

Exploiting Space and Location as a Design Framework for Interactive Mobile Systems

ALAN DIX

aQtive Limited

and

TOM RODDEN, NIGEL DAVIES, JONATHAN TREVOR, ADRIAN FRIDAY,
and KEVIN PALFREYMAN

Lancaster University

This article considers the importance of context in mobile systems. It considers a range of context-related issues and focus on location as a key issue for mobile systems. A design framework is described consisting of taxonomies of location, mobility, population, and device awareness. The design framework informs the construction of a semantic model of space for mobile systems. The semantic model is reflected in a computational model built on a distributed platform that allows contextual information to be shared across a number of mobile devices. The framework supports the design of interactive mobile systems while the platform supports their rapid development.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems—*Distributed applications*; H.1.2 [**Models and Principles**]: User/Machine Systems; H.4.3 [**Information Systems Applications**]: Communications Applications; H.5.3 [**Information Interfaces and Presentation**]: Group and Organization Interfaces—*Synchronous interaction*; *Theory and models*

General Terms: Design, Human Factors, Theory

Additional Key Words and Phrases: Mobile systems, context information, virtual space, awareness, platform support, shared interaction, design framework, location-sensitive applications

Authors' addresses: A. Dix, Birmingham Research Park, aQtive Limited, Birmingham B15 2SQ, UK; T. Rodden, N. Davies, J. Trevor, A. Friday, and K. Palfreyman, Department of Computing, Lancaster University, Lancaster LA1 4YR, UK.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2000 ACM 1073-0516/00/0900-0285 \$5.00

1. INTRODUCTION

1.1 The Emergence of Connected Mobility

The increased availability of communication facilities has seen a shift in the nature of mobile computers systems and applications. Prior to the widespread emergence of mobile communication services the work on mobile computer systems has tended to focus upon the development of small devices and the forms of interaction that they afford. In particular, we have seen considerable focus on the development of techniques and approaches to overcome display limitations and the ergonomics of new interactive devices [Bass et al. 1997]. However, the development of a new class of mobile devices that link with telecommunication services to offer connections to other systems means that we need to extend our consideration of design for mobile systems. These mobile systems are particularly problematic because of the various ways in which they break assumptions that are implicit in the design of fixed-location computer applications, leading to new design challenges for human-computer interaction.

One immediate issue arising from the nature of wireless communications are delays and outages, leading to slow and unpredictable temporal characteristics at the user interface. These issues have already been addressed in some detail in our previous work on pace of interaction and practical experience in building collaborative mobile applications [Davies 1994; Dix 1992; 1995]. They are also considered in the wider study of temporal issues in HCI [BCS HCI 1997; Howard and Fabre 1998; Johnson 1997; Johnson and Gray 1995].

Another issue that has been of considerable importance is the development of context-sensitive devices that have a distinct awareness of their location and other devices [Davies 1994; Fickas et al. 1997; Long et al. 1996; Want et al. 1995]. However, the treatment of these issues has, so far, been predominantly focused on the design of specific devices and applications. In this article we are aiming to produce a broader view of context and location-aware computation. In particular, we will present a framework to support the design of interactive mobile systems based on an understanding of location within these systems.

1.2 Context and Mobility

In mobile systems, as in other areas of human-computer interaction, it is not sufficient to focus on the specific interface of a device. The device operates within a broader context. This context includes the network and computational infrastructure, the broader computational system, the application domain, and the physical environment.

Each of the different contexts represents a different part of the design space within which mobile systems must be placed, and the features of infrastructure, system, domain, and environment all suggest trade-offs that developers must address in realizing mobile interactive systems. Currently,

designers undertake this trade-off with little support or guidance, as little is known about the design space into which mobile applications are placed.

The design framework developed in this article focuses on the context-sensitive nature of mobile devices in the design and development of cooperative mobile systems. We wish to chart the design space for mobile systems to allow developers to consider the properties of mobile systems under construction and how they may be related to other applications and systems. In developing this design framework we focus particularly on location, as ideas of space and location are of paramount importance in any consideration of the context of these systems.

In undertaking this task we wish to build upon the research lessons from research on temporal issues in interaction. This research community has successfully created frameworks to improve the general understanding of temporal aspects prior to design and construction of new forms of technology. We wish to apply a similar tactic by moving from a theoretical consideration of the context of mobile systems, through an examination of location and the development of a framework for the design of mobile systems, to the development of models to support these systems. Our particular interest is in the way in which the interactive nature of cooperative mobile systems drives the development of supporting infrastructures, and our endpoint is a computational infrastructure we have developed for these systems.

1.3 Structure of this Article

In developing our design framework we focus particularly on the importance of location and space in mobile systems. However, it would be unwise for us to suggest that location is the only manifestation of context in mobile systems. In order to place our framework within the broader picture of the design of interactive mobile applications we begin, in Section 2, by considering the nature of the context in which interaction with mobile applications takes place. This broader consideration of context acts as a general backdrop for our focus on location as a means of designing for context-sensitive mobile systems.

In Section 3, we address the need to consider location in terms both of the physical space in which a mobile device exists and of the virtual models of space exploited by applications. The importance of location underpins the development, in Section 4, of multiple taxonomies of location, mobility, population, and device awareness. This conceptual analysis allows us to exploit location as a means of understanding interactive mobile applications. In Section 5, we build upon these taxonomies to suggest a simple semantic model of space that allows us to more generally represent and reason about the location of devices. We then use this semantic model to develop a simple computational model and supporting infrastructure that allows the understanding of location to be conveyed across a number of devices working in tandem to realize cooperative mobile applications.

Our computational model of location is realized on top of a distributed architecture and infrastructure. In Section 6, we discuss this architecture

and general architectural issues for contextual systems. This section ends with a short description of the distributed architecture that instantiates many of the principles espoused in this article by realizing a computational model of space that is accessible from a range of mobile devices.

2. THE CONTEXTUAL NATURE OF MOBILE SYSTEMS

In considering the design of mobile systems we wish to focus particularly on the situation where mobile devices behave differently and offer different interaction possibilities depending on the particular context in which the system is being used. For example, in the development of mobile multimedia guides, such as the systems at Georgia Tech [Long et al. 1996] and the Lancaster GUIDE [Davies et al. 1998], the information presented to the user and the interaction possibilities are strongly linked to the location where the device is being used. In these cases interaction is no longer solely a property of the device but rather of the *device in context*. While location is often the principal determinant used to represent this context and the main focus of this article it is worth briefly considering the general importance of context in mobile systems and why understanding and modeling this context is important.

A considerable amount of research surrounding the development of mobile devices has obviously focused on the portable nature of these devices and the technical problems of implementation. Mobile computing devices represent real technical challenges and have always stretched the state of the art in terms of displays and interaction devices.

The emergence of mobile telecommunication standards such as GSM and the increased availability of these services have also led more recently to the development of devices that provide mobile access to on-line services (e.g., the Nokia communicator). This merging of computer and communication facilities allows the development of systems that provide immediate on-line access to information. These portable networked devices have also been combined with the use of GPS technologies to develop portable devices that are aware of their position [Long et al. 1996].

The current generation of portable devices has an awareness of their setting and an increased ability to access network resources. This means that we need to balance the current consideration of the interaction properties of individual devices with a broader consideration of the context of use. The importance of context to interactive systems is not unique to mobile devices and is already reflected in research in ubiquitous computing, wearable computers, and augmented reality [Aliaga 1997; Weiser 1991; 1993], as well as more recent work at MIT on the development of devices that exploit context to provide an ambient awareness of interaction [Ishii and Ullmer 1997].

The term “context” has become problematic for interactive systems development and is itself the subject of some debate. One aim of the growing focus on context is to allow the highly situated nature of interactive devices to be reflected in the design of systems. This focus on the

situated nature of these devices reflects their growing acceptance and the need to allow them to closely mesh with existing practices, and mirrors previous work in the development of interactive systems within CSCW [Hughes et al. 1995]. Rather than engage in this much broader discussion we wish to concentrate on location. However, in the remainder of this section we wish to briefly characterize some of the broader debate on context that is relevant to mobile systems.

2.1 Context in the Design of Mobile Systems

In order to reflect the broader role of context in our consideration of location we wish to unpack what we might mean by the term context and how we may exploit it to determine different interaction possibilities for mobile systems. In the following sections we consider some of the ways in which context has played a key design role in the development of distributed mobile applications. We outline some of the different forms of context that influence interaction with mobile systems before we consider the central role of location in Section 3. Our consideration of context moves from the nature of the underlying *infrastructure context* to consider the overall *system context*, the broader application *domain context*, and finally the actual *physical context*.

2.1.1 Infrastructure Context. The interaction offered by mobile applications is not solely dependent on the particular features of the mobile devices used. Rather it is a product of the device and the supporting infrastructure used to realize the application. The impact of the properties of the supporting distribution infrastructure on different styles of interaction has been discussed in CSCW and HCI [Greenberg and Marwood 1994]. In mobile systems the nature of the infrastructure is even more likely to change as the application is used, and the sort of service available may alter dramatically. This variability in the infrastructure can dramatically affect interaction, and it is essential that interaction styles and interfaces also reflect the state of the infrastructure.

In essence, the user interfaces to mobile applications must be designed to cope with the level of uncertainty that is inevitably introduced into any system that uses wireless communications. Consider our experiences in the development of an advanced mobile application used to support collaborative access to safety-critical information by a group of field engineers [Davies 1994]. If one of these engineers becomes disconnected from the group as a result of communications failure then it is vital that the remaining users' interfaces reflect this fact. For example, if an engineer is about to work on a cable it is important that the system either (a) correctly reflects the current state of the cable or (b) clearly shows that the information is not current. If this does not hold, the engineer could easily touch a live cable with potentially fatal consequences. Reflecting this information requires interaction between the application's user interface and the underlying infrastructure via which failures will be reported. In addition, if the information being manipulated is replicated by the distributed-

systems platform the validity of each replica will clearly be important to the engineers. In this case the user interface needs to reflect this platform information.

