Seamful Interweaving: Heterogeneity in the Theory and Design of Interactive Systems

Matthew Chalmers & Areti Galani
Computing Science, University of Glasgow, UK
matthew@dcs.gla.ac.uk, areti@dcs.gla.ac.uk

Abstract
Design experience and theoretical discussion suggest that a narrow design focus on one tool or medium as primary may clash with the way that everyday activity involves the interweaving and combination of many heterogeneous media. Interaction may become seamless and unproblematic, even if the differences, boundaries and 'seams' in media are objectively perceivable. People accommodate and take advantage of seams and heterogeneity, in and through the process of interaction. We use an experiment with a mixed reality system to ground and detail our discussion of seamful design, which takes account of this process, and theory that reflects and informs such design. We critique the 'disappearance' mentioned by Weiser as a goal for ubicomp, and Dourish's 'embodied interaction' approach to HCI, suggesting that these design ideals may be unachievable or incomplete because they underemphasise the interdependence of 'invisible' non-rationalising interaction and focused rationalising interaction within ongoing activity.

Categories & Subject Descriptors: H5.0. Information interfaces and presentation (e.g., HCI); H1.2. Models and Principles: User/machine systems.

INTRODUCTION
An important recent HCI text [11] drew upon philosophy in discussing the accommodation of new technology by users, and their appropriation of it as they find their own ways to use and understand it. Dourish suggested that everyday human interaction is embodied i.e. is non–rationalising, intersubjective and bodily activity. Traditional approaches to HCI offer many guidelines for system design, but do not take full account of embodiment, according to this view. They are not in accord with the activity they aim to support. He raises the issue of embodiment but draws back from offering specific principles and guidelines, favouring instead statements that help sensitise designers to the general issue, e.g. users, not designers, create and communicate meaning and users, not designers, manage coupling. This paper uses similar theory, but tries to move forward with regard to discussion and understanding of accommodation and appropriation. We apply that understanding in making specific design critiques, suggestions and guidelines, centred on the issue of heterogeneity—spatial, temporal and technological—as a catalyst of deeper understanding. Although we focus on interactive systems that most obviously consist of a mix of media, such as mixed reality and ubicomp systems, we suggest that the use of any interactive system involves a degree of interdependence with other media. Therefore, we believe that the issue of heterogeneity is relevant to the wider field of HCI and interactive systems design.

Mixed reality systems have their roots in collaborative virtual environments (CVEs). CVEs and VEs have primarily gained wide public acceptance in the form of computer games. The focused engagement in such games is designed to fit with the closed world of the virtual environment. A player can become immersed in a game—closed off from the 'real' world—by attention as much as by apparatus. A PC at home can be as engaging as the head–mounted displays and immersive projection technologies of research labs. However, even a single–player non–networked 3D game may be a resource for social interaction, e.g. played by one person while friends and family shout advice from the sofa, order pizza by phone and slip into the kitchen to get more drinks. A computer game is a resource for far more social interaction than the system's architecture may suggest. Games may be tightly interwoven into people's interaction, collaboration and culture, but in general the wider context of system use is hardly modelled or represented. In technological terms, such games and CVEs are decoupled from their users' wider context such as more traditional interactions of family members, the overall educational activities of the school and the business of the workplace.

In response, many researchers are working on systems that are more 'out in the world' than traditional CVEs, contextualising and connecting them to the other media that we use in everyday activity, and sensing and tracking users' wider context and activity beyond the computer. For example, in mixed reality (MR) systems, users of a virtual environment may see artefacts or images from a traditional workplace, and vice versa, e.g. [20]. In augmented reality (AR) systems, users may use the technologies of VEs combined with tangible artefacts in more traditional media, such as urban models and interaction devices made from wood, wire and plastic [33], or tiles and book pages made from toner and paper [2]. Here, as discussed in the early days of ubiquitous computing [34], designers do not focus on one digital or

DIS2004, August 1–4, 2004, Cambridge, Massachusetts, USA.

traditional 'space' as the primary or dominant medium. In MR, AR and ubicomp systems, the distinction between digital media and traditional media is clear if one looks for it, but the idea is that, effectively, one is not aware of it because one focuses on the overall experience: on tasks instead of tools. The new technology and the seams of where it joins to old media are, as Weiser put it [35], "literally visible, effectively invisible". With such interwoven or simultaneous use, the notion of each medium being a space itself becomes problematic, as has been discussed by authors such as Harrison and Dourish [20], Brown and Perry [4] and Chalmers [6].

One might go so far as to call computer games 'ubiquitous computing' in the sense that Mark Weiser discussed in [35], where he suggested that even a "glass TTY UI can be ubicomp," if its use is well woven into the fabric of people's collaboration and interaction. This may seem contradictory to the common notion of ubicomp, involving technologies such as location sensors, mobile displays and wireless communication, but Weiser was clear that it was not technology in itself that made for ubicomp. Instead he suggested that we should aim for the accommodation and appropriation of computing into everyday life, so that its use is non–rationalised, intersubjective and interwoven with the other media that we use. In good MR and AR (and ubicomp) design, according to Weiser, interaction using heterogeneous media is so tightly coupled in user activity that the obvious differences and boundaries—what he called 'seams'—between the parts of a system become less significant than the quality of interaction with the whole. The seams are perceivable—the technology is 'seamful'—but we can call the whole system a single, hybrid object because coupled use of the parts is so unproblematic in users' interaction i.e. interaction is non–rationalised and seamless. The MagicBook, for example, works when users get past a rationalising focus on each of the interconnected media, and instead start reading, learning and imagining.

