambridge, 1987.

34. D. Hindus et al., “Casablanca: Designing Social Communications Devices for the Home,” *Proc. 2001 ACM Conf. Human Factors in Computing Systems* (CHI 2001), ACM Press, New York, 2001, pp. 325–332.

35. E.D. Mynatt et al., “Digital Family Portraits: Providing Peace of Mind for Extended Family Members,” *Proc. 2001 ACM Conf. Human Factors in Computing Systems* (CHI 2001), ACM Press, New York, 2001, pp. 333–340.

36. W. Gaver, A. Dunne, and E. Pacenti, “Design: Cultural Probes,” *Interactions: New Visions of Human- Computer Interaction*, ACM Press, New York, 1999.

37. W. Gaver and A. Dunne, “Projected Realities: Conceptual Design for Cultural Effect,” *Proc. 1999 ACM Conf. Human Factors in Computing Systems* (CHI 99), ACM Press, New York, 1999, pp. 600–607.

38. Philips Corporate Design, *Vision of the Future*, 1996; www.design.philips.com/vof.

the AUTHORS

Gregory D. Abowd is an associate professor in the College of Computing and the GVU Center at Georgia Institute of Technology. His research interests involve applications research in ubiquitous computing, concerning both HCI and software engineering issues. He founded the Future Computing Environments Group at Georgia Tech and is currently director of the Aware Home Research Initiative. He received a DPhil in Computation from the University of Oxford, where he attended as a Rhodes Scholar. He is a member of the IEEE Computer Society and ACM. Contact him at abowd@cc.gatech.edu.

Elizabeth D. Mynatt is an associate professor in the College of Computing and the GVU Center at the Georgia Institute of Technology. She directs the research program in Everyday Computing within the Future Computing Environments Group, examining the implications of having computation continuously present in many aspects of everyday life. She previously worked for three years at Xerox Parc with Mark Weiser. Her research explored how to augment everyday places and objects with computational capabilities. Contact her at mynatt@cc.gatech.edu.

Tom Rodden is a professor of computing at the University of Nottingham. He currently codirects the Equator Interdisciplinary Research Collaboration, a multidisciplinary research program exploring the ubiquitous computing across a wide range of application domains (see www.equator.ac.uk). His research interests are in computer-supported cooperative work and the design of interactive systems that are sensitive to the needs of many users. Most recently, this has involved exploring how to combine physical and digital interaction to develop ubiquitous computing environments. Contact him at tar@cs.nott.ac.uk.

For more information on this or any other computing topic, please visit our Digital Library at http://computer.org/publications/dlib.