Seamful and Seamless Design in Ubiquitous Computing

Matthew Chalmers and Ian MacColl

Computing Science, University of Glasgow, United Kingdom

matthew@dcs.gla.ac.uk, ianm@dcs.gla.ac.uk

Abstract: In this paper, we question the assumption that seamless integration of computer system components is necessarily a design requirement for ubiquitous computing. We explore Mark Weiser's notions of seamlessness and 'seamfulness', and apply them in discussing system design and use. The physical nature of ubicomp systems reveals itself in, for example, uncertainty in sensing and ambiguity of representations. These 'seams' may be inevitable, and users perceive and appropriate them for their own uses. Users can benefit from them, and new opportunities for seamful design arise if we take fuller account of them. We offer some examples of seams and initial suggestions for seamful designs, drawing from previous work in ubiquitous computing, mixed reality systems and media spaces, but focusing on our own system that lets a visitor using a PDA in a museum exhibition co-visit with people using virtual reality and web versions of the same exhibition.

Keywords: ubiquitous computing, CSCW, uncertainty, customisation, adaptation.

1 Introduction

In Mark Weiser's vision of ubiquitous computing, one of the broader design goals is invisibility (Weiser 1994a):

> A good tool is an invisible tool. By invisible, I mean that the tool does not intrude on your consciousness; you focus on the task, not the tool.

In other words, one no longer needs to attend to the tool when using it. Ubiquity involves multiple, heterogeneous devices providing highly dispersed input, output and computational capabilities. These parts collectively form a tool for interaction that is "literally visible, effectively invisible" because their design, along with one's experience and understanding of them, lets one focus on interaction *through* the whole instead of *on* the parts—just as a carpenter does carpentry without a constant and conscious focus on his hammer.

In (Weiser 1994b), one example of invisible technology is electricity. There are often 20–30 electric motors in a car, and yet we do not need to attend to each one when driving. Another example of invisible technologies is "literary technology" such as books, wallboards and notepads. The famous 'dangling string' of Natalie Jeremijenko was also simple in technological terms but sophisticated in interaction design terms. Weiser suggests that even a "glass TTY UI can be ubicomp," if its use is well woven into the fabric of people's collaboration and interaction.

However, it appears that this notion of invisibility has been translated into requirements for seamless integration of computer system components, as well as the interactions supported by those components. Some easily found examples, with our emphasis added, include IEEE Pervasive Computing's FAQ—

> Many key building blocks needed for [Weiser's] vision are now viable commercial technologies. The challenge is to combine these technologies into a *seamless* whole.

IBM's Pervasive Computing site—

> We expect devices—personal digital assistants, mobile phones, office PCs and home entertainment systems—to access that information and work together in one *seamless*, integrated system.

and the EU's Disappearing Computer initiative:

> A world full of interacting artefacts could easily confuse people. Research is needed in order to make sure that environments are coherent and understandable. This could include [...] Approaches that ensure "*seamless* interaction", for example, for an activity that takes place across different locations and different stages in time.

Seamlessness is an attractive prospect, extending the ideas of metaphoric direct manipulation to make our interactions with computers more literal, reducing the distractions that such interactions currently introduce. However, Weiser describes seamlessness as a misleading or misguided concept. In his invited talks to UIST94 (Weiser 1994b) and USENIX95 (Weiser 1995) he suggests that making things seamless amounts to making everything the same, and he advocates *seamful* systems (with "beautiful seams") as a goal. Paraphrasing Weiser's talk slides only slightly, and retaining his emphasis: making everything the same is easy; letting everything be *itself*, *with* other things, is hard.

Making everything the same meant, to Weiser, reducing components, tools and systems to their 'lowest common denominator'. Seamlessness could mean sacrificing the richness of each tool in order to obtain bland compatibility. A similar danger arises when one tool is chosen as primary and the others are reduced and simplified so that they conform to it.