2.1.2 System Context. Most advanced mobile applications are distributed in nature. Rather than functionality residing solely within a single machine (or device) it is spread across the system as a whole. This means we need to consider the interactive properties of the system in terms of the distributed nature of the application. This is particular true when we consider issues of pace and interaction [Dix 1992]. For example, rapid feedback is an accepted premise of HCI design, and many applications provide direct-manipulation interfaces that rely on rapid feedback. The development of distributed applications and the impact of the delays inherent in the technical infrastructure have already seen a reconsideration of feedback [Dix 1995; Ramduny and Dix 1997]. The need to consider the overall functionality of the application and to design structures that provide appropriate access to different levels of functionality is amplified in the case of mobile applications where the infrastructure may vary dynamically as the application is in use.

Consider, for example, the development of caching strategies for field engineers who will only ever be examining or servicing units within a subregion of a particular area. The choice of the appropriate location to cache information will depend on the required feedback, the safety-criticality of the information, the speed and reliability of different parts of the network, and on who else is likely to be using and updating the cached information. A local writeable cache would improve the feedback for an individual user. However, it may also make it appear that the user's data changes have been reflected in the system as a whole, when, in fact, the connection between the cache server and the data server is broken and other users are seeing dangerously out-of-date information.

Another aspect of this system context is the extent to which a device is aware of other devices in its vicinity and, related to this, the extent to which an application is aware of other applications. This is important, partly because such devices can adversely affect one another as they contend for resources, but more importantly, because combinations of devices may be able to offer more advanced services to the user. This can lead to a planned or accidental emergent behavior of the devices as a group, which is not defined in any individual device. One example of this is onCue (see Section 6.3), which only offers certain services when particular software is available on the local machine.

2.1.3 Domain Context. As well as addressing the infrastructure and system issues discussed above, distributed mobile applications need to consider the semantics of the application domain. The situated nature of advanced multimedia applications is such that design needs to explicitly identify the nature of the work being supported and the practicalities of this work. In doing so, developers need to consider the relationship between the mobile devices and their users and how this can be used to determine

the nature of the interfaces presented. In the case of mobile applications the normal design considerations are amplified by the need to consider the limited interaction facilities of mobile devices.

Mobile devices are intended to be readily available and useful to the community of users being supported. As a consequence we need to consider the highly situated nature of this interaction. Developing a clear understanding of what people do in practice and the relationship with technology is essential to informing the development of these applications. The relationship between users and mobile technology is still unclear, and few studies have taken place that consider the development of mobile cooperative applications [Davies 1994].

For example, we may choose to exploit the personal nature of these devices to associate mobile devices with users. This allows us to tailor applications to allow them to be sensitive to the identity of the user of the device. This information may be exploited along with additional contextual information (e.g., location) to present appropriate information. One example of this would be a particular doctor visiting patients within a hospital. At a particular bed, who the doctor is and his or her relationship to the patient in the bed may determine the information presented. Contrast this situation with the development of a museum guide where the devices need to be considered as general purpose and where no information is available about the relationship between users and the artefact being described.

Another aspect of the domain is the level of trust and mutual awareness between participants in collaborative interactions. This is particularly important if devices are to be used to identify users and potentially make information about their location and what they are doing available to others. In this case a consideration of the issues of privacy and the need for some management of privacy is essential [Harper 1992].

2.1.4 Physical Context. Finally, mobile computer systems are likely to be aware of, or embedded into, their physical surroundings. Often this is because they are embedded in an application-specific device, e.g., in a mobile telephone or car. In these situations the computer system is mobile by virtue of being part of a larger mobile artefact. This context can and does affect the application interface, e.g., the telephone directory within a mobile telephone can be very different from one in an independent PDA. Another example is a car radio (now often computer-controlled) which has different design considerations to a static radio including the need to automatically retune as the car travels between local radio areas and transmitter zones. Because the computer systems are embedded into application-specific devices, they may also be aware of their *environmental context*, e.g., the speed of the car. Some of this sensory information may be used simply to deliver information directly to the user, but some may be used to modify interface behavior, e.g., in a tourist guide, increasing text size in poor lighting conditions, or, in a car system, limiting unimportant feedback during periods of rapid maneuvering.

Table I. Taxonomy of Context

Context	Relationship with	Issues
infrastructure	network bandwidth, reliability, and display resolution	variability of service, user awareness of service, liveness of data
system	other devices, applications, and users	distributed applications, pace of feedback and feedthrough, emergent behavior
domain	application domain, style of use, identification of user	situated interaction, personalization, task and work studies, privacy
physical	physical nature of device, environment, location	nature of mobility, location-dependent information, use of environmental sensors

2.1.5 Context in Context. Context is largely about relationship, and the four forms of context considered in this section have focused on the relationship between an interaction device and surrounding elements. These forms of context, the associated relationships, and the issues raised are summarized in Table I.

Although each of the different kinds of context discussed in this section are worthy of further study in their own right, this article is predominantly concerned with physical context and the use of location in determining this. In the following section we discuss the central role of location and physical context for mobile systems. This is partly because of our desire to produce a computational infrastructure to support location-aware applications, but also reflects the dynamic changes of location as a unique feature of mobile systems. However, as you would expect, these various forms of context are closely related, and we will see elements of various other kinds of context throughout this article.

3. THE IMPORTANCE OF LOCATION IN UNDERSTANDING CONTEXT

Clearly, the very idea of “mobility” demands an understanding of location, and one of the unique aspects of mobile devices is that they can have an awareness of the location within which they are being used. Furthermore, this location information may be exploited as a means of understanding the overall context within which the system is placed. Essentially, location becomes a useful indexing device from which to infer the overall context influencing the mobile application: in order to ask “what devices are near this device” (system context) we need to know the location of this device and others; it only makes sense to measure the environment (physical context) when a device is physically located in space. Furthermore, the need to be contextually aware in other ways depends to a large extent on the mobility in space of devices. If devices are spatially static, many aspects of their environment are also static, or at most slowly varying.




Fig. 1. A device situated in space.

3.1 Location and Space

Any notion of location puts the device within some form of space. The space within which a device is located may also contain other devices and users with which the device may interact. A device involved in a mobile system can be considered as

- having location in the space,
- having an effect on the space (and devices and users within it), and
- being subject to influencing events from the space (and devices and users within it).

Essentially, devices are situated and embedded within a space, and their interaction is mediated through this space (Figure 1). Consequently, understanding the nature of their location in that space is key to understanding the nature of the mobile system being designed and provides a means of reflection on the context.

If the device in Figure 1 and other devices were at fixed locations, then the nature of these interactions would be one of configuration. However, the interesting and challenging nature of mobile interfaces is the changing nature of these relationships. So, to have an overall model of spatially situated interaction, we need to understand the following:

- location in space (of the device and other bodies)
- mobility through space (of these)
- the kinds of bodies populating the space (which the device may interact with)
- the awareness (of the device) of these other bodies.

In Section 4, we will develop each of these, but before we consider the different relationships between a device and space it is worth reflecting on what we may mean by space.

3.2 Of Real and Virtual Worlds

Focusing on a device situated in space is only the start of any consideration of the importance of space to mobile systems. Consider a purely physical device like a lawnmower: it inhabits the real world and can be used to affect the real world. *It exists only in a single space*, and its relationship is only with the physical world it inhabits. Its location is unique to that space, and its influences and effects are only through the physical space within which it resides.

However, this is not the case for many mobile systems and our interaction with mobile systems. Computers and mobile devices are different in that we can consider their existence and presence in terms of many spaces. They can be thought of as simultaneously inhabiting a real world and some form of virtual world (or indeed multiple virtual worlds). A computational device's ability to exist in the physical world while also having an existence in an electronic or virtual worlds is significant to any consideration of mobility.

3.2.1 *The Emergence of Virtual Space.* Our consideration of an existence in an electronic world extends beyond the realm of virtual reality, and we would suggest is equally evident as we surf the Web, use FTP to access remote files, or even simply explore our own file system. In all of these cases we are in a sense inhabiting virtual space. Even the vocabulary we use reflects this: we “visit,” “explore,” “go to,” “navigate”; our Web browsers even have a button to go “back.” This turning to virtual spaces and spatial approaches generally grows from the use of spatial metaphors and techniques to represent information and action in electronic systems. One of the early examples of the use of spatial metaphors includes the use of a rooms metaphor to allow the presentation of information [Henderson and Card 1985]. From these early spatial approaches we have seen concepts of spatial arrangement exploited in the development of desktop conferencing systems such as Cruiser [Root 1988] and more generally in the work of Mediaspaces [Gaver 1992].

The recent development of cooperative systems in CSCW has also seen a growing application of concepts drawn from spatial arrangements. These include the development of *groupkit* to form *teamrooms* [Roseman and Greenberg 1996], the emergence of the *worlds* system [Fitzpatrick et al. 1996], and the use of a notion of places to support infrastructure [Patterson et al. 1996]. This exploitation of virtual spaces is most notable in the development of shared social worlds existing solely within the machine [Benford et al. 1995]. However, the use of space and virtual spaces has not been isolated to an existence solely within the computer, and a number of researchers have considered how space and location can be considered both virtually and physically within the development of applications. This is most evident in the augmenting of existing physical spaces to form digital spaces populated by electronically sensitive physical artefacts (or tangible bits) [Ishii and Ullmer 1997] that are sensitive to their position within both physical and virtual space.