The next section of this paper discusses principles and assumptions underlying ubicomp systems. It focuses on the process by which any new technology becomes woven into the fabric of people's collaboration and interaction, i.e. interwoven and interdependent with the use of the hetereogeneous media that are used in everyday life. A number of these theoretical issues were explored and refined in our system design work, and so a later section of the paper uses the paper's initial theoretical discussion in reviewing key aspects of one system design and the associated user experience. The paper then puts forward a number of general design suggestions for supporting the process of accommodation and appropriation in ubicomp and MR systems, and discusses further potential methods of design for appropriation, where systems are designed to show and support the change of how they are used and how they are structured.

HETEROGENEITY AND UBIQUITY

This section explores the process of experiencing and understanding how to weave a new system into one's everyday life. It emphasises the temporal, spatial and social patterns of use of all the media one has at hand, rather than treating a tool as an isolated 'thing in itself'.

It is relatively common for an interactive system to be designed for use in a relatively isolated way, so that the use of the digital 'space' or medium stands above or apart from others. However, studies of use consistently point out that accommodation and appropriation are key to the adoption of new technologies: users design their activity to fit 'our' technologies into the many and varied media that they use in their everyday lives, often changing or adapting the technology along the way. One influential paper on the long–term use of video communication points out "complex patterns of behaviour built up around the interactional details of the video medium […] When the medium changes, the mechanisms change too; but the communicative achievements remain" [12]. This process has also been observed in email [24], Lotus Notes [30] and workflow technologies [3].

It is normal for users to create new forms of interaction beyond those considered by designers and unlike face–to–face interaction. They accommodate the characteristic affordances of a new technology, but they also appropriate it to suit the practices and priorities of their own contexts and communities of use i.e. other, older tools and media, and other people. As they do so, the use of the new technology becomes everyday, in the sense that "the most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it" [34].

In laying the foundations for ubiquitous computing, Weiser put forward a design approach that relies on the fit and coupling of the system design with the context of use i.e. with the other tools and media used in everyday communication, activity and interaction. He also emphasised the contextual and social aspects of design to support this disappearance, e.g. "the unit of design should be social people, in their environment, plus your device" [35], and computational media as being embedded and embodied in social interaction. Social people, in their environment, continually mix and couple media in everyday communication: walking, gesturing and pointing while one talks, and referring to places and what people did in them as one writes.

Weiser's notion of disappearance, where a tool is "literally visible, effectively invisible" is from philosophical hermeneutics [19]. Weiser says that:

> 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.

An old example from Heidegger is the way that a skilled carpenter engaged in his work focuses on the use of the hammer, and how it changes and is combined with other tools and materials, rather than focusing on the hammer in itself. Heidegger called this practically engaged and non–rationalising use 'ready–to–hand', in contrast to the rationalising, objectifying and abstracting activity he categorised as 'present–at–hand'. He saw both modes or categories of use as being set within a circular process of

interpretation, in which one is influenced by one's understanding and experience of older tools and media when using any new tool or medium. One's use of the tool in the course of everyday, situated and social interaction, combining the new tool with the heterogeneous others used in everyday life, builds up new experience and understanding—that will affect how one uses and interprets another new tool. In time, this process of accommodation and appropriation lets one focus on the use of the tool, and not on the tool in itself, thus making the tool 'disappear' as Weiser also points out.

Influenced by Weiser but also drawing directly from similar philosophical sources, in [11] Dourish similarly called for a move towards the design of interactive systems which have a better fit with everyday human activity, understanding and interaction, and with the practically engaged and non–rationalising way that everyday activity takes place. Dourish draws upon Heidegger, as well as Schutz' elucidation of the social or intersubjective element of everyday perception and activity, Merleau–Ponty's discussion of the way that the body, through the interwoven senses, plays a vital role in everyday perception, and Wittgenstein's emphasis on the way that meaning and activity are based on the patterns of use of the heterogeneous mix of media that constitute language: "the meaning of a word is its use in the language".

Weiser and Dourish focus on raising our awareness of embodied interaction, i.e. the interpretation of a system by a user as ready–to–hand. They present traditional HCI design as being based on its opposite, i.e. as rationalising, objectifying, abstracting and present–at–hand interpretation and use. Dourish discusses the shift between these two categories of interpretation as varying the degree of coupling between the interpreter and the system. As he puts it [11, p. 139], the existence of both modes is critical to the effective use of technologies. However, Weiser and Dourish both swing from one extreme to the other, focusing almost entirely on design to support embodied or ready–to–hand interaction. They do not fully address the relationship between the two modes. In particular, how does a tool become invisible or ready–to–hand?

Heidegger, and his successors such as Gadamer and Ricoeur, held that situations where a tool becomes present–at–hand may be crucial to the individual's learning and to the differences between individuals. The ongoing 'feedback loop' of interpretation and understanding integrates these two modes, and affords variation in people's understanding as well as consistency in their behaviour. For example, creativity can be considered as the variation of an individual's subjective understanding from his or her prior understanding and from others'. The individual may then be very conscious of his or her own activity, rationalising it and very aware of it, i.e. the system, tool or symbol is present–at–hand. With experience of its use, however, it may become understood and familiar, i.e. more ready–to–hand and embodied. Similarly, as two people perceive one another's use, with each interpreting and reacting to each other, they can achieve intersubjective consistency of behaviour; consistent with each other, but not necessarily with the use expected by the designer. A use or activity that is new and present–at–hand for one of them can thus become learned and ready–to–hand for both. The circular process of interpretation, whereby perception and activity are influenced by understanding, but also feeding into and changing understanding, thus relies on the interplay between ready–to–hand and present–at–hand interpretation.