Around Xerox PARC, where many researchers worked on document tools, Weiser used an example of seamful integration of a paint tool and a text editor (Weiser, personal communication). He complained that seamless integration of such tools meant that the user was forced to use only one of them, or to crudely patch together components that were of one type or the other i.e. the seams were ugly. Seamfully integrated tools would maintain the unique characteristics of each tool, through transformations that retained their individual characteristics. This would let the user brush some characters with the paint tool in some artful way, then use the text editor to 'search and replace' some of the brushstroked characters, and then paint over the result with colour washes. Interaction would be seamless even though the features of each tool were apparent. Seamful integration is hard, but the quality of interaction can be improved if we let each tool 'be itself'.

Weiser suggests to the system designer that "the unit of design should be social people, in their environment, plus your device". A device that senses, models and lets the user take advantage of the context of 'other things', such as nearby people and the non–digital objects in their environment, is of course well–established within the ubicomp community. We suggest, however, that letting a ubicomp system be itself means accepting all its physical and computational characteristics—that may either be weaknesses or strengths. A user's activity is influenced by what they perceive and understand of sensors, transducers and other I/O devices, as well as the system's internal models and infrastructure.

We offer as an example the now ubiquitous mobile phone. Signal strength is a physical property of the system infrastructure that is sensed and made apparent to phone users in everyday use through the number of bars in the interface, but also through whether they can hear people clearly, make a call, get an Internet connection, and so on. In contrast, users usually do not want or need to know what cell their phone is using, or that their phone has been handed over to another cell. Cell handover remains deep in the infrastructure so that, to the user, cell handover is handled seamlessly. However, mobile phones can be set to display the current cell (if the service provider permits), and some people (including one of the authors of this paper) choose to enable this facility. This is an elegant ambient or peripheral presentation of potentially useful information that is characteristic of the phone as a physical sensor and the phone network as a cell structure. Users can choose whether this information is presented and, if so, what use to make of it. For example, one might seek a stronger signal by moving to a location that forces a perceptible handover to another cell.

The physical characteristics of ubicomp systems are often apparent as uncertainty and inaccuracy. For example, digital sensors such as cameras, GPS receivers and ultrasonic trackers have limited sampling rates and resolution, are prone to communication delays and disconnections, have finite data storage limits and have representational schemes of finite scope and accuracy. As designers, we can be defensive or negative about such seams in devices and infrastructure, and try to design them out. This approach is not always affordable or practical, but we might invest resources to improve phone signal strength and coverage, so that cell handovers can stay hidden in the infrastructure, or buy higher resolution cameras so that images are less coarse. However, we will always have finite resources for sensing, computation and storage. We can let the tool 'be itself' by accepting its characteristics and designing our systems more pragmatically. We can choose a more positive design approach by making seams a resource for users. Note that we do not claim that seamlessness is always bad, or that seamfulness is always good. Rather, they form a continuum or design space—with lots of room for new seamful work.

The following sections of this paper present more examples of seams and seamful design, with the next section dealing especially with uncertainty. We draw from systems we have developed in City, a project within the Equator interdisciplinary research collaboration (www.equator.ac.uk), as well as other Equator projects. We then go on to discuss another aspect of seamfulness related to the way that users 'design' their activity to take advantage of seams and uncertainties, and even appropriate them i.e. use them in ways that the designers may not have intended or imagined. Again, mobile phones offer an example. Variable signal strength has the advantage of giving people a credible excuse for not returning a call, or for ending a call early—sometimes even when the signal is strong. GPS drawing (www.gpsdrawing.com) offers a fascinating and whimsical use of GPS, e.g. drawing an elephant on a map by walking through city streets.

In the section on appropriation, we offer some examples of users of our systems taking account or advantage of uncertainty, and also discuss similar examples in other forms of computer–mediated communication. Since this appears to be an inevitable and useful part of system use, we consider the prospects for a positive design approach: *design for appropriation*. Ubicomp generally involves representing user context and making it a resource for users to use, control and change. However, if we accept that a ubicomp system's devices, interfaces and infrastructure are part of the user's context, then we are faced with the design issue of how to support users' control and change of deep system structure.