3.2.2 *Combining the Real and the Virtual.* The work in tangible bits undertaken by Ishii and Ullmer [1997] represents the start of a trend to interweave real and virtual spaces. This work exploits the combined use of a number of devices within a space so that their physical manipulation can be used to generate a computational (or virtual) effect. Various other strands of recent research have explored these boundaries between the physical and virtual including wearable computing and augmented reality. One example of this has been the explicit development of boundaries that span between the physical and the virtual [Benford et al. 1998]. We would suggest that this interplay between the real and the virtual is at the core of the design of cooperative mobile applications, as devices and users have a location and presence that is both virtual and physical and since each is available to the computer application.

This interplay between the real and the virtual provides a starting point for the development of our taxonomy. A direct result of the need to recognize this coupling is that many of the categories we will consider for understanding mobile and context-aware computation have counterparts in both the real physical world and the virtual electronic world. There are important differences—the virtual world does not always behave in ways we have come to expect from the physical world—and designers and developers often exploit these differences.

In particular, even the object of interest for mobile computation may have a physical or virtual existence depending on the nature of the application. At one extreme we have hand-held GPS systems that simply tell you where you are in physical space—perhaps these do not even rank as mobile computation. At the other extreme there are agents that only have an existence within the virtual world, e.g., Web crawlers or the components within CyberDesk [Wood et al. 1997]. Between these we have more complex physical devices, such as the PDA, which both have a real-world existence and serve as windows into virtual space (especially when combined with mobile communications).

To some extent context-awareness can be seen in hardware-configuration architectures such as Plug&Play. In both, the emphasis is on self-discovery and automatic reconfiguration of software to reflect the current hardware. The main difference between these and “real” context-aware applications is one of time—the rate of reconfiguration required by, say, Plug&Play on a personal computer may only be a few times during the lifetime of the device. However, as interdevice protocols and standards such as BlueTooth become more prevalent, hardware reconfiguration will become far more frequent and will be an important source of contextual information for higher levels of the interface.

As we consider taxonomies and then models of space and location in Sections 4 and 5, we explicitly consider both physical and virtual location and endeavor to construct theoretical and computational models which encompass both.

4. A DESIGN FRAMEWORK FOR MOBILE SYSTEMS

The core of our framework for understanding the design of mobile systems is a series of taxonomies that considers the relation between different devices and the spaces they inhabit using location as a starting point for this consideration. Rather than seek to understand all senses of mobility for all potential forms of space, we will particularly focus on the physical space of devices as a distinguishing feature. However, we will also draw on examples of the virtual where they are instructive to highlight the coexistence of these two forms of space and the issues of mobility that may exist in both. Although it is worth stressing that our understandings of this virtual space is still under development. We will also return to the issue of the real and the virtual when we consider the development of a model of space to support mobile systems.

As described in Section 3.1, the core of our framework is an understanding of

- location in space (of the device and other bodies)
- mobility through space (of the device and other bodies)
- the kinds of bodies populating the space (which the device may interact with) and
- the awareness (of the device) of these other bodies.

In this section we will develop taxonomies of location, mobility, and population, in turn, then finally a decomposition of types of device awareness.

4.1 A Taxonomy of Location

Mobility makes us think automatically about location, the way in which this sense of location can be understood in the system, as well as how changes in location can affect the system. Any simple mobile device will have a physical location in space. It is important to understand the nature of this location, and how the developers of interactive mobile applications may exploit this understanding. In this section we wish to consider what we might actually mean by location in space. This brief exploration is more than a mere issue of terminology, as developing an understanding of what we actually mean by location represents a consideration of one of the core design concepts in the production of mobile systems.

Looking at the spatial dimension, we can see there are some devices (e.g., GPS-based map systems) where the exact Cartesian position in 2D or 3D space is important in defining a sense of absolute physical location. For others a more topological idea of space is sufficient in understanding position, and in these cases location is considered not in an absolute sense but in relation to other objects or sensors. For example the Lancaster GUIDE system is based on radio cells roughly corresponding to rooms and sections of Lancaster Castle, and the CyberGuide [Long et al. 1996] system

Table II. Location in Different Kinds of Space

Space	Real	Virtual
Cartesian	GPS	VR
Topological	room	hypertext

at Georgia Tech. shows visitors around the GVU laboratory by altering its behavior, depending on what item of equipment is closest.

In Section 5 we will look at more formal models of space in greater detail. For now we will just use this two-way distinction between Cartesian space and topological space and consider location in these terms. As we discussed in Section 3.2, it is important to consider location in both a physical and a virtual sense. If we consider ideas of virtual location, e.g., position within a hypertext, we see that we may have similar ideas of space within the electronic domain. This consideration of location provides us with the simple taxonomy shown in Table II.

Note that these are not mutually exclusive categories: an item in a room also has a precise longitude and latitude, and a computational entity may have an existence in one or more virtual spaces as well as physical space. Indeed, possibly many of the most interesting interaction possibilities occur when these different ideas of location are linked. For example, moving a display up and down in physical space could be used to change the display of hypertext help system for the maintenance of a piece of machinery. Similarly, in an aircraft cockpit, setting the destination city (topological destination) instructs the autopilot to take an appropriate course in Cartesian space/time. This interplay between the real and the virtual is central to the development of augmented reality spaces where the movement of devices within a space may manifest in effects that are both real and virtual. These spaces only work because the location of the device can be controlled in virtual and physical space and because its effects provide alterations to either the physical or virtual space.

4.2 A Taxonomy of Mobility

Our core concern in the development of our design framework is the issue of mobility and its implications for how we understand human-computer interaction. In the previous section we considered how the issue of location can be unpacked to provide understanding in both a physical and a virtual sense and how the nature of the space affects our consideration of location. In this section we wish to focus on how to understand mobility and what potential design issues may emerge from a more detailed consideration of mobility.

Devices may be mobile for a number of reasons. They may be mobile because they are carried around by users (as with a PDA or a wearable computer), because they move themselves (robots), or because they are embedded within some other moving object (a car computer). Furthermore, a number of different devices may be spread within our environment so

that they become pervasive, as in the case of an active room such as the ambient room suggested by Ishii and Ullmer [1997]. The issue of pervasiveness is itself a rather thorny one in that it is not clear what constitutes pervasiveness in terms of devices and how this relates to previous discussions surrounding ubiquitous devices. Ubiquitous computing has focused on the backgrounding of the device and the computer essentially “disappearing” into the environment. For us the issue of pervasive devices has less to do with the devices fading into the environment and more to do with an expectation that devices are normally available. Pervasive computing is intimately bound up with the interrelationship between different devices and the expectation that these devices can work together to provide some form of shared functionality. An active room is active because it contains a number of devices which when they work in unison provide some function. Essentially, we are seeing a number of computing devices which in cooperation provide some functionality. Some of these devices may be mobile, but many are not. Consider, for example, the layout of radio-LAN based stations for the GUIDE tourist information system. These base stations have a fixed location in Lancaster, but are the source both of location and other information displayed on mobile devices. Neither the base stations nor the mobile devices can function by themselves, but together they allow the space to offer a pervasive computing facility.

We can disentangle the different levels of mobility into three dimensions that are used in Table III to classify examples of mobile systems.

First, we can consider the *level of mobility* within the environment, which divides into three main categories:

- fixed*: that is, the device is not mobile at all! (e.g., a base station fixed in a particular place)
- mobile*: may be moved by others (e.g., a PDA or wearable computer that is carried around)
- autonomous*: may move under its own control (e.g., a robot)

The devices *relation to other devices* or its environment provides our second dimension and can also be divided into three different categories:

- free*: the computational device is independent of other devices, and its functionality is essentially self-contained.
- embedded*: the device is part of a larger device
- pervasive*: the functionality provided by the device is essentially spread throughout the environment.

These separations do not consider the nature of the device and the sort of functions it may afford. The physical design of the device itself is an issue that needs to be considered carefully, especially in terms of existing traditions of aesthetic and practical design. The consideration of these features is beyond the scope of the framework and taxonomy we wish to present here, which focuses on the development of the device.

Table III. A Taxonomy of Different Levels of Mobility

		Personal	Group	Public
Free	Fixed	office PC	Liveboard	computer lab.
	Mobile	PDA	tour guides	tour guides
	Autonomous			factory robot
	Fixed		active fridge	ATM
Embedded	Mobile	wearable devices	car computer	shopping cart
	Autonomous		auto pilot	monorail
	Fixed		active room	active room
Pervasive	Mobile			Star Trek
	Autonomous	Web agent	HAL	Web crawler

As a final part of our taxonomy we can reflect the cooperative nature of advanced mobile applications by considering the extent to which the device is bound to a particular individual or group. We have three classes for this too:

- personal*: the device is primarily focused on supporting one person
- group*: the device supports members of a group such as a family
- public*: the device is available to a wide group

We do not suggest that these categories are absolute but rather provide them as sample equivalent cases of utility to designers. All the categories have gray cases, but perhaps this last dimension most of all. In particular we should really consider both the static and dynamic nature of how these categories are applied. For example, we could classify a computer laboratory as “public,” but of course, after logging in, each computer becomes personal. We will return to these dynamic aspects when we look at how devices can become aware of their users.

In fact, the “group” category really covers two types of device. Some, like a liveboard actually support a group working together. Others, like an active refrigerator (which allows messages to be left, email browsing, etc.), may primarily support one person at a time but are available to all members of a family. In-car computer systems exhibit both sorts of “groupness”: they may perform functions for the benefit of the passengers of the car as well as the driver, and the exact mix of people from within the family (or others) in the car may vary from trip to trip.

Some of the examples in Table III are clear, but some may need a little explanation. The “Star Trek” reference is to the computer in Star Trek that responds to voice commands anywhere in the ship, but does not actually control the ship’s movements. This pervasiveness of interaction is also evident in the work on ubiquitous environments developed by those at PARC [Want et al. 1995]. Here the computational infrastructure is considered to be continually available. In a similar vein, we put HAL (the computer from the 1960’s movie *2001: A Space Odyssey*) in the group

category, as it has a small crew, but this is exactly one of the gray distinctions in constructing a taxonomy of this form. Our reference to “shopping cart” refers to the development of smart supermarket trolleys that allow shoppers to scan the barcode of items as they are added to the trolley and keep track of your purchases to enable a fast checkout. Often these require the insertion of a shopper identification, in which case they become dynamically personalized.