Embodied interaction, as Dourish and Weiser made clear, is an aspect of human activity that is under–emphasised in HCI. Nevertheless, ready–to–hand embodied interaction and present–at–hand objectification are interdependent—and neither author addresses this. We have to expect that a new technology will be to some degree present–at–hand, no matter how well the designer aims towards embodied or present–at–hand interaction. This is most clearly the case when the technology is new, but other situations arise that neither Weiser nor Dourish fully address. One is *breakdown*, where the affordances of even the most familiar tool may significantly differ from those of everyday ready–to–hand use e.g. when the head of the carpenter's hammer becomes loose, so that he has to consciously concentrate on using it towards his task. Another example might be the breakdown that occurs with a mobile phone when it loses its network signal: one's attention may turn from a conversation 'through' the phone and its infrastructure to the tool itself. Another usefully present–at–hand situation is where the one can no longer work through the tool in a transparent way because the task 'is' the tool itself. This might happen because of breakdown: the carpenter may work on the hammer, to fix it, and the phone user may focus on the signal strength indicator, waiting or moving until he or she regains a signal. It also may occur as an act of conscious learning or analysis, e.g. a novice carpenter trying to improve his hammer swing, or a researcher studying how a new mobile technology works in use.

Activity continually combines and cuts across different media, building up the temporal patterns of coupling and interweaving that constitute experience and understanding. A person's work or activity may be influenced by a 3D computer graphics display in front of them, and the interactions that such a system affords, but also by books, telephones, hypermedia, furniture, buildings and so forth—and other people's use of all of these media. People act and work through the full range of media they have at hand. A narrow emphasis on one digital system or 'virtual space' as the paramount resource for activity underrates the influence of other media. Recent technological developments, such as mobile phones and email, heighten or highlight a phenomenon already familiar in the use of older media such as written text, maps and cinema, and well–explored in philosophy, semiotics [27] and linguistics [31]. It is hard to claim that any digital medium stands by itself, as users have preconceptions and expectations of how to use it, how it compares to other media and how it can be combined with them. More generally, a medium cannot be fully used or understood in an isolated or 'singular' way. For example, a city's meaning is not just in its bricks and

mortar, but also in our understanding and use of the information about it. At any time, one is likely to have symbols in a number of heterogeneous media available for interpretation and use. As I walk through a train station towards a city square, the map in my hand, the voice of a friend on my mobile phone, the signs informing me of exit routes and the posters advertising exciting shopping opportunities are all open for my interpretation and action. Temporally, symbols in an even broader range of media influence me, as my activity is influenced by my past experience and my expectations of the future. Past experience may include my previous visits to that city, my browsing of a web site with good maps to print out, and my experience of magazines, books and films about urban life, and so forth. My language and culture, spanning media old and new, affect me as much as the immediate perception of spatial form. The early decades of the 20th century saw dramatic advances in philosophy, linguistics and semiotics, as they took account of how activity and language is constituted by all the symbols and all the media one uses, with each symbol interpreted through immediate perception as well as past experience and social interaction. Contemporary neurophysiology is in strong accord with this view [9, 13], as is architecture and urban design [22], the field most obviously related to the theory and design of space.

The differences between media are usually very obvious. We can characterise media and treat each one as if it were an isolated individuated entity because of the senses we use in perceiving each one, and also because of our understanding of how to relate and to distinguish examples of each one. For example, it is easy to distinguish the spoken word "red" from the written word *red* because of the senses one uses in each case. Despite having the same letters, it is easy to distinguish *tar* from *rat* by looking at the order of letters within each written word. Simple rules about what one can immediately see, hear, etc. within a word begin to strain and then break when one considers, for example, how we distinguish homonyms such as *rose.* The written word *rose* can mean many things, including a flower and having risen. When spoken, the same syllables can also mean linear structures (rows), about or belonging to fish eggs (roe's), moving in a boat (rows), small deer (roes) and multiple occurrences of the Greek letter (rhos). Saussure [31] established that a word's usage is understood through understanding and experience of patterns of use i.e. of other symbols that generally co–occur with it in use in language—and not just through the perception of the word's syllables or letters. A digital system or tool also has this property: its meaning is its use in the heterogeneous mix of media that is language.

Context of use becomes progressively more important as we turn from thinking about the differences between media, and the distinction of symbols, to the similarities of media and the relatedness of symbols. This understanding is not solely dependent on the form or medium of each symbol, but also on experience and understanding of how we use each symbol in the context of other symbols—context that may include symbols in any or all media. For example, the spoken word "red" and the written word *red* are related because, based on past experience and current context, we can use either of them in the context of rose blooms, fresh blood, the former USSR and so forth. We understand, relate and differentiate symbols through experience of contexts of use within a culture.

Using this section's theoretical discussion, we can take another look at the design of ubiquitous computing and mixed reality systems in general. A typical 'context–aware' ubicomp system involves the coupling and interdependence of media for an individual user, but we often seem much keener to couple information to space than vice versa. A museum exhibition might be associated with a set of web pages, so that walking into a room on a particular architect triggers the display of text describing the life and work of that architect. Ubicomp systems rarely treat space as secondary, so that reading text about the architect triggers display of a map or visualisation of the museum room, and affords access to a structured collection of blueprints, design sketches and building models. Perhaps neither should be primary: each should be coupled to the other, and part of the context of the other, so that the space of the room and the text of the page are treated as peers.

If a system synchronously couples different media used by several people—rather by an individual—its support of social interaction may make it more likely to be called a mixed reality system than a ubicomp system. For example, a person walking into a museum room might be made aware of a friend's concurrent use of a VR model of the room, suggesting openness to conversation about the exhibition despite the two people being geographically remote from each other. Again, note a tendency to treat space as primary. Opening a museum web page might show images from the museum via a webcam, but it is rare to find video going the other way, from the reader back to the museum visitor.