2 Uncertainty

In this section we present examples from Equator involving uncertainty, in particular uncertainty in sensor technologies. The City project has developed a prototype system for the Mackintosh Interpretation Centre, an exhibition in Glasgow. Our system allows a visitor using a PDA in the exhibition room (Figure 1, top left) to co-visit with other people using VR and web versions of the same exhibition (MacColl 2002a). Design was informed by field studies of findings in studies of co-visiting in conventional museums (Galani 2002), and allowed visitors to talk over a shared audio channel, share location and interact around 'hybrid' museum exhibits.

In an extensive user trial (Brown 2003), we found that the system effectively supported many, but not all, of the social aspects of a shared visit to a traditional museum. The system was also seamful enough to let the visitors maintain the characteristics of the new media. The PDA user had as a resource the full visual and tactile richness of the exhibition room, which we could not offer the VR and web visitors, but the latter two could use online media unavailable on the PDA. The shared resources that the system provided were effective in allowing users to navigate together, collaborate around objects, share their experiences and talk about the exhibition. The individual resources ensured that each participant had a contribution that only he or she could make.

The PDA includes a sensor package that is part of an ultrasonic location system (Randell 2001), and an electronic compass for orientation information. The location and orientation are displayed on a map of the exhibition room, along with the locations and orientations of the other two visitors (Figure 1, bottom left).

The VR visitor's position was that of his or her avatar in the 3D model. Other visitors' positions were shown using avatars, as in (Figure 1, top right). Each exhibit is modelled in the VR at a coarse level that shows overall form but not fine detail such as textual labels, and each exhibit has a corresponding web page containing text and 2D images from the exhibition catalogue. By approaching an exhibit in the 3D model, the VR visitor triggers display of information about that exhibit in the web browser. The spatial location of the visitor is converted to a name that represents a spatial extent, or zone, in the room. A browser applet responds to a zone change by loading a new HTML page corresponding to main exhibit in the new zone.

The web visitor's position was determined from his or her use of a 2D map of the room, as shown in an applet (Figure 1, bottom right). Mouse clicks were interpreted as movements, with the direction from the old location to the new location treated as the new orientation. Moving to a new location triggered display of web information in much the same way as for the VR visitor, except that a menu of links is offered an alternative to map–based navigation.

The ultrasonic positioning, in common with all physical sensing, is subject to error, leading to uncertainty about the position of the PDA visitor. In addition to sensing error, the exhibition space is challenging, split into two large areas by a partial wall, with some display areas covered by roofs. The space includes surfaces that are acoustically very reflective. For aesthetics and coverage, ultrasonic transmissions are reflected off the ceiling. Testing indicates 50% accuracy of 0.52m, 95% accuracy of 1.83m, and an overall standard deviation of 1.29m.

Spatial uncertainty is problematic for several reasons. The aim of shared spatial awareness is mutual visibility, indicating to other visitors what a particular visitor might be viewing. Uncertainty about the actual position of a PDA visitor showed in the spatial awareness displays by apparent jumps of up to 2m. This sometimes made it difficult for trial participants to establish shared context although visitors did resolve some of this uncertainty through talk. They developed habits and strategies so that they could focus on the tasks they were set in the trial, rather than tools such as the handheld and the ultrasonics. Similarly, the