Notice there are various blank cells in this taxonomy reflecting our use of it as a means of charting the design space for interactive mobile devices. Some of these blanks represent difficult cases where there may not be any sensible device. For example, a fixed-pervasive-personal device would have to be something like an active hermit’s cell. In fact, the whole pervasive-personal category is problematic, and the items “Web agent” and “Web crawler” in the final row may be better regarded as virtual devices of the free-autonomous class.

Other gaps represent potential research opportunities. For example, what would constitute a free-mobile-group device? This would be a portable computational device that supports either different individuals from a group, or a group working together—possibly an electronic map that can be passed around and marked.

As we suggested at the outset of our discussion of the design framework most of the examples are of physical devices. Virtual devices may also be classified in a similar way; for example, Word macros are embedded-mobile (or even autonomous in the case of macro viruses!) as are Java applets. The only virtual devices in Table III are the items “Web agent” and “Web crawler” in the final row which, as we have said, may be regarded as virtual devices of the free-autonomous class. This ambiguity is because any virtual device or agent must be stored and executed upon a physical computational device and the attributes of the physical device and virtual device may easily differ. For example, a PDA may contain a diary application. This is mobile by virtue of being stored within the PDA (a virtual device embedded within a physical device). However, if the PDA is used as a Web browser it may execute a Java applet that is a form of virtual agent embedded within a Web page (a virtual embedding in a mobile artefact). That is, we have an embedded-mobile-public virtual agent temporarily executing on a free-mobile-personal device! This dual presence in multiple contexts is both the problem and the power of virtual environments and will require significant further research to resolve.

4.3 A Taxonomy of Population

In addition to having a location in a space and exhibiting some degree of mobility, devices also need to be aware that they populate a space and need to reflect the coupling with the space they inhabit depicted in Table I. This awareness may include both the physical nature of the space (light, temperature, weather) and the electronic environment (network state, available memory, current operating system). A simple example of virtual

Table IV. Examples of Bodies within the Environment

	Physical (e.g., car computer)	Virtual (e.g., active web page)
People	current driver of car	visitor at Web page
Devices	other cars	running applets
Objects	roadside fence	other pages on the site

devices understanding the space they inhabit are Javascript Web pages that run different code depending on the browser they are running on.

Spaces are normally populated with a range of different devices. Within the physical and virtual spaces of a device there may be other computational devices, people (including the user(s) of the device), and passive objects such as furniture. These may be used to modify the behavior of the device. For example, in CyberDesk “ActOn” buttons are generated depending on what other applications are available and the types of input they can accept [Wood et al. 1997]

We can consider the issue of population in terms of the sorts of *bodies* that populate a space and the different spaces they populate. Table IV gives examples of items in the environment that may be relevant for a mobile or context-aware device and the bodies that may populate the space. In order to illustrate the development of this taxonomy, we have taken a car computer and an active Web page as two simple running examples.

4.4 Measurement and Awareness

Each of the three taxonomies developed in the framework relies on devices having an awareness of the surrounding space and using this as a resource in interaction. The central role of awareness of the surrounding environment and how this awareness is conveyed to others is an issue of some sensitivity in design. For example, in the case of active badges the issue of awareness of users and how this may be applied became embroiled within a discussion of privacy [Harper 1992]. This may become even more problematic in the case of multiple devices that display an awareness of others. Consider the suggested “fun” interest badge devices offered by Philips in the development of its visions of the future design study or the “Meme tags” developed at MIT [Borovoy et al. 1998]. These badges are programmed with a set of profiles for people and are intended to light up when you meet someone else with a compatible profile. The social acceptability of this form of device may well become a significant issue in determining their success and the general acceptance of devices of this form. For example, they have proven successful in conference settings [Borovoy et al. 1998].

In order to have an awareness of their environment, devices must be able to detect or measure significant attributes (e.g., we have mentioned their location, environment, other devices, people, and things). The discrete nature of computation means that both specification and implementation of computer systems tends to focus on events that occur at specific times. However, most of the contextual information above is of a different kind:

Table V. Types of Measurement and Examples

		How	
		Direct—Own Sensors	Indirect—Told by Others
What	own attributes (self)	GPS	PDA—location set
	other bodies	proximity sensors	object registration
	other (third party)	proximity sensors	Recommender

status phenomena, i.e., they are things which constantly have a value that can be sampled. The translation of status phenomena into events is problematic and is often done “accidentally” within systems, with the consequent probability of errors. Status-event analysis is a collection of techniques that lay equal weight to events and status. In particular, one strand of previous work in this area has looked in detail at the ways in which an active agent can become aware of a status change [Dix and Abowd 1996; Ramduny et al. 1998]. In short, these reduce to finding out *directly* by its own sensors or *indirectly* via another agent (human or electronic). For example, a car with a built-in GPS sensor can detect its position directly and thus give directions to the driver, but a simple PDA may need to be told of the current location by its user in order to adjust time zones. Other computational agents may also be important sources of information about themselves (as in the case of onCue, where objects register themselves with the system) and about other parts of the environment (e.g., recommender systems, which say “others have visited here”).

This leads to the two-way table (Table V). The first axis is *how* a device finds out about contextual information: *directly* using own sensors measurement or *indirectly* told another device/user. The second axis is *what* is being discovered: the device’s own attributes (location, etc.), the attributes of another device, and, in the case when told indirectly, whether that device is telling of its own or a third-party’s attributes.

Items in the environment (people, devices, objects) are particularly difficult: not only may they change their attributes (position, etc.), but also the configuration of items may change over time (e.g., people may enter or leave an active room). This leads to three levels of awareness. We will look at these with the example of a car computer:

- presence*: someone has sat down in the driver’s seat, but all the car can tell is that the door has been opened then closed
- identity*: the driver enters a personal pin number, and the car can then adjust the seat position for the driver
- attributes*: the car detects from the steering behavior that the driver is getting drowsy and sounds a short warning buzzer.

Notice how, in this example, presence was not detected at all; identity was informed by the driver, but the sleepiness of the driver was detected directly. In other cases different combinations of detection or informing may be found. Security systems often have ultrasonic sensors to tell that

Table VI. Taxonomy of Device Awareness of Self and Other Bodies

	Self		Other Bodies	
	Direct	Indirect	Direct	Indirect
Presence	implicit	n/a	proximity sensor	announcement
Identity	Intel chip!	DHCP	Identifying attribute	announcement
Internal Attributes	implicit	n/a	only if explicitly represented in external attributes	announcement
External Attributes	autosensors, e.g., screen Pantone color matcher	Windows message “did that screen setting work?”	sensors	autosensing + announcement

someone is near (presence). Similarly, the car could be equipped with a pressure sensor in the driver’s seat. Active badges, video-based face recognition, or microphones matching footstep patterns can be used to tell a room who is there and hence play the occupant’s favorite music and adjust the room temperature.

These examples are all about detecting people, but the same things occur in other settings. In the virtual world an agent may need to detect the same levels of awareness: presence—whether any other applications are running; identity—if so what they are (e.g., Netscape); and attributes—what Web page is currently being viewed. Also, physical devices may detect one another, e.g., allowing several people with PDAs to move into “meeting” mode. In fact, awareness models that do just this form of detection within the virtual world abound [Rodden 1996].

Table VI summarizes these various factors laying out awareness levels against the “what” from Table V. Differences between direct/indirect measurement are drawn out where relevant. Perhaps most interesting is the “presence” row. There is no need for a device to measure its own presence—a computational equivalent of *cognito ergo sum*. Also, in the attributes, we have distinguished internal attributes (memory state of device) from external ones (sound coming from the speaker, position in space). Again this is because, a device is implicitly able to be aware of the former (although may not be in practice), whereas external attributes need some form of physical sensors. This state of affairs is reversed when looking at other bodies. Note that in this table the word “announcement” means some sort of directed or broadcast communication between the other body and the device.

In all the cases considered, detection and measurement may vary in accuracy: perhaps a box was put onto the car-seat pressure sensor, or the driver lied about her identity, or the ultrasonic sensor cannot tell whether there is one person or more. There will also typically be some delay, especially when indirect means are used, which is especially problematic if

the attribute being measured changes rapidly. Thus actual detection is a trade-off between accuracy, timeliness, and cost. Depending on the outcomes certain adaptations may be ill advised (a car wrongly identifies its driver and adjusts the seat thinking the driver is short, the real driver is quite tall and ends up squashed behind the steering wheel). The fidelity of awareness is very closely tied to the demands of the application and represents a genuine trade-off between the cost of measurement, the nature of the measurement, and the importance of accuracy in the awareness information.

In developing the framework in this section, we have explored the overall design space and suggested some ways in which we might characterize it. This characterization forms the basis of the development of a model of space that supports mobile devices in maintaining an awareness of others reported in the following section. This model builds upon the taxonomies of location (Table II) and bodies (Table IV). This focus is partly pragmatic and partly intrinsic: we cannot computationally model mobility until we model location; we cannot model awareness of other bodies in close locations until we have modeled bodies and their location. So, for the purposes of this article, the taxonomies in Tables III, V, and VI will inform our discussion, but will not be explicitly represented.

5. DEVELOPING SUPPORTING MODELS

The focus on the characterization in the previous section has been a sense of location and mobility in space. In developing this characterization we have concentrated on physical space while suggesting that a significant feature of the interactive nature of mobile systems is that they tie together different forms of virtual space without elaborating on the nature of these spaces. One reason for this is that while considerable agreement exists on the basic structure and nature of physical space a similar general model of electronic spaces has yet to emerge.