Ubicomp often focuses on context as based on immediately observable objective features, in a rather present–at–hand way, whereas theory and studies of use suggest that context also has temporal and intersubjective features that cross or interrelate media—and that these features are especially important in ready–to–hand use. For example, many Ubicomp systems rely on a simple 'walk up, pop up' approach whereby only one's current location triggers information display, but one's current information, e.g. the pages one has recently viewed in a web browser, rarely triggers location display. There are some partial counterexamples, of course, many of which have been applied in the museum domain. HIPPIE adapted the presentation of information in a museum, and of the attributes within the system's database, based on a record of what displays and related information a visitor had seen before, either in the museum or previously [29]. The system of [32] captured video images of paintings in a museum tour, and then would later automatically retrieve video recordings of the tour guide if one later came across the same paintings (or realistic enough reproductions). A rather simpler example was the HP Cooltown Rememberer system, which built up a visit record, consisting of a set of

web pages. Users left the museum with an artifact that was intended to remind the user of the visit and which contained a URL for the visit record, for example a fridge magnet with an embedded RFID tag [15].

Although they do reflect slightly broader notions of context, dealing with time and heterogeneity, these systems were essentially single–user systems and were relatively asymmetric or biased in terms of their coupling and use of media. For example, HIPPIE's representation of a person's interests and activity in the museum, or in the museum information application, was isolated within the application. For example, general web browsing about related exhibitions and artists would not affect HIPPIE. The system of Schiele et al. would not respond to an image of a previously seen painting on TV or a video playback via a VCR, to remind one or guide one to the museum. With specialised equipment, Rememberer could support retrieval of a visit record from home, but again this was rather one–way: more general use of the web would not directly trigger display of the visit record (or even of the fridge). When browsing the web, one rarely if ever gets reminders or peripheral awareness of the relevance of the current web page to places one has visited.

In our system design work, we are beginning to explore notions of context and interdependence that go further away from the spatial and synchronous approaches of these and other ubicomp and MR systems. The theoretical discussion in this section suggests that treating media as peers, and taking more account of the temporal and experiential aspects of context that cut across the boundaries of heterogeneous media in more symmetric ways, opens up new possibilities for technology design and for computer–mediated social interaction. More particularly, the aim of embodied interaction, and the fit with user activity, may be improved if system designers can better understand and support this rich but complex process of interweaving, accommodation and appropriation. As mobile computers become net–connected, and can be used to access other people and other computers, and as they gain rich sensor and interaction devices, these design possibilities become ever more feasible. Experimenting with these possibilities, and finding practical ways to design for the process of appropriation, is the subject of the following sections.

USER EXPERIENCE OF HETEROGENEITY

The City project, set within the Equator interdisciplinary research collaboration (www.equator.ac.uk), aims for mixed reality and ubicomp design that increases and takes advantage of this coupling and interdependence, and theory that lets us understand it. In our work we aim to treat digital media as peers, rather than treating any one space or tool as the primary focus or locus of activity. Our intention is to support social context and interaction, as well as individual activity and interpretation, through heterogeneous media. This blend of social and individual activity is familiar from traditional cultural institutions, where co–visitors use awareness of each other's interaction with exhibits as a resource for their interaction with each other, and use interaction with each other as a resource for their interpretation of the exhibits [14, 17]. City explores the process of coupling and contextualisation of digital and traditional media, and of different digital media, as users weave them together to form resources for their interaction and interpretation. This section mostly uses the Mack Room user experience to ground and exemplify the discussion of the previous section. The user experience was presented in [5], but here we discuss different fragments and issues from the Mack Room user trials. In this section and the next, we also begin to introduce other systems from City and from related Equator projects. The Mack Room system predominantly addressed issues of synchronous social interaction among co–visitors, and the discussion in this section will mainly focus on dealing with heterogeneity in the course of the visit. However, the excerpts of the trials also highlight aspects of the asynchronous communication among trial participants, which have fed into newer system designs.

The Mack Room system supported a shared visiting experience for three visitors, involving talk, spatial awareness and overlapping content. They could speak to and hear from each other, and the subsystem for talk was relatively homogeneous across visitors. In part, the system was designed around a scenario of users wishing to share a museum visit but being geographically remote from each other, but it was also intended to allow us to explore users' handling of heterogeneity. In order to explore their accommodation and interweaving of technologies, users' representations in the system and their spatial and content information were deliberately heterogeneous, i.e. the visitors used different spaces and tools.

Interaction among people in different locations and contexts, by definition, means people with different resources at hand. As remote collaborators discuss and refer to contextual information, some heterogeneity is inevitable: one person can use the non–digital resources of his or her location while others have only digital representations of that location at hand. A case that is more easily handled is audio: each person will hear his or her own voice and sounds from other nearby sources differently to others, because of the digitisation and transmission of audio, but we have become relatively accustomed to handling this. A much more challenging heterogeneity is that of people's position, orientation and gesture within rooms and buildings. For example, the Mack Room presents much greater visual and tactile richness than the room's digital representations e.g. maps and VR models. We addressed this inevitable heterogeneity by coupling media together, tracking activity in each medium and representing it in others, and by providing content that users might discuss and share in social interaction.

The 'on–site' visitor used a handheld or wearable computer in the Mack Room, with its location tracked or derived via an ultrasonic positioning system. The handheld showed an outline map of the room, but no web content—we considered that the traditional exhibition content was already rich enough. A second visitor used the World Wide Web (including a 2D map) on a laptop or PC in another room, with this visitor's 'location' in the exhibition

derived from his or her map marker. A third used 3D graphics on a similar machine in a third room, and had a location derived from his or her avatar. The two online visitors' movements, in map and VR respectively, were used to dynamically generate location–specific web content.

All three visitors' locations were used in each visitor's spatial representation i.e. the on–site visitor's PDA map showed three visitor icons, as did the web user's map applet, and the VR user's 3D graphical model had a first–person point of view and had an avatar for each of the other two visitors. However, note that each of the three visitor locations came from a different 'space' i.e. the room, the map applet and the VR respectively. Rather like the content drawn from the catalogue, that was close enough to the Mack Room artefacts and descriptions to support interaction, the spatial representations were similar but not identical.