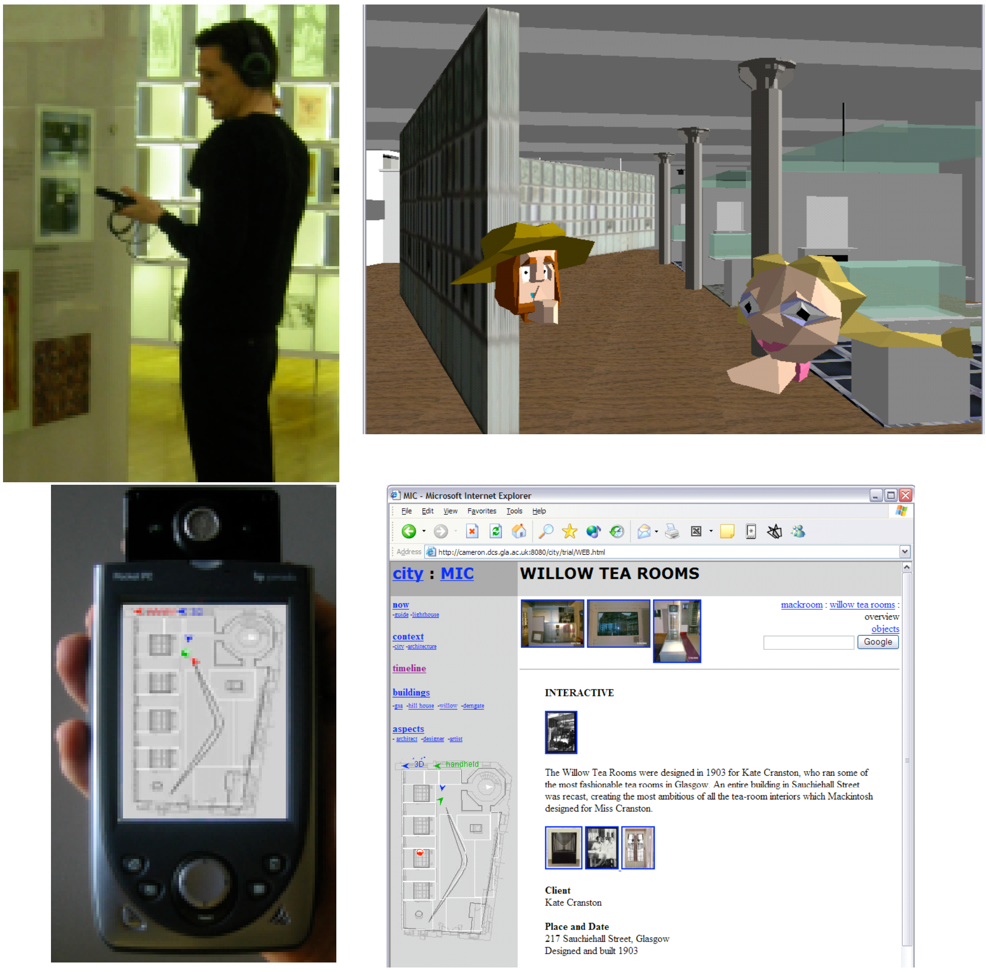

Figure 1. The image top left shows a visitor in the Mackintosh Interpretation Centre with a PDA, ultrasonic receiver, headphone and microphone. A close up view (bottom left) shows the PDA and receiver, with the PDA showing a map of the centre with all three visitors' locations. The image top right shows the VR visitor's 3D graphical display, with avatars for co–visitors. The web visitor's browser can be seen bottom right, with a map in an applet as well as a page of content describing the Willow Tea Rooms.

generation of location-sensitive web content for PDA visitors is difficult in the face of uncertainty.

A variety of technical solutions have been deployed to address the general issue of uncertainty, and others are under investigation. For user trials, the accuracy of the PDA visitor's modelled position was improved by small changes to hardware, e.g. radio frequency beacons in roofed areas, and software, e.g. limiting jumps and jitter. Dead reckoning, spatial and temporal inference, and probabilistic estimates are all being actively explored within Equator and, of course, in many other research groups

A variety of technical solutions have been deployed to address the general issue of uncertainty, and others are under investigation. For user trials, the accuracy of the PDA visitor's modelled position was improved by small changes to hardware, e.g. radio frequency beacons in roofed areas, and software, e.g. limiting jumps and jitter. Dead reckoning, spatial and temporal inference, and probabilistic estimates are all being actively explored within Equator and, of course, in many other research groups.