To address this problem we have developed a general model of space that supports mobile devices in maintaining an awareness of others. Not surprisingly, ideas of space and location are of critical importance in developing such a model. In order to adapt itself to its location, a mobile device needs to be able to ask:

- (1) Where am I?
- (2) What else is nearby?
- (3) How should I behave in the light of (1) and (2).

And in the case of autonomous devices, a mobile device needs to be able to ask

- (4) How do I get to where I want to be?

We will concentrate on the models of location needed to answer the first two questions, as the final two questions are highly application specific.

Actual Space and Represented Space. In practice, a device does not access the “actual” location of itself or other objects directly; instead it accesses some computational representation of location held by itself or some sort of location service. Some form of transducer relates the “actual” space and the representation of the space. This actual vs. representation is not just an issue for physical space, but also virtual space. An awareness mechanism on a Web server may tell you about other current visitors to the site based on a site login/logout. However, some of the current “visitors” may have simply omitted to logout before going to pages on another site. Thus, the server’s representation of the virtual location of those users is not their actual location in this virtual space.

To represent this separation between the actual space and representation of the space we need two kinds of model:

—a semantic model which can be used for both the actual space and the computational representation of the space and

—a computational model which is part of the run-time architecture.

The semantic model gives a common meaning to the actual and representational space and, thus, allows us to discuss issues in the mapping between them including the fidelity of that mapping.

Although we need to deal with different kinds of space, if they are suitable for questions of type (1) and (2), they must share some idea of location and some idea of *nearness*. There are several mathematical models of space that are informative as well as implicit models of space in various awareness models. In both types of model we will find explicit representations of *nearness*. We will briefly review these existing models and use these together with the taxonomies from Section 4 to inform our construction of a semantic model, which in turn, will be instantiated in our computational model.

5.1 Existing Models of Space and Awareness

We all feel we have some knowledge of ordinary physical space, and those with a scientific background are used to encoding this in the x,y,z-coordinates of Cartesian geometry. The Cartesian view of physical space allows a unique labeling of space and allows us to understand the relationships between locations in terms of their coordinates alone. Scientifically it has been of tremendous importance, and practically it enables global navigation and the civil construction. In virtual reality it is this Cartesian 3D space that is emulated and in desktop interfaces Cartesian 2D space. One of the requirements we have is to have a measure of nearness, and Cartesian geometry supplies this with the familiar Pythagorean (as the crow flies) distance:

$$\begin{aligned} \text{dist}^2 &= x^2 + y^2 && \text{—2D space} \\ \text{dist}^2 &= x^2 + y^2 + z^2 && \text{—3D space} \end{aligned}$$

The awareness model of Benford et al. [1995] was designed to deal with proximity and attention in shared virtual environments. It is thus formulated within a strongly Cartesian spatial framework. Some of the concerns driving this model were pragmatic: “how can we know when an object is not the center of a user’s attention and so render it in less detail?” and “how can we know to whom to transmit a particular user’s audio so as not to drown everyone in a uniform babble?”. The concepts of *aura*, *nimbus*, and *focus* (and in later work, *third-party objects*) introduced in this model capture a relative notion of “nearness”: “what can I see/hear?”. The fact that this is set within a Cartesian virtual reality environment means that there are already clear “nearness” clues given by the scaling of objects with distance.

Despite its influence and conceptual power, Cartesian geometry is not as universal in the physical world as first appears. Cartesian coordinates are themselves built upon Euclidean geometry, which for almost two millennia was seen as self-evident. It was only comparatively recently (17th century) that alternative regular geometries were discovered: spherical geometry (the surface of a sphere, where there is too little “space” as one moves farther away) and hyperbolic geometry (where there is too much “space” as one looks further away—cabbage leaf geometry!). Still more recently with general relativity it has become clear that large-scale space is neither Euclidean nor regular, but instead “curves” as it is influenced by anything and everything that has mass or energy. At the quantum level things are still worse, and it appears that space may become fractal.

In mathematics there are a number of fields of study aimed at understanding alternative kinds of space. Important historically was the study of the geometry of regular spherical and hyperbolic space, following in the same vein as traditional geometry with theorems about triangles, circles, etc. and a whole study of spherical trigonometry. More interesting for virtual environments are various kinds of “space” that are less regular and embody more abstract notions of “nearness.” Two common abstract mathematical models of space that capture aspects of nearness are Metric Spaces and Topological Spaces.¹ Both of these abstract mathematical spaces capture an idea of nearness. In the case of Metric Spaces this is a numerical measure of the distance between two points which satisfies the “triangle inequality”:

$$\text{dist}(a,c) + \text{dist}(c,b) \geq \text{dist}(a,b)$$

This effectively says that if you want to go from A to B, it is always as fast or faster to go directly rather than to stop of at some other place C on the way—a reasonable minimal property of distance.

¹Note that when we used the word “topological” in Section 4.1, it was in the weaker sense in which it is used in computing rather than the precise mathematical formulation of a “Topological Space” (which will be capitalized to avoid confusion).

In the case of Topological Spaces, the idea of nearness is captured by ever-decreasing “neighborhoods” which contain a point and all sufficiently close neighbors. In both these kinds of spaces, the main interest in mathematics is in the notion of series of points which get ever closer without reaching a given point (convergent sequences); they are treated like “rubber sheets” which can be stretched as much as you like so long as they are not torn (continuous mappings). Hence, even in the case of Metric Spaces which have a numeric measure of distance, the important factor for their mathematics is not the absolute measure of nearness, but the use of the numbers to see whether things are getting closer.

More abstract notions of space can be found in Rodden’s formalization [Rodden 1996] of the Benford et al. awareness model. The spatial model underlying Benford et al.’s work was clearly Euclidean, but largely implicit. Rodden’s work looked at awareness over a graph structure as is found in the Web and many other computational domains. Nearness in such a space can be measured by number of arcs traversed or similar weighted measures both of which yield Metric Spaces. However, the critical properties of nearness in this work do not depend on these particular properties of the underlying graph. This suggests that we need models of space which may be stronger in that we would like some absolute sense of nearness and weaker in that we do not need the complex mechanisms needed to discuss convergent sequences, etc.

A final form of mathematical “space” which is relevant is the Differential Manifold. This is used to model curved space-time in General Relativity. This is not directly relevant as a model of the kinds of location found in virtual space or much-slower-than-light-speed physical space. However, the ways in which relativity has challenged our understanding of “space” in the physical world has a lot to teach us about the challenges of “virtual” space. One particular point is the way general relativity models space using mathematical structures called Differential Manifolds. Because space curves may have “singularities” (such as black holes) and even distant linked points (wormholes), it is impossible to use a single coordinate system to refer to all points. Instead, the models consist of a number of patches, each of which has “ordinary” Cartesian coordinates. Where the patches overlap there is a gentle transition between the coordinate systems (in mathematical terms they are related by a smooth function). Virtual spaces, such as the Web, may similarly have no global map or model, but if we can establish patches with well-defined structure and clear transitions between them then there is some hope for lost users.

Not only is space in general relativity not flat, but its shape and “size” change in time. We have all heard of the expanding universe. This does not mean simply that the stars are flying apart through space, but instead that the *space itself* between the galaxies is stretching. This at first sounds as if it is only of interest to cosmologists. However, it is also precisely the experience of those using wireless communications when their connection is broken. Before the break in communication they have established a sense of “nearness” in virtual space with other people and things on the network.

Then when the connection breaks this virtual geometry suddenly changes—things that were near suddenly become far away. It is precisely the difference between this and “normal” space that makes such disconnections so disturbing, especially if a collaborative system engenders any sense of immersion.

The feature of space in General Relativity that is perhaps most well known (although not necessarily understood) is that time and space are dealt with on an equal and interlocked basis: the time-space continuum. This blending of time and space can also be found in more mundane areas of virtual environments and interface design.

The Aether model of Sandor et al. [1997] adopts a graphical network as its underlying space, very like Rodden’s model. However, whereas both Benford et al. and Rodden have declarative definitions of awareness, the Aether model adopts a more process-oriented mechanism whereby the influence of an object (*aura*, *nimbus*, and *focus*) percolates through the network, getting weaker as it passes from node to node. The choice of this mechanism was largely driven by implementation considerations of producing an “awareness engine,” but is, of course, very like the physical transmission of sound and light. The Aether model has an implicit measure of nearness given by the rate at which network links and nodes attenuate influence, but also the Aether model explicitly introduces time as part of its awareness model. Whether in physical space or virtual, as soon as one takes into account transmission delays, space and time become inseparably interlinked.

This interlinking of time and space also becomes important as we consider different sensory experiences [Dix 1996]. Different senses give us different “cuts” through time and space. For objects within sight, we can consider the speed of light as practically instantaneous. Hence a quick glance around tells you about an area of space at a particular instant in time. If you want to know where something was a few seconds ago, you need to have looked then and remember. Imagine, however, that you are a dog or mole and are working using a sense of smell. As you sniff at a particular location you get some idea of the various creatures that have passed and even recent weather conditions at that point, i.e., smelling tells you about recent time at a single point of space. If you want to know what happened at other locations you need to have smelled there and remember. Finally consider a creature that uses sonar such as a whale or bat. Because sound takes time to travel through water or air, the echoes heard at a single moment correspond to close things recently, but further things longer ago. Figure 2 shows how each of these give us a particular cut of space-time.

In virtual space, network delays mean that we have sonar-like mixing of time and space. But also computer systems embody a memory of interaction—traces of past commands, windows opened for a previous purpose and never closed, copies of possibly out-of-date information—more like the world of smell. To add to the confusion, these different cuts through time-space are typically all presented visually!




Fig. 2. Different cuts through space-time.