By presenting themselves to each other through talk and through maps and VRs, the visitors wove together heterogeneous media so as to increase the degree to which the heterogeneity was "literally visible, effectively invisible". In [5] we referred to the construction of 'hybrid objects', each of which was actually a set of heterogeneous representations tightly interwoven in visitors' interaction. Many of the artefacts and exhibits had corresponding descriptions and representations for each of the three visitors, and some of these sets of corresponding objects were used as hybrid objects—but other correspondences were constructed by the visitors themselves and some of the pre–designed correspondences were not used at all. In other words, the term 'hybrid object' is a way for us to describe a pattern of use and reference that is socially constructed by designers and users. The process of construction is central to or even constitutive of the shared visit—and hence central to or constitutive of a good system design.

There were many commonalities between the themes, images and descriptions that visitors had but, as pointed out in [5], there were also differences and heterogeneities. The on–site visitor had the richness of the traditional exhibition, including touch screens and video displays that on–line visitors lacked. The on–site visitor lacked the access to the web that the two online visitors had, and could not move and jump between Mack Room locations as quickly as the on–line visitors. The 3D view of the Mack Room VRML model had greater visual richness than the 2D maps, but also gave rise to visual occlusions that contrasted with the overview of a 2D map. Therefore, for example, the VR visitor could lose sight of his or her co–visitors. Incidentally, the images and furniture in the Mack Room are not all 'real', in that some of them are reproductions. It is interesting that we often call a printed copy of a painting 'real' but call a digital copy 'virtual'.

As vom Lehn, Heath and Hindmarsh [23] discussed in their studies, members of a group of visitors collaborate in the exploration of both galleries and displays by conversing with each other, animating displays for each other and so forth. We also observed that, during collaborative exploration of displays, members of a group contribute to the shared exploration of a display by volunteering information and highlighting interesting points. They engage in 'creative discussions' [14], i.e. conversations about and around displays and exhibitions, which give speakers the opportunity to share their knowledge and understanding, to develop their own ideas and interpretations, and to establish new shared understanding and use of terms and references. Here museum visitors refer to and use their previous experience, whether this was gained long before the visit or just a few minutes earlier. Creative discussions can be contrasted with 'functional conversations' comprised of descriptions of the environment as it appears now in objective terms: information about the look and feel of displays, and directions to where people or objects are. We liken creative discussions to readiness–to–hand with regard to use of the exhibition, and functional conversations to presence–at–hand. Creative discussions among museum visitors are perceived—by museum professionals—as an indicator of successful interaction. Functional conversations are seen to reflect a shallow or objectifying view of the museum and, sadly, are often associated with IT in museums. Nevertheless, the total novice or the student of museums may find this latter mode of interpretation useful—as do most people, occasionally.

Echoing our characterisation of Ubicomp in the last section, functional conversation is about synchronously observable objective features, whereas creative discussions are about past experience and intersubjectively established features. In the Mack Room system, these two categorisations of visitors' conversation overlapped more than in traditional museum visits. The look, feel and content of the exhibition varied among visitors, and so individual interpretations were brought to bear from the outset. Conversation and interaction served both creative and functional purposes. In establishing shared understanding, they articulated what they had in common and what was different. Shared content, shared terminology and common experiences support discussion and exchange of opinions among co–visitors, but differences do too. In the Mack Room, differences often led to greater engagement with the exhibition, and discovery of details that would probably have been unnoticed otherwise. Often this enhanced their engagement, as in the following example where the on–site visitor is motivated by her friend to explore the content of a touch screen display.

*In the excerpts, G is the on–site visitor, B the VR visitor, and R the hypermedia visitor. Square brackets show overlapping talk, underline shows speaker's emphasis, italic indicates text from museum labels, and numbers *1* show when images were captured.*

B: Is it something about Derngate? 78 Derngate.
G: Yeah, yeah, I think that's it.
B: [Is it something, yeah]
G: [yeah], I am not so sure where that was though, the house... *1*
B: Right, *it's a late Georgian terrace house in the middle of Northampton, which Mackintosh altered, decorated and furnished for*
G: Ahhh!
B: *for Wenman J. Bassett-Lowke,* Look, Lowk...
G: [Ohh, I got a little screen]

B: [*who engineered models] and made model railway engines*.
G: OK
B: Can you see the *dark lounge hall*?
G: No… I think I can find that, it's an interior presumably…
B: Yeah, it's quite, it's quite something!
G: (not audible) Would you like to click at the *guest bedroom*? *2*
B: Hmmm, I am not sure, I am not sure if I can. Oh wait, maybe I can. Is it two beds in it? Two single beds very close together?

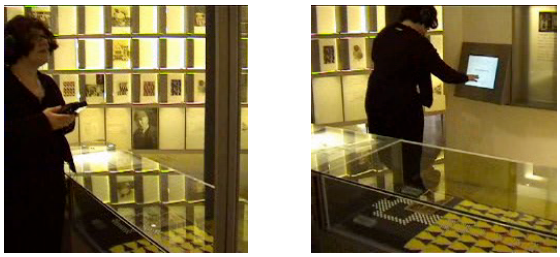

Figure 1. The 78 Derngate example

The visitors were unfamiliar with 78 Derngate Street, a house with an interior designed by Mackintosh, and their questions related both to the house and to each other's opinions on it. Both visitors repeated or reused phrases from the text in the labels, for example 'dark lounge hall'—behaviour common in groups of museum visitors [25]. Such phrases became established as shared references for their interaction, and were pivotal in establishing the relationship between the components of hybrid objects, but shared phrases were not always derived from the exhibition materials. For example, the exhibition designer might not be pleased to hear his elegantly curving 'time line wall', visible on the PDA map, called a 'boomerang':