Ultimately, however, there is a difference between the precise position of a web or VR visitor and the imprecise (sensed) position of a PDA visitor. Here, we were being excessively seamless, in the sense of choosing a precise VR–like representation as primary, and simplifying or reducing the other visitor representations to conform to it. Users' interaction soon revealed to them that this apparently precise representation was not as precise and accurate as it seemed, and so showing the PDA visitor as spatial extent or probability distribution might have been more useful for them.

Another form of uncertainty was apparent for the web visitor. In looking at a web page for an exhibit, and possibly following links to more detail about it, no map movement was apparent. Similarly, a web visitor reading about a painting might follow a link to a page about its past owners, or another link to a page about the style of brushstrokes the painter used. Again, showing the visitor as an extent might be adequate in this case, but more difficult cases exist. It would be difficult or uninformative to show a single sharp location for a visitor reading about a topic exemplified in many or all the artefacts in an exhibition, such as the development of a painter's style through his whole career. The same page may refer or be associated with many or all of the exhibits. We should not expect to discriminate thematic or referential differences in spatial ways.

These are, of course, isolated examples of the much larger problem (or opportunity) of uncertainty. Beyond the inaccuracy of physical sensing and the ambiguity of references, ubiquitous systems must increasingly deal with complex and dynamic technical problems related to bandwidth, power, latency, disconnection, and so forth. Non-technical aspects are also affected by uncertainty, such as awareness of others' locations and activity. These are often apparent through the patterns of social interaction more than interaction with devices and interfaces. Privacy, for example, can be seen as explicit control of the degree of certainty we permit others to have about us, e.g. by permitting others to know roughly, but not exactly, where we are.

Rather than fighting against the uncertainty or ambiguity, we could make a deliberate choice to present and use it. There are several presentation policies (suggested by our colleague, Steve Benford) that may be suitable:

- *pessimistic:* only show information that is known to be correct;
- *optimistic:* show everything as if it were correct;
- *cautious:* explicitly present uncertainty; and
- *opportunistic:* exploit uncertainty.

We are starting to develop cautious presentations to accommodate uncertainty due to ultrasonic- and GPS-based positioning, for example, showing a sensed position as a spatial extent, rather than as a point, and offering maps of estimated signal strength and positional accuracy in the exhibition room and city streets. The user could then have a resource to help change his or her non-digital context in order to change the digital aspects of his or her context. System design would help users 'design' their activity, in ways that take account of the seamful characteristics of the system. We are also considering opportunistic presentations that may be (in the words of our colleague, Bill Gaver), discordant, deliberately leading users to pause or reflect.

In City and in other Equator projects, we are actively exploring these design issues and opportunities. Another strand of our work is in studying how people handle or take advantage of uncertainty and seams in traditional settings, in other technologies and of course in our own systems. In the next section, we offer some examples of our own work in Equator as well as studies and design issues arising from other ubicomp and media space systems. Our goal here is to inform our own system design but also to develop and explore more generally applicable concepts such as 'design for appropriation'.

3 Appropriation

In this section we explore issues of appropriation raised by ubicomp and mixed reality systems as well as other technologies and approaches to computer supported collaborative work (CSCW) such as video and media spaces.

In their paper on the duality of space and place, Harrison and Dourish (1996) compare experiences of the use of video to link rooms at Xerox and at Bellcore. They note that the Bellcore system was reported as disappointing, while the Xerox system was "wonderful". They argue that a critical factor in these different results is the ability to "participate, adapt and appropriate". The Xerox system used inexpensive, easily manipulated, visible hardware. The Bellcore system used expensive, high-quality hardware to attempt to convincingly simulate face-to-face co-presence. Harrison and Dourish propose that the expensive, complex system couldn't be "owned" by its users, inhibiting adoption and enjoyment.

In a related paper reporting on studies of long–term use of video communication (Dourish & Adler 1996), it was argued that we should take a view which "emphasises emergent communicative practices, rather than looking for the transfer of face-to-face behaviours." They offered examples of users learning to accommodate the particular interactional resources that video offered, and also fitting them in to their interaction with people outside of the media space e.g. people off–camera but potentially audible through the media space. They also gave a number of examples of appropriation of media spaces, where prolonged use of the medium led to new and complex patterns of behaviour that took advantage of these characteristic interactional details: "when the medium changes, the mechanisms change too; but the communicative achievements remain."