In order to be able to talk about time-space interactions precisely, we need a semantic model of space (virtual and physical). Before we consider such a model it is worth highlighting that we will not be able to investigate all the aspects of time-space in this article and that the model presented in the following section is only one part of a much richer picture.

5.2 Developing a Semantic Model of Space

As discussed, existing awareness models are focused primarily on virtual space taking lessons from physical phenomena and are based on different underlying models of space. In order to support this range of approaches to awareness we need an abstract model of space that includes both Euclidean space and network space. We do not need the full richness and complexity of the mathematical spaces, but we do need an explicit formulation, as we need to be able to talk about several simultaneous spaces and their relationships.

Kinds of Space. The fundamental concepts we require are *location* and *nearness*. So we define a “Space-Kind” to include precisely these two plus some functions relating them:

Space-Kind = Location – set of elements representing “locations”
 Nearness – partially ordered set of “nearness” values
 dist: Location \times Location \rightarrow Nearness
 ...

The Location set depends on the precise kind of space. In 2D Cartesian space it would be the set of all (x,y) coordinate pairs; on the Web it would be the set of all URLs; in a building it would include locations such as “floor 3,” “room A 312,” or “northeast stairwell.”

The Nearness set will normally contain a minimal element “HERE” representing “at the same place as.” In the case of Cartesian space it is simply positive real numbers, and “dist” would be the normal as-the-crow-flies distance. In other spaces Nearness is a less precise concept and has values such as “on the same Web page,” “at the same site as,” “in the same room as,” and “in the same building as.”

The Nearness measure need not be a total order. For example, in a geographical information system we may be able to locate roads within towns, so “in the same road as” is obviously closer than “in the same town as.” However, if we look in the countryside it may simply be able to tell us “on the same mountain as.” Is this a closer measure than “in the same town as”? Although this Nearness set is intended to give some idea of absolute distance, we need to be careful. A very clear lesson from the mathematical studies of metric and topological spaces is that nonlocal measures of distance need to be treated with extreme caution. In our context we want to be able to conclude that if A is a device and X and Y are objects in a space such that

$$\text{dist}(A,X) < \text{dist}(A,Y)$$

then it is fair to make A’s behavior more dependent on X’s presence than that of Y. However, if A and B are devices such that

$$\text{dist}(A,X) < \text{dist}(B,Y)$$

then we should be extremely cautious about making any strong statements about the comparative strength of the relationships A-X and B-Y.

This does not mean we never use absolute judgments of distance. We have to be able to say things like the following:

*If object X is in the same room as device A,
then A shows a representation of X in its screen.*

However, when designing such rules we have to be aware that “in the same room as” could mean a broom cupboard or an auditorium.

Some kinds of location have a natural idea of containment. In an office complex “room A 315” may be on “floor 3” of “building A.” Similarly, the hierarchy of a Web site leads to a natural hierarchy of locations. In Cartesian spaces locations are mutually exclusive, but one can have regions of space, e.g., “all points within three miles of Great St. Mary’s Church Cambridge.” In order to capture both these uniformly we allow a space to have a set of regions which may either be a subset of locations (in the case of a hierarchy), or represent well-formed sets of points (in the case of a Cartesian space) or some other domain-specific concept:

Space-Kind = ...
 Region – sensible areas
 contains: Region \times Location \rightarrow Boolean

Spaces and Bodies. As we have already noted, even in the physical world we have several simultaneous ideas of space: longitude and latitude, town-street, etc. In the virtual world this multiplies further. So our model of the world has a number of spaces each of a particular kind and with other domain-specific attributes:

World = Spaces – set of elements representing “locations”
 kind: Spaces \rightarrow Space-Kind
 attr: Spaces \rightarrow Attributes
 ...

People, devices, and passive objects also inhabit the world as we have discussed previously. We call these collectively “Bodies.” These again have various domain-specific attributes, but the crucial question is the location of a specific body. This is not absolute, but defined relative to a specific space (e.g., GPS coordinates):

World = ...
 Bodies – (Bodies = People + Devices + Objects)
 attr: Bodies \rightarrow Attributes
 loc: Bodies \times Spaces \rightarrow Location

The “loc” function is partial, as there may be “spaces” for which a particular body has no clear relationship, e.g., there is no sensible answer to the question “what URL (Web location) is my tea cup at?”.

5.3 A Computational Model

Our semantic model of space is based on the concepts of *Space and Bodies* and a representation of location. We have developed a corresponding computational model to allow mobile applications to share a common awareness of a space and the bodies that inhabit that space. The general approach is to use an object-oriented model with a small number of simple objects that can be made shared across a distributed information space. This allows the state of defined objects to be accessed by a number of different devices.

The core of the model depends upon the definition of a virtual model of space in which the bodies relevant to the system are located and the ability to reason about the location of these in terms of a developed virtual space and the physical space of the real world. The central elements in the computational model are a *world object* and a *body object*. Each of these objects are intended to provide the root of two distinct specialization trees to represent the different forms of space and the bodies that exist within

the space. These two object hierarchies essentially instantiate the different kinds of space and bodies suggested in the semantic model. This development of a number of models of space mirrors the taxonomy we developed in the design framework and provides us way of reasoning about the location of devices in mobile systems in terms of both a real and virtual locations.

All the objects are realized on top of a distributed platform that allows the state of Java objects to be shared between different applications. The classes introduced below are therefore subclasses of *SharedEntity*, which is the root of all objects shared across the distributed platform.

The Space Object. The space object focuses on the ability of a space to act as a container of objects and on the way in which space can structure the world in terms of containment. The core space object has only two significant attributes: a set of bodies that it contains and a set of locations for these objects in the space. Location depends on a location object that can represent different senses of location.

The space object is provided to developers as a Java class that can be extended and specialized in order to represent different forms of space. Updates to the space object are propagated to all those that have registered an interest in the space. The key elements of the space object are reflected in Java as

```
Class Space extends SharedEntity {
    Kind kind; // The kind of space (e.g., Cartesian or Topology..)
    Vector bodies;
    Vector locations; // corresponding locations of bodies
    ... }
```

The core *Space* class includes methods for adding and removing bodies to a space, finding the location of a body within the space, and moving bodies in the space. Specializations drawn from this core object exploit the semantics of the space to provide more sophisticated views of proximity and distance.

The Location Object. A location class handles the location of bodies in space and is closely associated with the space object: each location object has an attribute to determine the kind of space it refers to, and each body in the space has an associated instance of a location object. Each location object basically represents the more general structure of the space within which bodies are placed. The attributes of the core Location class are

```
Class Location extends SharedEntity {
    Kind kind; // The kind of space (e.g., Cartesian or topological)
    Vector connected_entities; // The entities connected to this one
    Position position; // The position in the space.
    ... }
```

This basic location class has attributes that allow two different kinds of space to be represented: Cartesian spaces that define a location in terms of a position with reference to a fixed origin and topological spaces that consider the linkage between different spaces. This is also reflected in the

fact that the space provides two distinct methods *position* that returns a 3D position for a given object and *connected* which returns a list of spaces that a given space is connected to. The *connected* attribute allows us to represent a range of graph-like spaces. Although there are other forms of non-Cartesian space these graph-like spaces include those most commonly found in information systems, and the base class can easily be subclassed for other kinds of space. The *kind* attribute of the location should of course agree with the kind of space it is being used in, and the methods provided by the core classes maintain this consistency.

The Body Class. The definition of the location class allows us to represent and reason about the location and position of bodies within any space. The core of our design framework was the need to consider bodies as having both real and virtual locations and to manage interaction in terms of the correspondence between these. This means that our computational model needs to allow an interaction with bodies in terms of their position in multiple spaces. The link between the bodies representing the overall system and the spaces in which they reside are reflected in the definition of the class that provides the root of the bodies hierarchy making up the overall system. This is achieved in the computational model by having a Body class definition that allows state information about the spaces bodies are in to be externalized. The core attributes of the root Body class is a list of the spaces the body exists in. (Recall that a body may simultaneously exist in several physical and virtual spaces.) This is represented in the Java class as a simple vector:

```
Class Body extends SharedEntity
    Vector Spaces; // The spaces a body exists in
    ... }
```

As in the case of the definition of the Space class the Body class inherits from SharedEntity. This means that the state information can be shared and made available across the distributed platform. Each shared entity has a unique name and a set of optional tags that can be used to find objects of particular interest. This arrangement allows the different components making up a mobile system to be aware of the location of other entities in the various spaces in which they reside.

Using the Model. The computational model has been realized over a distributed infrastructure called Limbo (discussed in Section 6.2). This platform allows search facilities over the distributed shared object store, thus allowing any application to find out

- what spaces exist and what bodies are in those spaces,
- what other entities are in locations close to a body, and
- what other spaces is this body in and what is its location in these spaces.

This information can be used to infer distinct contextual cues about the nature of space. As an example consider a simple illustration of the use of

the platform drawn from our experiences of the GUIDE project. A portable notebook has facilities that allow it to know which cell it is in a cell-based radio infrastructure. This is a topological space with a close correspondence to the physical arrangement of the devices in the real world.

This complete space can be represented as an instance of (a subclass of) Space called “physical radio” which has a Kind attribute set to “Topological.” Each location in the space is named. Location names are based on the name of the base station supporting the cell. All the notebooks within the space are associated with a location, and we can ask the platform for the bodies in the space and their location. The physical movement of notebooks is reflected as changes to the associated location in the “physical radio” space.

Each notebook’s body definition also records the other spaces in which it is present, and this can be exploited or even coupled with the information about its location in the physical space. For example, each of the notebooks in the GUIDE project shows a Web page based on the radio cell that it is physically located in. This is achieved by putting the notebook in a virtual information space we shall call “guide space.” This space is actually a set of connected Web pages. The guide browser shows the appropriate page by finding its location within the virtual information space and updating this location as the location of the notebook in the “physical radio” space changes.