G: There is a big thing along, it looks like a boomerang shape, that's a big wall with glass with pictures on it.
R: Where is the big wall?
G: I am walking along that, is on my left hand side as I move up *1*
R: Yes on your left side. It's got a boomerang [shape]?
G: [Boomerang] shape, a–ha

Figure 2. The 'boomerang' example

In this example, the hypermedia visitor (R) asked the on-site visitor to clarify the shape and the size of the partition in the room. The on–site visitor, by describing visible aspects of the gallery, e.g. 'big wall' and 'boomerang', did not only offer an account of the environment as he saw it but also established an additional word to use later on in the visit. Referring to the previous section, as the trial participants perceived one another's use and activity, with each interpreting and reacting to each other, they achieved intersubjective consistency of naming—consistent with each other, but not with the designer. Generally, these elements were initially used by one person in one space or medium—not just the on–site visitor—and became collaboratively used by all participants in their interactions.

Social interaction also let the visitors accommodate variability of positioning. The web visitor moved between discrete map 'zones', and the tracking of the on–site visitor's position via ultrasonics varied in accuracy across the room. We did make the mistake of not showing the VR visitor the boundaries of the zones, leading to extra repair conversation, but—especially with regard to positioning—visitors did build up a shared understanding of these 'seams' in the infrastructure and how to handle them, suggesting again that that social interaction may be effectively supported through only approximate location, rather than highly precise position. A visitor's engagement is based not solely on his or her own precise position and gaze direction—the type of synchronous features that many ubicomp systems support—but also on the general areas where their co–visitors are and have been, and what they are showing and have shown an interest in.

Making suggestions or recommendations as to what to see next is one of the most important ways in which shared terminology and experience is used. This may be explicit, for example by asking one's friends to see something, or implicit e.g. through a visitor's engagement with a display being seen or remembered by others. In our studies we found that co–visitors were aware of where their friends were and what they were looking at, and often also of where they had been and had looked at. They took advantage of their friends' engagement in shaping their own actions. In many cases, the way one presents one's engagement operates as a recommendation for co–visitors to follow. In the case of the Mack Room system, visitors either achieved this through gestures specialised to the media involved, or through verbal description:

B: I am looking at the reconstruction of the guest bedroom in the Hunterian Art Gallery
G: Is what you are looking at?
B: I am. Quite stripy!
R: Oh, me too now.

In this example, the VE visitor (B) announced what she was looking at, and also expressed her personal opinion about the look of the specific room. The hypermedia visitor (R) joined her, but the on-site visitor (G) continued to look at another display in another part of the gallery. However, he was aware of what his friends were looking at and used it to inform his own exploration one minute later:

G: Did you see me passing? Do you see me go by?
B: I do, where are you going? I am going to follow you then.
G: Oh, are you? I was going to go to the bit you were looking at which was…
B: Oh, I was walking… oh where did you go again?
G: [I was looking at]
B: [Where…]
G: [Where did you]
B: [Not audible]
G: Ha, what was the exhibition you were looking at before?
B: It was the Hunterian Art Gallery, the guest bedroom.

This case serves in reiterating that interaction among companions is not strictly based on proximity and gaze, but, more generally, on awareness of each other's current and past activity. The participants appeared willing to follow their friends regardless of the media they were

using, passing the 'leading role' among them. Similarly, the on–site visitor did not dominate the generation of terminology and landmarks central to their social interaction. Although one might expect the on-site or 'real' exhibition to have primary impact on people's choices, participants regularly treated all media as equal resources for interaction as long as the media supported the social activity at hand.

Part of the way through the trial, we realised that our design for the web user had a weakness based on making the space of the map primary in an aspect of individual interaction—clicking on a web link about a part of the museum room did not change the visitor's location in the maps and VR. This was partly because of the cost and complexity of 'geo–referencing' each accessible page. For example, a web visitor reading about a painting might follow a link to a page about a topic exemplified by many or all the artefacts in an exhibition, such as the development of the artist's style throughout his career. More generally, it can be difficult to discriminate thematic or textual differences in spatial ways.

While the Mack Room design concentrated on relatively symmetric awareness across heterogeneous media, one limitation of our system highlighted by the earlier theoretical discussion was that, like a good proportion of ubicomp systems, it did not directly support awareness across time. One of our newer systems currently under trial, George Square, supports synchronous awareness rather as in the Mack Room, but it also supports asynchronous awareness via logs of activity and a collaborative filtering system based on Recer [7]. The system makes contextually specific recommendations of locations and URLs by comparing each person's recent activity (with any and all our tools) with similar sections of the past activity of selected others. Both spatial and informational recommendations are shown in each of the media in our system i.e. in 2D maps, in textual web pages, and in 3D VRs. This issue of representing the past, along with other guidelines arising from theory, design and user studies, is explored further in the next section.

GUIDELINES & DIRECTIONS FOR DESIGN

In this section we use four topics to summarise and structure some of our design–oriented findings and suggestions: variation and precision of positioning, heterogeneity and 'correctness' of content, making the past a resource for ongoing interaction, and exposing some of the limits and boundaries of the media we use.

In a way contrary to a great deal of work in ubicomp and MR, highly precise positioning may not always be necessary to support social exploration of artifacts and interaction. The issue of the degree of positioning accuracy being appropriate to the task or activity at hand, rather than an end in itself, has been raised in some earlier work e.g. in Equator's CityWide project [16]. The way that people accommodated significant spatial variation in accuracy in the Mack Room reinforces this point.