Studies of media spaces and of other collaborative technologies, such as email (Mackay 1990), Lotus Notes (Orlikowski 1992) and workflow technologies (Bowers 1995), consistently point out that accommodation and appropriation is key to the adoption of new technologies: users design their activity to accommodate the particular technologies we offer them. They create forms of interaction not considered by designers in order to fit technologies into the practices and priorities of their own contexts and communities of use.

Each user trial of our system in the Mackintosh Interpretation Centre only lasted an hour, on average, but we did observe some simple examples of users designing their activity to take account of our system. For example, a VR visitor developed a 'wiggle' gesture to show others using maps where he was. He moved his avatar to run back and forth, and initially also said that he was running back and forth. This would seem excessive or eccentric in everyday interaction, where smaller scale gestures would be appropriate, such as raising or waving a hand. Here, however, the scale of presentation on small maps meant that a larger movement was necessary and appropriate in producing the same 'communicative achievement'.

Another example of accommodation in our system involved trial participants shaping their talk to the different resources they each had. The PDA visitor could see all of a long 'timeline' wall of chronologically ordered panels that described an artist's life. The VR and web visitors could only see information for a single year at a time. The sensing of the position of the PDA visitor as well as the display of all the visitors' positions were both too coarse–grained for position to be used as a gesture or reference towards a particular panel. The visitors developed a way to guide each other by using talk to guide each other to particular images and panels. They verbally emphasised and repeated the year far more than would be necessary among traditional museum visitors, in ways that would seem odd among visitors to the existing Centre.

Another Equator project, CityWide, involved longer–term use of ubicomp technology for collaboration—or, more accurately, for competition. CityWide ran a mixed reality game called "Can You See Me Now?" (CYSMN), which involved the use of handheld computers and augmented reality displays while moving through the buildings and streets of Sheffield (Flintham et al 2003). CYSMN was live for six and a half hours during Friday 30th November and Saturday 1st December 2001. It was designed to be a fast-paced chase game in which up to twenty online players (members of the public using the Internet) were chased across a map of the city by three runners (professional performers) who were moving through the city streets. Runners and players shared an online map, but runners had a global view showing all avatars while players had only a local view showing avatars in their vicinity. The runners talked to one another over a shared audio channel using walkie-talkies. This audio channel was then digitally encoded and streamed to the players over the Internet.

The runners' positions were determined using GPS, whereas the players' positions were controlled through their map tools. Analysis of system logs shows that, from a technical perspective, GPS was indeed quite inaccurate. Estimated errors ranged from 4m to 106m with a mean of 12.4m and a standard deviation of 5.8m. Error varied according to location in the game area, with some of the more open spaces exhibiting typically only a few meters error while the more narrow built up streets suffered considerably more. The more extreme errors were most likely due to multi-path reflections or temporary losses of satellite visibility.

214 players took part over the Internet. The best 'score'—time without being caught—was 50 minutes. The worst was 13 seconds. These periods of involvement were much shorter than those of the runners. The runners were active for almost all of the event, and had time to talk with each other and develop tactics, as shown in this extract from a post–event interview of a runner by a member of the CityWide 'crew'—with our italics added:

> Runner: If they're in a place that I know it's really hard to catch them, I walk around a little bit a wait till they're heading somewhere where I can catch them.
> Crew: Ambush?
> Runner: Yeah, ambush.
> Crew: What defines a good place to catch them?
> Runner: A big open space, with *good GPS coverage*, where you can get quick updates because then every move you make is updated when you're heading towards them; one of the problems is if you're running towards them and you're in a place where it slowly updates, you jump past them, and that's really frustrating. *So you've got to worry about the GPS as much as catching them.*

The online players would move and react without taking GPS inaccuracy into account, but the runners exploited a growing knowledge of it over the course of the two days' play. They became increasingly aware of its effects and also where on the city streets it was most likely to be experienced. By the second day, they had begun to exploit this knowledge by waiting for players to enter areas with higher GPS accuracy, i.e. where a runner could move more quickly. On the first day of play, runners struggled to catch many players, but on the second day their appropriation of inaccuracy significantly changed the balance of the game.