The computational model allows us to represent a range of different models of space and location central to the contextual interaction underpinning mobile systems. This model can then be accessed by a range of mobile devices and shared between them. This shared computational model provides a higher-level representation, which allows the rapid development and alteration of interactive applications essential to most prototyping approaches.

However, note that this computational model needed to be realized over an underlying infrastructure and system architecture (in our case the extended Limbo platform [Palfreyman et al. 1999]). In the next section, we will discuss the various issues involved in designing and selecting such architectures and discuss how the Limbo platform meets these requirements.

6. FROM REQUIREMENTS TO ARCHITECTURE

The taxonomies we have developed in this article have highlighted a wide range of application niches and suggests many exciting design possibilities for specific applications exploiting the contextual nature of mobile devices. Although we are investigating some of these in a number of projects the primary aim of our current “infrastructure” project is to examine the generic requirements to emerge from taxonomies of this form. These requirements can then be exploited to develop the underlying toolkits, architecture and infrastructure needed for temporally well-designed, context-aware, collaborative mobile systems.

Mobile systems extend our considerations of interaction beyond the user interface to consider interaction in terms of the entire environment (human, physical, and computational). As our framework has highlighted, in mobile systems the relevant semantics includes issues of location in physical and virtual space; proximity of other devices and people; and capabilities of devices and communication infrastructure. The necessary information and functionality to exploit this context is typically widely distributed within the computational environment—spread over different devices, spread over different physical locations, and spread between different layers in the system. The individual application developer will simply not have the relevant information and functionality available unless the infrastructure is designed taking into account human interface requirements. Thus, in mobile systems—more than other areas of HCI—the design of infrastructure is a central and essential concern.

6.1 Requirements

Unfortunately, research has repeatedly demonstrated the shortcomings of existing infrastructure components for supporting adaptive mobile applications [Davies 1994; Joseph et al. 1995]. In more detail, existing components have two critical shortcomings. Firstly, they are often highly network specific and fail to provide adequate performance over a range of network infrastructures (e.g., TCP has been shown to perform poorly over wireless networks [Cáceres and Iftode 1994]). Secondly, existing components often lack suitable APIs for passing status information to higher levels. As a consequence of these shortcomings new systems are increasingly being developed using bespoke communications protocols and user interfaces. For example, the GUIDE system described in Davies et al. [1998] uses a broadcast-style protocol. This is appropriate in a location-based information system where it is likely that pages of information needed by one device will be useful to all. It also uses the presence of base stations as an indicator of location, a technique shared with several current location-aware systems.

As these devices become more widespread the need increases for generic application architectures for at least specific subclasses of the mobile domain. There is clear commercial pressure for this; in particular, Windows-CE is being promoted for use in embedded systems. However, if these are simply developed by modifying architectures and toolkits originally designed for fixed environments there is a danger that some of the rich interaction possibilities afforded by mobile devices may be lost.

There are some examples of generic frameworks on which we can build. In Georgia Tech., location-aware guides are being constructed using the CyberDesk/Cameo architecture [Wood et al. 1997]. Because the phenomena we are trying to model are largely status, there is a great advantage of the underlying architecture also reflects status and change in status. For example, Cameo is a software architecture based on the theoretical framework of status-event analysis (as discussed in Section 4.4).

Another major architectural issue for context-aware applications is the way in which contextual issues cut across the whole system design. This is reminiscent of other aspects of user interface where the structures apparent at the user interface often do not match those necessary for efficient implementation and sound software engineering [Dix and Harrison 1989]. In UI design this has led to a conflict between architectures which decompose in terms of user interface layers, such as the Seeheim and ARCH-Slinkey models [Gram and Cockton 1996] and more functionally decomposed object-oriented models. In fact the object- and agent-based architectures themselves usually include a layered decomposition at the object level as in the MVC (Model-View-Controller) model [Lewis 1995] and in the PAC (Presentation-Abstraction-Control) model [Coutaz 1987]. Although the display and input hardware may be encapsulated in a single object or group of objects, its effects are felt in the architectural design of virtually every user interface component. In a similar fashion the hardware that supplies contextual information may well be encapsulated within context-objects, but their effect will permeate the system. This requires a similar orthogonal matrix structure similar to that found in models such as PAC or MVC. We expect context-awareness mechanisms to emerge as structures cutting across application layers and interface components.

Reviewing our discussion, an architecture for supporting mobile, context-aware applications must be

- distributed*: as this is the nature of the devices over which it operates,
- capable of representing *location*,
- able to effectively deal with both *status and event phenomena*, and
- be *orthogonal* to other interface components.

6.2 Extending Limbo to Provide a Supporting Platform

Our computational model has been instantiated over a distributed platform that allows a number of devices to make state information accessible to each other and thus allows the creation of a community of devices. This framework builds directly on the authors' previous work on Limbo and the development of a shared interaction platform [Palfreyman et al. 1999]. The platform exploits a distributed tuple space to share state information between geographically remote clients allowing them to function as a single collaborating system. The developed infrastructure is constructed using a combination of C++ and Java and has four significant components:

- A distributed tuple space (Limbo) that allows tuples of data values to be shared between different devices and accessed across a range of communication facilities. Our particular tuple space is a mobile variant of the established Linda model [Gelernter 1985].
- A notification services that informs applications about changes in the tuple space.




Fig. 3. The platform architecture.

- An infrastructure that allows structured information to be mapped onto shared tuples.
- The set of specific Java context objects (as described in Section 5.3) that allow communal access to the shared information about bodies, spaces, and locations.

The general architectural arrangement is shown in Figure 3. The general platform interface is provided through a set of Java objects. The implementation of the distributed tuple space allows the rapid replication of these objects and for changes in state to be propagated.

We will examine this architecture using the requirements identified in Section 6.1.

Distribution. Limbo is one of a number of distributed platforms. Indeed, our use of Limbo has similarities to the recent emergence of JavaSpaces [Sun Microsystems 1998] which provides a tuple-based infrastructure for Java programs. However, there are significant differences in the implementation resulting from the requirements that have driven their respective designs. Our platform is specifically designed for rapid dissemination of events and context information, and this is reflected in the use of multicast within its underlying implementation. Limbo is thus not only distributed, but optimized for the kinds of transactions we envisage for context-aware applications.

Location. The underlying Limbo architecture supports the sharing of any kind of objects. This, combined with our computational model of space, allows the representation and sharing between devices of

- different elements of space,
- bodies and their locations, and
- other domain-specific objects.

Together these facilities allow applications to provide interaction possibilities sensitive to context information pooled from a number of distributed sources.

Status and Event Phenomena. Limbo, the distributed Linda tuple space, upon which we have built our computational model of space and location, is already a status-orientated representation. The raw Linda space is a passive status requiring polling to discover changes. However, the additional “Shared Universe” layer adds explicit event notification facilities allowing applications to react to status-change events and to use the platform for general event notification. Because of the ability of the platform to manage both status and event phenomena, we expect it to support other forms of context awareness as well as the location services explored in detail here.

Orthogonality. It is in recognition of the pervasive nature of contextual dependency within the interface that we have developed our shared space model as an underlying service rather than a widget or component. The distributed nature of Limbo means that the shared space model allows appropriate definitions of space to be shared between devices and for these devices to exploit this contextual information represented by this shared context. The orthogonality of the model to other infrastructure components means that it is capable of capturing location information from both low-level sources (such as a GPS data interface) or higher-level sources (such as a Web browser’s current page).

7. CONCLUSION

This article has considered the emergence and development of a new class of advanced cooperative application and the different forms of interaction that may need to be supported. The maturing of technology to allow the emergence of multiuser distributed applications that exploit mobile applications means that we can no longer focus the issues of interaction on the nature of the device. Rather we must explicitly consider impact of the context in informing the design of different interaction techniques.

We have focused on understanding the design space to emerge for this new class of application and the importance of location in mobile systems. The article has presented a characterization of this design space and, for a particular portion of the design space, described more detailed models and supporting platforms that reflect the general approach to understanding the design space.

The general approach to developing our characterization of the design space has been to focus on the central role of location in mobile systems. This focus represents only one of the potential contexts of relevance to mobile systems considered in Section 2. The importance of space and location described in Sections 3 and 4 respectively represent a fairly unique aspect of mobile systems and underpin the development of the more

detailed models described in Section 5 that are instantiated in the platform described in Section 6.

The use of location within this article represents only one approach to understanding and managing context for these systems, but as we said in Section 2 the issues of context are much broader than location. In addition to the location of the device the overall context needs to be considered in terms of the devices relationship with the technical infrastructure, the application domain, the socio-technical system in which it is situated, and the physical nature of the device. The interaction style supported by mobile applications is as dependant on this context as the properties of the device itself. As a result, it is essential that work on the nature of these devices is complemented by a broader consideration of the nature of interaction. Our consideration of location and the development of the taxonomies, models, and supporting platforms represents one step in this broader consideration.