A closely related point, but a more contentious one perhaps, is that precisely overlapping content can occasionally be detrimental to the user experience: small differences can serve as individual contributions and spurs to deeper engagement. We propose that overlap in content and presentation should be substantial but not necessarily total. Shared homogeneous facilities may aid what Aoki et al. called a "cohesive social experience" [1], but slightly varied resources may aid debate and discussion. Here we do not contradict the findings of our own earlier work, where we noted that differences can be problematic, e.g. in creating locational confusion or unclear referents. Nor do we contradict work with systems such as Aoki et al.'s *Sotto Voce,* which established coupled, shared presentation of audio information in a ubicomp system, so that users understood what had been shared, and how. Sotto Voce improved upon a state of little coupling or overlap in the moment–by–moment presentation of information, and we do not propose a retrograde step. Homogeneous content may often be useful, but we suggest three practical situations in which a limited degree of heterogeneity of content may be useful: when users have different past experiences to draw from, when they have different tools available and yet wish a shared experience, and when the designer's and the users' interest is in the ambiguous or contradictory.

As people become more engaged in a visit, and have established common references and landmarks, they are likely to draw from individual experience in the course of creative discussion. System support for this, as in the use of social recommender systems and also systems such as HIPPIE, may let each user access selected relevant information from his or her past activity. Such information may be different to others' and yet have much in common with others', and therefore may be both individual and worth sharing. In our George Square system we are beginning to explore this issue, especially with regard to learning and recommendation.

Another situation where heterogeneity of content may be productive is when users have different tools or media at hand. For example, geographic separation may force constraints on the technology available to different users, much as in the Mack Room scenario. Similarly, finite budgets of institutions or individuals may mean limited technological resources, and users may then not have bought, hired or borrowed the same equipment. Also, users may have varying perceptual abilities, as when a partially sighted person and a fully sighted person engage in a shared experience. Then, some information may be made accessible to only a subset of users, and yet the system may afford valuable shared insight and discussion.

Research such as [18] reminds us that ambiguous or contradictory information may potentially have a positive effect: "ambiguity can be frustrating, to be sure. But it can also be intriguing, mysterious, and delightful. By impelling people to interpret situations for themselves, it encourages them to start grappling conceptually with systems and their contexts, and thus to establish deeper and more personal relations with the meanings offered by those systems". The user may value and appreciate an experience that makes the familiar present–at–hand.

We emphasise that heterogeneity, like ambiguity, may be a resource for design, but it is not an end in itself. Such a feature still has to be well–designed, like any other design feature. We also emphasise that we do not suggest that shared homogeneous content is a bad idea. Instead, we suggest that small variations in content that complement a core of shared material may be an option to explore in future system designs.

Another design suggestion we offer is to make past activity across media a resource for ongoing or synchronous activity in each medium. Theory and user studies suggest that people use past activity, in all the media they have used, as a resource for interaction with each other. They also use interaction with each other as a resource for use, accommodation and appropriation of each medium. As mentioned above, recommendations may be done individually (i.e. heterogeneously with regard to users), but collective recommendations, drawing from all users' histories and presented homogenously, are feasible [28].

Lastly, we suggest that designers may consider selectively and carefully revealing differences and limitations of systems, in ways that support social interaction i.e. seamful design. We can show a person's sensed position as a spatial extent rather than as a point, for example. In recent work we have been exploring systems that let people use spatial representations such as maps and VRs to see and even take advantage of where wireless communication networks are (and are not), and where GPS positioning is poor [8]. We are particularly interested in seamful systems whose underlying infrastructural 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 infrastructure. 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 categorisation, and the processes that changed them [10]. Another potentially relevant approach is recombinant computing, as investigated in the Speakeasy project [26]. 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.

Since seams can be 'user context' too, we suggest that the way that we designers traditionally classify and isolate our system components e.g. as models of user activity, infrastructure, sensors, transducers, I/O devices, and so forth may be at odds with use and interaction. Similarly, we should not always rely on the traditional categorisation of error and uncertainty as features of the system to be hidden and reduced. In the long run, we should consider accommodation and appropriation as a process that designers contribute to by selectively revealing system structures, 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 that lets users focus on their task, contextualising a tool and interweaving it with others, even when that task involves changing the tool itself.

CONCLUSION

People often weave interactive systems into their social interaction and local environment without technological support. Designing a system so that it supports this process of accommodation and appropriation involves coupling and interconnection with other media, systems and spaces. As designers treat media more as coupled, interdependent peers, for example by supporting CVE-like remote collaboration via ubicomp systems, user interaction with heterogeneous media becomes not a new feature of interaction and use—just a more explicitly designed one.

Design often focuses on a circumscribed set of interface features, distinct from infrastructure, but people build up their understanding and use over time, relating a new system to their own use of heterogeneous systems and spaces, and the seams of infrastructure and connections between media often show through in interaction. Theoretical discussion and analysis of systems in use suggest that we move away from a treatment of one system or medium as primary. Also, a narrow focus on one mode of interaction, such as ready–to–hand or embodied interaction, may not reflect the way that present–at–hand interaction is an unavoidable part of this process of accommodation and appropriation.

In our design work we explored a combination of CVE technology with hypermedia, mobile computers and the architecture and exhibits of the Mack Room. We aimed to understand some of the detail of how people accommodate and even appropriate such heterogeneity, in and through the temporal process of social interaction, and thus weave together media into a more unified experience. A system's design and its designers influence the way that users manage the ongoing process of coupling, contextualisation and appropriation. We can support this process with rich synchronous and asynchronous awareness between users, via talk and spatial representations such as maps and VRs, with overlapping content that may vary in small but useful ways between users, and with seamful revealing of sensing, communication and structure.

More generally, we suggest that theory and design be closely linked, and can jointly feed into future designs for interactive systems. We can enrich our work, and aid the use and adoption of our systems, with design practice and theory that take fuller account of heterogeneity, seamfulness and the social and asynchronous aspects of context.