Users' appropriation of the particular seams and characteristics of our systems, and of related systems and technologies, seems to be vital to technology adoption. Users' interactions not only let them achieve their moment–by–moment tasks and goals, but also let them build up a shared understanding of how to resolve interactional problems and how to take advantage of the particular features of our system in their particular context. Dynamic emergence of new patterns of interaction seems common and often necessary, and so we are becoming progressively more interested in designing to support this process, i.e. *design for appropriation.*

Extending the analysis of Harrison and Dourish, one approach to designing for appropriation is to aim for systems whose underlying mechanisms are "literally visible, effectively invisible" in that everyday interaction does not require attention to these mechanisms' representations—but one can selectively focus on and reveal them when the task is to understand or even change the tool. These mechanisms and their representations must be robust, simple and flexibly manipulable. Using these ideas, Dourish used computational reflection to offer manipulable 'accounts' of deep system structure and categorization, and the processes that changed them. Another potentially relevant approach is recombinant computing, as investigated in the Speakeasy project (Newman 2002). Speakeasy explores distributed computing patterns and possible user experiences for ubiquitous computing. Rather than supporting seamless connection and access of devices and services, their approach is to enable users to discover and manipulate devices, services and their interconnections.

Here we also find a useful parallel with design for privacy and awareness among users of ubicomp technology. Bellotti & Sellen (1993) put forward a framework for the design of mechanisms for feedback on and control over the system's models and representations of a user's activity, and how those representations were used by other people. The design of feedback and control mechanisms was based not just on the media involved, but on their effects and uses in interaction. Another design principle was that such mechanisms should be interconnected to allow graceful changes to the degree of engagement (Gaver 1992), so that we support gradual shifts between peripheral awareness and engaged interaction.

Since seams are context too, we suggest that these approaches should be extended so that they are not based solely on how we designers traditionally classify our system components e.g. as models of user activity, infrastructure, sensors, transducers, I/O devices, and so forth. Similarly, we should not always rely on the traditional categorization of error and uncertainty as features of the system to be hidden and reduced. Instead, we suggest that we might offer feedback and control over whatever system components they find useful as they achieve their moment–by–moment tasks and goals, resolve problems in social as well as system interaction, and build a shared understanding of the system as a part of their own overall user experience.

Over time, designers may find patterns and correlations that describe which aspects of system structure, sensing and categorization to reveal, and in what form—but where should we start looking for them? We may be able to begin the process of finding which components these are through sociological methods, such as field studies, and technological methods, such as instrumentation of system components and user activity, to track which components are used, where, how and when. We may be able to find correlations, and offer recommendations, but explanations will be harder to find. In the long run, we must consider deep customization to be something that designers contribute to by revealing system structures and seams, and affordances for their *potential* use, but it is users who through their interactions with our system and with each other choose what to use and why. The ultimate design goal here is a good tool lets users focus on their task—even when that task involves changing the tool itself.

4 Conclusion

It might be considered heretical to suggest that ubiquitous computing might be invisible, but not seamless. On the other hand, there is the danger of treating Mark Weiser's words as gospel truth. However, we do not see our work in either of these ways. We have tried to understand Weiser's discussion, how he drew his ideas from fields such as philosophy, psychology and sociology, and the systems that arose from his design proposals. Our experience with system design and user experience, as well as our understanding of other studies and systems, lead us to suggest that seamfulness is a useful and practical issue for system designers.