REFERENCES

- ALIAGA, D. G. 1997. Virtual objects in the real world. *Commun. ACM* 40, 3, 49–54.
- BASS, L., KASABACH, C., MARTIN, R., SIEWIOREK, D., SMALAGIC, A., AND STIVORIC, J. 1997. The design of a wearable computer. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (CHI '97, Atlanta, GA, Mar. 22–27), S. Pemberton, Ed. ACM Press, New York, NY, 139–146.
- BCS HCI. 1997. *Proceedings on British HCI Group Workshop on Time and the Web*. British Computer Society Workshop Series. British Computer Society, Swinton, UK.
- BENFORD, S., BOWERS, J., FAHLEN, L. E., MARIANI, J., AND RODDEN, T. 1994. Supporting co-operative work in virtual environments. *Computer J.* 37, 8.
- BENFORD, S., GREENHALGH, C., REYNARD, G., BROWN, C., AND KOLEVA, B. 1998. Understanding and constructing shared spaces with mixed-reality boundaries. *ACM Trans. Comput. Hum. Interact.* 5, 3, 185–223.
- BOROVY, R., MARTIN, F., VEMURI, S., RESNICK, M., SILVERMAN, B., AND HANCOCK, C. 1998. Meme tags and community mirrors: moving from conferences to collaboration. In *Proceedings of the 1998 ACM Conference on Computer-Supported Cooperative Work* (CSCW '98, Seattle, WA, Nov. 14–18), S. Poltrock and J. Grudin, Chairs. ACM Press, New York, NY, 159–168.
- CÁCERES, R. AND IFTODE, L. 1994. The Effects Of Mobility on Reliable Transport Protocols. In *Proceedings of the 14th International Conference on Distributed Computer Systems* (ICDCS, Prozen, Poland, June 22-24). 12–20.
- COUTAZ, J. 1987. PAC: an object oriented model for dialogue design. In *Proceedings on Human Computer Interaction* (INTERACT'87), H.-J. Bullinger and B. Shackel, Eds. Elsevier North-Holland, Inc., New York, NY, 431–436.
- DAVIES, N., BLAIR, G., CHEVERST, K., AND FRIDAY, A. 1994. Supporting Adaptive Services in a Heterogeneous Mobile Environment. In *Proceedings of the Workshop on Mobile Computing Systems and Applications* (Mobile '94, Santa Cruz, CA, Dec.), L.-F. Cabrera and M. Satyanarayanan, Eds. IEEE Computer Society Press, Los Alamitos, CA, 153–157.
- DAVIES, N., MITCHELL, K., CHEVERST, K., AND BLAIR, G. S. 1998. Developing a Context Sensitive Tourist Guide. University of Lancaster, Lancaster, UK. Technical Report Computing Department.
- DIX, A. 1992. Pace and interaction. In *Proceedings on People and Computers VII: Proceedings of HCI '92* (HCI '92, York, UK, Sept.), A. Monk, D. Diaper, and M. D. Harrison, Eds. Cambridge University Press, New York, NY, 193–207.
- DIX, A. J. 1995. Cooperation without (reliable) Communication: Interfaces for Mobile Applications. *Distrib. Syst. Eng.* 2, 3, 171–181.

- DIX, A. J. 1996. Closing the Loop: modelling action, perception and information. In *Proceedings on Advanced Visual Interfaces (AVI'96, Gubbio, Italy)*, T. Catarci, M. F. Costabile, S. Levialdi, and G. Santucci, Eds. ACM Press, New York, NY, 20–28.
- DIX, A. AND ABOWD, G. 1996. Modelling status and event behaviour of interactive systems. *Softw. Eng. J.* 11, 6, 334–346.
- DIX, A. J. AND HARRISON, M. D. 1989. Interactive systems design and formal development are incompatible?. In *The Theory and Practice of Refinement*, J. McDermid, Ed. Butterworth Publishers, Stoneham, MA, 12–26.
- FICKAS, S., KORTUEM, G., AND SEGALL, Z. 1997. Software issues in wearable computing. In *Proceedings of the CHI Workshop on Research Issues in Wearable Computers* (Atlanta, GA). ACM, New York, NY.
- FITZPATRICK, G., KAPLAN, S., AND MANSFIELD, T. 1996. Physical spaces, virtual places and social worlds: a study of work in the virtual. In *Proceedings of the 1996 ACM Conference on Computer-Supported Cooperative Work (CSCW '96, Boston, MA, Nov. 16–20)*, M. S. Ackerman, Ed. ACM Press, New York, NY, 334–343.
- GAVER, W. W. 1992. The affordances of media spaces for collaboration. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4)*, M. Mantel and R. Baecker, Chairs. ACM Press, New York, NY, 17–24.
- GELERNTER, D. 1985. Generative communication in Linda. *ACM Trans. Program. Lang. Syst.* 7, 1 (Jan.), 80–112.
- GRAM, C. AND COCKTON, G., EDS. 1997. *Design principles for interactive software*. Chapman and Hall, Ltd., London, UK.
- GREENBERG, S. AND MARWOOD, D. 1994. Real time groupware as a distributed system: concurrency control and its effect on the interface. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '94, Chapel Hill, NC, Oct. 22–26)*, J. B. Smith, F. D. Smith, and T. W. Malone, Chairs. ACM Press, New York, NY, 207–217.
- HARPER, R. H. R. 1992. Looking at ourselves: an examination of the social organisation of two research laboratories. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '92, Toronto, Canada, Oct. 31–Nov. 4)*, M. Mantel and R. Baecker, Chairs. ACM Press, New York, NY, 330–337.
- HENDERSON, D. A. AND CARD, S. 1986. Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Trans. Graph.* 5, 3 (July), 211–243.
- HOWARD, S. AND HOWARD, J., EDS. 1998. Temporal Aspects of Usability: The relevance of time to the development and use of human-computer systems. *Interact. Comput.*
- HUGHES, J., KING, V., RODDEN, T., AND ANDERSEN, H. 1995. The role of ethnography in interactive systems design. *interactions* 2, 2 (Apr.), 56–65.
- ISHII, H. AND ULLMER, B. 1997. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97, Atlanta, GA, Mar. 22–27)*, S. Pemberton, Ed. ACM Press, New York, NY, 234–241.
- JOHNSON, C. AND GRAY, P. 1996. Workshop Report: Temporal Aspects of Usability. *SIGCHI Bull.* 28, 2.
- JOHNSON, C. W. 1997. The impact of time and place on the operation of mobile computing devices. In *Proceedings of the conference of HCI '97 on People and Computers XII (HCI '97, Bristol, UK)*. 175–190.
- JOSEPH, A., DELSPINASSE, A., TAUBER, J., GIFFORD, D., AND KAASHOEK, M. F. 1995. Rover: A Toolkit for Mobile Information Access. In *Proceedings of the 15th ACM Symposium on Operating System Principles (SOSP, Copper Mountain Resort, Colorado, U.S., 3–6 Dec.)*. ACM Press, New York, NY, 156–171.
- LEWIS, S. 1995. *The art and science of Smalltalk*. Prentice Hall International (UK) Ltd., Hertfordshire, UK.
- LONG, S., KOOPER, R., ABOWD, G. D., AND ATKESON, C. G. 1996. Rapid prototyping of mobile context-aware applications: the Cyberguide case study. In *Proceedings of the 2nd Annual International Conference on Mobile Computing and Networking (MOBICOM '96, Rye, NY,*

- Nov. 10–12), H. Ahmadi, R. H. Katz, I. F. Akyildz, and Z. J. Haas, Chairs. ACM Press, New York, NY, 97–107.
- PATTERSON, J. F., DAY, M., AND KUCAN, J. 1996. Notification servers for synchronous groupware. In *Proceedings of the 1996 ACM Conference on Computer-Supported Cooperative Work (CSCW '96, Boston, MA, Nov. 16–20)*, M. S. Ackerman, Ed. ACM Press, New York, NY, 122–129.
- RAMDUNY, D. AND DIX, A. 1998. Why, What, Where, When: Architectures for Co-operative work on the WWW. In *Proceedings of the primary European Annual Conference on Human Computer Interaction (HCI'98, Sheffield Hallam University, Bristol, UK, Sept. 1-4)*. Springer-Verlag, Vienna, Austria, 283–301.
- RAMDUNY, D., DIX, A., AND RODDEN, T. 1998. Getting to Know: the design space for notification servers. In *Proceedings on CSCW*.
- RODDEN, T. 1996. Populating the application: a model of awareness for cooperative applications. In *Proceedings of the 1996 ACM Conference on Computer-Supported Cooperative Work (CSCW '96, Boston, MA, Nov. 16–20)*, M. S. Ackerman, Ed. ACM Press, New York, NY, 87–96.
- ROOT, R. W. 1988. Design of a multi-media vehicle for social browsing. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW '88, Portland, OR, Sept. 26–28)*, I. Greif, Chair. ACM Press, New York, NY, 25–38.
- ROSEMAN, M. AND GREENBERG, S. 1996. TeamRooms: Network places for collaboration. In *Proceedings of the 1996 ACM Conference on Computer-Supported Cooperative Work (CSCW '96, Boston, MA, Nov. 16–20)*, M. S. Ackerman, Ed. ACM Press, New York, NY, 325–333.
- SANDOR, O., BOGDAN, C., AND BOWERS, J. 1997. Aether: An Awareness Engine for CSCW. In *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work (ECSCW'97)*, J. Hughes, Ed. Kluwer Academic, Dordrecht, Netherlands, 221–236.
- SUN MICROSYSTEMS. 1998. Java Distributed Computing White Papers. <http://java.sun.com/products/javaspaces/whitepapers/index.html>.
- WANT, R., SCHILIT, B. N., ADAMS, N. I., GOLD, R., PETERSEN, K., GOLDBERG, D., ELLIS, J. R., AND WEISER, M. 1995. An Overview of the ParcTab Ubiquitous Computing Experiment. *IEEE Personal Commun.* 2, 6 (Dec.), 28–43.
- WEISER, M. 1991. The computer of the 21st century. *Sci. Am.* 265, 3 (Sept.), 66–75.
- WEISER, M. 1993. Some computer science issues in ubiquitous computing. *Commun. ACM* 36, 7 (July), 75–84.
- WOOD, A., DEY, A., AND ABOWD, G. D. 1997. CyberDesk: automated integration of desktop and network services. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '97, Atlanta, GA, Mar. 22–27)*, S. Pemberton, Ed. ACM Press, New York, NY, 552–553.
- PALFREYMAN, K., TREVOR, J., AND RODDEN, T. 1999. PSI: A Platform for Shared Interaction. In *Proceedings on ECSCW (ECSCW99, Copenhagen, Denmark, 12-16 Sept.)*. Kluwer Academic, Dordrecht, Netherlands, 351–371.