ACKNOWLEDGEMENTS

We offer thanks to our Equator colleagues past and present, especially Steve Benford, Barry Brown, Bill Gaver, Ian MacColl, Cliff Randell, Tom Rodden and Anthony Steed, to Paul Dourish, and also to the anonymous referees. Equator is funded by the UK EPSRC (GR/N15986/01). Matthew would also like to thank Mogwai (Happy Songs for Happy People), Metallica (St. Anger) and The Mars Volta (De-Loused in the Comatorium).

REFERENCES

1. Aoki, P. et al., Sotto Voce: Exploring the Interplay of Conversation and Mobile Audio Spaces, *Proc ACM CHI*, 431–438, 2002.
2. Billinghurst M et al. The MagicBook: A Transitional AR Interface. *Computer Graphics*, 25:745-753, 2001.
3. Bowers, J et al. Workflow from Within and Without: Technology and Cooperative Work on the Print Industry Shopfloor, *Proc. Euro. Conf. on Computer Supported Cooperative Work (ECSCW),* 51-66, 1995.
4. Brown, B. & Perry M. Of Maps and Guidebooks: Designing Geographical Technologies, *Proc. ACM DIS,* 246–254, 2002.
5. Brown, B. et al. Lessons from the Lighthouse: Collaboration in a Shared Mixed Reality System. *Proc ACM CHI,* 577–584, 2003.
6. Chalmers M Awareness, Representation and Interpretation, *J. CSCW* 11:389–409, 2003.
7. Chalmers M., Rodden K., Brodbeck D., The Order of Things: Activity–Centred Information Access. *Proc. World Wide Web (WWW98)*. Published as *Computer Networks and ISDN Systems*, 30:359-367, 1998.
8. Chalmers et al., Seamful Design: Showing the Seams in Wearable Computing, *Proc. IEE Eurowearable*, Birmingham, UK, 11-17, 2003.
9. Churchland P.M., Churchland P.S. *On the Contrary: Critical Essays 1987–1997*, MIT Press, 1998.
10. Dourish, P. Developing a Reflective Model of Collaborative Systems, *ACM Trans. CHI,* 2(1), 40–63, 1995.
11. Dourish, P. *Where the Action Is: The Foundations of Embodied Interaction.* MIT Press, 2001.
12. Dourish, P. et al. Your place or mine? Learning from long-term use of audio-video communication, *J. CSCW* 5(1), 33-62, 1996.
13. Edelman G, Tononi G *Consciousness: How Matter Becomes Imagination*. Penguin, 2000.
14. Falk H, Dierking L *The Museum Experience*. Whalesback Books, Washington, 1992.
15. Fleck, M. et al., Rememberer: A Tool for Capturing Museum Visits, *Proc. Ubicomp 2002*, 48–55, 2002.
16. Flintham, M, et al., Where on-line meets on-the-streets: experiences *with mobile mixed reality games. Proc. ACM CHI,* 569–576, 2003.
17. Galani A., Chalmers, M. Can You See Me?: Exploring co-visiting between physical and virtual visitors *Proc. Museums and the Web,* pp 31-40. Archives and Museum Informatics, 2002.
18. Gaver, W. et al. Ambiguity as a Resource for Design, *Proc. ACM CHI*, 233–240, 2003.
19. Grondin J. *Introduction to Philosophical Hermeneutics*. Trans. J. Weinsheimer. Yale University Press, 1994.
20. Harrison S. & Dourish P. Re-Place-ing Space: The Roles of Place and Space in Collaborative Systems, *Proc. ACM CSCW*, 67–76, 1996.
21. Koleva, B. et al. Traversable Interfaces Between Real and Virtual Worlds. *Proc. ACM CHI,* pp 233–240, 2000.
22. Leach N. (ed) *Rethinking Architecture: A Reader in Cultural Theory*, Routledge, 1997.
23. vom Lehn, D. et al. Exhibiting Interaction: Conduct and Collaboration in Museums and Galleries. *Symbolic Interaction,* 24(2):189-216, 2001.
24. Mackay, W. Patterns of Sharing Customizable Software, *Proc. ACM CSCW*, 209-221, 1990.
25. McManus, P. M. Oh yes, they do: how museum visitors read labels and interact with exhibit texts. *Curator* 32(3):174-189, 1989.
26. Newman, M. et al. Designing for serendipity: Supporting end-user configuration of ubiquitous computing environments. *Proc. ACM DIS*, 147-156, 2002.
27. Nöth W. *Handbook of Semiotics*, Indiana University Press, 1995.
28. O'Connor, M. et al. PolyLens: A Recommender System for Groups of Users. *Proc. Euro. Conf. on Computer Supported Cooperative Work (ECSCW)*, 199-218, 2001.
29. Oppermann R., Specht M. & Jaceniak I. Hippie: A Nomadic Information System, *Proc Handheld & Ubiquitous Computing (HUC)*, 330–333, 1999.
30. Orlikowski, W. Learning from Notes: Organizational Issues in Groupware Implementation, *Proc. ACM CSCW*, 362-369, 1992.
31. de Saussure, F. *Course in General Linguistics*. Trans. W. Baskin. McGraw–Hill, 1983. (Original published in 1906.)
32. Schiele B., Jebara T. & Oliver N. Sensory-Augmented Computing: Wearing the Museum's Guide, *IEEE Micro*, 21(3):44–52, 2001.
33. Underkoffler J, Ishii H. Urp: A Luminous-Tangible Workbench for Urban Planning and Design. *Proc. ACM CHI*, 386–393, 1999.
34. Weiser, M. The Computer for the Twenty-First Century, *Scientific American*, 94-110, Sept. 1991.
35. Weiser M. Creating the invisible interface (invited talk). *ACM Conf on User Interface Software and Technology (UIST)*, 1, 1994.