We do not see seamlessness as always bad and seamfulness as always good. Similarly, supporting appropriation may be a bad design choice in some situations, e.g. where consistent interaction is desirable for legal, medical or educational reasons, and a good choice in others e.g. where personalization, adaptation and exploration are required. We should treat these design approaches as different tools, to be understood and used in ways that suit the settings, technologies and users we design for.

5 Acknowledgements

This paper grew out of a short workshop paper (MacColl, Chalmers, Rogers & Smith 2002) which itself resulted from discussions at an Equator workshop on Understanding Interaction. This paper has also been influenced by discussions with Steve Benford (U. Nottingham), Barry Brown (U. Glasgow), Bill Gaver (RCA) and Cliff Randell (U. Bristol).

The City project is supported by a donation from the Hewlett-Packard Art & Science programme.

References

Bellotti, V., Sellen A. (1993) Design for Privacy in Ubiquitous Computing Environments. *Proc. Euro. Conf. on Computer Supported Cooperative Work (ECSCW'93),* pp. 77-92. Kluwer.

Flintham, M. et al. (2003) Where On-Line Meets the Streets: Experiences with Mobile Mixed Reality Games. To appear in *Proc ACM CHI 2003.*

Bowers, J., Button, G., Sharrock, W. (1995) Workflow from Within and Without: Technology and Cooperative Work on the Print Industry Shopfloor, *Proc. Euro. Conf. on Computer Supported Cooperative Work (ECSCW'95),* pp.51-66. Kluwer.

Brown, B. et al. (2003) Lessons from the Lighthouse: Collaboration in a Shared Mixed Reality System. To appear in *Proc ACM CHI 2003.*

Dourish, P. (1995) Developing a Reflective Model of Collaborative Systems, *ACM Trans. CHI*, 2(1), 40–63.

Dourish, P., et al. (1996) Your place or mine? Learning from long-term use of audio-video communication, *Computer Supported Cooperative Work* 5(1), 33-62.

Galani A., Chalmers, M. (2002) Can You See Me?: Exploring co-visiting between physical and virtual visitors *Proc. Museums and the Web 2002,* pp 31-40. Archives and Museum Informatics.

Gaver, W., et al. (1992) Realising a Video Environment: EuroPARC's RAVE System, *Proc. ACM Conf. on Human Factors in Computer Systems (CHI'92),* pp.27-35.

Harrison, S. & Dourish, P. (1996), Re-place-ing space: The roles of place and space in collaborative systems. *In Proc. ACM Computer Supported Collaborative Work* (CSCW96), pp.67-76.

MacColl, I., Chalmers, M., Rogers, I. & Smith, H. (2002) Seamful ubiquity: Beyond seamless integration, *Proc. Ubicomp 2002 Workshop on Models and Concepts for Ubiquitous Computing.*

MacColl, I., Millard, D., Randell, C., Steed, A. (2002) Shared visiting in EQUATOR City. *Proc. ACM Collaborative Virtual Environments (CVE2002),* pp. 88-94.

Mackay, W. (1990) Patterns of Sharing Customizable Software, *Proc. ACM Computer Supported Collaborative Work* (CSCW90), pp. 209-221.

Newman, M. et. al. (2002) Designing for serendipity: Supporting end-user configuration of ubiquitous computing environments. *Proc. ACM Conf. Designing Interactive Systems (DIS 2002), pp. 147-156.*

Orlikowski, W. (1992) Learning from Notes: Organizational Issues in Groupware Implementation, *Proc. ACM Conf. Computer Supported Cooperative Work (CSCW'92)*, pp.362-369.

Randell C., Mueller, H. (2001) Low cost indoor positioning system. *Proc. Ubiquitous Computing* 2001, pp. 42-48.

Weiser, M. (1994a) The world is not a desktop, *ACM Interactions* 1(1), 7-8.

Weiser M. (1994b) Creating the invisible interface: (invited talk). *ACM Conf on User Interface Software and Technology (UIST94)*, p.1, 1994.

Weiser M. (1995) Ubiquitous computing: (invited talk). *Proc. USENIX Conference*, 1995.